
Self-Adaptive Isogeometric Discretisations of the Second-Order Forms of the Neutron Transport Equation with Dual Weighted Residual or Goal-Based Error Measures

Academic Supervisors:

Dr Matthew Eaton

Dr Michael Bluck

Industrial Supervisors:

Dr Paul Warner

Professor Alan Copestake

Author:

Charles Latimer

Department of Mechanical Engineering

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in
Nuclear Engineering of Imperial College London and the Diploma of Imperial College London

December 2019

Abstract

In this thesis several second-order forms of the neutron transport equation (NTE) are spatially discretised with isogeometric analysis (IGA). IGA allows for the exact representation of geometries that are produced using computer aided design (CAD) software. Finite element (FE) spatial discretisation methods are incapable of exactly representing all the geometries that are produced by CAD software. Furthermore, the NURBS basis functions allow for high-order continuity within a NURBS patch, whereas FE basis functions are typically C^0 continuous between adjacent FEs. The advantages and disadvantages of NURBS based IGA will be investigated by comparisons to FE based spatial discretisations.

The second-order forms of the NTE investigated in this thesis are: the self-adjoint angular flux (SAAF) equation, the least-squares (LS) equation, and the weighted least-squares (WLS) equation. The discrete ordinate (S_N) method will be used to angularly discretise these equations. A number of verification benchmark problems will be used to determine the numerical accuracy, convergence, and computational efficiency of both the second-order forms of the NTE, and the IGA spatial discretisation method.

One major challenge associated with continuous IGA spatial discretisations is performing local refinement of the IGA discretisation. Therefore, a constraint based local adaptive mesh refinement (AMR) algorithm is developed to overcome this challenge. Both heuristic and dual weighted residual (DWR) or goal-based error measures are developed in order to determine where local refinement needs to be performed. The DWR error measures enable the numerical error in both global (K_{eff}) and local (reaction rate) quantities of interest to be determined rigorously. Once again a number of verification benchmark problems are used to analyse the numerical accuracy, convergence and computational efficiency of the IGA AMR algorithm.

Declaration of Originality

I hereby declare that the work contained in this thesis is my own unless stated otherwise. In such cases the work will be appropriately referenced.

Copyright Declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Acknowledgements

I would like to thank my principal academic supervisor Dr Matthew Eaton and my co-supervisor Dr Mike Bluck for their support and advice throughout the duration of my PhD studies. I would like to thank Dr Jozsef Kópházi and Professor Ryan McClarren for helping improve my understanding of theoretical and computational neutron transport. I would also like to acknowledge the financial support of EPSRC (Engineering and Physical Sciences Research Council) and Rolls-Royce. Finally, I would like to thank my Rolls-Royce industrial supervisors Dr Paul Warner and Professor Alan Copestake for their support and advice during my PhD.

Contents

Abstract	2
1 Introduction	19
1.1 Project Scope	22
1.2 Thesis Structure	23
2 Computational Modelling of Neutron Transport	25
2.1 The First-Order Form of the Neutron Transport Equation (NTE)	26
2.1.1 Solution Methodologies	27
2.2 Integral Transport Methods	28
2.2.1 Collision Probability Method (CPM)	28
2.2.2 Method of Characteristics (MoC)	29
2.3 Angular Discretisation Methods	30
2.3.1 Discrete Ordinate Method (S_N)	30
2.3.2 Spherical Harmonics Method (P_N)	32
2.3.3 Simplified Spherical Harmonics Method (SP_N)	33
2.4 The Discretisation of the Energy Domain	35
2.5 Spatial Discretisation Methods	38
2.5.1 Finite Difference (FD) Method	38
2.5.2 Finite Element (FE) Method	39
2.6 Second-Order Forms of the Neutron Transport Equation (NTE)	41
3 Isogeometric Analysis (IGA)	45
3.1 Introduction & Motivation	45
3.2 Basis Functions	47
3.2.1 Bernstein Polynomials and Bézier Curves	47
3.2.2 B-splines	49
3.2.3 Non-Uniform Rational B-splines (NURBS)	51
3.3 Refinement Strategies	52
3.3.1 Knot Insertion	52
3.3.2 Order Elevation	53
3.3.3 k -refinement	54
3.4 Geometry Creation	54
3.5 Mathematical Spaces and Numerical Quadrature	59

3.5.1	Quadrature	60
3.6	Neutron Diffusion Simulation Results	60
3.6.1	Two Group Bare Nuclear Fuel Pin	60
3.6.2	Mono-Energetic Lattice Calculation	64
3.6.3	ANL Two Dimensional Nuclear Reactor Physics Benchmark	66
3.6.4	OECD/NEA C5G7 Nuclear Reactor Physics Benchmark	68
3.7	Conclusion	71
4	The Self Adjoint Angular Flux Equation (SAAF)	73
4.1	Algebraic Derivation of the Self-Adjoint Angular Flux NTE and its Associated Weak and Variational Forms	75
4.1.1	Algebraic Derivation of the SAAF Form of the NTE	76
4.1.2	The Weak Form of the SAAF NTE	77
4.1.3	Source Iteration Compatible Variational Formulation of the SAAF NTE	77
4.1.4	Equivalence of Algebraic and Variational Derivation of the SAAF Form of the NTE	81
4.2	Spatial and Angular Discretisation of the Weak Form of the SAAF NTE	82
4.2.1	Bubnov-Galerkin IGA Approximation	82
4.2.2	Angular (Ω) Discretisation of the SAAF NTE	83
4.2.3	IGA Spatial Discretisation of the Weak Form of the SAAF NTE	83
4.3	Numerical Results	85
4.3.1	MMS: Gaussian Modulated Tensor Product of Trigonometric Functions	85
4.3.2	The IAEA Swimming Pool Nuclear Reactor Physics Verification Benchmark Test Case	86
4.3.3	Seven-Group, Bare UO_2 Fuel Pin Nuclear Reactor Physics Verification Benchmark Test Case	89
4.3.4	OECD/NEA, Seven-Group, Two-Dimensional (2D) C5G7 Nuclear Reactor Physics Verification Benchmark Test Case	91
4.4	Conclusion	97
5	Constraint Based Locally Refined IGA Applied to the SAAF Equation	99
5.1	Constraint Based Local Refinement	100
5.1.1	The Two Patch Case	100
5.1.2	The Three Patch Case	102
5.2	The Continuous (Physical) Adjoint of the SAAF Equation	105
5.2.1	Derivation of the Continuous (Physical) Adjoint	106
5.3	Adaptive Mesh Refinement (AMR) Methods	107
5.3.1	The Forward Error Indicator (FEI)	108
5.3.2	The Weighted Error Indicator (WEI)	110
5.4	Results	110
5.4.1	Method of Manufactured Solutions Test Case	111
5.4.2	Pincell Lattice Calculation	112
5.4.3	Adaptive Mesh Refinement Applied to the IAEA Swimming Pool Problem	114

5.4.4	Lathouwers Radiation Shielding Problem	117
5.5	Conclusions	122
6	The Weighted Least Squares Equation	123
6.1	Derivation of the Weighted Least Squares (WLS) Equation	125
6.1.1	Choice of the Weighting Operator $\mathcal{W}$	128
6.2	Spatial and Angular Discretisation of the Weak Form of the WLS Neutron Transport Equation	129
6.2.1	Discrete Ordinate (S_N) Discretisation of the WLS Neutron Transport Equation	129
6.2.2	IGA Spatial Discretisation of the Weak Form of the WLS Neutron Transport Equation	131
6.3	Causality in Second-Order Forms of the Neutron Transport Equation	132
6.4	Numerical Results	135
6.4.1	Bare Fuel Pin Extraneous (Fixed) Source Nuclear Reactor Physics Verification Benchmark Test Case	136
6.4.2	Seven Group, C5G7 UOX Pincell Lattice Fission Source Nuclear Reactor Physics Verification Benchmark Test Case	138
6.4.3	BWR Supercell Extraneous (Fixed) Source Nuclear Reactor Physics Verification Benchmark Test Case	140
6.4.4	Dog-Leg Duct Extraneous (Fixed) Source Radiation Shielding Verification Benchmark Test Case	142
6.5	Conclusion	149
7	Constraint Based Locally Refined IGA Applied to the WLS equation	151
7.1	The Physical (Continuous) Adjoint of the WLS Equation	151
7.2	Generation of the Constraint Operator Using Bézier Extraction	153
7.2.1	Methodology	153
7.3	Results	155
7.3.1	Timing Comparison of the Cox DeBoor and Bézier Extraction Methodologies .	155
7.3.2	KAIST Pincell	155
7.3.3	Reed Cell Problem	158
7.3.4	Cartesian ANL Quarter Core Nuclear Reactor Physics Benchmark Test Case .	163
7.4	Conclusion	169
8	Conclusions	171
8.1	Conclusion	171
8.2	Future Work	172
8.2.1	Trigonometric and Exponential Splines	172
8.2.2	Geometric Spatial Multigrid with Energy Dependent Meshes and Adaptivity .	173
8.2.3	Spatial Domain Decomposition on Shared and Distributed Memory Systems .	173
8.2.4	Spherical NURBS for Angular IGA	174
8.2.5	Boundary Element IGA	174
8.2.6	T-splines	175

8.2.7	Multiscale IGA	175
8.2.8	Reduced Order Models (ROM)	175

List of Figures

2.1	A two dimensional pincell in a Cartesian coordinate system can be approximated by a circular pincell in one dimensional polar coordinates due to rotational symmetry. The radius of the circular moderator region is chosen such that the volume of the moderator is preserved. This is known as the Wigner-Seitz approximation.	29
2.2	Examples of the ray effect for a source in a strongly absorbing media. Oscillations in the surrounding media are qualitatively present.	31
2.3	Microscopic neutron fission cross-section of U^{238} . (From: http://www.nndc.bnl.gov/) .	38
2.4	Two FE mesh approximations of a circular region. On the far right is the circular region being approximated. The approximations are formed from triangular elements. A large number of elements are necessary to get a good qualitative representation. However, even in the case where a large number of elements are used, both the surface area and the volume of the approximations can not be exact at the same time.	40
3.1	Bernstein polynomials for $n = 3$. Also corresponds to a B-spline basis generated over the knot vector $\Xi = \{0, 0, 0, 0, 1, 1, 1, 1\}$	48
3.2	A fourth order Bézier curve. The red crosses are the control points and the dotted bounding box shows the convex hull property. Note that at the start and end of the Bézier curve the curve intersects with the control points.	48
3.3	An example of one dimensional B-spline functions with the quadratic open knot vector $\Xi = \{0, 0, 0, 0.25, 0.25, 0.5, 0.75, 1, 1, 1\}$. Notice the C^0 continuity at $\xi = 0.25$ caused by the repeated knot whilst there is C^1 continuity at $\xi = 0.5$ and $\xi = 0.75$	50
3.4	A selection of the two dimensional B-spline basis functions generated over a unit square with linear open knot vectors $\Xi_\xi = \Xi_\eta = \{0, 0, 0.25, 0.5, 0.75, 1, 1\}$	51
3.5	NURBS basis functions for a quadratic open knot vector Ξ before and after insertion of a knot at $\xi = 0.5$	53
3.6	NURBS basis functions for a knot vector Ξ where order elevation has been performed twice. Notice all functions at $\xi = 0.4$ maintain C^0 continuity after order elevation. . .	53
3.7	NURBS basis functions showing that knot insertion and order elevation do not commute.	55
3.8	Mapping from the parametric domain $\hat{V}$ to the physical domain V	57

3.9	Scalar neutron flux plots along the line $y = 2\text{cm}$. Global high-order continuity of basis functions can cause propagation of spurious oscillations in the solution. Notice for the cases C^1, C^2 , and C^3 with 64 elements that there are oscillations near the domain boundary. In the case of 256 elements the flux near the boundary is lower than the C^0 case. The first-order solution is included for qualitative comparison. All other results are from a neutron diffusion solution. The C^0 results are closer to the neutron transport solution near the boundary in all cases.	58
3.10	A five patch pincell. The requirement that patches have compatible, consistent discretisation means that when knots are inserted into one patch, this insertion propagates through the domain. In the second picture a knot has been inserted in each parametric dimension of the southern patch. In the third picture two knots have been inserted in each parametric dimension in the western patch, and one along the horizontal axis of the southern patch.	59
3.11	Error in the K_{eff} for the bare nuclear fuel pin problem with radius $r = 1$ and zero flux boundary conditions prescribed on all external boundaries.	61
3.12	Error in the K_{eff} for the bare pin problem with radius $r = 1$ and vacuum boundary conditions prescribed on all external boundaries.	63
3.13	Condition number of the bare nuclear fuel pin plotted against number of dof and number of elements. It can be seen that κ is lower for the IGA system than the FE system when plotted against number of elements.	65
3.14	Pincell geometry for the mono-energetic lattice calculation. Reflective boundary conditions are prescribed on all external boundaries.	65
3.15	Error in the disadvantage factor for the mono-energetic lattice calculation.	66
3.16	Geometry for the ANL two dimensional nuclear reactor physics benchmark testcase. The grey square regions correspond to material III.	67
3.17	Error in the K_{eff} for the 2D IAEA quarter core nuclear reactor physics benchmark verification test case. The NDE has been solved for this problem. It can be seen that IGA spatial discretisation is uniformly better than FE. This increase in accuracy per dof is due to the higher inter-element continuity of the IGA basis functions within a NURBS patch.	68
3.18	Layout of dof within a C5G7 fuel pin for the FE code CRONOS2 and the IGA code ICARUS. Both codes use quadratic basis functions. CRONOS2 is using isoparametric parabolic triangular elements. Notice CRONOS2 has higher spatial fidelity than the ICARUS geometry.	70
4.1	Example MMS function, $C = 10, m = 20, n = 10$	85
4.2	Spatial convergence plots of MMS verification test case	86
4.3	Geometry and region numbers for the IAEA swimming pool nuclear reactor physics verification test case. Vacuum boundary conditions are prescribed on all boundaries of the solution domain.	87

