

Spatial adaptivity of the SAAF and Weighted Least Squares (WLS) forms of the neutron transport equation using constraint based, locally refined, isogeometric analysis (IGA) with dual weighted residual (DWR) error measures

C. Latimer^a, J. Kópházi^a, M. D. Eaton^a, R. G. McClarren^b

^a*Nuclear Engineering Group
Department of Mechanical Engineering, Imperial College London
City and Guilds Building, Exhibition Road, South Kensington Campus, SW7 2AZ*

^b*Department of Aerospace and Mechanical Engineering University of Notre Dame
Fitzpatrick Hall, Notre Dame, IN 46556*

Abstract

This paper describes a methodology that enables NURBS (Non-Uniform Rational B-spline) based Isogeometric Analysis (IGA) to be locally refined. The methodology is applied to continuous Bubnov-Galerkin IGA spatial discretisations of second-order forms of the neutron transport equation. In particular this paper focuses on the self-adjoint angular flux (SAAF) and weighted least squares (WLS) equations. Local refinement is achieved by constraining degrees of freedom on interfaces between NURBS patches that have different levels of spatial refinement. In order to effectively utilise constraint based local refinement, adaptive mesh refinement (AMR) algorithms driven by a heuristic error measure or forward error indicator (FEI) and a dual weighted residual (DWR) or goal-based error measure (WEI) are derived. These utilise projection operators between different NURBS meshes to reduce the amount of computational effort required to calculate the error indicators. In order to apply the WEI to the SAAF and WLS second-order forms of the neutron transport equation the adjoint of these equations are required. The physical adjoint formulations are derived and the process of selecting source terms for the adjoint neutron transport equation in order to calculate the error in a given quantity of interest (QoI) is discussed. Several numerical verification benchmark test cases are utilised to investigate how the constraint based local refinement affects the numerical accuracy and the rate of convergence of the NURBS based IGA spatial discretisation. The nuclear reactor physics verification benchmark test cases show that both AMR algorithms are superior to uniform refinement with respect to accuracy per degree of freedom. Furthermore, it is demonstrated that for global QoI the FEI driven AMR and WEI driven AMR produce similar results. However, if local QoI are desired then WEI driven AMR algorithm is more computationally efficient and accurate per degree of freedom.

Email address: c.latimer15@imperial.ac.uk (C. Latimer)

1. Introduction

Isogeometric analysis (IGA) is a numerical technique for the spatial discretisation of partial differential equations (PDEs). The IGA method is a generalisation of traditional finite element (FE) spatial discretisation techniques that has been shown to: produce improved numerical accuracy per degree of freedom for neutron diffusion and transport problems [1, 2, 3], correctly retrieve modes in crash simulations [4], obtain physically accurate contact stresses [5], and generate accurate frequency spectra for structural vibration problems [6]. The two main features of IGA that lead to these advantages are the ability to exactly represent geometries generated by commercial CAD (Computer Aided Design) software, such as conic sections (2D) and quadric surfaces (3D), and high-order continuity of the NURBS (Non-Uniform Rational B-spline) basis functions within a NURBS patch. FE basis functions are typically C^0 over element boundaries, whereas IGA basis functions can have any level of continuity within a NURBS patch, depending upon the order in which refinement operations are performed [7].

There are many types of basis function that can be used to perform IGA such as THB-splines (Truncated Hierarchical B-splines) [8], LR B-splines (Locally Refined B-splines) [9], and analysis suitable T-splines [10]. The most widely used basis functions by commercial CAD software packages are the NURBS functions [11]. The fundamental aim of IGA is to streamline the interface between CAD and computer aided engineering (CAE) analysis. Therefore, the vast majority of the IGA literature utilises the NURBS basis functions that are typically used within CAD software [12]. One disadvantage of NURBS based IGA is the inability to perform local spatial refinement. This inability to perform local spatial refinement of the NURBS mesh manifests itself in two ways. First, control points on the interface between NURBS patches are typically assumed to be in one-to-one correspondence. This assumption simplifies the logic of IGA numerical algorithms, and also minimises the computational effort required to assemble the global stiffness matrix [11]. Second, two-dimensional (2D) and three-dimensional (3D) NURBS basis functions are generated using a tensor product structure. Native local refinement within a NURBS patch is not supported when basis functions are constructed in this manner.

Hierarchical splines [13] are an option for remedying the second local refinement issue. These hierarchical splines could also be utilised in the standard multipatch NURBS based IGA framework. In order to overcome the first local refinement issue a constraint based IGA methodology is investigated in this paper. The constraint based local refinement approach was first developed to overcome the issue of hanging nodes in the FE spatial discretisation methods [14, 15], and has also been applied to IGA before [11].

Once local spatial refinement can be performed, a methodology for deciding where to apply local spatial refinement is necessary. These methods can be divided into two categories: *a priori* methods, and *a posteriori* methods. *A priori* methods provide error estimates based on the underlying PDE being solved and the spatial discretisation being used. However, these normally only yield information about the asymptotic behaviour of the solution and often are only valid for solutions that are sufficiently regular [16]. *A posteriori* methods provide error estimates for the solution based upon the solution itself [17, 18]. Furthermore, *a posteriori* methods can be tailored to estimate the error in specific quantities of interest (QoI) by using dual-weighted residual (DWR) or goal-based error measures [19]. However, *a posteriori* can be significantly more computationally demanding than *a priori* methods. This is because they require both an exact solution and a calculated solution, or approximations of these quantities, in order to evaluate numerical errors. As a consequence the use of *a posteriori* can be computationally demanding. The use of *a posteriori* methods enables spatial mesh refinement to be performed only in those spatial

regions that contribute most to the improvement of the numerical accuracy of prescribed local and global QoI. Therefore, in principle *a posteriori* methods can lead to computationally efficient AMR schemes. However, there is always a balance between the improvements in numerical accuracy achieved using *a posteriori* methods and the computational effort required to apply these methods to practical problems in nuclear reactor physics and radiation shielding.

Nuclear reactor core design, reactor shield design, and radiation protection are fields where both global and local QoI are of great importance. For example, in criticality calculations, the K_{eff} of a nuclear reactor is an important global QoI. The K_{eff} indicates whether the number of neutrons in a system will decrease (sub-critical), increase (super-critical), or stay constant (critical) as time progresses. However, nuclear reactor physics calculations are also concerned with local QoI such as the power generated in specific nuclear fuel pins. This is of interest in order to make sure that burn-up/depletion within a nuclear fuel assembly is reasonably uniform and that the power output stays relatively constant [20]. On the other hand, reactor shielding and radiation protection calculations are often concerned with highly localised QoI and sometimes QoI that consist of a single point. For example, the dose that an operator might receive, or a detector response is of interest to ensure that radiation levels in a certain area are under the recommended limit [21]. The neutron fluence (time-integrated scalar neutron flux) is important in understanding the radiation damage, which is measured in displacement per atom (DPA), to a reactor pressure vessel (RPV) [22].

In Section 2 the NURBS basis functions are introduced and the reasons for the lack of local spatial refinement in standard NURBS based IGA are expanded upon. In Section 3 a constraint based local spatial refinement methodology for NURBS based IGA is described. In Section 4, a heuristic error indicator (FEI) and a goal-based error measure (WEI) are derived. The abbreviation FEI stands for Forward Error Indicator, due to the fact that areas are flagged for refinement purely due to errors in the forward solution. The abbreviation WEI stands for Weighted Error Indicator, due to the fact that the forward error is weighted by the error in the adjoint solution. Both of these are *a posteriori* methods that require a projection operator between spatial meshes of differing fidelity. In Section 5 two second-order forms of the neutron transport equation are introduced. These are the self-adjoint angular flux (SAAF) and weighted least squares (WLS) neutron transport equations. Section 6 derives the mathematical adjoint of both of these equations and discusses how the goal functional can be chosen in the adjoint equation to allow the use of the WEI for a given local QoI. Finally, in Section 7, the effect of the constraint based locally refined NURBS based IGA on the numerical accuracy, and the rate of convergence, of the solution is investigated. Furthermore, comparisons are made between the relative numerical accuracy of FEI driven AMR and WEI driven AMR for various QoI and for various nuclear reactor physics and radiation shielding verification benchmark test cases.

2. NURBS based Isogeometric Analysis (IGA)

Isogeometric analysis (IGA) is a methodology for the spatial discretisation of partial differential equations (PDEs) that was first developed in 2005 [7]. Since then it has been applied to many areas such as structural mechanics [6], nuclear reactor physics [23], optimisation problems [24], crash simulations [4], and medical modelling [25]. In this section the NURBS basis functions will be explained in order to illustrate how issues regarding local refinement arise. IGA is not explained in detail in this paper and if the reader is interested in more information about refinement techniques, the generation of geometry, Galerkin IGA spatial discretisations, or

NURBS basis functions, in general, then the following references are suggested [11, 26, 27, 28]. To define the NURBS basis functions it is first necessary to define B-splines.

2.1. B-splines

B-splines are a set of piecewise polynomials that require a knot vector in order to be generated. A knot vector $\Xi = \{\xi_1, \xi_2, \dots, \xi_{n-1}, \xi_n\}$ is a series of monotonically increasing real numbers. The B-spline basis functions over Ξ are defined as follows:

$$N_{i,0}(\xi) = \begin{cases} 1, & \text{for } \xi \in [\xi_i, \xi_{i+1}), \\ 0, & \text{otherwise,} \end{cases}$$

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

where p is the order (or degree) of the B-splines. Entries in Ξ can appear multiple times consecutively and the number of times a knot ξ_i appears in Ξ is referred to as the multiplicity m_i of knot ξ_i . The B-spline basis functions are C^{p-m_i} over knot ξ_i and C^∞ at all other $\xi \in \Xi$. A B-spline curve is defined by:

$$\mathbf{C}(\xi) = \sum_{i=1}^n N_{i,p}(\xi) \mathbf{P}_i, \quad \text{for } \xi \in \Xi, \quad (2)$$

where the $\mathbf{P}_i \in \mathbb{R}^3$ are the control points of the B-spline curve. Furthermore, in a similar manner to FEs, a function can be represented by the product of the B-spline basis functions and *control variables* as follows:

$$f(\xi) = \sum_{i=1}^n N_{i,p}(\xi) d_i, \quad \text{for } \xi \in \Xi, \quad (3)$$

where the control variable $d_i \in \mathbb{R}$ is associated with control point \mathbf{P}_i . In order to generate B-spline objects in higher dimensions a tensor product structure is used. That is, given two knot vectors Ξ_ξ and Ξ_η , the B-spline surface is defined as:

$$\mathbf{S}(\xi, \eta) = \sum_{i=1}^N \sum_{j=1}^M N_{i,p}(\xi) N_{j,p}(\eta) \mathbf{P}_{i,j}, \quad \text{for } \xi \in \Xi_\xi \text{ and } \eta \in \Xi_\eta. \quad (4)$$

An expression similar to equation (3) that represents a 2D function in terms of control variables and basis functions can be written as:

$$f(\xi, \eta) = \sum_{i=1}^N \sum_{j=1}^M N_{i,p}(\xi) N_{j,p}(\eta) d_{i,j}, \quad \text{for } \xi \in \Xi_\xi \text{ and } \eta \in \Xi_\eta, \quad (5)$$

where the control variable $d_{i,j}$ is associated with control point $\mathbf{P}_{i,j}$. The extension to higher dimensions follows a similar tensor product structure.

2.2. NURBS

NURBS basis functions are a set of rational piecewise polynomials that are constructed using the B-spline basis functions. To define the NURBS basis functions a set of weights must be defined $\{w_i\}_{i=1}^N$. Then the NURBS basis functions of order p are given by:

$$R_{i,p}(\xi) = \frac{N_{i,p}(\xi)w_i}{\sum_{j=1}^N N_{j,p}(\xi)w_j}, \text{ for } i = 1, 2, \dots, N. \quad (6)$$

A NURBS curve is defined in a similar manner to equation (2), that is:

$$\mathbf{C}(\xi) = \sum_{i=1}^N R_{i,p}(\xi)\mathbf{P}_i, \text{ for } \xi \in \Xi, \quad (7)$$

and a NURBS surface is constructed using the same tensor product structure:

$$\mathbf{S}(\xi, \eta) = \sum_{i=1}^N \sum_{j=1}^M R_{i,p}(\xi)R_{j,p}(\eta)\mathbf{P}_{i,j}, \text{ for } \xi \in \Xi_\xi \text{ and } \eta \in \Xi_\eta. \quad (8)$$

A NURBS patch is defined as a collection of knot vectors, weights, and control points. Therefore, each NURBS patch has its own set of NURBS basis functions. NURBS patches are topologically equivalent to a hypercube due to the tensor product structure of the NURBS basis functions. In 2D, NURBS patches are topologically equivalent to a square. By combining NURBS patches complex, multiply connected, heterogeneous domains can be modelled.

The tensor product formulation of the B-spline and NURBS surfaces in equations (4) and (8) is responsible for the inability of NURBS based IGA to locally refine within a NURBS patch. If the basis functions in one parametric dimension are refined then this refinement propagates throughout all other parametric dimensions. This is illustrated in two dimensions in Figure 1. The aim of this paper is not to tackle the inability of NURBS based IGA to locally refine on this level. As mentioned before, degrees of freedom (dof) on the interfaces between NURBS patches are typically assumed to be in one-to-one correspondence. This reduces the complexity and the amount of computational effort required to form global stiffness matrices. However, this assumption precludes the ability for NURBS patches to be locally refined spatially, and also for neighbouring NURBS patches to have different polynomial orders. In order for local spatial refinement of multipatch domains to be possible a constraint based local spatial refinement methodology is described. The effect of spatial refinement on standard NURBS based IGA, and the effect the constraint based local refinement will have, are illustrated in Figure 2.

3. Constraint Based Local Spatial Refinement

Constraint based local spatial refinement is a method that was originally used in FE to restore the continuity of the FE space when hanging nodes are present [14, 15]. The methodology has been extended to the IGA framework [11] in order to restore continuity of the NURBS basis function space for arbitrary order basis functions between NURBS patches when hanging nodes are present.

Two cases are considered in this section. The first case is two square NURBS patches side by side (see Figures 3 and 4). This case will illustrate how a constraint is applied and also how the global stiffness matrix is formed. The second case is three square NURBS patches in a corner configuration (see Figure 5). This case details how to deal with a point being constrained multiple times. These two cases can be combined to deal with most 2D geometries.




Figure 1: Illustration of the inability of NURBS based IGA to be locally refined within a NURBS patch. The whole domain represents a NURBS patch with internal lines referring to knot insertion within the NURBS patch. The aim is to insert degrees of freedom (dof) into the lower left of the NURBS patch. On the left is the case of NURBS based IGA. When a knot is inserted into the x direction the refinement propagates in the y direction resulting in five dof being created. This is due to the tensor product structure of the NURBS basis functions. On the right is the case of local refinement within a NURBS patch. When a knot is inserted only two dof are created.

3.1. The Two NURBS Patch Case

The two NURBS patch case is sufficient to deal with geometries where each dof is constrained at most once. The geometry for this case, with an example mesh superimposed onto it, is presented in Figure 3. The more refined NURBS patch is referred to with the superscript F for fine, and the less refined NURBS patch with superscript C for coarse. It is necessary that, at the coarsest level, these two NURBS patches form a compatible discretisation meaning that on the coarsest mesh the mappings and parameterisations between neighbouring patches are identical [11]. Furthermore, it is assumed that the control points on the shared interface of the fine NURBS patch are the result of knot insertion (spatial refinement) of the control points from the coarse NURBS patch. However, this assumption is not a requirement. In the case where the NURBS patches are not a direct refinement of one another the greatest common knot vector can be found. The knot vectors from both NURBS patches are then constrained to match the greatest common knot vector.