4.4	Error in the K_{eff} for the eigenvalue IAEA swimming pool reactor physics benchmark. Both the IGA and FE use quadratic basis functions and a S_{10} angular quadrature set. Both solutions were computed using ICARUS where a triangular Legendre-Chebyshev discrete ordinate angular quadrature has been used. The results are compared to the reference solution provided by an Inferno S_{20} solution.	89
4.5	Error in the volume averaged scalar neutron flux for the extraneous (fixed) source IAEA swimming pool reactor physics benchmark in each region computed by ICARUS using quadratic IGA and quadratic FEs with an S_{10} angular quadrature. The reference solutions were produced using the discontinuous Galerkin IGA S_N code Inferno.	90
4.6	Geometry for the seven-group, bare UO_2 fuel pin eigenvalue (K_{eff}) nuclear reactor physics verification benchmark test case. Vacuum boundary conditions have been prescribed on the boundary of the bare UO_2 fuel pin.	91
4.7	Error in the K_{eff} for the bare fuel pin problem for varying levels of spatial and polynomial refinement within a NURBS patch. All results calculated using ICARUS	92
4.8	Geometry specification for the 2D OECD/NEA C5G7 benchmark.	93
4.9	Scalar neutron flux profiles for the 2D OECD/NEA C5G7 quarter core nuclear reactor physics verification test case for several energy groups. The scalar neutron flux is normalised according to the C5G7 benchmark specification [210].	94
5.1	The geometry for the two patch case. The coarse patch is on the left and its control points are red. The fine 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 $\Xi^C = \{0, 0, 0.5, 1, 1\}$ and $\Xi^F = \{0, 0, 0.25, 0.5, 0.75, 1, 1\}$	100
5.2	The layout of the control points for quadratic basis functions. The coarse patch is on the left and its control points are red. The fine 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 $\Xi^F = \{0, 0, 0, 0.5, 1, 1, 1\}$	102
5.3	The geometry for the three patch case. The coarse patch to the west is referred to as C_w and its control points are red. The coarse patch to the south is referred to as C_s and its control points are yellow. The fine 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, $\Xi^{C_e} = \Xi^{C_s} = \{0, 0, 0.5, 1, 1\}$ and $\Xi^F = \{0, 0, 0.25, 0.5, 0.75, 1, 1\}$ in both parametric directions.	103
5.4	Geometry for the MMS verification test case.	111
5.5	Plot of the error in the MMS verification test case. The colour scale in both plots is the same.	112
5.6	Spatial convergence plots for the MMS verification test case. The rates of convergence for the coarse patches and the uniformly refined case match exactly for all orders of basis function.	113

5.7	Geometry of the pincell row problem. All boundaries have reflective boundary conditions applied to them.	113
5.8	Discretised geometry of the pincell row problem.	114
5.9	Error in ζ for the pincell row problem.	115
5.10	Geometry of the IAEA swimming pool nuclear reactor physics test case.	116
5.11	Error in the K_{eff} calculated from the S_4 -SAAF equation solved for the IAEA swimming pool nuclear reactor physics test case. An AMR scheme driven by the WEI is compared to uniform refinement.	116
5.12	Geometry specification for shielding problem I.	118
5.13	Error flux integral quantities. The WEI η^W and the FEI η^F are used to drive AMR. The decision to refine a 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.	119
5.14	Mesh produced from the FEI strategy with 8469 control points.	120
5.15	Mesh produced from the WEI strategy with 7581 control points.	120
5.16	Error flux integral quantities. The WEI η^W and the FEI η^F are used to drive AMR. 10% of the patches that form the domain are flagged for refinement in each AMR step. Therefore, the total number of refinements is equal for FEI and WEI. This does not imply that the number of control points is the same.	121
6.1	The algorithm for determining a value of w_{max} based upon a solution of the first-order neutron transport equation.	130
6.2	Geometry for the void wall problem.	132
6.3	First-order and WLS solutions of the void wall problem. S_{10} angular quadrature has been used and $w_{\text{max}} = 100.0$ has been chosen.	134
6.4	Geometry for the bare fuel pin extraneous (fixed) source nuclear reactor physics verification benchmark test case. A vacuum boundary condition is prescribed on the outer boundary of the bare fuel pin.	136
6.5	Errors in the integral of the scalar neutron flux over the domain for the bare fuel pin extraneous (fixed) source nuclear reactor physics verification test case. S_{10} angular quadrature has been used.	137
6.6	Geometry for the seven group, C5G7 UOX pincell nuclear reactor physics verification benchmark test case. All of the boundary conditions are prescribed to be reflective.	138
6.7	Error in the K_{eff} for the UOX C5G7 pincell fission source nuclear reactor physics verification test case angularly discretised with S_8 discrete ordinate angular quadrature.	139
6.8	Geometry of the BWR supercell extraneous (fixed) source nuclear reactor physics verification benchmark test case.	140
6.9	Error in the QoI for the BWR supercell extraneous (fixed) source nuclear reactor physics verification benchmark test case discretised using an S_6 LCT angular quadrature set.	141
6.10	Geometry of the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.	143
6.11	Error in various QoI for the two-dimensional S_{14} dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.	144

6.12	Scalar neutron flux profiles for the S_{14} solution of the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.	145
6.13	Error in I_n with $w_{\max} = 16.8$ and an S_{14} angular quadrature set for the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.	146
6.14	Error in various QoI for the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case for values of $w_{\max}$ selected using the algorithm in Figure 6.1. An S_{14} angular quadrature set was used for all calculations.	147
7.1	The elements e and $\bar{e}$ used for the generation of T^f using Bézier extraction. As the fine patch is assumed to be a refinement of the coarse patch both elements are illustrated on the same patch to signify that $\bar{e} \subset e$	154
7.2	Timings for the generation of the T_f matrix run 1000 times. Four different orders of basis function are displayed. For each order the timing of the Cox-DeBoor algorithm (CDB), Bézier Extraction algorithm (BE), and the time taken to invert the Bézier element extraction operator C^e required in the BE algorithm are plotted.	156
7.3	Geometry of the KAIST pincell nuclear reactor physics verification benchmark test case. Reflective boundary conditions are prescribed on all surfaces. The pincell consists of fuel surrounded by a moderator. There is a layer of cladding between the fuel and moderator and a small air gap is present. This air gap is included in real reactors to allow fuel pellets to swell without causing cracking in the cladding.	157
7.4	Element boundaries of the KAIST pincell nuclear reactor physics verification benchmark test case as defined by the knot spans. This is the refinement referred to in the local case. The thin void region is unrefined whilst the neighbouring fuel pin and cladding are highly refined.	158
7.5	Error in the disadvantage factor ζ for two refinement schemes. ‘Uniform’ refines all areas uniformly, and ‘local’ refines everywhere except the void. The WLS equation spatially discretised with quadratic IGA, angularly discretised with S_8 LCT discrete ordinates and with $w_{\max} = 1.0$ was used to generate results.	159
7.6	Geometry and boundary conditions for the one dimensional Reed cell radiation shielding benchmark test case.	160
7.7	Plot of the scalar neutron flux for the reed cell problem with an S_8 discrete ordinate quadrature. All problems have been uniformly refined. A larger value of $w_{\max}$ results in a more qualitatively accurate solution.	161
7.8	L_2 error in the neutron scalar flux for several refinement schemes. Two AMR schemes are used, FEI which is a heuristic based scheme and WEI which is a goal-based error measure. The goal for the WEI has been set as $\int_3^5 \phi(x)dx$. The L_2 error is calculated using an analytical S_8 solution.	162
7.9	L_2 error in the neutron scalar flux for several refinement schemes. Two AMR schemes are used, FEI which is a heuristic based scheme and WEI which is a goal-based error measure. The goal for the WEI has been set as $\int_3^5 \phi(x)dx$	164
7.10	Condition number of the global stiffness matrix for various refinement cases.	165

7.11	Geometry for the IAEA problem. Quadrilateral areas coloured grey correspond to region III	165
7.12	Error in the K_{eff} for several refinement schemes. The WLS equation with $w_{\text{max}} = 1$ has been spatially discretised with linear IGA and angularly discretised with S_8 discrete ordinates. The goal for the WEI method has been set as the K_{eff} . The no void description refers to the ANL problem with the void in region V not modelled.	167
7.13	Mesh for the ANL problem for the third iteration of the FEI driven AMR scheme. . .	167
7.14	Mesh for the ANL problem for the third iteration of the WEI driven AMR scheme. The goal functional has been set as the K_{eff}	168

List of Tables

2.1	Description of dependent variables, independent variables, and neutron macroscopic cross-section data for the first-order neutron transport equation.	27
3.1	Macroscopic neutron cross-sections for the bare pin nuclear reactor physics test case. If a quantity is not specified then it is assumed to be zero.	61
3.2	Spectral condition numbers of the unpreconditioned global stiffness matrix for the bare nuclear fuel pin problem.	64
3.3	Spectral condition numbers of the SSOR preconditioned global stiffness matrix for the bare nuclear fuel pin problem.	64
3.4	Macroscopic neutron cross-sections for the mono-energetic lattice nuclear reactor physics benchmark verification test case. Superscripts F and M denote quantities in the fuel pin and moderator regions respectively	65
3.5	Macroscopic neutron cross-sections for the two dimensional ANL nuclear reactor physics benchmark verification test case. Group 1 represents fast neutrons and group 2 represents thermal neutrons.	67
3.6	Various QoI for several C5G7 solves. The results generated by PIRANA are taken from this paper [175] and the results generated by CRONOS2 are taken from the C5G7 benchmark specification [210]	69
4.1	K_{eff} for various discrete ordinate (S_N) orders as calculated by Inferno for the IAEA swimming pool nuclear reactor physics benchmark.	87
4.2	One-group macroscopic neutron cross-section data for the eigenvalue (K_{eff}) IAEA swimming pool nuclear reactor physics verification test case.	88
4.3	Extraneous (fixed) neutron source strengths for the fixed source IAEA swimming pool nuclear reactor physics verification test case.	88
4.4	Reference volume averaged scalar neutron flux solutions for the extraneous (fixed) source IAEA swimming pool nuclear reactor physics verification test case. These solutions were generated using the Inferno code.	88
4.5	Error in CRONOS2 solution compared to the Inferno C5G7 reference result. The K_{eff} is converged to 10^{-6} . The particular S_8 angular quadrature set is unknown.	95
4.6	Error in ICARUS solution compared to Inferno C5G7 reference result. The K_{eff} has been converged to 10^{-6} and a triangular Legendre-Chebyshev S_8 angular quadrature set has been used.	95
4.7	Error in ICARUS solution compared to Inferno C5G7 reference result. The K_{eff} has been converged to 10^{-6} and a level-symmetric S_8 angular quadrature set has been used.	95

4.8	Error in ICARUS solution compared to the Monte-Carlo C5G7 reference solution found in the benchmark specification [210]. The K_{eff} has been converged to 10^{-7} and a S_8 triangular Legendre-Chebyshev angular quadrature set has been used.	95
5.1	Macroscopic neutron cross-section data for the pincell row test cases. A F superscript refers to the fuel and a M superscript refers to the moderator.	114
5.2	One-group macroscopic neutron cross-section data for the eigenvalue (k_{eff}) IAEA swimming pool nuclear reactor physics verification test case.	115
5.3	Macroscopic neutron cross-section data for shielding problem I.	117
6.1	Macroscopic neutron cross-sections and source strengths for the void wall extraneous (fixed) source radiation shielding test case.	133
6.2	Macroscopic neutron cross-sections and source strengths for the bare fuel pin extraneous (fixed) source nuclear reactor physics verification test case.	136
6.3	Macroscopic neutron cross-sections and source strengths for the BWR supercell extraneous (fixed) source nuclear reactor physics verification test case.	140
6.4	Macroscopic neutron cross-sections and source strengths for the dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.	142
6.5	Values of w_{max} produced by the algorithm in Figure 6.1 for different coarse meshes applied to the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.	146
7.1	Condition number κ for the KAIST pincell. The condition number was converged to 1×10^{-9} and the system was preconditioned with a SSOR scheme. If it took more than 1000 GMRES iterations then the value at the 1000 th iteration is taken and is coloured red.	159
7.2	Macroscopic neutron cross-section data for the one dimensional Reed cell radiation shielding benchmark test case.	160
7.3	Macroscopic neutron cross-section data for the two dimensional ANL nuclear reactor physics benchmark test case.	163

Chapter 1

Introduction

Around the world there is growing interest in the use of low-carbon, renewable power generation. This is being driven largely by international climate change commitments, such as the 1997 Kyoto protocol [1], and the 2016 Paris agreement [2], which is leading to decommissioning of old fossil fuelled power plants and the introduction of low carbon, renewable, power generation technology such as solar, wind, and tidal power. The UK has passed into law the 2008 climate change act which provides legally binding requirements to reduce carbon dioxide (CO₂) emissions by 80% by 2050 compared to base-line 1990 levels [3]. These challenging commitments are requiring the UK to invest heavily in both onshore and offshore wind power technology as well as tidal and solar. However, one issue associated with these types of power generation technology is their intermittent nature, which means that they cannot be used to provide a constant “base-load” of electrical power to the national grid.

This has forced the UK to invest in gas-fired power stations to provide back-up electrical power generation to account for the intermittent nature of most low-carbon, renewable power generation. Longer term, low-cost, high-power density energy storage is needed but this remains very much at the research level and indeed, may not be feasible. The UK’s fleet of nuclear power stations comprises 15 operating nuclear power plants: 14 Advanced Gas Cooled Reactors (AGRs) and 1 Pressurised Water Reactor (PWR) which provide approximately 20% of the UK’s electrical energy demands. However, the AGRs are all due to decommission by 2030 which will leave a very significant gap in the UK’s total base-load energy requirements [4]. Therefore, the UK government has kept the option of developing new nuclear power as a means of providing a base-load of low-carbon power generation. The current nuclear new build programme in the UK revolves around conventional generation III and III+ Light Water Reactor (LWR) [5] and Boiling Water Reactor (BWR) [6] designs such as the AREVA Evolutionary Power Reactor (EPR) which is being constructed at the new Hinkley Point C site in Somerset, England, UK [7, 8].

In the defence sector, the UK is renewing its fleet of Submersible Ship Nuclear (SSN) attack submarines, with the latest generation being the Astute class. Seven Astute class submarines are being built and commissioned with the final submarine, Agincourt, being in service around 2024. These are powered by the latest core H design Rolls-Royce PWR2 nuclear steam raising plant (NSRP) [9]. The core H PWR2 removes the need for refuelling allowing a submarine to avoid two reactor refits in its service life [9]. The PWR2 NSRP was originally developed for the Vanguard class of Submersible Ship Ballistic Missile Nuclear (SSBN) submarines. In addition to the Astute class submarines, the UK is renewing its continuous at sea deterrent with a new fleet of four SSBN submarines being designed,

built, and constructed by BAE systems and Rolls-Royce. These four submarines will be called the Dreadnought class and will be powered by a new Rolls-Royce PWR3 NSRP. The PWR3 NSRP is designed to be simpler and safer with a longer service life and lower maintenance requirements than the PWR2 NSRP. The first of class will be HMS Dreadnought and is due in service in the early 2030s [10].

The developments in both the civil and defence nuclear sectors is driving innovations in nuclear reactor technology and design. The primary reason is to reduce costs and improve reliability, performance, and safety. As a consequence the UK is investing in new small modular reactors (SMRs) [11] and generation IV [12] nuclear power plants (NPPs) such as molten salt and high temperature reactors (HTRs). These investments in new nuclear technologies are being funded through the Business Enterprise and Industrial Strategy (BEIS) department of Her Majesty’s government (HMG) [13]. In order to reduce the cost and speed-up the design and regulatory approval of next generation nuclear reactors, HMG is investing in new approaches to modelling and simulation via the BEIS digital reactor design (DRD) programme. The aim is to develop digital twins of next generation nuclear reactors in order to reduce the time needed to design such NPPs, as well as reduce the time needed to gain regulatory approval [14]. A further aim is to reduce the manufacturing time of components as well as the final assembly of the NPP. These are very ambitious aims and will require the development of large-scale, high-fidelity, multiscale, multiphysics modelling and simulation methods [15, 16, 17]. In this regard high-fidelity, exact geometry, nuclear reactor physics and neutron transport simulations may have a role to play in providing a digital twin of a NPP.

However, one of the main issues in developing a digital twin of a NPP is the multiscale nature of most nuclear reactor physics and neutron transport problems which provides a significant computational challenge [18]. On the level of the pincell, there are small, highly detailed features in the design. The nuclear fuel is kept separate from the moderator by a thin layer of cladding only millimetres thick. Between the cladding and the fuel there is a gas gap in order to enable the fuel to expand as it heats up without causing damage to the surrounding clad. This gas gap is a small fraction of 1mm thick. Inside the NPP the nuclear fuel pins are combined together to form nuclear fuel assemblies that are typically 20cm $\times$ 20cm in size. These nuclear fuel assemblies are combined together to form the entire nuclear reactor core which will be on the scale of metres. If a quantity of interest (QoI) is outside of the NPP then the surrounding areas must also be modelled requiring more degrees of freedom in a deterministic code, or more particle histories in a Monte Carlo code. In a submarine the NPP takes up a relatively large part of the overall volume that must be shielded so the increase in the size of computational domain may not be too large. Conversely, in civil NPPs, the NPP may comprise a smaller percentage of the overall area that must be shielded.

Different types of nuclear reactor physics and neutron transport simulations require different aspects of the physics to be accurately resolved. For example, nuclear reactor physics calculations require macroscopic and microscopic neutron cross-sections that capture the effects of spatial and energy self-shielding, Doppler broadening, and the resonance structures of certain nuclei [19]. Macroscopic neutron cross-section data is typically homogenised over some region or volume of the nuclear reactor core such as a nuclear fuel pin or fuel assembly. This nuclear data is also only specified for a coarse energy group structure rather than the pointwise continuous nuclear data that is available and utilised in Monte Carlo neutron transport codes. The homogenised macroscopic neutron cross-section

data for a variety of nuclear fuel temperatures, moderator temperatures, and moderator densities, amongst other variables, are generated using highly refined, fine group, lattice physics calculations [20]. These lattice physics calculations may have hundreds of groups [21, Ch. 9]. The homogenised, coarse group cross-sections are then used in nodal diffusion codes, which may have as few as two energy groups [22]. Whilst nodal diffusion codes provide good estimates to quantities that relate to homogenised regions, retrieving accurate information about certain heterogeneities that have been homogenised can be difficult [23]. One approach to reconstructing the fine scale detail within a nuclear fuel assembly is to utilise flux reconstruction algorithms or de-homogenisation methods. These are types of multiscale algorithms that use the fine scale neutron scalar flux distributions from single nuclear fuel assembly lattice physics calculations (assuming periodic boundary conditions) and combine these with the coarse scale neutron scalar flux distributions from whole core neutron diffusion calculations using nodal methods, in order to generate approximate small scale data [23, 24].

Reactor shielding calculations require a well-resolved gamma-ray photon fluence in order to determine the radiation dose rate both outside and inside the reactor pressure vessel (RPV). Microscopic neutron and gamma-ray photon cross-section libraries for reactor shielding calculations typically have many energy groups. For example, the BUGLE library has 67 energy groups, 47 for neutron energies and 20 for photon energies [25]. A detailed knowledge of the energy-dependent neutron and gamma-ray photon fluence is necessary in order to ensure that the dose received by civilian workers and submarine crew is below the required statutory levels. Furthermore, the distribution of the radiation fluence throughout the NPP is key to accurately predicting the radiation damage to the RPV; as well as any secondary systems and components. Radiation damage is a significant factor in determining the expected service life of components within a NPP. The gamma-ray photon fluence can be determined by utilising a similar equation to the neutron transport equation, only the cross-section data is different [26]. This means the same transport code can be utilised to determine both the neutron and gamma-ray photon fluence. The main difference between the modelling of neutron and gamma-ray photons, aside from the cross-section data, is that fast neutrons are produced by fission events whereas activation gamma-ray photons are produced by the decay of activated isotopes. Therefore, the initial source term and its distribution in space is thus different.

Determining the steady-state or time-dependent solution of the neutron transport equation is only part of the challenge posed in modelling the behaviour of a NPP. The behaviour of the nuclear reactor core is also influenced by many other factors such as: nuclear fuel burn-up/depletion, generation of neutron poisons, fluid dynamics, gamma-ray photon transport, and thermal hydraulic effects. These physical phenomena must all be modelled in order to obtain a detailed understanding of the behaviour of the NPP. Early on in the development of computer technology, all these physical phenomena were modelled by separate computer programs or modules. The numerical results of each computer program were used as an input to the other programs or modules to produce a calculational chain of results that could model the behaviour of the NPP. It is important to note that, due to computational constraints and limitations in the early days of computer development, this required significant simplifications to the mathematical and computational models assumptions. As discussed previously, a move towards the use of high-fidelity, integrated, multiscale and multiphysics software (IMMS) has occurred [15, 16]. The use of IMMS has removed many of the pessimisms and approximations that are prevalent in the current generation of commercial nuclear reactor analysis software. One important numerical approximation

that is used in commercial nuclear reactor analysis software, is the inexact representation of the geometry of the problem domain. In reality most engineering design of components and systems is performed using computer aided design (CAD) software where the geometry is described using non-uniform rational B-spline (NURBS) surfaces and most commercial numerical software is unable to represent such geometries exactly.

The use of IGA based spatial discretisation enables the direct use of the underlying CAD geometry description within the engineering analysis process. It is estimated that up to 77% of the analysis time within large-scale engineering projects is associated with geometry description, manipulation of CAD, mesh generation, and mesh optimisation [27]. The use of IGA could significantly reduce the time required to perform engineering design and analysis whilst significantly reducing any errors within this process. Furthermore, in strongly coupled multiphysics simulations, traditional finite element (FE) based approaches require the numerical solution on one mesh to be interpolated onto another mesh. For instance time-dependent computational neutron transport calculations might be performed on a coarser mesh than the computational fluid dynamic (CFD) simulations. The particular refinement of each mesh will be governed by the physics of the particular phenomena being analysed. In neutron transport this might be associated with the gradient in the scalar neutron flux or the mean free path of neutrons. In CFD simulations this would be associated with the turbulent length scales and boundary layers. Mesh-to-mesh interpolation is incredibly complex for overlapping, unstructured, higher-order, FE meshes. In addition, adaptive mesh refinement requires the use of an ancillary adaptive mesh generation algorithm. Moreover, refinement and coarsening of FE meshes often does not preserve the underlying geometry of the computational domain. Exact geometry IGA methods overcome many of these limitations as they do not require ancillary adaptive mesh generation algorithms. They can naturally incorporate higher-order mesh-to-mesh interpolation algorithms between two different IGA refinements and they also naturally preserve the underlying geometry even on the coarsest resolution. The use of IGA analysis could significantly streamline both CAD to computer aided engineering (CAE) analysis as well as large-scale IMMS.

1.1 Project Scope

The aim of this project is to investigate the effect of an exact geometry IGA discretisation on second-order forms of the neutron transport equation (NTE) for both nuclear reactor physics and reactor shielding. The FE method is a widely used numerical discretisation technique that has proven itself to be both flexible and reliable, and is typically used for the spatial discretisation of both first and second-order forms of the NTE. It has been demonstrated that the application of IGA to the first-order and even-parity neutron transport equations is superior to the same systems discretised with FEs [28, 29]. In this project IGA will be applied to the self-adjoint angular flux (SAAF) and weighted least squares (WLS) forms of the neutron transport equation and comparisons with the FE method will be performed. The WLS equation allows for solutions to be generated for problems containing void regions. The advantages in accuracy of the IGA method come from its ability to exactly represent geometry and the high-order continuity of its basis functions. These two properties will be explained in Chapter 3.

The novel parts of this thesis are as follows:

1. A variational derivation of the source iteration compatible self-adjoint angular flux (SAAF)

- equation is presented and discretised with NURBS based IGA (Chapter 4).
2. NURBS based IGA is applied to the weighted least squares (WLS) equation and an algorithm for choosing the weighting operator used in void regions is described (Chapter 6).
 3. Constraint based local refinement is applied to the SAAF and WLS equations discretised with NURBS based IGA (Chapters 5 and 7).
 4. The physical adjoint of the SAAF and WLS equations are derived. These adjoint formulations, discretised with NURBS based IGA, are used to drive goal-based adaptive mesh refinement algorithms (Chapters 5 and 7).

1.2 Thesis Structure

The chapters of this thesis are as follows. Chapter 2 discusses the various formulations of the neutron transport equation and many of the computational techniques used to solve it. Chapter 3 explains isogeometric analysis (IGA) and illustrates its properties by discretising the neutron diffusion equation using IGA. Chapter 4 derives a variational formulation of the source iteration compatible self-adjoint angular flux (SAAF) equation and presents several nuclear reactor physics test cases solved using the SAAF equation discretised with IGA. Comparisons between the IGA-SAAF and the FE-SAAF formulations are made. In Chapter 5 the constraint based local refinement methodology is introduced. A heuristic error indicator and a dual weighted residual (DWR) or goal-based error indicator are introduced and an adaptive mesh refinement (AMR) algorithm is proposed. The physical adjoint of the SAAF equation is derived and is used to develop a DWR or goal-based measure for the constraint based IGA AMR scheme. In Chapter 6 the weighted least squares (WLS) equation is derived and discretised using an IGA based spatial discretisation. This formulation is applied to several nuclear reactor physics and radiation shielding problems that contain void regions. The effect of the strength of the weighting factor upon the solution is discussed. Furthermore, an algorithm for choosing the weighting factor is presented. In Chapter 7 the derivation of the physical adjoint of the WLS equation is presented and is used to develop a DWR or goal-based error measure for an AMR algorithm for the WLS equation. This AMR algorithm is applied to several nuclear reactor physics and radiation shielding problems containing voids. The effect of local refinement on void regions, as well as the efficacy of the AMR algorithm, is discussed. Finally, Chapter 8 proposes some avenues for future work, as well as providing a summary of the work in this thesis.

Chapter 2

Computational Modelling of Neutron Transport

The computational analysis of nuclear reactor physics and reactor shielding problems can, in general, be achieved by solving one of several partial integro-differential equations (PIDEs) that describe the migration and density of neutrons in a prescribed host media. Whilst there are several different PIDEs that can be solved to determine the scalar neutron flux, each one has different mathematical properties that make them amenable to different discretisation methods and algorithms. Together, these PIDEs are referred to as the forms of the neutron transport equation (NTE).