For simplicity this explanation will continue assuming that the fine NURBS patch is the direct result of knot insertion being performed on a coarse NURBS patch with the knot vector Ξ^C . Therefore, the control points for this patch before and after knot insertion can be related in the following way:

$$\mathbf{P}^F = T\mathbf{P}^C, \quad (9)$$

where the matrix operator T represents knot insertion, and \mathbf{P}^F and \mathbf{P}^C are the control points in the fine and coarse patches respectively. As T represents knot insertion it is a unitary matrix. The relationship in expression (9) can be expanded into:

$$\begin{pmatrix} \mathbf{P}_n^F \\ \mathbf{P}_f^F \end{pmatrix} = \begin{pmatrix} T_n & 0 \\ 0 & T_f \end{pmatrix} \begin{pmatrix} \mathbf{P}_n^C \\ \mathbf{P}_f^C \end{pmatrix}, \quad (10)$$




Figure 2: Illustration of the inability of standard NURBS based IGA to be locally refined over a multipatch domain. The domain consists of six NURBS patches and spatial refinement of the parameter space within a NURBS patch is indicated with a dashed line. The aim is to insert dof into the lower left of the domain. On the left is the case of standard NURBS based IGA. When knots are inserted into the lower left NURBS patch the refinements propagate through neighbouring NURBS patches resulting in eight dof being created. On the right is the case of constraint based local refinement covered in this paper. When knots are inserted into the lower left NURBS patch only five dof are created. This is due to the fact that dof on the interface between NURBS patches do not need to be in one-to-one correspondence.




Figure 3: The geometry for the two NURBS patch case. The coarse NURBS patch is on the left and its control points are red. The fine NURBS patch is on the right and its control points are blue. If control points are in the same location then they are coloured purple. Black lines represent refinement in the parametric domain. The linear knot vectors are given by $\Xi^C = \{0, 0, 0.5, 1, 1\}$ and $\Xi^F = \{0, 0, 0.25, 0.5, 0.75, 1, 1\}$ in both parametric directions.

where the control points in the NURBS patch have been partitioned into two sets: control points that are on an interface with another NURBS patch, denoted by a f subscript, and control points that are not, denoted by an n subscript. In particular, from equation (10), it can be found that:

$$\mathbf{P}_f^F = T_f \mathbf{P}_f^C, \quad (11)$$

where T_f represents knot insertion performed on the dof on the interface between NURBS patches. Due to the fact that T is a unitary matrix, T_f is also a unitary matrix. Following from this it is apparent that the control variable associated with each control point on the interface between patches must satisfy a similar relationship, that is:

$$\mathbf{d}_f^F = T_f \mathbf{d}_f^C. \quad (12)$$

where \mathbf{d} is a vector representing the control variables used as the coefficients for the NURBS basis functions in the finite expansion of the solution in equation (3). Control variables with a F superscript belong to the fine patch, and control variables with a C superscript belong to the coarse patch.

The system of linear equations that represents the problem on the fine NURBS patch can be written in the following form:

$$K^F \mathbf{d}^F = \begin{pmatrix} K_{nn}^F & K_{nf}^F \\ K_{fn}^F & K_{ff}^F \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^F \\ \mathbf{d}_f^F \end{pmatrix} = \begin{pmatrix} S_n^F \\ S_f^F \end{pmatrix} = S^F, \quad (13)$$

and on the coarse NURBS patch:

$$K^C \mathbf{d}^C = \begin{pmatrix} K_{nn}^C & K_{nf}^C \\ K_{fn}^C & K_{ff}^C \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^C \\ \mathbf{d}_f^C \end{pmatrix} = \begin{pmatrix} S_n^C \\ S_f^C \end{pmatrix} = S^C, \quad (14)$$

where K^F and K^C are the stiffness matrices and S^F and S^C are the source vectors for the fine and coarse patches respectively over the sets of control variables implied by the subscripts. So

K_{mn}^F is the stiffness matrix for the fine patch that relates the dof not on the interface to the dof not on the interface. Expression (12) is then substituted into equation (13) giving:

$$\begin{pmatrix} K_{mn}^F & K_{nf}^F \\ K_{fn}^F & K_{ff}^F \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^F \\ T_f \mathbf{d}_f^C \end{pmatrix} = \begin{pmatrix} S_n^F \\ S_f^F \end{pmatrix}. \quad (15)$$

In order to assemble the patch-wise systems of linear equations in equations (14) and (15) into one global system that represents the problem over the entire two NURBS patch domain, it is necessary for both systems of equations to operate on the same set of control variables on the shared interface. Therefore, the unknowns on the shared interface of the fine NURBS patch system, equation (15), must be constrained to the same set as the coarse NURBS patch system, equation (14). This is achieved by operating on the equations for the control variables in the fine NURBS patch that are on the interface, $T_f \mathbf{d}_f^C$, with T_f^T . That is:

$$K_{fn}^F \mathbf{d}_n^F + K_{ff}^F (T_f \mathbf{d}_f^C) = S_f^F \rightarrow T_f^T K_{fn}^F \mathbf{d}_n^F + T_f^T K_{ff}^F T_f \mathbf{d}_f^C = T_f^T S_f^F. \quad (16)$$

Substituting this into equation (15) gives:

$$\begin{pmatrix} K_{mn}^F & K_{nf}^F T_f \\ T_f^T K_{fn}^F & T_f^T K_{ff}^F T_f \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^F \\ \mathbf{d}_f^C \end{pmatrix} = \begin{pmatrix} S_n^F \\ T_f^T S_f^F \end{pmatrix}. \quad (17)$$

Now there is an expression for the stiffness matrix over the fine NURBS patch that operates on the same set of control variables on the interface as the coarse NURBS patch and that is not over or under determined. In other words, there is now one equation for each control variable in equation (17). Therefore, the two linear systems of equations, (14) and (17), can be assembled into a global system of linear equations. In this particular case yielding:

$$\begin{pmatrix} K_{mn}^C & K_{nf}^C & 0 \\ K_{fn}^C & K_{ff}^C + T_f^T K_{ff}^F T_f & T_f^T K_{fn}^F \\ 0 & K_{nf}^F T_f & K_{nn}^F \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^C \\ \mathbf{d}_f^C \\ \mathbf{d}_n^F \end{pmatrix} = \begin{pmatrix} S_n^C \\ S_f^C + T_f^T S_f^F \\ S_n^F \end{pmatrix}. \quad (18)$$

In Figure 3, due to the linear basis functions along the interface, the control points that belong to the coarse NURBS patch are a proper subset of the control points that belong on the interface. This will not be the case for non-linear basis functions. The case of quadratic basis functions is illustrated in Figure 4. It can be seen that there is a control point on the interface that belongs only to the coarse NURBS patch. The constraint based methodology detailed above deals with this case by design as all of the refinement information is contained in T_f .

3.2. The Three NURBS Patch Case

In order to fully describe the constraint based local refinement methodology the case of three NURBS patches must be considered. The reason this case is illustrative is that it now includes a control variable that is subject to more than one constraint. A schematic of the geometry for this case is displayed in Figure 5 and the point that will have more than one constraint applied to it is coloured black.

The methodology begins as in the two NURBS patch case where it is noted that control variables along interfaces can be written as a refinement of one another. The same steps can be

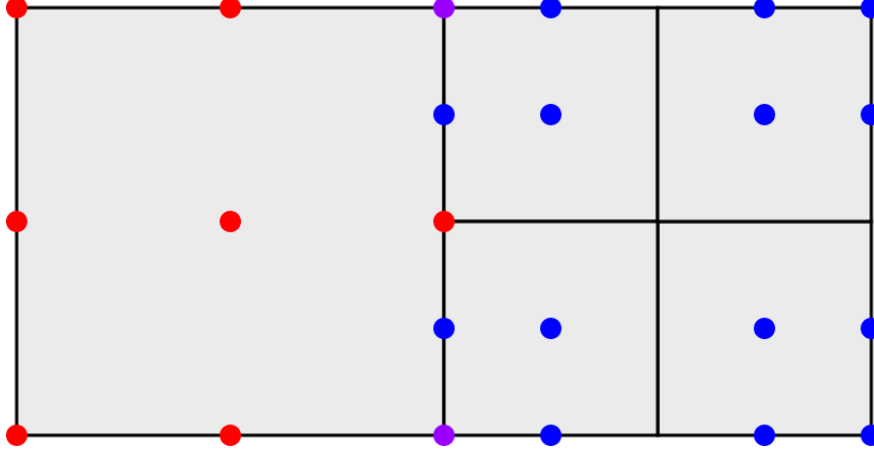


Figure 4: The layout of the control points for quadratic basis functions. The coarse NURBS patch is on the left and its control points are red. The fine NURBS patch is on the right and its control points are blue. If control points are in the same location then they are coloured purple. Black lines represent refinement in the parametric domain. Therefore $\Xi^C = \{0, 0, 0, 1, 1, 1\}$ and $\Xi^F = \{0, 0, 0, 0.5, 1, 1, 1\}$ in both parametric directions.

followed until three NURBS patch-wise stiffness matrices are formed, two for the coarse NURBS patches which are unconstrained:

$$\begin{pmatrix} K_{nn}^{C_w} & K_{nf_w}^{C_w} \\ K_{f_w n}^{C_w} & K_{f_w f_w}^{C_w} \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^{C_w} \\ \mathbf{d}_{f_w}^{C_w} \end{pmatrix} = \begin{pmatrix} S_n^{C_w} \\ S_{f_w}^{C_w} \end{pmatrix}, \quad (19)$$

$$\begin{pmatrix} K_{nn}^{C_s} & K_{nf_s}^{C_s} \\ K_{f_s n}^{C_s} & K_{f_s f_s}^{C_s} \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^{C_s} \\ \mathbf{d}_{f_s}^{C_s} \end{pmatrix} = \begin{pmatrix} S_n^{C_s} \\ S_{f_s}^{C_s} \end{pmatrix}, \quad (20)$$

and one for the fine NURBS patch which has had two sets of constraints applied to, namely T_{f_s} and T_{f_w} :

$$\begin{pmatrix} K_{nn}^F & K_{nf_w}^F T_{f_w} & K_{nf_s}^F T_{f_s} \\ T_{f_w}^T K_{f_w n}^F & T_{f_w}^T K_{f_w f_w}^F T_{f_w} & T_{f_w}^T K_{f_w f_s}^F T_{f_s} \\ T_{f_s}^T K_{f_s n}^F & T_{f_s}^T K_{f_s f_w}^F T_{f_w} & T_{f_s}^T K_{f_s f_s}^F T_{f_s} \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^F \\ \mathbf{d}_{f_w}^{C_w} \\ \mathbf{d}_{f_s}^{C_s} \end{pmatrix} = \begin{pmatrix} S_n^F \\ T_{f_w}^T S_{f_w}^F \\ T_{f_s}^T S_{f_s}^F \end{pmatrix}. \quad (21)$$

The subscripts f_w and f_s refer to the interface between the central NURBS patch and west-facing NURBS patch and the interface between the central patch and the south-facing NURBS patch respectively.

What is important to consider here is the set $\mathbf{d}_{f_w}^{C_w} \cup \mathbf{d}_{f_s}^{C_s}$. Referring to Figure 5 this is the union of the purple and the green points. Therefore, the set $\mathbf{d}_{f_w}^{C_w} \cup \mathbf{d}_{f_s}^{C_s}$ is the singular black point on the south-west corner of the fine NURBS patch. This point is referred to as d_l . Since d_l has been constrained twice it appears twice in equation (21), once in the set $\mathbf{d}_{f_w}^{C_w}$ and once in the set $\mathbf{d}_{f_s}^{C_s}$.

It can be noted that d_l is located at the corner of the parametric domain, is interpolatory, and will always be present regardless of the amount of refinement. Therefore, the structure of T_f will always constrain d_l to obey $d_l = d_l$. So when equations (19), (20), and (21) are assembled into a global system the row that relates all of the control variables to d_l will be added twice in this case. Therefore, to get the correct answer the coefficients in the rows corresponding to d_l must be combined and divided by two.




Figure 5: The geometry for the three NURBS patch case. The coarse NURBS patch to the west is referred to as C_w and its control points are red. The coarse NURBS patch to the south is referred to as C_s and its control points are yellow. The fine NURBS patch in the centre is referred to as F and its control points are blue. If control points are in the same location then they are coloured a mixture of the relevant colours. Black lines represent refinement in the parametric domain. Therefore, $\Xi^{C_w} = \Xi^{C_s} = \{0, 0, 0.5, 1, 1\}$ and $\Xi^F = \{0, 0, 0.25, 0.5, 0.75, 1, 1\}$ in both parametric directions.

In three-dimensions (3D) this repeated assembly of the row corresponding to d_l could happen more than twice depending on the geometry. Therefore, in three dimensions the number of times that d_l is constrained must be monitored so that the rows corresponding to d_l can be combined and divided by the correct value and the correct system of linear equations can be retrieved.

There are two more points to address. First, all of the above cases only deal with B-splines as the NURBS patches are all square with uniform weights. The extension to NURBS is accounted for through the generation of the T_f operator. In the case of NURBS, T_f will be generated for the four-dimensional B-splines in homogeneous coordinates which can be mapped back to the NURBS basis functions. T_f then represents the correct operator to be applied in the case of the constraint based local refinement methodology applied to NURBS based IGA.

Second, the above two cases have been performed with unit square NURBS patches. This means that the Jacobian of the geometry map is the identity matrix. As mentioned, NURBS patches are topologically equivalent to a hypercube. Due to the requirement that the unrefined geometry must form a compatible discretisation it is guaranteed that the geometry can be mapped back into a collection of squares that have common interfaces between them. Therefore, the extension of this method to non-square NURBS patches simply requires that the patch-wise stiffness matrices (K_{ff}^C, K_{nf}^F , etc) should have the Jacobian of their geometry map included in the integrals of the basis functions that form the stiffness matrices. This already happens within IGA algorithms so the extension to general NURBS geometries is naturally included in IGA spatial discretisations.

4. Adaptive Mesh Refinement (AMR) Methods

Adaptive mesh refinement (AMR) algorithms indicate specific areas of the problem domain for refinement based upon some underlying error metric or indicator. In this paper, two *a posteriori* error indicators will be considered. The refinement that is desired could be spatial, [29, 30], polynomial order [31], or a hybrid of the two [32, 33, 34, 35]. In this work only spatial refinement strategies are investigated. There is no reason that this methodology could not be extended to polynomial or hybrid refinement. All that is needed is an operator that represents the mapping between control points on the interface between neighbouring NURBS patches. Hybrid (spatial and polynomial) refinement has been shown to be the optimal refinement strategy. By spatially refining where the solution is non-smooth and increasing the order of basis functions where the solution is smooth, exponential convergence rates can be achieved [34, 36, 37].

The methods for indicating NURBS patches for refinement described in this section could be applied in lieu of a local refinement algorithm. Refinement of a specific NURBS patch can still be performed in standard NURBS based IGA. However, without the constraint based local refinement methodology, the refinements would not be truly local, as refinement of a NURBS patch would propagate to neighbouring NURBS patches and through subsections of the domain as was illustrated in Figure 2. Therefore, these error indicators, and their application to AMR algorithms, are discussed in isolation from the constraint based local refinement methodology detailed in Section 3. However, in Section 7, these AMR algorithms will be applied to locally refinable NURBS based IGA.