The derivation of each form of the NTE involves a different set of techniques and assumptions that are applied in order to reduce the complexity of the mathematics or underlying physics. Reductions in complexity can lead to easier implementation of methods, as well as to less computational work, and therefore time, being necessary to arrive at the same answer. Each form of the neutron transport equation can then be discretised to produce a matrix system of linear equations that can be solved using matrix solution algorithms. Despite the myriad of different forms, all forms of the NTE are still challenging to solve due to two major factors. First, the phase space describing the neutron population is seven dimensional (7D) and consists of: position $\mathbf{r} = (x, y, z)$, energy E , angle $\boldsymbol{\Omega} = (\theta, \varphi)$, and time t . In order to accurately represent the neutron angular and scalar flux each of these variables must be discretised with suitable fidelity. This leads to a large computational problem that requires the application of preconditioning and iterative matrix solution algorithms so that it can be solved efficiently.

Second, typical problem geometries over which the NTE must be solved are challenging to model accurately and efficiently. These geometries are typically multiscale, ranging from cladding air gaps (sub-mm) to fuel pins (cm) to fuel assemblies (m), and highly heterogeneous, with neighbouring materials having varying macroscopic neutron cross-section data that can depend on space, temperature, angle, and energy of a given neutron. Furthermore, nuclear reactor geometries are generally non-Cartesian, almost all reactors contain cylindrical fuel pins and some contain non-Cartesian fuel assemblies, such as cylindrical and hexagonal nuclear fuel assemblies seen in CANDU and high-temperature gas-cooled reactors (HTGR) respectively [30]. These non-Cartesian components introduce significant geometric errors into the solution unless discretised with a large number of geometric primitives.

2.1 The First-Order Form of the Neutron Transport Equation (NTE)

One of the most commonly used forms of the NTE is the first-order form of the NTE and was first derived by L. S. Ornstein and G. E. Uhlenbeck [31]. This form of the NTE can be derived in many different ways [32, 33]. The conventional approach to deriving the equation is by considering the conservation of a large number of neutrons over an infinitesimal volume of phase space [34]. The first-order NTE is a Boltzmann equation as it describes the average statistical behaviour of a large population of neutrons [35]. The time dependent first-order neutron transport equation is presented in equation (2.1).

$$\begin{aligned} \frac{1}{v(E)} \frac{\partial}{\partial t} \psi(\mathbf{r}, \boldsymbol{\Omega}, E, t) + \boldsymbol{\Omega} \cdot \nabla \psi(\mathbf{r}, \boldsymbol{\Omega}, E, t) + \sigma_t(\mathbf{r}, E, t) \psi(\mathbf{r}, \boldsymbol{\Omega}, E, t) = \\ \int_0^\infty \int_{4\pi} \sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}, E' \rightarrow E, t) \psi(\mathbf{r}, \boldsymbol{\Omega}', E', t) d\boldsymbol{\Omega}' dE' + \\ \frac{\chi(\mathbf{r}, E, t)}{4\pi} \int_0^\infty \nu(\mathbf{r}, E', t) \sigma_f(\mathbf{r}, E', t) \phi(\mathbf{r}, E', t) dE' + Q(\mathbf{r}, \boldsymbol{\Omega}, E, t). \end{aligned} \quad (2.1)$$

The terms appearing in equation (2.1) are explained in Table 2.1. The key assumption made in the derivation of the first-order NTE is the neglect of neutron-neutron interactions. This is justified by the relatively small size of a neutron compared to a target nuclei and the relatively high density of target nuclei in a host material. These two properties make neutron-neutron interactions highly unlikely compared to neutron-nucleus collisions [19]. Therefore, equation (2.1) is a linearised form of the original non-linear Boltzmann transport equation.

Several other assumptions are typically made in order to simplify equation (2.1). First, it is assumed that all homogeneous materials are isotropic. This means that the material properties do not depend on the incoming and outgoing angles of neutrons, only the relative change in angle [36, p. 225]. Second, when determining the steady state solution, if it exists, the material properties are assumed to not change with regard to time (e.g. burn-up/depletion, temperature feedback through Doppler broadening etc). For non-multiplying media in which the prescribed, non-negative, extraneous (or external) source and boundary conditions are time-independent there will always be non-negative solutions to the steady-state NTE. For multiplying media there only exists a steady-state solution to the NTE for sub-critical systems with a prescribed, time-independent, extraneous source or, a just critical source-free system. Third, it is assumed that all neutrons are produced immediately after a fission event, that is $\chi(\mathbf{r}, E)$ is the prompt fission spectrum. In reality, a fraction of the neutrons are released as the results of decaying fission products [30] and are referred to as delayed neutrons. The release of prompt neutrons from fission is too fast for the reaction to be controlled by an operator. The delayed neutrons increase the time taken for the neutron population in a nuclear reactor to change. Therefore, they enable the nuclear reactor to be controlled by mechanical means, such as control rods. The modelling of delayed neutron precursors is not considered in this thesis and so the prompt fission spectrum, $\chi(\mathbf{r}, E)$, is used. These three assumptions lead to the time-independent or steady-state

Symbol	Description
$\mathbf{r}$	Spatial location defined by a position vector
$\boldsymbol{\Omega}$	Direction of neutron travel
E	Neutron energy
t	Time
$\psi(\mathbf{r}, \boldsymbol{\Omega}, E, t)$	Neutron angular flux
$\phi(\mathbf{r}, E, t)$	Neutron scalar flux given by: $\phi = \int_{4\pi} \psi(\boldsymbol{\Omega}) d\boldsymbol{\Omega}$
$v(E)$	Neutron velocity
$\sigma_t(\mathbf{r}, E, t)$	Macroscopic total cross-section
$\sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}, E' \rightarrow E, t)$	Macroscopic differential scattering cross-section
$\chi(\mathbf{r}, E, t)$	Energy spectrum of prompt neutrons produced by fission
$\nu(\mathbf{r}, E, t)$	Average number of prompt neutrons produced per fission
$\sigma_f(\mathbf{r}, E, t)$	Macroscopic fission cross-section
$Q(\mathbf{r}, \boldsymbol{\Omega}, E, t)$	Extraneous (fixed) neutron source

Table 2.1: Description of dependent variables, independent variables, and neutron macroscopic cross-section data for the first-order neutron transport equation.

first-order form of the NTE:

$$\begin{aligned}
& \boldsymbol{\Omega} \cdot \nabla \psi(\mathbf{r}, \boldsymbol{\Omega}, E) + \sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, \boldsymbol{\Omega}, E) = \\
& \int_0^\infty \int_{4\pi} \sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \cdot \boldsymbol{\Omega}, E' \rightarrow E) \psi(\mathbf{r}, \boldsymbol{\Omega}', E') d\boldsymbol{\Omega}' dE' + \\
& \frac{\chi(\mathbf{r}, E)}{4\pi} \int_0^\infty \nu(\mathbf{r}, E') \sigma_f(\mathbf{r}, E') \phi(\mathbf{r}, E') dE' + Q(\mathbf{r}, \boldsymbol{\Omega}, E),
\end{aligned} \tag{2.2}$$

where the dependence of the macroscopic neutron differential scattering cross-section on angle has been transformed: $\sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}, E' \rightarrow E) \equiv \sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \cdot \boldsymbol{\Omega}, E' \rightarrow E)$.

Further assumptions can be made in order to solve equation (2.2) analytically such as: infinite media [37], diffusion approximations, symmetries [38], isotropic scattering [39], or low dimensional geometries [34]. However, in almost all real-world applications, equation (2.2) must be solved computationally.

2.1.1 Solution Methodologies

The methodologies for solving neutron transport problems using a computer can be broadly categorised as either deterministic or Monte Carlo methods. Deterministic methods are where a system of linear equations that can be written in matrix form is assembled. The matrix can then be inverted to provide a solution. For example:

$$A\mathbf{x} = \mathbf{b} \implies \mathbf{x} = A^{-1}\mathbf{b}, \tag{2.3}$$

where $\mathbf{x}$ is the solution vector, $\mathbf{b}$ is the source vector, and A is the matrix representation of the system of linear equations.

The Monte Carlo method involves simulating large numbers of neutron histories undergoing interactions in the domain. The events that a neutron undergoes in its history are determined by sampling random numbers from probability distributions. Using this method, scattering, capture, and fission events can be simulated. By simulating large numbers of neutron histories the variance in a given

quantity of interest (QoI) can be reduced.

There are two primary advantages of the Monte Carlo method. First, the geometry of the domain can be represented exactly and is well suited to complex geometries [35]. Second, because of the fact that neutron paths are independent of one another, Monte Carlo is amenable to parallel execution on large distributed memory high performance computing (HPC) systems [40]. The main disadvantage of the Monte Carlo method is that the variance in the solution reduces $O(\frac{1}{\sqrt{n}})$, meaning that large numbers of neutron histories may have to be considered in order to arrive at a satisfactory answer. Whilst variance reduction techniques, such as adjoint weighting [41], can be used, they can still not provide a large enough increase in performance to outperform deterministic methods.

In order to solve equation (2.2) using deterministic methods each dimension of the phase space must be numerically discretised. A detailed discussion of the different approaches to numerically discretising these phase space variables, as well as their relative advantages and disadvantages, is now presented.

2.2 Integral Transport Methods

Integral transport methods are based upon the characteristic form of the NTE. It is derived by considering the integration of the streaming operator over a straight line in direction Ω , referred to as the characteristic direction. In this approach the first-order form of the NTE is transformed into an ordinary differential equation along Ω [42]. The backward characteristic form of the NTE is written as:

$$\frac{d}{ds}\psi(\mathbf{r} + s\Omega, \Omega) + \sigma_t(\mathbf{r} + s\Omega)\psi(\mathbf{r} + s\Omega, \Omega) = Q(\mathbf{r} + s\Omega, \Omega), \quad (2.4)$$

where $\mathbf{r}$ is the starting point of a particle and s is the distance the particle has travelled along the characteristic [21, Ch. 8]. The source Q consists of the scattering term, fission term, and any extraneous (fixed) sources. There are two main numerical methodologies that can be derived from the characteristic form of the neutron transport equation: the collision probability method (CPM), and the method of characteristics (MoC).

2.2.1 Collision Probability Method (CPM)

The collision probability method (CPM) can be derived from equation (2.4) by applying the integrating factor method where the integrating factor is the exponential of the optical path length of a neutron. By assuming the neutron scattering is isotropic, Peierls' equation for the scalar neutron flux can be derived [21, Ch. 5]. The problem domain is then divided into N regions. In order to determine the scalar neutron flux, an $N \times N$ dense matrix must be inverted. Due to the fact that a dense matrix must be inverted to solve the system of linear equations, the CPM does not scale well for large problems [21, Ch. 9]. This scaling is so intractable in higher dimensions that the Wigner-Seitz approximation is often applied to simple two dimensional (2D) pincells in order to reduce them to one dimensional (1D) infinite cylindrical problems. A 2D Cartesian pincell may be converted to a 1D problem in polar coordinates by converting the square moderator region into a circle and preserving the area of the moderator [21, Ch. 9]. This introduces rotational symmetry to the geometry and is known as the Wigner-Seitz approximation [19, 43]. Figure 2.1 demonstrates this procedure.




Figure 2.1: A two dimensional pincell in a Cartesian coordinate system can be approximated by a circular pincell in one dimensional polar coordinates due to rotational symmetry. The radius of the circular moderator region is chosen such that the volume of the moderator is preserved. This is known as the Wigner-Seitz approximation.

Although the CPM can treat general geometry computational domains, the inversion of dense matrices makes the CPM unsuited to large neutron transport problems. Rather than inverting large dense matrices, larger computational domains can be dealt with by the coupled current collision probability (CCCP) method, where CPM is applied to small subdomains and these subdomains are coupled via their currents and using response matrices. This means that CPM can be used for small problems where it is best suited. However, an assumption must be made about the currents leaving and entering subregions. Any assumption may not be valid in all circumstances throughout the problem. For example, in the presence of strong absorbers, such as control rods or fuel rods containing burnable poisons, the angular and spatial variations of surface currents will be very different from the behaviour of surface currents between homogeneous subdomains [21, Ch. 9][44].

Finally, anisotropic scattering and void regions are not dealt with naturally by the CPM or CCCP method. To deal with void regions reduced collision probabilities must be calculated [45] whilst linearly anisotropic scattering is included in the equation using a transport cross-section [19]. This makes the CPM and CCCP unsuitable for coupled neutron gamma shielding calculations as high orders of anisotropy will be present in the gamma ray scattering. For certain calculations, such as fine-mesh lattice calculations, the disadvantages of the CPM are so great that MoC is the method of choice [21, Ch. 9].

2.2.2 Method of Characteristics (MoC)

The method of characteristics (MoC) integrates equation (2.4) over a series of characteristic directions or “tracks” given by a prescribed quadrature set. MoC schemes can be broken down into long and short characteristic methods. Long characteristic methods typically provide more accuracy as they integrate along each characteristic until the problem boundary is reached. The long characteristic method has no numerical dispersion [46]. On the other hand, short characteristic methods only integrate along the characteristic over a few zones. This means that short characteristic methods are computationally faster and require less memory than long characteristic methods but suffer from numerical dispersion effects [47].

In MoC, each zone is assumed to have constant material properties [48]. Typically the source is also

assumed constant within a zone, known as the flat source or step characteristic (SC) approximation, which requires a very fine mesh to reduce the discretisation error in the neutron scalar flux [49]. The SC scheme is a strictly positive scheme (no negative fluxes) and is also a conservative discretisation scheme along a prescribed characteristic [21, Ch. 8][50]. By solving equation (2.4) over a collection of characteristic directions the scalar neutron flux can be determined using an angular quadrature rule. This is very similar to a discrete ordinate (S_N) methodology and can suffer from the ray-effect for analogous reasons. Choosing the directions of the characteristics in order to mitigate the ray-effect and minimise discretisation error is an open problem. Characteristics can be chosen to avoid geometric discontinuities using the macroband method [51] - or to preserve transmission probabilities [52].