The *a posteriori* error measures used to indicate regions for spatial refinement can broadly be split into two groups: heuristic and goal-based error measures or indicators. Heuristic methods refine based on some general property of the solution. For example, the gradient of the solution, the error in the solution compared to some reference solution, or the jump in the solution between neighbouring elements in discontinuous Galerkin methods [38]. Heuristic error indicators can be

applied to any type of spatial discretisation method or solution. In this regard they are “black box” type AMR algorithms. Dual weighted residual (DWR) or goal-based error measures are tailored towards the specific outputs or “goals” such as the degree of criticality (K_{eff}) of a nuclear system, or the reaction rate within certain regions. In this work this is achieved by solving the adjoint neutron transport equation. The solution of the adjoint equation can broadly be considered to give a measure of a neutron’s importance to the scalar neutron flux [39]. By correctly setting the source terms of the adjoint neutron transport equation, the solution can be considered to give a measure of a neutron’s importance to a prescribed quantity of interest (QoI).

Both a heuristic error indicator and a goal-based error indicator are derived in the following section. Both of these *a posteriori* methods are based upon the error between a reference solution on a fine NURBS mesh and an approximate solution over a coarse NURBS mesh. The approximate solution can either be acquired by solving the neutron transport equation over the coarse NURBS mesh or by projecting the reference solution from the fine mesh to the coarse NURBS mesh.

IGA spatial discretisations enable spatial refinement within a NURBS patch to be performed without an ancillary mesh generation algorithm. Furthermore, IGA based AMR algorithms preserve the exact geometry of the spatial domain during the spatial refinement process. Indeed, even on the coarsest spatial refinement the geometry is exact. This is unlike traditional FE AMR algorithms, where a new FE mesh must be generated by an ancillary mesh generation algorithm at each AMR iteration.

4.1. The Forward Error Indicator (FEI)

The forward error indicator (FEI) is based upon calculating the error in a given NURBS patch with respect to some norm. This error is written as:

$$\|\psi - \psi_h\|_N, \quad (22)$$

where N is represents some functional norm, ψ is the exact solution, and ψ_h is the approximate solution. Some examples of N are the L^2 norm:

$$\|f(\mathbf{r})\|_{L^2} = \int_V f(\mathbf{r})^2 d\mathbf{r}, \quad (23)$$

or the H^1 norm:

$$\|f(\mathbf{r})\|_{H^1} = \|f(\mathbf{r})\|_{L^2} + \|\nabla f(\mathbf{r})\|_{L^2} = \int_V f(\mathbf{r})^2 + \nabla f(\mathbf{r}) \cdot \nabla f(\mathbf{r}) d\mathbf{r}. \quad (24)$$

The H^1 seminorm can also be considered:

$$|f(\mathbf{r})|_{H^1} = \|\nabla f(\mathbf{r})\|_{L^2} = \int_V \nabla f(\mathbf{r}) \cdot \nabla f(\mathbf{r}) d\mathbf{r}, \quad (25)$$

which is a seminorm due to the fact that $|f(\mathbf{r})|_{H^1} = 0 \not\Rightarrow f(\mathbf{r}) \equiv 0$.

In this work the H^1 norm is considered as this accounts not only for the magnitude of the function through the $\|f(\mathbf{r})\|_{L^2}$ term but also how oscillatory the function is through the $\|\nabla f(\mathbf{r})\|_{L^2}$ term. Therefore, the heuristic error indicator considered here is:

$$e^F = \|\psi - \psi_h\|_{H^1}, \quad (26)$$

where ψ is the exact scalar neutron flux and ψ_h is the calculated scalar neutron flux. The FEI e^F is calculated in each NURBS patch yielding e_p^F . All the e_p^F 's are then formed into an ascending list and a method for indicating which NURBS patches need refinement is selected. Two options are considered here.

In the first option, a NURBS patch is indicated for refinement if the error in that NURBS patch is within a certain threshold of the maximum error:

$$e_p^F > \epsilon \max\{e_p^F\}_{p=1}^N, \quad (27)$$

where $0 \leq \epsilon \leq 1$ and N is the number of patches. This is referred to as the threshold method. The second option is called the list method. A NURBS patch is indicated for refinement if it is in the top given percentile in the sorted list $\{e_p^F\}_{p=1}^N$. For example, if there are 100 NURBS patches and $\epsilon = 0.3$ then the top 30 NURBS patches are indicated for refinement.

The second strategy will ensure that the number of NURBS patches indicated for refinement in each AMR iteration will be the same. The first strategy should produce a more suitable NURBS mesh for a given problem but could require more AMR iterations.

In practice the error indicator from expression (26) is not suited to all problems as it contains the exact solution ψ . The exact solution for neutron transport is only available ahead of time for a small subset of problems that can be solved analytically [40]. Therefore, in order to approximate the exact solution, the substitution $\psi = \psi_{h+1}$ is made. ψ_{h+1} represents the solution of the neutron transport equation on a NURBS mesh where every NURBS patch has been uniformly refined to a higher level of spatial resolution than the NURBS mesh used to calculate ψ_h . Any highly refined solution could be considered, for example the solution $\psi_{h+1,p+1}$ where the solution has been solved over a once spatially refined NURBS mesh with basis functions that are one order higher [41]. A more flexible FEI is then given by:

$$\mu^F = \|\psi_{h+1} - \psi_h\|_{H^1}. \quad (28)$$

The process of calculating the FEI μ^F now involves solving two neutron transport problems: one over a coarse NURBS mesh to calculate ψ_h and one over a fine NURBS mesh to calculate ψ_{h+1} . One more step is taken to produce a FEI that only requires one neutron transport equation to be solved. This is achieved by letting $\psi_h = \Pi_h^{h+1}\psi_{h+1}$ where Π_h^{h+1} is a projection operator that projects the solution ψ_{h+1} to the coarse NURBS mesh. This choice can be justified by a mathematical argument from Wang and Ragusa [34]. Therefore, a flexible, computationally efficient, heuristic FEI is given by:

$$\eta^F = \|\psi_{h+1} - \Pi_h^{h+1}\psi_{h+1}\|_{H^1}. \quad (29)$$

4.2. The Weighted Error Indicator (WEI)

The weighted error indicator (WEI) is a DWR or goal-based error measure that builds upon the FEI η^F given in expression (29) in a manner similar to other work [29, 30, 42]. This is done by weighting η^F with the error in the adjoint solution. This is expressed mathematically as:

$$e^W = \eta^F \|\psi^\dagger - \psi_h^\dagger\|_{H^1} = \|\psi_{h+1} - \Pi_h^{h+1}\psi_{h+1}\|_{H^1} \|\psi^\dagger - \psi_h^\dagger\|_{H^1}, \quad (30)$$

where ψ^\dagger is the exact adjoint scalar neutron flux and ψ_h^\dagger is the calculated adjoint scalar neutron flux. Computing the WEI e^W requires the exact solution of the adjoint neutron transport equation as well as the numerical solution of one primal (forward) and one adjoint neutron transport

equation. An analogous series of approximations to the ones performed in Section 4.1 can be made to arrive at a WEI that requires the numerical solution of one primal neutron transport equation and one adjoint neutron transport equation. A flexible, goal-based WEI is then given by:

$$\eta^W = \eta^F \|\psi_{h+1}^\dagger - \Pi_h^{h+1} \psi_{h+1}^\dagger\|_{H^1} = \|\psi_{h+1} - \Pi_h^{h+1} \psi_{h+1}\|_{H^1} \|\psi_{h+1}^\dagger - \Pi_h^{h+1} \psi_{h+1}^\dagger\|_{H^1}. \quad (31)$$

Furthermore, the computational cost of this procedure could be halved by solving a special neutron transport formulation that yields the primal scalar neutron flux as its solution which can be differentiated to retrieve the adjoint scalar neutron flux in a manner similar to the even/odd parity neutron flux equations [43].

5. Second-order Forms of the Neutron Transport Equation

In this paper the FEI and WEI driven AMR algorithms will be applied to nuclear reactor physics and radiation shielding problems. In particular, two different second-order forms of the neutron transport equation are considered which are the self-adjoint angular flux (SAAF) equation:

$$-\mathbf{\Omega} \cdot \nabla \frac{1}{\sigma_t(\mathbf{r}, E)} \mathbf{\Omega} \cdot \nabla \psi(\mathbf{r}, \mathbf{\Omega}, E) + \sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, \mathbf{\Omega}, E) = Q(\mathbf{r}, \mathbf{\Omega}, E) - \left(\mathbf{\Omega} \cdot \nabla \frac{Q(\mathbf{r}, \mathbf{\Omega}, E)}{\sigma_t(\mathbf{r}, E)} \right), \quad (32)$$

and the weighted least squares (WLS) neutron transport equation:

$$\begin{aligned} & -\mathbf{\Omega} \cdot \nabla \mathcal{W} \mathbf{\Omega} \cdot \nabla \psi(\mathbf{r}, \mathbf{\Omega}, E) - \mathbf{\Omega} \cdot \nabla [\mathcal{W} \sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, \mathbf{\Omega}, E)] + \sigma_t(\mathbf{r}, E) \mathcal{W} \mathbf{\Omega} \cdot \nabla \psi(\mathbf{r}, \mathbf{\Omega}, E) \\ & + \mathcal{W} \sigma_t^2(\mathbf{r}, E) \psi(\mathbf{r}, \mathbf{\Omega}, E) = (\sigma_t(\mathbf{r}, E) - \mathbf{\Omega} \cdot \nabla) \mathcal{W} Q(\mathbf{r}, \mathbf{\Omega}, E), \end{aligned} \quad (33)$$

where \mathbf{r} is the position of a neutron in 3D space, E is the energy of a neutron, $\mathbf{\Omega}$ is the direction in which a neutron is travelling in steradians, $\sigma_t(\mathbf{r}, E)$ is the total macroscopic neutron cross-section, and $\psi(\mathbf{r}, \mathbf{\Omega}, E)$ is the angular neutron flux. \mathcal{W} is a weighting operator defined by

$$\mathcal{W} = \min\left(\frac{1}{\sigma_t}, w_{\max}\right). \quad (34)$$

The weighting operator defined above allows the WLS equation to solve neutron transport problems that contain void regions due to the removal of the singularity at $\sigma_t = 0$. The SAAF equation cannot model void regions explicitly due to the fact that the factor $\frac{1}{\sigma_t} \rightarrow \infty$ in a void region. A void treatment for the SAAF equation does exist, but the formulation does not produce symmetric positive definite (SPD) matrices in void regions [44]. Other void treatments for second-order forms of the neutron transport equation have been proposed [45, 46] but these methods also lose the SPD character of the global system. The source term $Q(\mathbf{r}, \mathbf{\Omega}, E)$ is defined as follows:

$$Q(\mathbf{r}, \mathbf{\Omega}, E) = Q_{\text{scatter}}(\mathbf{r}, \mathbf{\Omega}, E) + Q_{\text{fission}}(\mathbf{r}, E) + Q_{\text{ext}}(\mathbf{r}, \mathbf{\Omega}, E), \quad (35)$$

where

$$\begin{aligned} Q_{\text{scatter}}(\mathbf{r}, \mathbf{\Omega}, E) &= \int_0^\infty \int_{4\pi} \sigma_s(\mathbf{r}, \mathbf{\Omega}' \rightarrow \mathbf{\Omega}, E' \rightarrow E) \psi(\mathbf{r}, \mathbf{\Omega}', E') d\mathbf{\Omega}' dE', \\ Q_{\text{fission}}(\mathbf{r}, E) &= \chi(\mathbf{r}, E) \int_0^\infty \nu \sigma_f(\mathbf{r}, E') \phi(\mathbf{r}, E') dE', \end{aligned} \quad (36)$$

and $Q_{\text{ext}}(\mathbf{r}, \boldsymbol{\Omega}, E)$ represents an extraneous (fixed) source of neutrons. $\sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}, E' \rightarrow E)$ is the macroscopic differential neutron scattering cross-section, $\sigma_f(\mathbf{r}, E)$ is the macroscopic neutron fission cross-section, ν is the average number of prompt neutrons produced per fission, $\chi(\mathbf{r}, E)$ is the prompt fission spectrum, and $\phi(\mathbf{r}, E)$ is the scalar neutron flux. The choice of w_{max} in expression (34) can have a significant effect on the quality of the WLS solution. More information on the choice of w_{max} can be found in the following references [47, 48].

6. Derivation of the Physical Adjoints of the SAAF and WLS equation

The adjoint of the neutron transport equation has several uses in nuclear reactor physics and radiation shielding simulations. The solution of the adjoint neutron transport equation can be used as a variance reduction technique for Monte Carlo methods [49] and is also useful for calculating the effect of local perturbations in material coefficients on the reactivity of a system [50]. The solution of the adjoint neutron transport equation can be thought of as describing the importance of a neutron to the primal scalar neutron flux [39]. The adjoint is derived here to be used as part of a DWR or goal-based error measure. In particular, it will be used with the goal-based WEI derived in Section 4.2.

The adjoint of an operator H is defined as follows:

$$\langle Hu, v \rangle = \langle u, H^\dagger v \rangle, \quad (37)$$

where H^\dagger is the adjoint operator, $\langle \cdot, \cdot \rangle$ denotes some inner product, and u, v are two functions in the inner product space. The name self-adjoint angular flux implies that the SAAF equation would be its own adjoint. This is only the case when the term from equation (32), $\boldsymbol{\Omega} \cdot \nabla \frac{Q(\mathbf{r}, \boldsymbol{\Omega}, E)}{\sigma_t} \equiv 0$, which is satisfied for mono-energetic, homogeneous problems with no scattering. Therefore, in the general case, the adjoint of the SAAF equation must be formulated and solved separately.

The actions of discretising the individual dimensions of the phase space of the neutron transport equation $(\mathbf{r}, \boldsymbol{\Omega}, E)$ and taking the adjoint according to expression (37) do not, in general, commute. These adjoints are commonly referred to as the continuous-adjoint (physical adjoint) or the discrete-adjoint (mathematical adjoint) depending on the order in which these operations are performed [51]. However, when taking the adjoint of the second-order forms of the neutron transport equation there are several possibilities for the continuous-adjoint. Following from the notation used in the paper by Schunert et al. [52] a vector is used to signify the order in which discretisation and the action of taking the adjoint are applied. For example, the vector $(E, \boldsymbol{\Omega}, \dagger, \mathbf{r})$ means that first the multigroup approximation is applied, then the discrete ordinate method for angular discretisation, then the adjoint is taken, and then IGA is used to discretise the spatial domain. There are three main options to consider:

1. True adjoint: $(\dagger, E, \boldsymbol{\Omega}, \mathbf{r})$
2. Physical adjoint: $(E, \dagger, \boldsymbol{\Omega}, \mathbf{r}) \equiv (E, \boldsymbol{\Omega}, \dagger, \mathbf{r})$
3. Mathematical adjoint: $(E, \boldsymbol{\Omega}, \mathbf{r}, \dagger)$

In this work the physical adjoint is used for the following reasons. First, if the true adjoint is used then adjoint flux weighted neutron cross-sections are necessary. In this work we rely heavily on neutron cross-section data being provided from test cases or from nuclear data libraries. The neutron cross-sections from nuclear data libraries have usually been collapsed from a fine group structure to a coarse group structure using lattice calculations and are weighted by the primal

scalar neutron flux [53]. Since adjoint weighted neutron cross-section data are not generally available this rules out the true adjoint.