The MoC scales much better than the CPM and three dimensional (3D) computations can be performed in principle. However, one of the main issues associated with MoC is the development of scalable, general geometry ray-tracing algorithms in 3D. Therefore, the majority of MoC algorithms use either axial buckling approximations or, more recently, a hybrid 2D MoC method coupled to a 1D pseudo-transport approximation in the axial direction such as the SP_N method [48, 53]. The magnitude of the error of the hybrid 2D MoC/1D axial pseudo-transport method is still an open problem in the field of neutron transport [54]. Unlike the CPM, MoC can deal with void regions naturally as ray-tracing in a void is relatively straightforward and consists of determining the uncollided flux. However, MoC cannot easily deal with arbitrary order anisotropic scattering and an augmented transport cross-section is typically used to approximate linearly anisotropic scattering [55, 56].

Finally, as with all spatial discretisation techniques mentioned in the rest of this chapter, the MoC is not able to exactly calculate areas and volumes of the geometry. The geometry can be represented exactly using combinatorial geometry [57] and CAD software [58] which provide high accuracy for ray-tracing algorithms. However, the volumes of the zones can not, in general, be calculated exactly due to the rectangular nature of the volumes associated with each characteristic direction [21, Ch. 9][59]. This could lead to errors calculating reaction rates, as well as in criticality problems, unless suitably adjusted.


2.3 Angular Discretisation Methods

2.3.1 Discrete Ordinate Method (S_N)


The discrete ordinate (S_N) method is an angular collocation method which discretises the angular variable of the NTE into a discrete set of directions. These directions are chosen from a quadrature set where each direction has a corresponding weight. The solution of the discrete ordinate equations yields the angular flux solution in the particular directions associated with the angular quadrature set. The collection of angular flux solutions, in particular directions, can be used to determine a discrete approximation of integration over angle using the following expression:

$$\int_{4\pi} f(\boldsymbol{\Omega}) d\boldsymbol{\Omega} \approx \sum_{i=1}^N w_i f(\boldsymbol{\Omega}_i). \quad (2.5)$$

The weights and quadrature directions are given by w_i and $\boldsymbol{\Omega}_i$ respectively, where $\{w_m, \boldsymbol{\Omega}_m\}_{m=1}^N$ form the angular quadrature set. Some typically used angular quadrature sets are the level symmetric set, triangular Legendre-Chebyshev, and square Legendre-Chebyshev. In one dimension the scalar neutron



(a) S_2 discretisation, the NTE is solved in 4 directions.



(b) S_4 discretisation, the NTE is solved in 12 directions.

Figure 2.2: Examples of the ray effect for a source in a strongly absorbing media. Oscillations in the surrounding media are qualitatively present.

flux can be formed using equation (2.5):

$$\phi(x) = \int_{-1}^1 \psi(x, \mu) d\mu \approx \frac{1}{2} \sum_{n=1}^N w_n \psi(x, \mu_n) = \frac{1}{2} \sum_{n=1}^N w_n \psi_n(x), \quad (2.6)$$

where μ is the angular variable, and the factor $\frac{1}{2}$ is required for normalisation purposes [19, 60].

The S_N method was first developed for astrophysics modelling [61] and is now used in radiative heat transfer [62] as well as to determine reference solutions for other radiation transport problems where ray-effects do not dominate [63, 64]. However, the S_N method suffers from two major drawbacks. First, if the problem domain contains optically thick media, where the media are many mean free paths thick and scattering dominates [65], this causes an increase in the computational time due to slow convergence of the scattering source [66, 67]. Second, the solution of the governing equation in a discrete number of streaming directions introduces the phenomenon known as the ray-effect.

The ray-effect manifests itself as spurious oscillations in the scalar neutron flux. This is illustrated in Figure (2.2) for two different levels of angular refinement. It is inherent in the weak coupling between angles in the S_N method and therefore, is independent of the level of spatial refinement [66]. If the NTE is formulated in spherical coordinates then the angles can be more strongly coupled eliminating the ray-effect but at the cost of having to solve all angles simultaneously which increases memory requirements and reduces parallel efficiency [68].

The magnitude of the ray-effect is problem dependent but can be easily observed in problems such as a source in a purely absorbing medium, or a scattering media filled with localised sources [69]. There are a number of algorithms that mitigate the ray-effect, such as Lathrops ‘fictitious source method’ [69]. However, they tend to suffer from a number of drawbacks such as dramatic increase in computational time, being highly problem dependent, or causing a loss of numerical convergence [70].

Only the spherical harmonics (P_N) method, where the basis functions are rotationally invariant on the unit sphere, completely eliminates the ray-effect [70]. However, the P_N method suffers from another numerical issue which is the Gibbs phenomena, which are spurious numerical artefacts in the angular neutron flux [71]. In 1-D slab geometry the S_N method with a Gauss-Legendre angular quadrature set can be shown to be equivalent to a P_{N-1} spherical harmonic angular discretisation with

Mark boundary conditions [72].

2.3.2 Spherical Harmonics Method (P_N)

Whilst the S_N method resembles a collocation method, the P_N method can be compared to Ritz's method as the angular dependence of the angular neutron flux on the unit sphere is expanded into a finite sum of known spherical harmonic basis functions. The method of spherical harmonics was first considered by J. H. Jeans [73] and was applied to the 1D NTE by J. C. Mark [74, 75]. The spherical harmonic basis functions form a rotationally invariant basis over the unit sphere. Therefore, any function that depends on Ω can be written as an infinite sum over the spherical harmonics:

$$f(\Omega) = \sum_{l=0}^{\infty} \frac{2l+1}{4\pi} \sum_{m=-l}^l f_l^m R_l^m(\Omega), \quad (2.7)$$

where the R_l^m are the real spherical harmonic functions and f_l^m is given by:

$$f_l^m = \int_{4\pi} f(\Omega) R_l^m(\Omega) d\Omega. \quad (2.8)$$

Since the angular neutron flux is a function of Ω the above method can be applied giving:

$$\psi(\mathbf{r}, \Omega, E) = \sum_{l=0}^{\infty} \frac{2l+1}{4\pi} \sum_{m=-l}^l \phi_l^m(\mathbf{r}, E) R_l^m(\Omega), \quad (2.9)$$

with ϕ_l^m given by:

$$\phi_l^m(\mathbf{r}, E) = \int_{4\pi} \psi(\mathbf{r}, \Omega, E) R_l^m(\Omega) d\Omega. \quad (2.10)$$

In order for this to be used to form a linear system of equations, the infinite sum must be replaced by a finite sum. The typical method for this is to truncate the series at the N^{th} order spherical harmonic:

$$\psi(r, \Omega, E) \approx \sum_{l=0}^N \frac{2l+1}{4\pi} \sum_{m=-l}^l \phi_l^m(\mathbf{r}, E) R_l^m(\Omega). \quad (2.11)$$

When choosing where to truncate the outer sum odd values of N are chosen. This is because the truncated equations supply $(N+1)$ boundary conditions. If N is even then the number of incoming and outgoing boundary conditions on a surface can not be equal which introduces an artificial asymmetry into the problem [19, p. 38].

The P_N method is free of the ray-effect due to the spherical harmonic basis functions being rotationally invariant on the unit sphere [64], and as $N \rightarrow \infty$ the P_N solution converges to the exact solution of the NTE [76].

When the P_N method is used in multiple dimensions, the number of equations to solve grows as $O(N^2)$ [77, p. 79] and as N increases, the coupling of terms in the streaming operator becomes much more complicated. Due to the fact that these complications are especially prevalent in higher dimensions the S_N method was favoured in the early development of angular discretisation methods due to its relative simplicity and lower computational cost [78, p. 147]. Furthermore, the time-independent or steady-state first-order form of the NTE in a vacuum region does not have a unique solution [79].

Indeed, there will be more unknowns than equations, in other words the linear system of equations that are formed would be under-determined [79].

Due to the fact that the spherical harmonic basis functions are defined over the whole angular domain (4π Steradians), there are issues when trying to apply exact boundary conditions that only apply to half of the angular domain (2π Steradians). For example, the vacuum boundary condition can not be represented exactly by the truncated spherical harmonic expansion used in equation (2.11) [78, p. 97]. Therefore, the half angle boundary conditions are applied approximately by using either the Marshak [50, p. 95] or the Mark [78, 80] boundary conditions.

The Mark boundary condition stipulates that the vacuum boundary condition holds in a finite number of angular directions:

$$\psi(\mathbf{r}, \boldsymbol{\Omega}_m, E) = 0, \quad (2.12)$$

and the Marshak boundary condition approximately sets the net neutrons entering the boundary to zero in a way that is consistent with the P_N expansion [81]. Furthermore, because any finite sum of spherical harmonic functions is continuous there can be issues with representing discontinuities in angular fluxes that tend to occur near interfaces. This is another contributing factor to the issues with vacuum boundary conditions [78]. Remedies such as the double P_N approximation have been suggested [35], but cannot be applied to geometries of arbitrary shape [82].

2.3.3 Simplified Spherical Harmonics Method (SP_N)

The SP_N method was first introduced by Gelbard in 1960 [83, 84, 85] at a time when solving the 3D NTE was computationally infeasible. The original derivation begins with the 1D first-order form of the NTE angularly discretised using the P_N method:

$$\begin{aligned} \frac{d\phi_1}{dx} + \sigma_0\phi_0 &= Q, \\ \frac{n}{2n+1} \frac{d\phi_{n-1}}{dx} + \frac{n+1}{2n+1} \frac{d\phi_{n+1}}{dx} + \sigma_n\phi_n &= 0, \text{ for } n = 1, \dots, N, \end{aligned} \quad (2.13)$$

where $\sigma_n = \sigma_t - \sigma_{s,n}$, $\sigma_{s,n}$ is the macroscopic neutron cross-section for the n^{th} order of anisotropic scattering, and the angular dependence of the angular neutron flux has been expanded in Legendre polynomials given as:

$$\psi(\mathbf{r}, \boldsymbol{\Omega}) = \sum_{n=0}^N \frac{2n+1}{4\pi} P_n(\boldsymbol{\Omega}) \phi_n(\mathbf{r}), \quad (2.14)$$

which is the one dimensional analogue of equation (2.11). The following transform is then applied to the equation set described in equation (2.13) in order to transform the equation from 1D into 3D:

$$\begin{aligned} \frac{d\phi_n}{dx} &\rightarrow \nabla \cdot \underline{\phi}_n, \text{ for } n \text{ even}, \\ \frac{d\phi_n}{dx} &\rightarrow \nabla \phi_n, \text{ for } n \text{ odd}. \end{aligned} \quad (2.15)$$

where in the case of n being even, ϕ_n has been expanded as a vector so that its divergence can be calculated. Combining equations (2.13) with the transform in expression (2.15) and manipulating the

result yields $\frac{N+1}{2}$ second-order, weakly coupled equations in 3D which are given as:

$$\begin{aligned}
& -\nabla \cdot \frac{1}{3\sigma_1} \nabla \phi_0 - \nabla \cdot \frac{2}{3\sigma_1} \nabla \phi_2 + \sigma_0 \phi_0 = Q, \\
& -\nabla \cdot \left(\frac{n(n-1)}{(2n+1)(2n-1)\sigma_{n-1}} \right) \nabla \phi_{n-2} - \nabla \cdot \left(\frac{(n+1)(n+2)}{(2n+1)(2n+3)\sigma_{n+1}} \right) \nabla \phi_{n+2} \\
& -\nabla \cdot \left(\frac{n^2}{(2n+1)(2n-1)\sigma_{n-1}} + \frac{(n+1)^2}{(2n+1)(2n+3)\sigma_{n+1}} \right) \nabla \phi_n + \sigma_n \phi_n = 0 \\
& \text{for } n = 2, 4, \dots, N-1.
\end{aligned} \tag{2.16}$$

The Mark and Marshak boundary conditions can be extended using the transforms described in expression (2.15) [77]. This derivation means that the SP_N method applied to a 1D slab is identical to the P_N method as each transform becomes the identity transform [86]. It can also be shown that $\text{SP}_1 \equiv \text{P}_1$ in 1D, 2D, and 3D [50].

Although the initial derivation was heuristic, mathematically rigorous derivations of the SP_N equations have been performed. In 1993, Larsen et al. derived the $\text{SP}_{1,2,3}$ equations for regimes with isotropic scattering using asymptotic analysis [64]. This derivation was valid for weakly multi-dimensional, heterogeneous problems. Weakly multi-dimensional means that if strong heterogeneities exist, they must look one dimensional locally. They showed that in these cases the SP_N method was significantly more accurate than P_1 whilst being much faster than S_4 methods. Pomraning provided another asymptotic derivation of the SP_N equations [87]. He assumed that the geometry is locally planar but that the scattering may be highly-peaked, demonstrating another domain of validity. A variational derivation was also performed by Larsen [88]. However, this variational form is limited to time-independent or steady-state homogeneous problems.

The computational performance and accuracy of the SP_N method has been investigated for problems with differing levels of optical thickness. It has been shown that the SP_N method is computationally more efficient than the P_N method depending on the optical thickness of the material [86]. Larsen et al. [88] expanded upon their previous work providing a derivation of the SP_N equations for optically thick domains that are dominated by non-forward peaked scattering. They demonstrated that the SP_3 and SP_5 equations provide an improvement in accuracy when compared to the P_1 , or diffusion theory, discretisation and a significant improvement in computational efficiency compared to the S_4 discretised first-order form of the NTE. The SP_3 and SP_5 equations were also shown to produce reasonably accurate solutions in void and highly absorbing regions although they produce more accurate solutions in the latter case. The SP_N method is rotationally invariant and therefore doesn't suffer from the ray-effect. This can lead to SP_N being preferable in anisotropic scattering regions compared to the S_N method [63, 88]. However, in cases where the S_N method does not suffer from the ray-effect, S_N is seen to be more accurate; albeit at a higher computational cost.