The comparison of the mathematical and physical adjoints is more subtle. In the case of the first-order neutron transport equation with a linear spatial discretisation, the physical and mathematical adjoints are the same so the comparison between the two adjoint formulations is not often considered [52]. When the mathematical adjoint is used the discretisation errors in the primal and adjoint calculations are the same. Therefore, the eigenvalues of the primal equation and mathematical adjoint equation agree exactly as they both include the same discretisation error [52]. This is not the case when the physical adjoint is used [54]. However, in the asymptotic limit, as the discretisation error goes to zero, the eigenvalues of the primal equation, mathematical adjoint equation, and the physical adjoint equation are identical [52].

It was mentioned earlier that the adjoint can be thought of as a measure of a neutron's importance to the system [39] which makes the adjoint solution suited to a goal-based error measure. This physical interpretation is the case for the physical adjoint. It is unclear what the physical significance of the mathematical adjoint is outside of the asymptotic limit [55], but it is assumed that, since the solutions of the physical and mathematical adjoints differ only by discretisation error that the mathematical adjoint would also be suitable for use in goal-based error measures. However, despite the literature suggesting that the mathematical adjoint is preferred to the physical adjoint in terms of accuracy and computational time [52, 54, 55], the physical adjoint is used here, mainly due to the fact that it is consistent with the forward SAAF equation [52].

In order to apply the WEI derived earlier in Section 4.2, the physical adjoints of the SAAF and WLS equations must be derived. They are derived in their multigroup form to avoid the neutron transport cross-sections being weighted by the adjoint scalar neutron flux [52], referred to as the true adjoint above. In order to derive the physical adjoints, the multigroup first-order form of the neutron transport equation is given:

$$\boldsymbol{\Omega} \cdot \nabla \psi_g(\mathbf{r}, \boldsymbol{\Omega}) + \sigma_{t,g}(\mathbf{r})\psi_g(\mathbf{r}, \boldsymbol{\Omega}) = Q_g(\mathbf{r}, \boldsymbol{\Omega}), \text{ for } g = 1, \dots, G \quad (38)$$

with a prescribed incoming boundary condition given by:

$$\psi_g(\mathbf{r}, \boldsymbol{\Omega}) = T_g(\mathbf{r}, \boldsymbol{\Omega}) \text{ for } \mathbf{r} \in S \text{ and } \boldsymbol{\Omega} \cdot \mathbf{n} < 0, \quad (39)$$

where S is the surface of the problem domain, \mathbf{n} is the outward unit normal vector, and $T_g(\mathbf{r}, \boldsymbol{\Omega})$ is a function representing the prescribed incoming angular neutron flux. It can be used to represent all relevant prescribed boundary conditions, for example:

- $T_g(\mathbf{r}, \boldsymbol{\Omega}) = 0 \implies$ vacuum boundary condition.
- $T_g(\mathbf{r}, \boldsymbol{\Omega}) = F_g(\mathbf{r}, \boldsymbol{\Omega}) \implies$ incoming boundary source condition.
- $T_g(\mathbf{r}, \boldsymbol{\Omega}) = \psi_g(\mathbf{r}, \boldsymbol{\Omega}_r) \implies$ reflective boundary condition, where $\boldsymbol{\Omega}_r$ is the angle of specular reflection of $\boldsymbol{\Omega}$.

Q_g corresponds to equation (35) where the multigroup method has been used to discretise energy. That is:

$$Q_g(\mathbf{r}, \boldsymbol{\Omega}) = \sum_{g'=1}^G \int_{4\pi} \sigma_{s,g' \rightarrow g}(\mathbf{r}, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}) \psi_{g'}(\mathbf{r}, \boldsymbol{\Omega}') d\boldsymbol{\Omega}' + \chi_g(\mathbf{r}) \sum_{g'=1}^G \nu \sigma_{f,g'}(\mathbf{r}) \phi_{g'}(\mathbf{r}) + Q_{\text{ext},g}(\mathbf{r}, \boldsymbol{\Omega}). \quad (40)$$

Furthermore, the adjoint of the multigroup first-order neutron transport equation, from equation (38), is given by

$$-\mathbf{\Omega} \cdot \nabla \psi_g^\dagger(\mathbf{r}, \mathbf{\Omega}) + \sigma_{t,g}(\mathbf{r}) \psi_g^\dagger(\mathbf{r}, \mathbf{\Omega}) = Q_g^\dagger(\mathbf{r}, \mathbf{\Omega}), \text{ for } g = 1, \dots, G \quad (41)$$

with

$$Q_g^\dagger(\mathbf{r}, \mathbf{\Omega}) = \sum_{g'=1}^G \int_{4\pi} \sigma_{s,g \rightarrow g'}(\mathbf{r}, \mathbf{\Omega} \rightarrow \mathbf{\Omega}') \psi_{g'}^\dagger(\mathbf{r}, \mathbf{\Omega}') d\mathbf{\Omega}' + \nu \sigma_{f,g}(\mathbf{r}) \sum_{g'=1}^G \chi_{g'}(\mathbf{r}) \phi_{g'}^\dagger(\mathbf{r}) + Q_{\text{ext},g}^\dagger(\mathbf{r}, \mathbf{\Omega}), \quad (42)$$

and the prescribed boundary conditions for equation (41) is given by:

$$\psi_g^\dagger(\mathbf{r}, \mathbf{\Omega}) = T_g(\mathbf{r}, \mathbf{\Omega}) \text{ for } \mathbf{r} \in S \text{ and } \mathbf{\Omega} \cdot \mathbf{n} > 0. \quad (43)$$

6.1. Derivation of the Physical Adjoint of the SAAF Equation

The starting point in the derivation of the physical adjoint of the SAAF equation is the multigroup adjoint first-order neutron transport equation. It then proceeds in the same fashion as the algebraic derivation of the primal SAAF equation [28, 56]. Equation (41) is rearranged to form the following expression:

$$\psi_g^\dagger = \frac{1}{\sigma_{t,g}} (Q_g^\dagger + \mathbf{\Omega} \cdot \nabla \psi_g^\dagger), \text{ for } g = 1, \dots, G. \quad (44)$$

Expression (44) is then substituted into the streaming term of equation (41) giving:

$$-\mathbf{\Omega} \cdot \nabla \frac{1}{\sigma_{t,g}} \mathbf{\Omega} \cdot \nabla \psi_g^\dagger + \sigma_{t,g} \psi_g^\dagger = Q_g^\dagger + \mathbf{\Omega} \cdot \nabla \frac{Q_g^\dagger}{\sigma_{t,g}}, \text{ for } g = 1, \dots, G. \quad (45)$$

The prescribed outgoing boundary conditions for equation (45) can be derived from equation (43) as:

$$\frac{1}{\sigma_{t,g}} (Q_g^\dagger + \mathbf{\Omega} \cdot \nabla \psi_g^\dagger) = T_g^\dagger(\mathbf{r}, \mathbf{\Omega}), \text{ for } \mathbf{\Omega} \cdot \mathbf{n} > 0. \quad (46)$$

For the prescribed incoming boundary condition it is assumed that the adjoint angular neutron flux should satisfy equation (41) yielding:

$$\frac{1}{\sigma_{t,g}} (Q_g^\dagger + \mathbf{\Omega} \cdot \nabla \psi_g^\dagger) = \psi_g^\dagger, \text{ for } \mathbf{\Omega} \cdot \mathbf{n} < 0. \quad (47)$$

Equations (45), (46), and (47) form the physical adjoint of the SAAF equation. This derivation is equivalent to the transform suggested by Schunert et al. [52]. The transform is repeated here for clarity:

- The scattering matrix is transposed. This is equivalent to the substitution: $\sigma_{s,g' \rightarrow g}(\mathbf{r}, \mathbf{\Omega}' \cdot \mathbf{\Omega}) \rightarrow \sigma_{s,g \rightarrow g'}(\mathbf{r}, \mathbf{\Omega}' \cdot \mathbf{\Omega})$.

- The product of the average number of prompt neutron per fission and the fission macroscopic neutron cross-section and the prompt neutron fission spectrum are switched. This is equivalent to: $\nu\sigma_{f,g'}(\mathbf{r}) \rightarrow \chi_{g'}(\mathbf{r})$ and $\chi_g(\mathbf{r}) \rightarrow \nu\sigma_{f,g}(\mathbf{r})$.
- The forward (primal) scalar and angular neutron fluxes are replaced with the adjoint (dual) scalar and angular neutron fluxes: $\psi_g(\mathbf{r}, \mathbf{\Omega}) \rightarrow \psi_g^\dagger(\mathbf{r}, \mathbf{\Omega})$, $\phi_g(\mathbf{r}) \rightarrow \phi_g^\dagger(\mathbf{r})$.
- Solutions are reinterpreted in the opposite direction. This is equivalent to replacing every instance of $\mathbf{\Omega}$ with $-\mathbf{\Omega}$.

6.2. Derivation of the Physical Adjoint of the WLS Equation

The adjoint of the WLS equation can be derived in two different ways. The first method, which will be detailed here, uses the adjoint first-order form of the neutron transport equation to begin the derivation of the primal WLS equation. In the second method, mentioned at the end of Section 6.1, the transform proposed by Schunert et al. [52] is applied to the primal WLS formulation in equation (33).

To begin the derivation equation (41) is written in operator form:

$$\mathcal{L}_g^\dagger \psi_g^\dagger(\mathbf{r}, \mathbf{\Omega}) = Q_g^\dagger(\mathbf{r}, \mathbf{\Omega}). \quad (48)$$

Equation (48) is then operated on by a within group weighting operator $\mathcal{W}_g = \min\left(\frac{1}{\sigma_{t,g}(\mathbf{r})}, w_{\max}\right)$ with $w_{\max} \in \mathbb{R}$, followed by the within group transport operator \mathcal{L}_g giving:

$$\mathcal{L}_g \mathcal{W}_g \mathcal{L}_g^\dagger \psi_g^\dagger(\mathbf{r}, \mathbf{\Omega}) = \mathcal{L}_g \mathcal{W}_g Q_g^\dagger(\mathbf{r}, \mathbf{\Omega}), \quad (49)$$

where the within group transport operators are given by:

$$\begin{aligned} \mathcal{L}_g &= \mathbf{\Omega} \cdot \nabla + \sigma_{t,g}(\mathbf{r}), \\ \mathcal{L}_g^\dagger &= -\mathbf{\Omega} \cdot \nabla + \sigma_{t,g}(\mathbf{r}). \end{aligned} \quad (50)$$

Expanding the operators in equation (49) yields the strong form of the physical adjoint of the WLS neutron transport equation:

$$\begin{aligned} & -\mathbf{\Omega} \cdot \nabla \mathcal{W}_g \mathbf{\Omega} \cdot \nabla \psi_g^\dagger + \mathbf{\Omega} \cdot \nabla \left(\mathcal{W}_g \sigma_{t,g} \psi_g^\dagger \right) - \sigma_{t,g} \mathcal{W}_g \mathbf{\Omega} \cdot \nabla \psi_g^\dagger + \sigma_{t,g}^2 \mathcal{W}_g \psi_g^\dagger \\ & = \left(\sigma_{t,g} + \mathbf{\Omega} \cdot \nabla \right) \mathcal{W}_g Q_g^\dagger. \end{aligned} \quad (51)$$

6.2.1. Choice of the Goal Functional

In order to have an effective WEI the goal functional for the goal-based error measure must be set. This is done by choosing an expression to represent the fixed source term, $Q_{\text{ext},g}^\dagger$ in equation (42). To set the goal as some neutron reaction rate in the system, the choice:

$$Q_{\text{ext},g}^\dagger(\mathbf{r}) = \sigma_{d,g}(\mathbf{r}) \phi_g(\mathbf{r}), \quad (52)$$

is made. For example, the rate of neutron absorption in the system can be set as the goal functional by choosing $\sigma_{d,g}(\mathbf{r}) = \sigma_{a,g}(\mathbf{r})$. If the K_{eff} is desired as the QoI then the choice $\sigma_{d,g}(\mathbf{r}) = \nu\sigma_{f,g}(\mathbf{r})$ is made. This is based upon the fact that the K_{eff} is calculated by:

$$K_{\text{eff}}^{(i)} = \frac{\int_V \int_0^\infty \nu\sigma_f(\mathbf{r}, E) \phi^{(i)}(\mathbf{r}, E) dE d\mathbf{r}}{\int_V \int_0^\infty \nu\sigma_f(\mathbf{r}, E) \phi^{(i-1)}(\mathbf{r}, E) dE d\mathbf{r}} K_{\text{eff}}^{(i-1)} \quad (53)$$

where $\phi^{(i)}$ is the scalar neutron flux from the i^{th} power iteration [57]. Having a goal-based error measure with the fission rate as the goal functional would also mean that K_{eff} could be considered a goal functional.

7. Results

In this section the continuous constraint based, locally refinable NURBS based IGA algorithm is applied to several nuclear reactor physics and radiation shielding verification benchmark test cases. The aim of these verification test cases is to illustrate the effect that constraint based local refinement has on the convergence rates of the solution. In addition, these test cases also illustrate how errors introduced by local refinement propagate throughout the problem domain. Furthermore, the heuristic error measure (FEI) and goal-based error measure (WEI) derived in Section 4 are applied to several mono-energetic nuclear reactor physics verification benchmark test cases. The results of applying both the FEI and WEI AMR algorithms are compared with the aim of identifying cases when the extra computational effort required for WEI is appropriate for a given QoI.

All IGA- S_N solutions of the second-order forms of the neutron transport equation have been performed using a modified version of ICARUS (Isogeometric Continuous self-Adjoint Radiation Using Splines) written in Modern Fortran [58]. Where necessary, reference solutions have been generated using the neutron transport code Inferno. Inferno is a first-order, discontinuous IGA discrete ordinate (S_N) neutron transport code with DWR or goal-based local spatial refinement capabilities [3, 38, 59].

7.1. Method of Manufactured Solutions (MMS) Verification Benchmark Test Case

The method of manufactured solutions (MMS) is a standard mathematical method for the numerical verification of numerical discretisation methods [60, 61]. It is used here to investigate the effect that constraint based local refinement has on the rate of convergence and errors in the constrained (fine) and unconstrained (coarse) NURBS patches.

The manufactured solution chosen to generate source terms for the SAAF equation is given by:

$$\psi_M(x, y) = x(x - 10)y(y - 10) \left(\cos\left(\frac{4\pi x}{5}\right) \cos\left(\frac{4\pi y}{5}\right) + 2 \right). \quad (54)$$

The problem domain chosen for the MMS verification test case is a $10\text{cm} \times 10\text{cm}$ square divided into four equally sized NURBS patches. The bottom left and top right NURBS patches are one level more refined than the bottom right and top left NURBS patches. In this case a level of refinement means that every non-zero knot span in the knot vector has had a knot inserted at its midpoint. The geometry for the coarsest level of refinement can be seen in Figure 6. Figure 7 presents the error between the manufactured solution and the discrete solution for locally refined and uniformly NURBS refined geometries.

Plotted in Figure 8 are the errors in the MMS test in the fine NURBS patches, in the coarse NURBS patches, and in the case of uniformly refined computational domain. The rate of convergence of a solution is indicated by the gradient of the line joining the solution points. The first point of note is that the rate of convergence and level of error in the solution in the coarse NURBS patches is identical to the rate of convergence and level of error in the solution for the uniformly refined computational domain. This is to be expected as the local constraints apply to the dof in the fine NURBS patch not in the coarse NURBS patch. This also explains why, for




Figure 6: Geometry for the MMS verification benchmark test case.

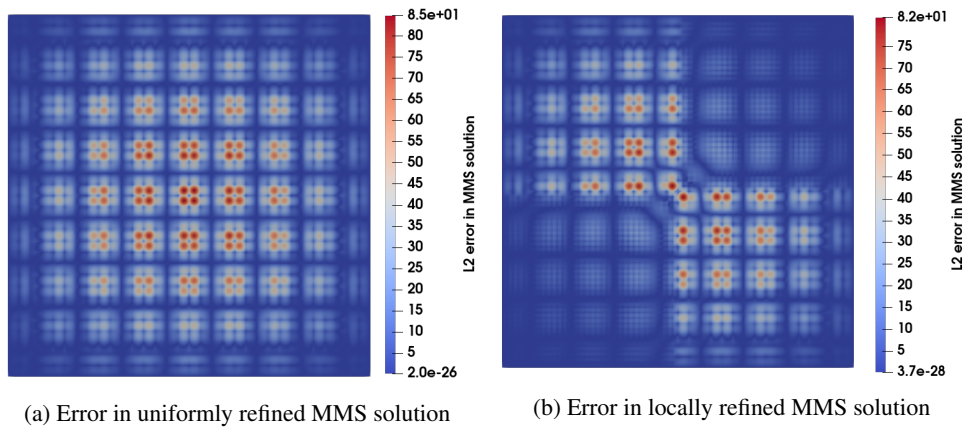


Figure 7: Plot of the error in the MMS verification benchmark test case.

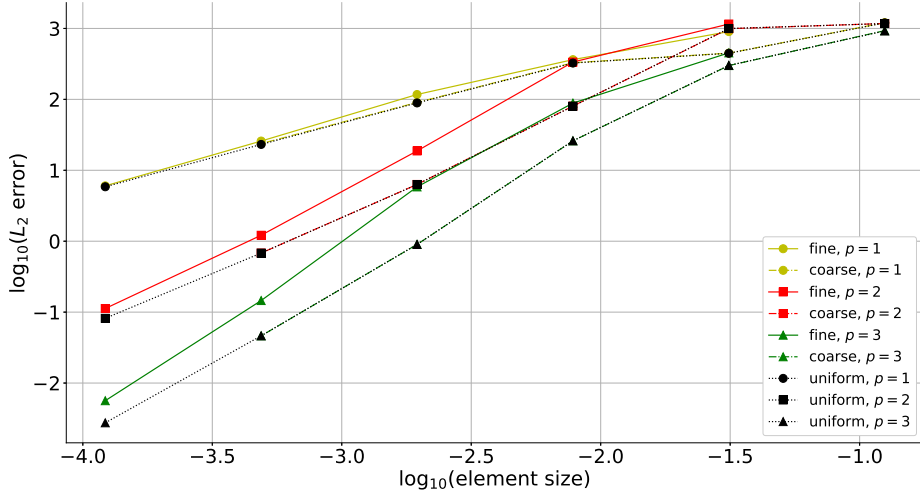


Figure 8: Spatial convergence plots for the MMS verification benchmark test case. The rates of convergence and L^2 error for the coarse NURBS patches and the uniformly refined NURBS patches match exactly for all orders of basis function.

the same element size, the error in the fine NURBS patches is greater than in the coarse NURBS patches. The dof in the fine NURBS patches are being constrained to match the dof in the coarse NURBS patches along the shared boundary, introducing an error into the fine NURBS patch solution.

Second, it can be observed that, whilst the error in the fine NURBS patches is greater than in the coarse NURBS patches for the same level of refinement, the rate of convergence of the solution in the fine NURBS patches is greater than the rate of convergence in the coarse NURBS patches. Furthermore, the error in the fine NURBS patches is tending towards the error in the coarse NURBS patches. It is expected that the level of error in the fine NURBS patches will decrease toward the level of error in the coarse NURBS patches as the problem is refined. In the fine NURBS patches, the fraction of dof being constrained is $O(\frac{n}{n^2}) = O(\frac{1}{n})$, where n is the number of dof on a side. As more dof are inserted into the geometry $n \rightarrow \infty$, $\frac{1}{n} \rightarrow 0$. Therefore, the proportion of dof being constrained, and the magnitude of the error induced by these constraints, will decrease. Once the level of error in both sets of NURBS patches match, it is also expected that the rates of convergence will be the same. The rates of convergence for each order of basis function are not commensurate with the expected theoretical rate of $O(h^{p+1})$, in fact they seem to be $O(h^p)$. This is due to the strong imposition of the zero flux boundary conditions caused by the big-spring method that is affecting the rate of convergence of the solution [62].

7.2. 5×1 Nuclear Fuel Pincell Lattice Reactor Physics Verification Benchmark Test Case

The aim of the next verification benchmark test case is to investigate how errors introduced by constraints and areas of low spatial resolution propagate through the domain. The geometry for this test case consists of a 5×1 array of nuclear fuel pincells with prescribed reflective boundary conditions. All nuclear fuel pincells have the same macroscopic neutron cross-sections and are




Figure 9: Geometry of the 5×1 nuclear fuel pincell row verification benchmark test case. All boundaries have prescribed reflective boundary conditions.




Figure 10: Discretised geometry of the 5×1 nuclear fuel pincell row verification benchmark test case.

geometrically identical. Therefore, this verification benchmark test case is similar to an infinite nuclear fuel pincell lattice calculation. The geometry is shown in Figure 9 and the specific spatial refinement configuration for this verification benchmark test case is shown in Figure 10. All nuclear fuel pincells, except the far left one, are formed of 5 bi-quadratic NURBS patches, each consisting of 256 elements. The far left nuclear fuel pincell is formed of 5 bi-quadratic NURBS patches, each one consisting of one element. The problem is discretised with quadratic NURBS basis functions and angularly discretised with a S_{12} discrete ordinate (S_N) quadrature set.

The QoI considered for this verification benchmark test case is the disadvantage factor in each nuclear fuel pincell. This is defined as the ratio of the average flux in the moderator to the average flux in the fuel. For the given geometry the cell disadvantage factor is given by:

$$\zeta = \frac{\pi}{16 - \pi} \frac{\int_{\text{moderator}} \phi(\mathbf{r}) d\mathbf{r}}{\int_{\text{fuel}} \phi(\mathbf{r}) d\mathbf{r}}. \quad (55)$$

Six different sets of material properties are considered and are given in Table 1. In each set of material properties the scattering ratio $c = \frac{\sigma_s}{\sigma_t} = 0.95$ is kept constant whilst σ_t and σ_s are varied. The ratio of the neutron cross-sections between fuel and moderator is also kept constant. The solution was converged until the residual of ϕ was less than 1×10^{-10} . Figure 11 displays the errors in the cell disadvantage factor for each nuclear fuel pincell. The error is calculated by

| Case | Macroscopic neutron cross-section data (cm ⁻¹) | | | | Extraneous source data (neutrons.cm ⁻³ .s ⁻¹) | |
|------|--|--------------|--------------|--------------|--|-------|
| | σ_t^F | σ_s^F | σ_t^M | σ_s^M | S_F | S_M |
| 1 | 2 | 1.9 | 1 | 0.95 | 1 | 1 |
| 2 | 5 | 4.75 | 2.5 | 2.375 | 1 | 1 |
| 3 | 10 | 9.5 | 5 | 4.75 | 1 | 1 |
| 4 | 20 | 19 | 10 | 9.5 | 1 | 1 |
| 5 | 200 | 190 | 100 | 95 | 1 | 1 |
| 6 | 500 | 475 | 250 | 237.5 | 1 | 1 |

Table 1: Macroscopic neutron cross-section data for the 5×1 nuclear fuel pincell row verification benchmark test case. A F superscript refers to the fuel and a M superscript refers to the moderator.

comparing the locally refined ICARUS solution with an uniformly refined ICARUS solution.

As σ_t^F and σ_t^M increase it can be observed that the distance the error propagates through the domain decreases. This is because the mean free path of a neutron is given by $\lambda = \frac{1}{\sigma_t}$. In the cases where $\sigma_t^F \geq 200$ the error introduced by the low resolution of nuclear fuel pincell one and by the constraint between nuclear fuel pincells one and two no longer effects the value of ζ in nuclear fuel pincell three due to a small λ . In the cases where $\sigma_t^F < 10$ the error can be seen to propagate throughout the whole domain due to a larger λ . This manifests as an error greater than 1×10^{-11} in ζ in nuclear fuel pincell five.

Another trend apparent from Figure 11 is that the error in ζ for nuclear fuel pincells one and two increases with σ_t . This is assumed to be due to the fact that the gradient of the solution over the boundary between the fuel and the moderator becomes steeper as σ_t^F increases. The low spatial resolution of nuclear fuel pincell one will lead to greater error in ζ as this gradient gets steeper.

7.3. KAIST Nuclear Fuel Pincell Reactor Physics Verification Benchmark Test Case

In order to investigate the interaction between the constraint based local refinement methodology and the impact of voids on the solution of the WLS equation the seven group KAIST nuclear fuel pincell verification benchmark test case is considered. The geometry of the nuclear fuel pincell is displayed in Figure 12. The macroscopic neutron cross-section data for all seven groups can be found in the following reference [63]. This nuclear reactor physics nuclear fuel pincell verification benchmark test case is more realistic when compared to the nuclear fuel pincell problem considered in Section 7.2. A section of cladding and a small air-gap are modelled as part of the domain. In a nuclear reactor cladding is used to contain the fuel and keep it separate from the moderator. A small gap is left between the fuel pellets and the cladding so that when the fuel pellets expand as they heat up the cladding does not begin to crack. The low value of σ_t in the air-gap can cause the SAAF equation to fail to solve this problem, due to the problem becoming ill-conditioned. The WLS equation is able to cope with the small values of σ_t and void regions due to the weighting operator \mathcal{W} from equation (34).

The disadvantage factor for each energy group is considered. For this particular geometry

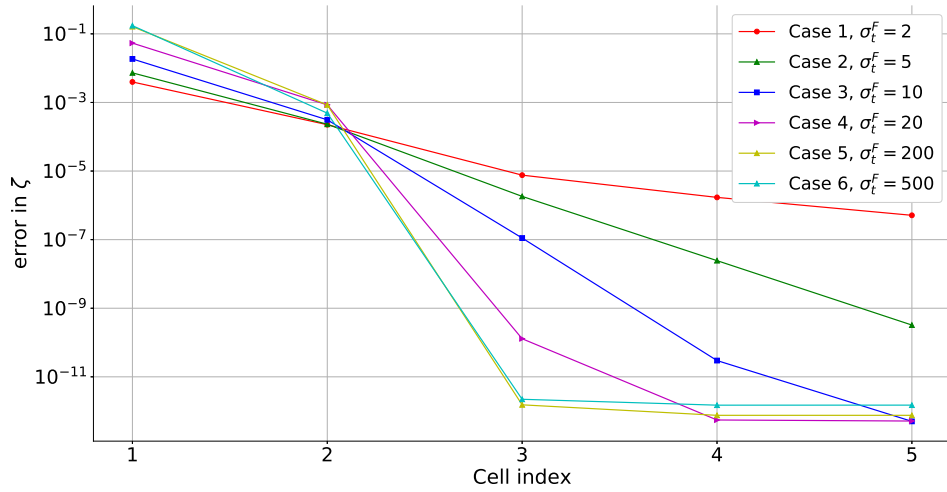


Figure 11: Error in ζ for the 5×1 nuclear fuel pincell row verification benchmark test case.



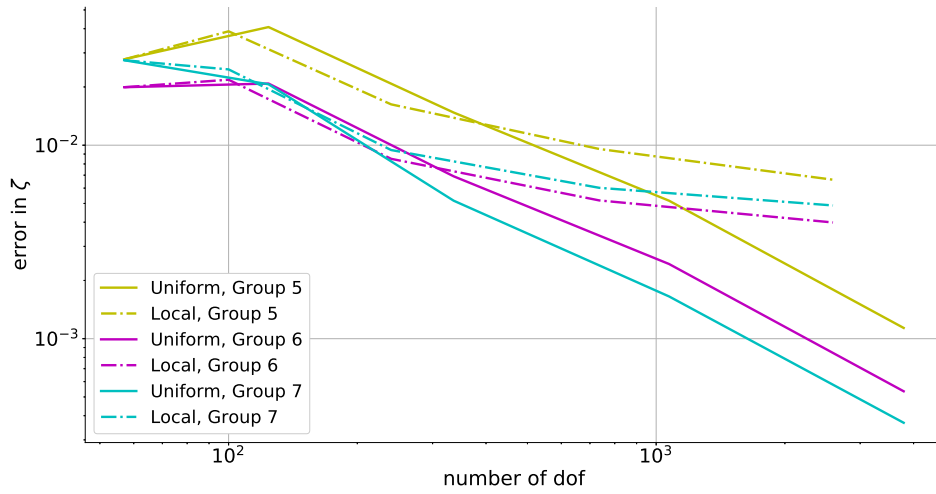


Figure 12: Geometry of the KAIST nuclear fuel pincell nuclear reactor physics verification benchmark test case. Reflective boundary conditions are prescribed on all external boundaries. The nuclear fuel pincell consists of fuel surrounded by a moderator with a layer of cladding between the fuel and moderator and a small air-gap between cladding and fuel. This air-gap is included in real nuclear reactors to allow fuel pellets to swell without causing cracking of the cladding.



(a) Error in ζ_g for groups one, two, three, and four.



(b) Error in ζ_g for groups five, six, and seven.

Figure 13: Error in the disadvantage factor ζ_g for two refinement schemes for the KAIST nuclear fuel pincell verification benchmark test case. 'Uniform' refines all areas uniformly, and 'local' refines everywhere except the air-gap (void region). The WLS equation spatially discretised with quadratic IGA and angularly discretised with S_8 LCT discrete ordinates (S_N) and with $w_{\max} = 1.0$ was used to generate results.




Figure 14: Element boundaries of the KAIST pincell nuclear reactor physics verification benchmark test case as defined by the knot spans. This is the refinement scheme referred to in the local refinement case. The thin void region is unrefined whilst the neighbouring fuel pin and cladding are highly refined.

the disadvantage factor for each energy group is given by the following expression:

$$\zeta_g = \frac{\int_{\text{moderator}} \phi_g(\mathbf{r}) d\mathbf{r}}{0.878778} \frac{0.5268814}{\int_{\text{fuel}} \phi_g(\mathbf{r}) d\mathbf{r}}. \quad (56)$$

The error in ζ_g for all energy groups is plotted in Figure 13 and has been generated by solving the WLS equation spatially discretised with quadratic NURBS based IGA and angularly discretised with S_8 Legendre-Chebyshev Triangular (LCT) discrete ordinates and with $w_{\max} = 1$. Two refinement schemes have been used: for uniform refinement all areas have been refined equally, and in the local case everywhere but the void has been refined as is displayed in Figure 14.

From Figures 13a and 13b it can be observed that across all energy groups the error in ζ_g is greater for the locally refined case. Furthermore, it is clear that there is a drop in the rate of convergence of ζ_g as indicated by the smaller gradient. However, this increase in error and drop in convergence rate does not show the entire picture.