One downside of the SP_N method is that as $N \rightarrow \infty$ the solution of the SP_N method is not guaranteed to converge to the solution of the NTE. Therefore, the angular order can not necessarily be increased to obtain better results unlike in P_N and S_N angular discretisation methods [63]. However, the SP_N method is much more computationally efficient than the P_N method. Indeed, the SP_N method can be as much as twenty times more efficient than the P_N method for certain problems [89]. This is due to the complex coupling between the P_N equations, the additional cross-derivative terms in the even-parity and self-adjoint angular flux (SAAF) P_N equations, and the greater number of equations.

Furthermore, Larsen et al. [88] demonstrated that the SP_2 method can be worse than P_1 , or diffusion theory, in heterogeneous media. This is caused by the lack of spatial continuity for the odd-order SP_N methods.

Alternative forms of the SP_N equations do exist such as the A_N equations [90]. The A_N formulation is derived in a similar manner to Gelbard's original derivation. However, it uses the first-order form of the NTE angularly discretised with the S_N method instead of the P_N method. The A_N equations take the form of coupled diffusion equations with coupling achieved through the scattering term [77]. For isotropic scattering the A_N equations are equivalent to the SP_{2N-1} equations. When anisotropic scattering is considered the A_N equations are also referred to as the canonical form of the SP_N equations [88].

2.4 The Discretisation of the Energy Domain

One simple way to deal with the energy dependence of the neutron flux and neutron cross-sections is to assume that neutrons scatter elastically from nuclei of infinite mass and experience no change in energy or velocity [91]. This assumption characterises the mono-energetic NTE as the neutron fluxes and cross-sections become independent of energy.

The mono-energetic NTE can be useful for treating neutron distributions that can be characterised as either very fast or very close to thermal [92, p. 66]. Furthermore, the mono-energetic NTE enables simpler derivations of analytical solutions for verification purposes. In addition the mono-energetic NTE is used by mathematicians to provide proofs of existence, completeness, and uniqueness of solutions to the NTE [21, p. 453].

However, in reality the neutron cross-sections do vary with energy. In most neutron transport problems this variation is characterised by concentrated peaks and troughs in the epi-thermal energy region (approximately 1 eV to 100 keV). These are referred to as the resonance structures of a nuclei and are a consequence of the internal nuclear structure. One important example is the microscopic fission cross-section of U^{238} which is plotted in Figure 2.3. In order to resolve the energy variation of neutron interactions the multigroup approximation is often used. In energy ranges where resonances are present many energy groups may be needed.

The multigroup approximation divides the energy spectrum into G discrete groups $[E_g, E_{g-1}]$, $g = 1, \dots, G$ where E_0 is the highest energy group. It was first developed by H. Hurwitz Jr. and R. Ehrlich for slowing down theory using the neutron age diffusion equation [93]. The multigroup approximation leads to the following identity:

$$\int_0^\infty f(E)dE = \int_0^{E_0} f(E)dE = \sum_{g=1}^G \int_{E_g}^{E_{g-1}} f(E)dE. \quad (2.17)$$

Generally neutrons will lose energy during a scattering event via one of two processes. The first process is elastic scattering, where a portion of the incident neutron's kinetic energy will be converted into kinetic energy for the target nuclei. The other process is inelastic scattering where, in addition to the transfer of kinetic energy, some of the incident neutron's kinetic energy may cause the target nuclei to enter an excited state. This excited state causes the target nuclei to produce gamma-ray photons as it returns to a ground state. If a neutron is thermalised then it may gain energy during a scattering

interaction. This is due to the fact that the velocity of the target nuclei caused by thermal excitation is relatively large compared to the magnitude of the velocity of the incident neutron [30, p. 48].

The multigroup form of the time-independent or steady-state first-order form of the NTE is derived by integrating equation (2.2) over each energy group giving:

$$\begin{aligned}
& \mathbf{\Omega} \cdot \nabla \psi_g(\mathbf{r}, \mathbf{\Omega}) + \sigma_{t,g}(\mathbf{r}) \psi_g(\mathbf{r}, \mathbf{\Omega}) = \\
& \sum_{g'=1}^G \int_{4\pi} \sigma_{s,g' \rightarrow g}(\mathbf{r}, \mathbf{\Omega}' \cdot \mathbf{\Omega}) \psi_{g'}(\mathbf{r}, \mathbf{\Omega}') d\mathbf{\Omega}' + \\
& \chi_g(\mathbf{r}) \sum_{g'=1}^G \nu_{g'}(\mathbf{r}) \sigma_{f,g'}(\mathbf{r}) \phi_{g'}(\mathbf{r}) + Q_g(\mathbf{r}, \mathbf{\Omega}), \text{ for } g = 1, \dots, G.
\end{aligned} \tag{2.18}$$

The group-dependent quantities, which are given by the subscript g or g' , are defined by:

$$\begin{aligned}
\psi_g(\mathbf{r}, \mathbf{\Omega}) &= \int_{E_g}^{E_{g-1}} \psi(\mathbf{r}, \mathbf{\Omega}, E) dE, \\
\phi_g(\mathbf{r}) &= \int_{E_g}^{E_{g-1}} \phi(\mathbf{r}, E) dE, \\
\chi_g(\mathbf{r}) &= \int_{E_g}^{E_{g-1}} \chi(\mathbf{r}, E) dE, \\
\sigma_{t,g}(\mathbf{r}, \mathbf{\Omega}) &= \frac{\int_{E_g}^{E_{g-1}} \sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, \mathbf{\Omega}, E) dE}{\int_{E_g}^{E_{g-1}} \psi(\mathbf{r}, \mathbf{\Omega}, E) dE}, \\
\sigma_{s,g' \rightarrow g}(\mathbf{r}, \mathbf{\Omega}', \mathbf{\Omega}) &= \frac{\int_{E_g}^{E_{g-1}} \int_{E_{g'}'}^{E_{g'-1}} \sigma_s(\mathbf{r}, \mathbf{\Omega}' \cdot \mathbf{\Omega}, E' \rightarrow E) \psi(\mathbf{r}, \mathbf{\Omega}, E') dE' dE}{\int_{E_g'}^{E_{g'-1}} \psi(\mathbf{r}, \mathbf{\Omega}', E') dE'}, \\
\nu \sigma_{f,g}(\mathbf{r}) &= \frac{\int_{E_g}^{E_{g-1}} \nu \sigma_f(\mathbf{r}, E) \phi(\mathbf{r}, E) dE}{\phi_g(\mathbf{r})}, \\
Q_g(\mathbf{r}, \mathbf{\Omega}) &= \int_{E_g}^{E_{g-1}} Q(\mathbf{r}, \mathbf{\Omega}, E) dE, \\
&\text{for } g = 1, \dots, G.
\end{aligned} \tag{2.19}$$

It can be seen that an angular dependence has been introduced into the total cross-section. The angular dependence is removed by assuming the neutron angular flux is separable with respect to angle. In other words, the neutron angular flux can be written in the following form:

$$\psi(\mathbf{r}, \mathbf{\Omega}, E) \approx \psi(\mathbf{r}, E) f(\mathbf{r}, \mathbf{\Omega}), \tag{2.20}$$

where $\psi(\mathbf{r}, E)$ is a specified neutron spectrum [21, Ch. 5]. When ansatz (2.20) is substituted into the macroscopic neutron differential scattering and total cross-sections from equation (2.19) they can be

written as:

$$\begin{aligned}\sigma_{t,g}(\mathbf{r}) &= \frac{\int_{E_g}^{E_{g-1}} \sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, E) dE}{\int_{E_g}^{E_{g-1}} \psi(\mathbf{r}, E) dE}, \\ \sigma_{s,g' \rightarrow g}(\mathbf{r}, \boldsymbol{\Omega}' \cdot \boldsymbol{\Omega}) &= \frac{\int_{E_g}^{E_{g-1}} \int_{E_{g'}}^{E_{g'-1}} \sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \cdot \boldsymbol{\Omega}, E' \rightarrow E) \psi(\mathbf{r}, E') dE' dE}{\int_{E_{g'}}^{E_{g'-1}} \psi(\mathbf{r}, E') dE'},\end{aligned}\tag{2.21}$$

as the function $f(\mathbf{r}, \boldsymbol{\Omega})$ cancels out of both of these expressions. These then constitute the multigroup macroscopic neutron cross-sections.

In order to evaluate the multigroup macroscopic neutron cross-sections an expression for the neutron spectrum $\psi(\mathbf{r}, E)$ is needed, which is ultimately related to the unknown scalar and angular neutron flux. Therefore, a ‘good enough’ guess for the neutron spectrum $\psi(\mathbf{r}, E)$ is required in order to generate the multigroup macroscopic neutron cross-section data [94]. The neutron spectrum $\psi(\mathbf{r}, E)$ is assumed to be separable in space and energy, that is:

$$\psi(\mathbf{r}, E) \approx \psi_{\mathbf{r}}(\mathbf{r}) \psi_E(E).\tag{2.22}$$

The spatial and energy dependence of the neutron spectrum can then be dealt with separately. The energy dependence of the neutron spectrum can be chosen from many options. For example, in LWRs, the fast neutron prompt fission spectrum may be approximated by the Watt spectrum [95], the epi-thermal region can use a slowing down calculation performed with a $\frac{1}{E}$ energy spectrum [96], and the thermal region typically uses a thermal Maxwellian spectrum [30]. Therefore, the neutron spectrum described above would be of the following form [94]:

$$\psi_E(E) = \begin{cases} \propto \exp(-1.036E) \sinh \sqrt{2.29E} & \text{(fast prompt fission spectrum),} \\ \propto \frac{1}{E} & \text{(epi-thermal spectrum),} \\ \propto E \exp\left(\frac{-E}{k_B T}\right) & \text{(thermal/Maxwellian spectrum).} \end{cases}\tag{2.23}$$

The neutron spectra in equation (2.23) are used to group condense (or collapse) the pointwise continuous microscopic data from nuclear data libraries such as JEFF (Joint Evaluated Fission and Fusion) [97] and ENDF/B-VII (Evaluated Nuclear Data File) [97] to produce multigroup microscopic cross-section data for use in lattice physics calculations. The nuclear data processing code NJOY performs the group condensation using these energy spectra and also performs a slowing down calculation in the epi-thermal energy region to determine the detailed spectrum in the epi-thermal region which varies approximately as $\frac{1}{E}$. An example of the multigroup cross-section library produced using the nuclear data processing code NJOY is the 172 energy group WIMS XMAS library [98]. WIMS then performs lattice physics calculations in 172 energy groups to determine the detailed spatial variation in the angular and scalar neutron flux. WIMS can also perform spatial homogenisation and further group condensation for use in whole core reactor physics codes such as PANTHER [99].

In reality the neutron cross-sections depend on temperature, time, and various other factors due to processes that occur in the nuclear reactor such as: fuel burn-up/depletion, moderator density variations, Doppler broadening, and varying fuel enrichment levels. In these cases the neutron cross-sections are typically homogenised and collapsed (or group condensed) into relatively few energy




Figure 2.3: Microscopic neutron fission cross-section of U^{238} . (From: <http://www.nndc.bnl.gov/>)

groups using infinite domain, time-independent or steady-state, lattice calculations and then passed to a program or module that solves the multigroup NTE [20]. The neutron cross-sections for various levels of fuel burn-up/depletion, and at different temperatures are retrieved by interpolating between the results from the lattice calculations.

In this thesis it is assumed that the appropriately calculated multigroup neutron cross-section data has been supplied for the calculation and the physical factors mentioned above are not explicitly modelled.

2.5 Spatial Discretisation Methods

2.5.1 Finite Difference (FD) Method

The finite difference (FD) method was one of the first numerical methods to be applied to the NTE in order to discretise the spatial domain. It involves defining a collection of mesh or grid points over the spatial domain and making an assumption about how the solution varies between these points. Some simple examples of this are the forward and backward FD formulas, which involve truncating a Taylor expansion of the solution at the required order of accuracy, and the diamond difference (DD) scheme.

The DD scheme is a second-order FD method which “leads to simple recursion formulas that generally give high accuracy in relation to the computing time required” [100], and corresponds to assuming that the flux varies linearly over the cell [101]. It is referred to as the diamond difference scheme as the FD stencil looks like a diamond in 2D. One major problem with the DD scheme is that it can produce negative fluxes. Whilst these negative fluxes only appear where the real flux would

be near zero, they can still slow the convergence of solvers and can lead to the failure of acceleration schemes [101]. Some alternatives designed to eliminate negative fluxes were developed such as the step method [19], the weighted-diamond difference (WDD) scheme [102], or the zero-and-rebalance method [103]. However, it was found that all schemes designed to enforce positivity of the flux led to a longer convergence time when compared to the second-order DD scheme [101].

In modern neutron transport codes, the finite element (FE) method is the more popular option for spatial discretisation due to its ability to model more general non-Cartesian geometries. However, some modern codes still include the option for these older FD schemes for use in the first-order NTE [104].