Table 2 displays the spectral condition number (κ) of a symmetric successive over-relaxation (SSOR) preconditioned spatial system matrix for a given angle and group in both the uniformly and locally refined cases. The spectral condition number of a normal matrix A is given by the following formula:

$$\kappa(A) = \frac{|\lambda_{\max}(A)|}{|\lambda_{\min}(A)|}, \quad (57)$$

where $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ are the maximum and minimum eigenvalues of the matrix A . A GMRES solver within PETSc with a restart of 1000 is used to estimate the eigenvalues. The option `-ksp_monitor_singular_value` must be used. More information about PETSc solution algorithm options can be found in the PETSc manual [64]. The global stiffness matrices for this problem were normal and so equation (57) is valid. The spectral condition number κ is not a

| Uniform refinement | | Local refinement | |
|--------------------|--------------------|------------------|--------------------|
| number of dof | κ | number of dof | κ |
| 124 | 8.09×10^4 | 100 | 4.61×10^2 |
| 336 | 1.21×10^5 | 240 | 5.72×10^2 |
| 1072 | 3.12×10^5 | 736 | 1.19×10^2 |
| 3792 | 3.73×10^5 | 2592 | 5.71×10^4 |

Table 2: Condition number κ for the KAIST nuclear fuel pincell. The condition number was converged to 1×10^{-9} and the system was preconditioned with a symmetric successive over-relaxation (SSOR) scheme. If it took more than 1000 GMRES iterations then the value at the 1000th iteration is taken and is coloured red.

direct measure of how long it will take to solve a system of linear equations but is an indication of the computational time required. A larger value of κ implies the matrix system is harder to invert.

It can be seen in Table 2 that κ is between one and three orders of magnitude less for the locally refined case. For example, in the second row of Table 2 the ratio $\frac{\kappa_{\text{uniform}}}{\kappa_{\text{local}}} = 211.5$ whilst the uniformly refined case only ran 2.36 times longer than the locally refined case. However, the increase in κ and runtime versus the potential increase in error is worth considering when determining regions for refinement.

In this case, the error introduced in ζ_g by not refining the gas region is quite severe. However, ζ_g is a global quantity in this problem. If the QoI was a quantity that was further away from the void region then it is reasonable to assume that the QoI would be relatively unaffected by the error introduced by not refining the void region. This, coupled with the large increase in κ , may make it desirable to not refine void regions in certain cases. Preconditioners could also be designed to enable efficient solution of the resulting matrices. For example, a non-overlapping domain decomposition solution algorithm combined with effective preconditioners for solving the domains that result in a higher spectral condition number, the void region in this case, could result in an efficient and parallelisable preconditioned solution algorithm [65, 66].

7.4. Adaptive Mesh Refinement (AMR) Applied to the IAEA Swimming Pool Nuclear Reactor Physics Verification Benchmark Test Case

To illustrate how constraint based local refinement would be used in practice, an IGA AMR algorithm has been implemented. Here it is applied to the IAEA swimming pool nuclear reactor physics verification benchmark test case [67]. The geometry and macroscopic neutron cross-section data for the verification benchmark test case can be seen in Figure 15 and Table 3 respectively. The IAEA swimming pool verification benchmark test case has been chosen as a benchmark due to the difference in the order of magnitude of the scalar neutron flux in each region. The region averaged scalar neutron fluxes vary by up to three orders of magnitude between regions [68]. Therefore, the problem will have strong localised gradients that can be efficiently resolved by local refinement.

The AMR algorithm uses the WEI η^W detailed in equation (31). The projection operation Π_h^{h+1} is performed using the energy dependent mesh interpolation methods developed by Welch [23] and Owens [59] that minimises the error between $f(\mathbf{r})$ and $\Pi_h^{h+1} f(\mathbf{r})$ with respect to the L^2 norm. Once η_p^W has been computed for each NURBS patch, a NURBS patch is indicated for refinement if $\eta_p^W > \epsilon \max\{\eta_p^W\}_{p=1}^N$ with $\epsilon = 0.3$, referred to in Section 4 as the threshold method.




Figure 15: Geometry of the IAEA swimming pool nuclear reactor physics verification benchmark test case.

This AMR step is repeated and, due to fact that in IGA knot insertion (spatial refinement) is performed at run time, no interfacing with an ancillary mesh generator is necessary between AMR iterations.

| Region | Macroscopic neutron cross-section data (cm^{-1}) | | |
|--------|---|------------|---------------|
| | σ_t | σ_s | $\nu\sigma_f$ |
| 1 | 0.6 | 0.53 | 0.079 |
| 2 | 0.48 | 0.20 | 0.0 |
| 3 | 0.70 | 0.66 | 0.043 |
| 4 | 0.65 | 0.50 | 0.0 |
| 5 | 0.90 | 0.89 | 0.0 |

Table 3: One-group macroscopic neutron cross-section data for the eigenvalue (K_{eff}) IAEA swimming pool nuclear reactor physics verification benchmark test case.

All results are calculated using the SAAF equation and the angular domain is discretised using a S_4 LCT discrete ordinate (S_N) quadrature set for all calculations, including the generation of the reference solution. This will eliminate the effect that the angular discretisation error will have between results, as identical angular quadrature sets are used for all results. The QoI chosen is the K_{eff} and is set as the goal for the WEI as set out in Section 6.2.1. A reference value of $K_{\text{eff}} = 1.0087598988$ was generated using Inferno with fourth-order NURBS basis functions and 468,571 elements.




Figure 16: Error in the K_{eff} calculated from the S_4 -SAAF equation solved for the IAEA swimming pool nuclear reactor physics verification benchmark test case. Vacuum boundary conditions are prescribed on all external boundaries. An AMR scheme driven by the WEI is compared to uniform refinement.

The results of the IGA AMR algorithm applied to the IAEA swimming pool verification benchmark test case are displayed in Figure 16. It can be seen that the error in the locally refined geometry is about half an order of magnitude less than the error in the uniformly refined geometry for the same number of dof. The WEI AMR algorithm is computationally more efficient and accurate than the uniformly refined solution with respect to accuracy per dof. However, the primal and adjoint SAAF equations have been solved for the WEI case. Consequently, for a given number of dof, the computational effort required to calculate the WEI is roughly double that required to calculate the uniformly refined solution. If the methodology laid out in work by Hanus and McClarren was implemented then the WEI results could potentially be acquired at no extra cost [43]. The effects of using the FEI driven AMR algorithm and comparing it to the WEI driven AMR algorithm will be investigated in subsequent verification benchmark test cases.

7.5. Reed Cell Radiation Shielding Verification Benchmark Test Case

In this section comparisons between FEI and WEI will begin to be made. The Reed cell problem is a 1D fixed source nuclear radiation shielding verification benchmark test case that contains a void region [69]. The geometry for the problem is displayed in figure 17 and the material properties are prescribed in Table 4. The problem is solved using the WLS equation spatially discretised with IGA and angularly discretised with a S_8 discrete ordinate (S_N) method. A value of $w_{\text{max}} = 100$ is chosen to solve the WLS equation. This choice of w_{max} is justified based on the qualitative appearance of the scalar neutron flux. The scalar neutron flux profiles for the first-order solution, analytical solution, and several values of w_{max} are presented in Figure 18. It can be seen that the higher the value of w_{max} the more qualitatively accurate the solution, particularly at material interfaces. However, there is still a quantitative difference between the



Figure 17: Geometry and boundary conditions for the one dimensional Reed cell 1D radiation shielding verification benchmark test case.

| Region | Macroscopic neutron cross-section data (cm ⁻¹) | | Extraneous source data (neutrons.cm ⁻³ .s ⁻¹) |
|--------|--|------------|--|
| | σ_t | σ_s | S |
| I | 50.0 | 0.0 | 50.0 |
| II | 5.0 | 0.0 | 0.0 |
| III | 0.0 | 0.0 | 0.0 |
| IV | 1.0 | 0.9 | 1.0 |
| V | 1.0 | 0.9 | 0.0 |

Table 4: Macroscopic neutron cross-section data for the Reed cell 1D radiation shielding verification benchmark test case.

$w_{\max} = 100$ and the analytical solution. The S_8 analytical solution of the Reed cell problem is available in the following reference [70].

Two AMR algorithms have been applied to the 1D Reed cell problem to illustrate the potential of local refinement when applied to neutron transport problems that contain voids. The FEI and the WEI have been used to determine the error indicators for the AMR algorithm in each case. The physical adjoint of the WLS equation has been solved in order to calculate the WEI in each patch. During each AMR iteration one NURBS patch is selected for refinement. Each patch represents a material, so there are five NURBS patches. The L^2 error in the scalar neutron flux for the locally refined cases and uniformly refined case are displayed in Figure 19.

Figure 19a shows the L^2 error in the scalar neutron flux over the entire domain for various levels of spatial refinement. It can be seen that both of the AMR schemes are worse than the uniformly refined solution for certain numbers of dof. This is suspected to be caused by the interaction of the void region with the constraint based local refinement. During the AMR steps the void region stays relatively unrefined, whilst its neighbours are highly refined. The influence of the coarsely refined void region on the highly refined non-void regions, combined with the constraints applied on material boundaries, may be what causes the AMR iterations to have a higher L^2 error per dof. Neither the FEI nor the WEI have been tuned to reduce the functional $\int_D |\phi_{\text{ref}}(x) - \phi(x)|^2 dx$ where D indicates integration over the whole domain $[0, 8]$. This integral represents the global L^2 error plotted in Figure 19a. The heuristic FEI scheme refines based on minimising the global L^2 error in the solution with respect to the H^1 norm whilst the goal for the WEI scheme has been set to minimise the error in the quantity $\int_{\text{III}} \phi(x) dx$ with respect to the H^1 norm.

The values of the L^2 error in the void region ($\int_{\text{III}} |\phi_{\text{ref}}(x) - \phi(x)|^2 dx$) have been plotted in Figure 19b. It can be seen that the WEI and FEI are more accurate per dof than the uniformly





Figure 18: Plot of the scalar neutron flux for the Reed cell 1D radiation shielding verification benchmark test case from solving the WLS equation with a S_8 discrete ordinate (S_N) quadrature. All problems have been uniformly refined. A larger value of w_{\max} results in a more qualitatively accurate solution.


refined solution in almost all cases. As the mesh becomes highly refined the error in the FEI solution begins to plateau. This is caused by the aforementioned fact that the FEI solution is refining to reduce the global H^1 error. This means that the L^2 error in any one NURBS patch is not necessarily minimised as the solutions in neighbouring NURBS patches affect one another. Despite the fact that the goal for the WEI solution is the integral quantity $\int_{\text{III}} \phi(x) dx$ the WEI solution still performs reasonably well as the mesh becomes highly refined and does not experience a plateau in the L^2 error.

Figure 20 shows the error in the integral of the scalar neutron flux. Figure 20b shows the error in $\int_{\text{III}} \phi(x)^2 dx$ and it can be seen that WEI outperforms FEI and uniform refinement. This is expected as this is precisely the goal of the WEI methodology. It can also be seen from Figure 18 that the $w_{\max} = 100$ solution is very similar to the analytical solution for $3 \leq x \leq 5$. Figure 20a shows the error in $\int_D \phi(x) dx$. The behaviour of the uniformly refined solution implies that the integral quantity is not converging to the analytical value. From Figure 18 it can be seen that the WLS solution over and under estimates the analytical solution in various regions of the computational domain. These errors do not cancel out and result in the convergence of the integral of the scalar neutron flux to the wrong value and a constant error level. The convergence of the integral of the scalar neutron flux to the wrong value means that not much can be concluded from the results of the FEI and WEI AMR algorithms in this case.

Further to the results in section 7.3, Figure 21 displays the condition number (κ) of the global stiffness matrix. The first point of note is that in the case when region 3 (void) is unrefined κ is at least half an order of magnitude smaller than the uniformly refined case and up to one order of magnitude smaller than when region 2 is unrefined. Refinement causes an increase in condition number [71] and in cases where low values of σ_t have been used to fill void or near void regions an increase in solution time, and therefore condition number, can be observed [72]. Therefore,




(a) L^2 error in the scalar neutron flux over the entire domain. That is the L^2 error in $\phi(x)$ for $x \in [0, 8]$.




(b) L^2 error in the scalar neutron flux in the void region. That is the L^2 error in $\phi(x)$ for $x \in [3, 5]$.

Figure 19: L^2 error in the scalar neutron flux for several refinement algorithms for the Reed cell 1D radiation shielding verification benchmark test case. Two AMR algorithms are used, the FEI (heuristic) AMR algorithm and the WEI (goal-based)AMD algorithm. The goal for the WEI has been set as $\int_{\text{III}} \phi(x) dx$. The L^2 error is calculated using an analytical S_8 solution.



(a) L^2 error in the scalar neutron flux over the entire domain. That is the L^2 error in $\phi(x)$ for $x \in [0, 8]$.



(b) L^2 error in the scalar neutron flux in the void region. That is the L^2 error in $\phi(x)$ for $x \in [3, 5]$.

Figure 20: L^2 error in the scalar neutron flux for several refinement algorithms for the Reed cell 1D radiation shielding verification benchmark test case. Two AMR algorithms are used, the FEI (heuristic) AMR algorithm and the WEI (goal-based) AMR algorithm. The goal for the WEI AMR algorithm has been set as $\int_{\text{III}} \phi(x) dx$.

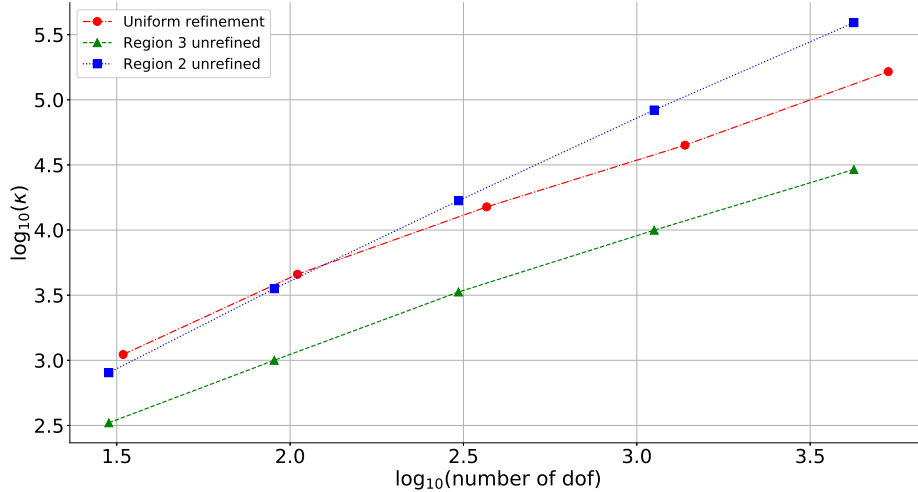


Figure 21: Condition number of the global stiffness matrix for various refinement cases for the Reed cell 1D radiation shielding verification benchmark test case.

it seems that refinement of the void region even in the case where $\mathcal{W} = w_{\max}$ causes a large increase in condition number of the resulting global matrix system. This is further supported by the results in section 7.3.

In the case when region 2 is unrefined, κ is larger than the uniformly refined case when the number of dof is greater than 100. This implies that κ is affected by the difference in refinement between neighbouring NURBS patches. That is, the more dof that have to be constrained between NURBS patches, the greater the increase in κ . However, in the case where region 3 is unrefined, the increase in κ by having to apply constraints between neighbouring NURBS patches is much less than the decrease in κ by not refining the region 3. In other words, it seems that leaving void regions unrefined leads to smaller values of κ for this problem.

Whilst an increase in condition number may seem minor, in neutron transport problems the global matrix system might be required to be solved with thousands of different source vectors. This is due to the necessity of solving for different angles, within different energy groups, converging scattering sources, and performing power iterations. Therefore, it is of great importance to minimise the condition number of the linear system of equations wherever possible.

7.6. Cartesian Geometry ANL Quarter Core Nuclear Reactor Physics Verification Benchmark Test Case

A one group 2D nuclear reactor physics verification benchmark test case has been selected to compare the effectiveness of the FEI and WEI AMR algorithms for a mono-energetic void problem. The geometry and macroscopic neutron cross-section data are presented in Figure 22 and Table 5 respectively. The problem is taken from a paper by D. Lathouwers [29] and is based on the “Multi-dimensional (x-y-z) LWR Model” benchmark from Argonne National Laboratory (ANL) [73].

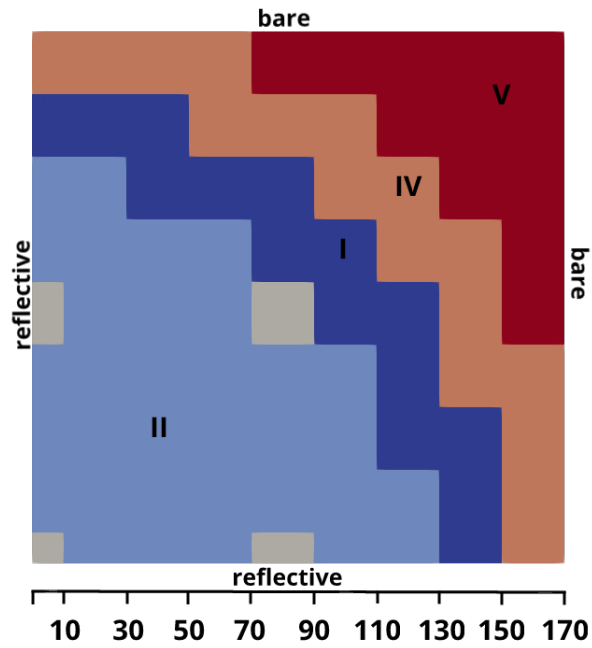


Figure 22: Geometry for the 2D ANL quarter core nuclear reactor physics verification benchmark test case. Quadrilateral areas coloured grey correspond to region III.

| Region | Macroscopic neutron cross-section data (cm^{-1}) | | |
|--------|---|------------|---------------|
| | σ_t | σ_s | $\nu\sigma_f$ |
| 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 | 0.0 |
| V | 0.0 | 0.0 | 0.0 |

Table 5: Macroscopic neutron cross-section data for the 2D ANL quarter core nuclear reactor physics verification benchmark test case.

The WLS equation with $w_{\max} = 1$ has been solved for this problem and is spatially discretised with linear NURBS and angularly discretised with S_8 LCT discrete ordinate (S_N) angular quadrature scheme. A reference solution has been taken from A. Owens Thesis [74]. Results are presented for the same refinement algorithms as in Section 7.5 that is: uniform refinement, the FEI AMR algorithm, and the WEI AMR algorithm. The goal for the WEI is the K_{eff} . Both AMR algorithms use the list method where 10% of the NURBS patches that make up the computational domain are indicated for refinement. The error in the K_{eff} for four cases is displayed in Figure 23. There are two cases for uniform refinement, one solving the WLS equation over the computational domain shown in Figure 22, and the other solving the SAAF equation over the same domain but with region V (the void region) not modelled. This means that the effect of the weighting operator on the solution and its convergence rate can be investigated. As $w_{\max} = 1$ the weighting operator will be active in regions I, II, III, and V.

It can be seen that the modelling of the void region and solving the WLS equation does not affect the convergence rate of the K_{eff} when compared against the K_{eff} from the SAAF solution. However, the case where the void region is not modelled and the SAAF equation is solved produces a slightly smaller error in the K_{eff} . The disadvantage of not modelling the external void region is that a QoI inside the external void region can not be calculated, for example, if the radiation dose in the void regions was desired, it would be unavailable in the case where the void is not modelled.

Figure 23 also shows that the error in K_{eff} is almost identical in the FEI and WEI cases. This is not unexpected as the FEI AMR algorithm is trying to reduce the error in the scalar neutron flux in the H^1 norm. Since the K_{eff} is a global QoI that depends on the accuracy in the scalar neutron flux between power iterations, it makes sense that the FEI AMR algorithm would accurately calculate the K_{eff} . The goal for the WEI AMR algorithm was set as the K_{eff} . However, it requires roughly twice the computational effort as the FEI AMR algorithm as it solves the forward and adjoint WLS equation. The meshes for the third AMR iteration for the FEI AMR algorithm and the WEI AMR algorithm are presented in Figures 24 and 25 respectively. It can be seen that the two meshes are incredibly similar. This result implies that for global QoI, such as the K_{eff} or the integral of the flux over the domain, that the FEI AMR algorithm is computationally more efficient than the WEI AMR algorithm. This is because it produces approximately the same level of accuracy for half of the required computational effort.

7.7. Lathouwers' Radiation Shielding Verification Benchmark Test Case I

The following verification benchmark test case is a radiation shielding problem from work by Lathouwers [30] and is referred to as radiation shielding verification benchmark test case I. The geometry of the computational domain is displayed in Figure 26. Region I contains a unit isotropic extraneous (fixed) source of neutrons and the QoI is a volumetric detector response in region III. The macroscopic neutron cross-section data for the problem is given in Table 6. Two AMR strategies will be compared, one driven by the FEI η^F , and one driven by the WEI η^W . In order to specify the goal for the WEI AMR scheme as the volumetric detector response in region III σ_d has been specified in Table 6. This is used to defined the adjoint source in expression (52). The SAAF equation, used for all calculations here, has been angularly discretised with an S_4 discrete ordinate (S_N) quadrature set to match the calculation performed by Lathouwers [30].

Figure 27 shows the error in the FEI and WEI AMR solutions for two quantities. Both the WEI and FEI AMR algorithms had NURBS patches selected for refinement using the error threshold method defined in expression (27). Figure 27a shows the error in the quantity $\int_D \phi(\mathbf{r}) d\mathbf{r}$, the integral of the scalar flux over the entire domain. It can be seen that both the FEI and the WEI




Figure 23: Error in the K_{eff} for several refinement schemes for the 2D ANL quarter core nuclear reactor physics verification benchmark test case. The SAAF and WLS equation with $w_{\text{max}} = 1$ has been spatially discretised with linear IGA and angularly discretised with an S_8 discrete ordinate (S_N) angular quadrature scheme. The goal for the WEI AMR algorithm has been set as the K_{eff} . When the SAAF equation has been solved the domain does not contain the void region.




Figure 24: Computational mesh for the 2D ANL quarter core verification benchmark test case problem for the third iteration of the FEI AMR algorithm.




Figure 25: Computational mesh for the 2D ANL quarter core verification benchmark test case problem for the third iteration of the WEI AMR algorithm. The goal functional has been set as the K_{eff} .

AMR algorithms lead to an improvement in the quantity $\int_D \phi(\mathbf{r}) d\mathbf{r}$ with respect to accuracy per dof. Both AMR solutions require about five times fewer dof to reach the same error tolerance as the uniformly refined solution. The WEI AMR error is also slightly less than the FEI AMR error for large numbers of dof. This is to be expected as the WEI requires the solution of both a primal and an adjoint neutron transport equation so more computational effort is required to get a more accurate solution.

Figure 27b shows the error in the quantity $\int_{III} \phi(\mathbf{r}) d\mathbf{r}$. It is clear that the FEI AMR algorithm led to a mesh that was unable to calculate the QoI accurately. The mesh generated from the WEI strategy is adequate as the error in the QoI is one order of magnitude less than the FEI AMR error. Indeed, the FEI AMR error is greater than the uniformly refined solution. This implies that the FEI AMR mesh failed to sufficiently resolve area III. The FEI AMR mesh can be seen in Figure 28. It can be seen that region III is relatively unrefined. The WEI AMR mesh is shown in Figure 29. The extra computation effort of solving the adjoint equation can be seen as the interface between region II and III has been significantly refined and this is reflected in the accuracy of the local QoI.

In order to check that the threshold method was not biased towards the WEI AMR algorithm the same calculation was run but this time the list method was used to select NURBS patches for refinement. NURBS patches that were within the top 10% of the list were indicated for refinement and Figure 30 displays the results. Similar trends to the ones in the threshold case can be observed. One unexpected behaviour in the FEI error is that at one refinement step the error in $\int_D \phi(\mathbf{r}) d\mathbf{r}$ increases whilst the error in $\int_{III} \phi(\mathbf{r}) d\mathbf{r}$ decreases by roughly the same amount. This unpredictable behaviour is undesirable in any AMR algorithm. The error of the WEI AMR algorithm solution is monotonically decreasing in Figures 27 and 30, implying that the WEI AMR algorithm is more stable than the FEI AMR algorithm for this particular verification benchmark test case.





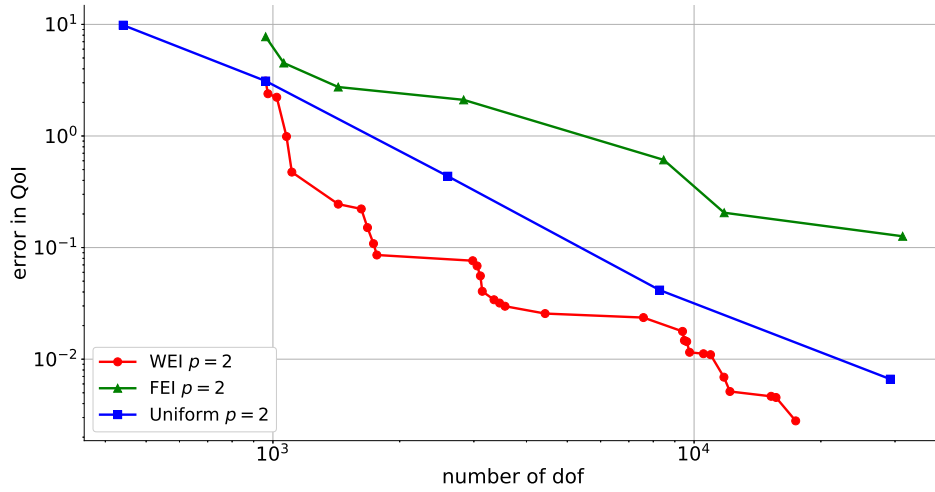
Figure 26: Geometry specification for Lathouwers' radiation shielding verification benchmark test case I.

| Region | Macroscopic neutron cross-section data (cm ⁻¹) | | | Extraneous source data (neutrons.cm ⁻³ .s ⁻¹) |
|--------|--|------------|------------|--|
| | σ_t | σ_s | σ_d | S |
| I | 1.0 | 0.5 | 0 | 1 |
| II | 0.1 | 0.01 | 0 | 0 |
| III | 0.3 | 0.1 | 1 | 0 |
| IV | 0.3 | 0.1 | 0 | 0 |

Table 6: Macroscopic neutron cross-section data for Lathouwers' radiation shielding verification benchmark test case I.



(a) Error in the quantity $\int_D \phi(\mathbf{r}) d\mathbf{r}$



(b) Error in the quantity $\int_{III} \phi(\mathbf{r}) d\mathbf{r}$

Figure 27: Error in the scalar neutron flux integral quantities for Lathouwers' radiation shielding verification benchmark test case I. The WEI η^W and the FEI η^F are used to drive the AMR. The decision to refine a NURBS patch is based on the size of its error indicator relative to the largest error indicator. The total number of refinements is not necessarily equal for FEI and WEI AMR algorithms.




Figure 28: Computational mesh for Lathouwers' radiation shielding verification benchmark test case I produced from the FEI AMR algorithm with 8469 control points.






Figure 29: Computational mesh for Lathouwers' radiation shielding verification benchmark test case I produced from the WEI AMR algorithm with 7581 control points.



(a) Error in the quantity $\int_D \phi(\mathbf{r}) d\mathbf{r}$



(b) Error in the quantity $\int_{III} \phi(\mathbf{r}) d\mathbf{r}$

Figure 30: Error in the scalar neutron flux integral quantities for Lathouwers' radiation shielding verification benchmark test case I. The WEI η^W and the FEI η^F are used to drive the AMR. 10% of the NURBS patches that form the computational domain are indicated for refinement in each AMR iteration. Therefore, the total number of refinements is equal for FEI and WEI AMR algorithms. This does not imply that the number of control points is the same.

8. Conclusions

In this paper a constraint based method for the local adaptive mesh refinement (AMR) of NURBS based isogeometric analysis (IGA) has been described and applied to the self-adjoint angular flux (SAAF) and weighted least squares (WLS) neutron transport equations. Two error indicators have been derived, a heuristic error indicator (FEI), and a dual-weighted residual (DWR) or goal-based error indicator referred to as the weighted error indicator (WEI). In addition, the physical adjoint of the SAAF and WLS equations have been derived so that WEI AMR algorithm can be applied to the SAAF and WLS equations.

The computational efficacy and accuracy of the constraint based local refinement algorithm has been investigated using the method of manufactured solutions (MMS) and a nuclear fuel pincell lattice reactor physics verification benchmark test case. It was shown that the errors introduced by the constraint based local refinement algorithm affect neither the convergence rates, nor the error in the solution significantly for the MMS verification benchmark test case. It can be observed that as the problem is further refined the rate of convergence of the locally refined solution approaches the rate of convergence of the uniformly refined solution. In the nuclear fuel pincell lattice reactor physics verification benchmark test case the magnitude of the error introduced by the constrained degrees of freedom, as well as the under resolution of certain parts of the domain, was demonstrated. It was noted that the propagation of this error depends upon the mean free path of neutrons within the computational domain.

The FEI AMR (heuristic) and WEI AMR (goal-based) were applied to several nuclear reactor physics and radiation shielding verification benchmark test cases. Through these verification benchmark test cases, it was demonstrated that both error indicators performed well for global quantities of interest (QoI). However, if a local QoI is considered, then WEI AMR algorithm, with a suitable goal functional, was more computationally efficient and accurate than the FEI AMR algorithm.

Further research is required to find a way to remove the requirement that the coarsest geometric representation must form a compatible discretisation for NURBS based IGA. In other words, that neighbouring NURBS patches have identical mappings and parameterisations when the NURBS geometry is initially generated. This would simplify the initial generate of geometries and could allow for easier processing of 'dirty' CAD geometries [11]. It could also allow for fewer areas of C^0 continuity between patches. High continuity of the basis functions has been shown to increase the accuracy per degree of freedom for smooth solutions in NURBS based IGA. Furthermore, whilst the application of the locally refinable NURBS based IGA discretisation scheme described in this paper to two-dimensional (2D) problems is of interest, the neutron transport problems that are of real interest are three-dimensional (3D) problems. The effective solution of fully heterogeneous multi-group 3D nuclear reactor physics problems is still an open problem in computational neutron transport. The ability of IGA to exactly represent volume, for fissile mass preservation, and surface area, to accurately represent neutron leakage effects, combined with the ability to only refine where necessary for a given QoI, may reduce the problem sizes enough to allow them to be tractable with more modest computing resources.

Finally, development of an error indicator that takes into account the effect of the neutron mean free path, as well as other material properties, when deciding where the system should be refined is desirable. It was demonstrated that refining certain regions can cause large increases in the spectral condition number of the system. These increases lead to an increase in solution time of the system of linear equations and could be avoided without sacrificing the accuracy if certain local QoI are of interest.

Acknowledgements

Dr C. Latimer would like to acknowledge the Engineering and Physical Sciences Research Council (EPSRC) through the Doctoral Training Partnership (DTP) PhD scheme (EPSRC Grant No.: EP/M507878/1). Dr C. Latimer also acknowledges the industrial and financial support of Rolls-Royce. Dr C. Latimer and Dr M.D. Eaton would like to acknowledge EPSRC for their support of the impact acceleration knowledge transfer secondment (KTS) grant “High Performance, scalable, self-adaptive, exact geometry isogeometric methods, for radiation shielding and nuclear reactor physics analysis of small modular reactors (SMRs) and Naval Nuclear Propulsion” (EPSRC Grant No.: EP/R511547/1). Dr M.D. Eaton and Dr J. Kophazi would like to acknowledge EPSRC for their support through the following grants: Adaptive Hierarchical Radiation Transport Methods to Meet Future Challenges in Reactor Physics (EPSRC Grant No.: 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 No.: EP/K503733/1). The authors would like to thank Mr M. Harvey, who is the research computing services (RCS) manager, and his team at Imperial College London. Finally, the authors would like to thank the anonymous reviewers for their suggestions.