2.5.2 Finite Element (FE) Method

The Finite Element (FE) method was first developed in the 1950s. One of the first papers was “Stiffness and Deflection Analysis of Complex Structures” by Turner et al. [105], which was developed in order to aid the field of structural mechanics. Mathematically, FE is based on work done by Walther Heinrich Wilhelm Ritz in 1909 [106] which is also sometimes credited to Lord Rayleigh (John William Strutt, 3rd Baron Rayleigh), although this claim can be a point of contention [107]. Ritz’s method requires the derivation of a variational problem. This variational functional is then turned into a finite problem and minimised by expanding the function that minimises the functional as a finite sum of basis functions. However, later on it was found by Boris Galerkin that the continuous operator could be turned directly into a discrete problem without the need to generate a variational problem. This is done by forcing the residual of the weighted continuous operator to be orthogonal to a chosen set of basis functions [108]. The Galerkin method therefore applies to a larger class of problems, as a variational problem does not need to exist [109]. Early work applying the Galerkin method to the neutron transport equation was done by Vasily Sergeyevich Vladimirov [110]. He applied the Bubnov-Galerkin method to the homogeneous, 1D NTE and derived the even parity neutron transport equation. It was not until the 1970s that FE was applied to the computational solution of reactor physics problems due to the expense of FE methods for large, multidimensional problems.

In addition to using the Galerkin method to discretise the continuous operator, FE requires the decomposition of the problem domain into a series of geometrically primitive subdomains - commonly line segments (1D), triangles & squares (2D), and tetrahedra & hexahedra (3D). The basis functions are then chosen to have local support over these subdomains and the global solution can be formed as a summation of these basis functions, similar to Ritz’s method.

Kaper et al. [111] first tackled the two-group homogeneous neutron diffusion equation (NDE) by comparing FE calculations to the FD method. A follow up paper published a year later by Kang et al. [112] contained theoretical analysis of FE with Hermite polynomials. It included results for time dependent diffusion criticality problems where piecewise polynomials were used to discretise space, energy, and time. One of the first results of FE applied to neutron transport problems was presented by Miller et al. in both 1D and 2D [113, 114], albeit for mono-energetic isotropic scattering. Martin & Duderstadt demonstrated that the first-order form of the NTE was not stable when discretised using continuous FE. The spurious numerical artefacts in the solution could be removed by using discontinuous FE methods [115]. This was demonstrated using several 1D nuclear reactor physics and radiation shielding problems. There were a lot of issues with regards to applying FE to the 3D NTE




Figure 2.4: Two FE mesh approximations of a circular region. On the far right is the circular region being approximated. The approximations are formed from triangular elements. A large number of elements are necessary to get a good qualitative representation. However, even in the case where a large number of elements are used, both the surface area and the volume of the approximations can not be exact at the same time.

as the computers available at the time had problems regarding the storage and solution of the large linear systems of equations that were generated [116].

Ackroyd developed a series of functionals that could be used to derive extreme variational principles that could be solved using FE for the even and odd parity equations (EP/OP) [82, 117, 118], least-squares (LS) formulations [119], and even including the self-adjoint angular flux (SAAF) equation [119, 120]. The FE method has since been applied to anisotropic scattering problems [121], neutron transport LS equations [122], the SAAF equation [123], and many other modern second-order formulations [124, 125, 126, 127].

The representation of complex geometries by geometric primitives inherent in FEs can lead to issues. The Babuska paradox, which is when the boundary conditions in the original problem and the limit problem are different, can occur when considering the discretised elliptic PDEs [128]. This is mostly seen in elasticity problems, but could affect the reflective boundary conditions for S_N neutron transport angular discretisation methods. Also, in general, the FE method approximates curved geometries using polygonal or polyhedral representations. For example, the FE method is incapable of exactly representing conic sections in 2D and quadric surfaces in 3D. In general, for nuclear reactor physics problems, it is desirable to maintain a domains volume and surface area in order to preserve both fissile mass and conserve neutron leakage. However, for conic sections and quadric surfaces it is impossible for FEs to exactly conserve both these quantities at the same time. The issues associated with modelling curved geometries, such as circles, is demonstrated in Figure 2.4. In the context of nuclear reactor physics lack of conservation of volume or surface area may lead to errors in determination of the reactivity, scalar neutron flux, or neutron leakage from regions within the computational domain [129].

Isoparametric FE methods provide a way of modelling curved geometries. However, the method is not general nor particularly robust. For example, linear functions are accurately interpolated when the isoparametric method is employed, but when quadratic or higher order basis functions are used, this interpolation becomes approximate [130]. Furthermore, isoparametric FEs are only C^0 continuous

between adjacent FEs and high-order continuity is required for some structural analysis problems. This led to the development of many different types of specialised FEs [131]. However, these specialised FEs were found to be too complicated and computationally expensive so were never pursued further [27]. Indeed, Lautard & Moller concluded that the use of NURBS is the next step in trying to represent geometry more precisely [129]. Therefore, this provides the motivation for the development and application of the isogeometric analysis (IGA) method, which will be presented in Chapter 3.

2.6 Second-Order Forms of the Neutron Transport Equation (NTE)

The time-independent or steady-state first-order form of the NTE, given by equation (2.2), is a hyperbolic partial integro-differential equation (PIDE). Despite the fact that this formulation can be derived based on the conservation of neutrons it is not necessarily the form that is most amenable to computational solution.

First, the first-order form of the NTE can be numerically unstable when discretised using certain methodologies. For example, a continuous Galerkin spatial discretisation of the first-order form of the NTE will not be stable, and may cause spurious numerical artefacts or oscillations to appear. Therefore, stabilising techniques, such as streamlined-upwind Petrov-Galerkin (SUPG), which can dramatically slow convergence [132, 133], or discontinuous FE must be employed [115]. The use of discontinuous Galerkin FE discretisations can lead to many more degrees of freedom than an equivalent continuous Galerkin FE discretisation [134, 135].

Second, the weak form of the first-order form of the NTE discretised using a Bubnov-Galerkin FE method does not produce a system of linear equations that is symmetric positive definite (SPD). SPD systems of linear equations can be solved using the memory efficient conjugate gradient (CG) algorithm [125] and the powerful preconditioners that have been developed for it [136]. Therefore, the first-order form of the NTE cannot utilise these types of iterative matrix solution algorithms. When discretised using discontinuous Galerkin FE methods, the system of linear equation is equivalent to a lower triangular system and so sweep based solvers are used [78, pg. 222]. Furthermore, the fact that the first-order form of the NTE is not self adjoint means that extremum variational principles cannot be derived, only saddle point principles [137]. Extremum variational principles allow for rigorous estimates of upper and lower bounds for the scalar neutron flux, and neutron current [138, 139].

A remedy to the above issues is to utilise one of the second-order forms of the NTE. Second-order forms of the NTE have been investigated since the early 1960s [110, 137]. They are characterised by the presence of a second-order derivative of the dependent variable. Some examples of second-order forms include the diffusion equation [37], simplified spherical harmonics (SP_N) equation [83, 84, 85], self-adjoint angular flux equation (SAAF) [137, 140], even parity (EP) & odd parity (OP) equations [82, 110], and the weighted least squares (WLS) [127] formulation. All of the aforementioned second-order forms can be derived from the time-independent or steady-state first-order form of the NTE either by manipulating the first-order NTE, or making certain angular approximations.

To illustrate some properties of the second order forms, as well as to show how they may be derived, a mono-energetic, isotropic scattering form of the EP NTE is derived. It begins with the mono-energetic, time-independent or steady-state, first-order form of the NTE with isotropic scattering and

an isotropic extraneous source:

$$\mathbf{\Omega} \cdot \nabla \psi(\mathbf{r}, \mathbf{\Omega}) + \sigma_t(\mathbf{r})\psi(\mathbf{r}, \mathbf{\Omega}) = \sigma_s(\mathbf{r})\phi(\mathbf{r}) + Q(\mathbf{r}). \quad (2.24)$$

The above equation can be evaluated at $\mathbf{\Omega} = -\mathbf{\Omega}$ giving:

$$-\mathbf{\Omega} \cdot \nabla \psi(\mathbf{r}, -\mathbf{\Omega}) + \sigma_t(\mathbf{r})\psi(\mathbf{r}, -\mathbf{\Omega}) = \sigma_s(\mathbf{r})\phi(\mathbf{r}) + Q(\mathbf{r}). \quad (2.25)$$

Equations (2.24) & (2.25) can be added together and subtracted from one another yielding:

$$\mathbf{\Omega} \cdot \nabla \psi^-(\mathbf{r}, \mathbf{\Omega}) + \sigma_t(\mathbf{r})\psi^+(\mathbf{r}, \mathbf{\Omega}) = \sigma_s(\mathbf{r})\phi(\mathbf{r}) + Q(\mathbf{r}), \quad (2.26)$$

$$\mathbf{\Omega} \cdot \nabla \psi^+(\mathbf{r}, \mathbf{\Omega}) + \sigma_t(\mathbf{r})\psi^-(\mathbf{r}, \mathbf{\Omega}) = 0, \quad (2.27)$$

where the even (ψ^+) and odd (ψ^-) parity fluxes are defined in the following way:

$$\psi^\pm(\mathbf{r}, \mathbf{\Omega}) = \frac{\psi(\mathbf{r}, \mathbf{\Omega}) \pm \psi(\mathbf{r}, -\mathbf{\Omega})}{2}. \quad (2.28)$$

Equation (2.27) can be rearranged to find an expression for the odd parity flux:

$$\psi^-(\mathbf{r}, \mathbf{\Omega}) = -\frac{1}{\sigma_t(\mathbf{r})}\mathbf{\Omega} \cdot \nabla \psi^+(\mathbf{r}, \mathbf{\Omega}), \quad (2.29)$$

which is then substituted into equation (2.26) giving the mono-energetic, steady-state, EP NTE with isotropic scattering and isotropic extraneous source:

$$-\mathbf{\Omega} \cdot \nabla \frac{1}{\sigma_t(\mathbf{r})}\mathbf{\Omega} \cdot \nabla \psi^+(\mathbf{r}, \mathbf{\Omega}) + \sigma_t(\mathbf{r})\psi^+(\mathbf{r}, \mathbf{\Omega}) = \sigma_s(\mathbf{r})\phi(\mathbf{r}) + Q(\mathbf{r}). \quad (2.30)$$

The above equation contains a second order derivative of the even parity flux, is elliptic, self-adjoint, and produces a symmetric Bubnov-Galerkin weak form. Therefore, it addresses all of the aforementioned negatives of the first-order neutron transport equation. In other words, PCG iterative matrix solution algorithms can be used and extremum variational principles can be derived [82]. Furthermore, it is stable under a continuous Bubnov-Galerkin spatial discretisation and can be discretised using P_N [141] and S_N [142] angular discretisations.

Use of the conjugate gradient (CG) solver is important due to the fact that the algorithm is readily parallelisable. A CG iteration can be thought of as one matrix-vector product and three vector dot products. Both of these operations have well known algorithms for parallel execution that are implemented in several libraries for different matrix storage structures [143]. The PETSc numerical matrix solution algorithm library is an example of one such library [144]. An efficient parallel implementation of the matrix-vector product is one of the dominating factors of a parallel CG implementation [145]. A series of CG iterations will approximate the action of inverting the matrix. Whilst the sweep based algorithms used to solve the first-order form of the NTE also involve a series of matrix inversions, the main difference in terms of parallel execution is that, while the matrices to be inverted in sweep based solvers are small, there are many of them. This means that, compared to one global matrix, more communication between processors may be necessary in a sweep based solve.

Because of the logically lower triangular structure of the discretised first-order NTE, CG solvers

cannot be used. If a Krylov subspace method is desired then a more general one must be used, for example, the biconjugate stabilised gradient method (BiCGSTAB) which also requires matrix-vector products. However, because of the lower triangular structure of the discretised first-order NTE the matrix-vector product will require more floating point operations (FLOPS) to complete, due to the fact that the discretised first-order NTE is less sparse than the discretised second-order NTEs. Furthermore, the more sparse the matrix system, the less data communication that is needed to compute the matrix-vector product. More of the data is kept local on a processor and so less communication is required. Therefore, any parallel algorithm for computing matrix-vector products will require a longer execution time to invert the global matrix from the discretised first-order NTE, than to invert the global matrix from the discretised second-order NTEs.

However, equation (2.30) is not without its disadvantages. The major disadvantage is the presence of the total cross-section in the denominator of the diffusion like term of equation (2.30). This causes equation (2.30) to tend towards a singularity in the presence of near void regions, and become singular in void regions. This singularity is an issue shared by most second-order forms of the NTE, as can be seen in the SP_N equations derived earlier, see equation (2.16). Void treatment methods have been developed, for example the iteration and extrapolation methods [146, 147]. The solution of problems containing void regions will be discussed in Chapter 6.

Chapter 3

Isogeometric Analysis (IGA)

3.1 Introduction & Motivation

Isogeometric Analysis (IGA) was first presented in the 2005 paper by Hughes with the aim of unifying the computer aided design (CAD) with computer aided engineering (CAE) [148]. The key difference between IGA and the FE method is the choice of the prescribed set of test and shape functions used within the weighted residual discretisation method. In the FE method, test and shape functions with convenient numerical properties are selected to approximate the unknown solution and are also used to approximate the known geometry. In IGA, the functions that are used to generate the known geometry by CAD software, are used as test and shape functions to approximate the unknown solution. This means that in IGA the geometry produced by CAD software is represented exactly at all times.