Data Statement

In accordance with EPSRC funding requirements all supporting data used to create figures in this paper may be accessed at the following URL: <https://doi.org/10.5281/zenodo.4108666>.

References

- [1] S. K. Hall, M. D. Eaton, and M. M. R. Williams. The Application of Isogeometric Analysis to the Neutron Diffusion Equation for a Pincell Problem with an Analytic Benchmark. *Annals of Nuclear Energy*, 49:160–169, 2012.
- [2] J. A. Welch, J. Kópházi, A. R. Owens, and M. D. Eaton. Isogeometric analysis for the multigroup neutron diffusion equation with applications in reactor physics. *Annals of Nuclear Energy*, 101:465–480, 2017.
- [3] A. R. Owens, J. A. Welch, J. Kópházi, and M. D. Eaton. Discontinuous isogeometric analysis methods for the first-order form of the neutron transport equation with discrete ordinate (S_N) angular discretisation. *Journal of Computational Physics*, 315:501–535, 2016.
- [4] Y. Chen, S. Lin, O. Faruque, J. Alanoly, M. El-Essawi, and R. Baskaran. Current Status of LS-DYNA Iso-geometric Analysis in Crash Simulation. In *14th International LS-DYNA Users Conference*, pages 1–24, Detroit, MI, USA, 2016.
- [5] I. Temizer, P. Wriggers, and T. J. R. Hughes. Contact treatment in isogeometric analysis with NURBS. *Computer Methods in Applied Mechanical Engineering*, 200:1100–1112, 2011.
- [6] J. A. Cottrell, A. Reali, Y. Bazilevs, and T. J. R. Hughes. Isogeometric analysis of structural vibrations. *Computer Methods in Applied Mechanics and Engineering*, 195:5257–5296, 2006.
- [7] T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs. Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. *Computer Methods in Applied Mechanics and Engineering*, 194:4135–4195, 2005.
- [8] C. Giannelli, B. Juttler, and H. Speleers. THB-splines: The Truncated Basis for Hierarchical Splines. *Computer Aided Geometric Design*, 29:485–498, 2012.
- [9] K. A. Johannessen, T. Kvamsdal, and T. Dokken. Isogeometric analysis using LR B-splines. *Computer Methods in Applied Mechanics and Engineering*, 269:471–514, 2014.
- [10] M. A. Scott, X. Li, T. W. Sederberg, and T. J. R. Hughes. Local refinement of analysis suitable T splines. *Computer Methods in Applied Mechanics and Engineering*, 213-216:206–222, 2012.
- [11] T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs. *Isogeometric Analysis: Toward Integration of CAD and FEA*. John Wiley and Sons, 2009.
- [12] V. P. Nguyen, C. Anitescu, S. P. A. Bordas, and T. Rabczuk. Isogeometric analysis: An overview and computer implementation aspects. *Mathematics and Computers in Simulation*, 117:89–116, 2015.
- [13] A. V. Vuong, C. Giannelli, B. Juttler, and B. Simeon. A hierarchical approach to adaptive local refinement in isogeometric analysis. *Computer Methods in Applied Mechanics and Engineering*, 200:3554–3567, 2011.

- [14] W. C. Rheinboldt and C. K. Mesztenyi. On a data structure for adaptive finite element mesh refinements. *ACM Transactions on Mathematical Software*, 6:166–187, 1980.
- [15] W. Bangerth, I. Kim, D. Sheen, and J. Yim. On Hanging Node Constraints for Nonconforming Finite Elements using the Douglas-Santos-Sheen-Ye Element as an Example. *SIAM Journal on Numerical Analysis*, 55(4):1719–1739, 2017.
- [16] R. Verfürth. *A Posteriori Error Estimation Techniques for Finite Element Methods*. Oxford Science Publications, 2013.
- [17] M. Ainsworth and J. T. Oden. *A Posteriori Error Estimation in Finite Element Analysis*. John Wiley & Sons, 2000.
- [18] T. Gratsch and K. Bathe. A posteriori error estimation techniques in practical finite element analysis. *Computers and Structures*, 83:235–265, 2005.
- [19] W. Bangerth and R. Rannacher. *Adaptive Finite Element Methods for Differential Equations, Lectures in Mathematics*. Birkhäuser Verlag, 2003.
- [20] N. Tsoulfanidis. *The Nuclear Fuel Cycle*. American Nuclear Society Inc., La Grange Park, Illinois, USA, 2013.
- [21] R. E. Faw and J. K. Shultis. *Radiological Assessment: Sources and Doses*. American Nuclear Society Inc., La Grange Park, Illinois, USA, 1999.
- [22] G. S. Was. *Fundamentals of Radiation Materials Science: Metals and Alloys*. Springer-Verlag, 2nd edition, 2017.
- [23] J. Welch, J. Kópházi, A. R. Owens, and M. D. Eaton. A geometry preserving, conservative, mesh-to-mesh isogeometric interpolation algorithm for spatial adaptivity of the multigroup, second-order even-parity form of the neutron transport equation. *Journal of Computational Physics*, 347:129–146, August 2016.
- [24] A. J. Herrema, J. Kiendl, and M. Hsu. A framework for isogeometric-analysis-based optimization of wind turbine blade structures. *Wind Energy*, 22:153–170, 2019.
- [25] G. Lorenzo, M. A. Scott, K. Tew, T. J. R. Hughes, Y. J. Zhang, L. Liu, G. Vilanova, and H. Gomez. Tissue-scale, personalized modeling and simulation of prostate cancer growth. *Proceedings of the National Academy of Sciences of the United States of America*, pages E7663–E7671, 2016.
- [26] L. Piegl and W. Tiller. *The NURBS Book 2nd edition*. Springer-Verlag, 1997.
- [27] G. Beer and S. Bordas, editors. *Isogeometric Methods for Numerical Simulation*. Springer-Verlag, 2015.
- [28] C. Latimer, J. Kópházi, M. D. Eaton, and R. G. McClarren. A geometry conforming isogeometric method for the self-adjoint angular flux (SAAF) form of the neutron transport equation with a discrete ordinate S_N angular discretisation. *Annals of Nuclear Energy*, 136:107049, 2020.
- [29] D. Lathouwers. Spatially adaptive eigenvalue estimation for the S_N equations on unstructured triangular meshes. *Annals of Nuclear Energy*, 38:1867–1876, 2011.
- [30] D. Lathouwers. Goal-oriented spatial adaptivity for the S_N equations on unstructured triangular meshes. *Annals of Nuclear Energy*, 38:1373–1381, 2011.
- [31] J. S. Warsa and A. K. Prinja. p -adaptive numerical methods for particle transport. *Transport Theory and Statistical Physics*, 28(3):229–270, 1999.
- [32] Y. Wang. HP-mesh adaptation for 1-D multigroup neutron diffusion problems. Master’s thesis, Texas A&M, 2006.
- [33] L. Demkowicz. *Computing with hp-adaptive finite elements - volume 1: one and two dimensional elliptic and Maxwell problems*. Chapman and Hall, 2007.
- [34] Y. Wang and J. Ragusa. Application of hp Adaptivity to the Multigroup Diffusion Equations. *Nuclear Science and Engineering*, 161:22–48, 2009.
- [35] N. Nahavandi, A. Minuchehr, A. Zolfaghari, and M. Abbasi. Spatially adaptive hp refinement approach for P_N neutron transport equation using spectral element method. *Annals of Nuclear Energy*, 85:1066–1076, 2015.
- [36] B. Guo and I. Babuška. The $h - p$ version of the finite element method: Part 1: The basic approximation results. *Computational Mathematics*, 1:21–41, 1986.
- [37] B. Guo and I. Babuška. The $h - p$ version of the finite element method: Part 2: General results and applications. *Computational Mathematics*, 1:203–220, 1986.
- [38] A. R. Owens, J. A. Welch, J. Kópházi, and 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. *Computer Methods in Applied Mechanics and Engineering*, 318:215–241, 2017.
- [39] J. Lewins. *Importance The Adjoint Function: The physical basis of variational and perturbation theory in transport and diffusion problems*. Pergamon Press, 1965.
- [40] B. D. Ganapol. *Analytical Benchmarks for Nuclear Engineering Applications - Case Studies in Neutron Transport Theory*. Nuclear Energy Agency, 2008.
- [41] Y. Wang. *Adaptive Mesh Refinement Solution Techniques for the Multigroup S_N Transport Equation Using a High-Order Discontinuous Finite Element Method*. PhD thesis, Texas A&M, 2009.
- [42] Y. Wang and J. C. Ragusa. Standard and goal-oriented adaptive mesh refinement applied to radiation transport on 2D unstructured triangular meshes. *Journal of Computational Physics*, 230:763–788, 2011.
- [43] M. Hanus and R. G. McClarren. On the Use of Symmetrized Transport Equation in Goal-Oriented Adaptivity. *Journal of Computational and Theoretical Transport*, 0(0):1–20, 2016.

- [44] Y. Wang, H. Zhang, and R. C. Martineau. Diffusion acceleration schemes for self-adjoint angular flux formulation with a void treatment. *Nuclear Science and Engineering*, 176:201–225, 2014.
- [45] V. M. Laboure, R. G. McClarren, and Y. Wang. Globally Conservative, Hybrid Self-Adjoint Angular Flux and Least-Squares Method Compatible with Voids. *Nuclear Science and Engineering*, 185:294–306, February 2017.
- [46] W. Zheng, R. G. McClarren, and J. E. Morel. An Accurate Globally Conservative Subdomain Discontinuous Least-Squares Scheme for Solving Neutron Transport Problems. *Nuclear Science and Engineering*, 189:259–271, 2018.
- [47] H. R. Hammer, J. E. Morel, and Y. Wang. A Weighted Least-Squares Transport Equation Compatible with Source Iteration and Voids. *Nuclear Science and Engineering*, 193:388–403, 2018.
- [48] C. Latimer, M. D. Eaton, J. Kópházi, and R. G. McClarren. A Geometry Conforming, Isogeometric, Weighted Least Squares (WLS) Method for the Neutron Transport Equation with Discrete Ordinate (S_N) Angular Discretisation. *Progress in Nuclear Energy*, 121:103238, 2020.
- [49] K. Ueki and P. N. Stevens. Variance Reduction Techniques Using Adjoint Monte Carlo Method in Shielding Problem. *Journal of Nuclear Science and Technology*, 16(2):117–131, 1979.
- [50] J. R. Lamarsh. *Introduction to Nuclear Reactor Theory*. American Nuclear Society Inc., La Grange Park, Illinois, USA, 2002.
- [51] J. E. Hicken and D. W. Zingg. Dual consistency and functional accuracy: a finite-difference perspective. *Journal of Computational Physics*, 256:161–182, 2014.
- [52] S. Schunert, Y. Wang, R. Martineau, and M. D. DeHart. A new mathematical adjoint for the modified SAAF- S_N Equations. *Annals of Nuclear Energy*, 75:340–352, 2015.
- [53] Y. Oka, editor. *Nuclear Reactor Design, An Advanced Course in Nuclear Engineering*. Springer-Verlag, 2010.
- [54] Q. Wu, F. Xu, X. Peng, Y. Yu, Q. Li, and K. Wang. Comparison of sensitivity calculation between mathematical and physical adjoint in the 2-D/1-D transport solver KYADJ. *Annals of Nuclear Energy*, 130:357–364, 2019.
- [55] E. Müller. On the non-uniqueness of the nodal mathematical adjoint. *Annals of Nuclear Energy*, 64:333–342, 2014.
- [56] J. E. Morel and J. M. McGhee. A self-adjoint angular flux equation. *Nuclear Science and Engineering*, 132:312–325, 1999.
- [57] E. E. Lewis and W. F. Miller. *Computational Methods of Neutron Transport*. American Nuclear Society, Inc., La Grange Park, Illinois USA, 1993.
- [58] M. Metcalf, J. Reid, and M. Cohen. *Modern Fortran Explained: Incorporating Fortran*. Oxford University Press, 2018.
- [59] A. R. Owens, J. Kópházi, J. A. Welch, and M. D. Eaton. Energy dependent mesh adaptivity of discontinuous isogeometric discrete ordinate methods with dual weighted residual error estimators. *Journal of Computational Physics*, 335:352–386, 2017.
- [60] K. Salari and P. Knupp. Code verification by the method of manufactured solutions. Technical Report SAND2000-1444, Sandia National Laboratory, 2000.
- [61] P. J. Roache. Code verification by the method of manufactured solutions. *Journal of Fluids Engineering*, 124:4–10, 2002.
- [62] Y. K. Cheung, S. H. Lo, and A. Y. T. Leung. *Finite Element Implementation*. Wiley-Blackwell, 1995.
- [63] N. Z. Cho. Benchmark Problem 2B: MOX Fuel-Loaded Small PWR Core. Technical report, Korean Advanced Institute of Science and Technology (KAIST), 2000.
- [64] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, D. A. May, L. Curfman McInnes, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang. *PETSc Users Manual*. Argonne National Laboratory, 2017.
- [65] L. Beirao da Veiga, L. F. Pavarino, S. Scacchi, O. B. Widlund, and S. Zampini. Isogeometric BDDC preconditioners with deluxe scaling. *Journal of Scientific Computation*, 36(3):1118–1139, 2014.
- [66] H. H. Kim, E. Chung, and J. Wang. BDDC and FETI-DP algorithms with adaptive coarse spaces for three-dimensional elliptic problems with oscillatory and high contrast coefficients. *Journal of Computational Physics*, 349:191–214, 2017.
- [67] IAEA. Transport Theory and Advanced Reactor Calculations: final report of the six-year co-ordinated research programme by the international atomic energy agency. Technical Report IAEA-TECDOC-254, IAEA, 1981.
- [68] J. Wood and M. M. R. Williams. Recent Progress in the Application of the Finite Element Method to the Neutron Transport Equation. *Progress in Nuclear Energy*, 14(1):21–40, 1984.
- [69] W. H. Reed. New Difference Schemes for the Neutron Transport Equation. *Nuclear Science and Engineering*, 46:309–314, 1971.
- [70] J. S. Warsa. Analytical S_N solutions in heterogeneous slabs using symbolic algebra computer programs. *Annals of Nuclear Energy*, 29:851–874, 2002.
- [71] K. P. S. Gahalaut, J. K. Kraus, and S. K. Tomar. Multigrid methods for isogeometric discretization. *Computer Methods in Applied Mechanics and Engineering*, 253:413–425, 2013.
- [72] C. R. E. de Oliveira, M. D. Eaton, A. P. Umpleby, and C. C. Pain. Finite Element-Spherical Harmonics Solutions

- of the 3D Kobayashi Benchmarks with Ray-Tracing Void Treatment. *Progress in Nuclear Energy*, 39(2):243–261, 2001.
- [73] ANL. Argonne Code Center: Benchmark Problem Book. Technical Report ANL-7416 Supplement 2, Argonne National Laboratory, 1977.
- [74] Alex Owens. *Discontinuous Isogeometric Analysis Methods for the First Order Form of the Neutron Transport Equation with Discrete Ordinate Angular Discretisation*. PhD thesis, Imperial College London, April 2017.