In the original paper, IGA was described using Non-Uniform Rational B-Splines (NURBS). This is because NURBS are typically used to represent geometries in commercial CAD software; albeit they are not the only types of basis functions that exist. Some other examples of suitable basis functions are LR B-splines [149], T-splines [150], hierarchical B-splines [151] and U-splines [152]. NURBS functions are a superset of B-splines, which in turn are a superset of Bernstein polynomials. As mentioned, because the NURBS functions are used to generate the geometry, using them as basis functions during analysis allows the analysis program to exactly represent CAD geometries, including conic sections in 2D and quadric surfaces in 3D. NURBS functions have a plethora of properties that make them highly suited for analysis such as forming a partition of unity, being highly differentiable, strict positivity, and they obey the convex hull property [153, 154]. Highly efficient computational algorithms exist for generating NURBS surfaces, volumes, and derivatives as well as other geometrical and numerical properties [155, 156, 157].

The ability of the analysis program to exactly represent CAD geometries is one of the main motivations behind using NURBS based IGA. The use of NURBS basis functions within the IGA spatial discretisation method provides two main advantages over conventional FE based spatial discretisation methods. First, many popular commercial CAD software packages have been developed [158, 159] and they typically use NURBS objects to build their models meaning IGA can be performed directly on CAD files such as STEP (Standard for the Exchange of Product Data) [160], IGES (Initial Graphics Exchange Specification) [161] and Parasolid XT [162] files. This eliminates the entire meshing process since, a volume does not need to be approximated by a collection of geometric primitives. This saves time due to not needing to generate an unstructured FE mesh. It also circumvents all the analysis

steps required in between such as checking the mesh is suitable for analysis, and fixing the mesh if it is not suitable. This is estimated to reduce the total time spent in analysis by 77% in some cases [27]. In addition, large mesh files do not have to be stored and generated for each case. Adaptive mesh refinement (AMR) may be performed at run time using dual weighted residual (DWR) or goal-based local error indicators where the computational resolution can be refined precisely where necessary.

The second advantage of exact geometry is the elimination of geometric discretisation errors associated representing surface areas and volumes with polygonal (2D) and polyhedral (3D) meshes. Whilst volume or surface area conservation algorithms can be applied to FE meshes, even in relatively simple cases, such as a circle, it is impossible to exactly conserve both the volume and the surface area in a given polygonal or polyhedral mesh. In a nuclear reactor physics context the choice between preserving fissile mass and surface area leads to a trade off in the accuracy of the K_{eff} or the neutron leakage [163]. With IGA the geometry is represented exactly on the coarsest mesh and at all further levels of refinement, eliminating the above issue.

Another major advantage of IGA lies in the basis functions themselves. The NURBS functions have high, readily controllable, levels of continuity over element boundaries. Through a certain combination of *hp* refinement, NURBS curves with any level of continuity can be created, down to C^0 continuity. The procedure of refining the basis functions whilst they maintain the maximum possible continuity is known as *k*-refinement and will be discussed in more detail later in this chapter. As the parameterisation does not change with refinement there is no need to regenerate the geometry when refinement is performed. Therefore, *hpk*-refinement, without any ancillary mesh generation algorithm, is possible leading to a more natural implementation of AMR, and easier generation of high-order basis functions. High-order basis functions are much more common and easy to apply in IGA computations. Therefore, many areas such as multigrid methods [164, 165, 166], direct matrix solution algorithms [167, 168], and domain decomposition solvers [169, 170] have been redeveloped for use with non-linear NURBS basis functions.

Having a higher level continuity between basis functions has been shown to increase the accuracy per degree of freedom (dof) when compared to FE [171, p. 31]. This has been shown for turbulent flow problems with varying Reynolds numbers [172], calculation of natural modes of free vibration of structures [173], and in applications to the first and second order NTEs [163, 174, 175, 176, 177]. A paper on Bézier extraction by Thomas et. al [178] allows for the easy incorporation of IGA into extant FE codes. This has been demonstrated in several cases [179, 180], including in the well known FE code FEniCs [181, 182]. The use of Bézier extraction for constraint based local refinement will be investigated in Chapter 7.

NURBS based IGA has been applied to many fields: structural mechanics [173], fluid dynamics [172], nano-structures [183], thermal-hydraulics [179], shape optimisation [184], crack propagation and the boundary element method (BEM) [180], to name a few. Collocation methods [185] have also been examined. The areas of nuclear reactor physics to which IGA has been applied are fairly limited compared to the equivalent FE literature. Hall et al. studied the application of IGA to the mono-energetic NDE and compared it to FE discretisations [163]. Welch et al. applied NURBS based IGA to the multi-group NDE [175] as well as to the even parity (EP) NTE [176]. An implementation of group-dependent meshes applied to the second-order even-parity NTE discretised with IGA was also studied. IGA allows for relatively simple generation of independent meshes for each energy group.

The mesh in each group will represent the same exact geometry so projecting solutions between group meshes is greatly simplified. Owens et al. applied discontinuous IGA to the first-order NTE [174] and studied the effects of a hanging-node discontinuous IGA formulation. Furthermore, a heuristic based AMR scheme was presented [186]. Energy dependent meshes in IGA neutron transport have also been applied to the first-order form of the NTE [187]. All of the above papers demonstrated the numerical advantages of IGA when compared to conventional FE based spatial discretisation methods.

3.2 Basis Functions

The basis functions that are chosen to represent the geometry will also be used to represent the function we are interested in. That is:

$$f(x) = \sum_{i=1}^N d_i F_i(x), \quad (3.1)$$

where F_i is some basis function from a finite set $\{F_i\}_{i=1}^N$ and the d_i 's are referred to as the *control variables*. Therefore, it is also important to consider the approximation properties of the basis functions, their ability to represent geometry, and their computational efficiency when deciding which basis functions to use.

It was mentioned earlier that NURBS are a superset of B-splines, which are a superset of the Bernstein polynomials. In order to fully describe NURBS basis functions and convey their beneficial properties, Bernstein polynomials and B-splines will also be described in this section.

3.2.1 Bernstein Polynomials and Bézier Curves

The creation of Bézier curves using Bernstein polynomials presents an intuitive method for the creation and manipulation of geometric objects. Furthermore, the algorithms that existed for their generation in the 1960's were fast, efficient, and didn't suffer from round-off error unlike the power basis form of curves [154]. The Bézier curve is named after Pierre Bézier who used them as part of Renault's UNISURF CAD system in the early 1960's. However, the curves were first developed by Paul de Casteljau after whom the popular recursive algorithm for the evaluation of Bézier curves is named [188].

An n th degree Bézier curve is defined as follows:

$$\mathbf{C}(\xi) = \sum_{i=0}^n B_{i,n}(\xi) \mathbf{P}_i, \quad \text{for } \xi \in [0, 1], \quad (3.2)$$

where the $\mathbf{P}_i \in \mathbb{R}^3$ are the *control points* in real space, and the functions under the summation are the n th degree Bernstein polynomials given by:

$$B_{i,n}(\xi) = \frac{n!}{i!(n-i)!} \xi^i (1-\xi)^{n-i}, \quad \text{for } \xi \in [0, 1]. \quad (3.3)$$

The Bernstein polynomials are non-negative, form a partition of unity, that is $\sum_{i=0}^n B_{i,n}(\xi) = 1 \forall \xi$,




Figure 3.1: Bernstein polynomials for $n = 3$. Also corresponds to a B-spline basis generated over the knot vector $\Xi = \{0, 0, 0, 0, 1, 1, 1, 1\}$.




Figure 3.2: A fourth order Bézier curve. The red crosses are the control points and the dotted bounding box shows the convex hull property. Note that at the start and end of the Bézier curve the curve intersects with the control points.

and can be recursively defined using the following formula:

$$B_{i,n}(\xi) = (1 - \xi)B_{i,n-1}(\xi) + \xi B_{i-1,n-1}(\xi). \quad (3.4)$$

Furthermore, the derivatives of Bernstein polynomials can be easily found. The third degree Bernstein polynomials are displayed in Figure 3.1.

Bézier curves have two properties that allow them to be easily manipulated, which is key during the engineering design process. First, the end points of the curve are interpolatory, that is $\mathbf{C}(0) = \mathbf{P}_0$ and $\mathbf{C}(1) = \mathbf{P}_n$. Bézier curves are non interpolatory at control points that are not the end point of a curve. The second important property is the convex hull property, meaning that a Bézier curve is bounded by the net formed by the control points used to define it. Both of these properties are illustrated in Figure 3.2.

Surfaces and volumes can be represented using a tensor product structure. An n th degree surface

is given by:

$$\mathbf{S}(\xi, \eta) = \sum_{i=0}^n \sum_{j=0}^n B_{i,n}(\xi) B_{j,n}(\eta) \mathbf{P}_{i,j}, \quad (3.5)$$

and an n th degree volume by:

$$\mathbf{V}(\xi, \eta, \zeta) = \sum_{i=0}^n \sum_{j=0}^n \sum_{k=0}^n B_{i,n}(\xi) B_{j,n}(\eta) B_{k,n}(\zeta) \mathbf{P}_{i,j,k}. \quad (3.6)$$

However, the Bézier curve formalism still leaves a lot to be desired. In order for a single Bézier curve to pass through n control points means that the curve must be $n - 1$ th degree. The algorithms for the generation of high degree Bézier curves can become unstable [189]. Furthermore, manipulating one control point value along the curve will affect the rest of the curve. In order to ameliorate these two concerns B-spline basis functions are introduced.

3.2.2 B-splines

In order to define Basis splines (B-splines) the concept of a knot vector must be introduced. A *knot vector* is a series of non-decreasing real numbers. A generic knot vector is given by:

$$\Xi = \{\xi_1, \xi_2, \xi_3, \dots, \xi_{n-1}, \xi_n\} \text{ such that } \xi_i \leq \xi_{i+1}. \quad (3.7)$$

The knot vector Ξ can span any range. However, without loss of generality, and for continuity with regard to the Bernstein polynomials, all knot vectors will be bounded such that $\Xi \subset [0, 1]$. Any knot vector that obeys this is referred to as a *unit knot vector*. Given Ξ the B-spline functions can be defined as:

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

$$N_{i,p}(\xi) = \frac{\xi - \xi_i}{\xi_{i+p} - \xi_i} N_{i,p-1}(\xi) + \frac{\xi_{i+p+1} - \xi}{\xi_{i+p+1} - \xi_{i+1}} N_{i+1,p-1}(\xi), \quad (3.8)$$

where if the expression $\frac{0}{0}$ occurs that term is assumed to be 0. The p th order B-splines are a set of piecewise polynomials of degree p that are infinitely differentiable at $\xi \neq \xi_i$ and $p - m_i$ times differentiable at $\xi = \xi_i$. The integer m_i represents the number of times that ξ_i appears in Ξ and is referred to as the *multiplicity* of knot ξ_i . The above statement can be written formally as:

$$N_{j,p}(\xi) \in \begin{cases} C^{p-m_i}, & \text{if } \xi = \xi_i, \\ C^\infty, & \text{otherwise.} \end{cases} \quad (3.9)$$

When the B-splines are referred to as having high variable continuity over element boundaries the property in expression (3.9) is what is being referenced. An example of this variable continuity is displayed in Figure 3.3. Traditional FE basis functions have C^0 continuity at element boundaries. Through careful refinement the continuity of B-spline functions can be readily controlled and FE basis functions can be mimicked if $m_i = p \forall i$.




Figure 3.3: An example of one dimensional B-spline functions with the quadratic open knot vector $\Xi = \{0, 0, 0, 0.25, 0.25, 0.5, 0.75, 1, 1, 1\}$. Notice the C^0 continuity at $\xi = 0.25$ caused by the repeated knot whilst there is C^1 continuity at $\xi = 0.5$ and $\xi = 0.75$.

Analogous to Bézier curves defined in equation (3.2), a p th degree B-spline curve is defined as:

$$\mathbf{C}(\xi) = \sum_{i=0}^N N_{i,p}(\xi) \mathbf{P}_i, \quad \text{for } \xi \in \Xi, \quad (3.10)$$

where $\mathbf{P}_i$ are the control points of the B-spline curve and N is the number of B-spline functions.


It can be seen that B-splines curves are a generalisation of Bézier curves by considering knot vectors of the following form:

$$\Xi = \{\underbrace{0, \dots, 0}_{p+1}, \underbrace{1, \dots, 1}_{p+1}\},$$


where the 0 and 1 have multiplicity $m = p + 1$. Due to this fact, the B-spline curves inherit most of the properties from Bézier curves for example, the convex hull property and forming a partition of unity. In order for the end points of a B-spline curve to be interpolatory ($\mathbf{C}(0) = \mathbf{P}_0$ and $\mathbf{C}(1) = \mathbf{P}_n$) open knot vectors must be considered. Similar to Bézier curves, B-spline curves are non interpolatory away from the end points.

An *open knot vector* of degree p is a knot vector where the first and last knots have multiplicity $p + 1$. This means that 0 and 1 appear $p + 1$ times in Ξ . This also enforces that the B-spline functions form a partition of unity for all $\xi \in \Xi$ which allows for the easy enforcement of strong Dirichlet boundary conditions. Open knot vectors are the *de facto* standard for all B-spline based IGA [27]. An open knot vector Ξ of degree p with length α has $\beta = \alpha - (p + 1)$ B-spline functions defined over it. From this point on, all knot vectors will be open unit knot vectors.

The fact that B-splines are a generalisation of Bernstein polynomials means that B-splines curves can represent a superset of the objects that Bézier curves can represent. This fact is illustrated in Figure 3.1 where the Bernstein polynomials can be defined as B-splines over a specific knot vector. Due to the piecewise polynomial nature of B-splines, a B-spline curve can contain many control points without needing to have a high polynomial order. This is due to the fact that the B-spline functions



(a) The 8th B-spline basis function



(b) The 16th B-spline basis function

Figure 3.4: A selection of the two dimensional B-spline basis functions generated over a unit square with linear open knot vectors $\Xi_\xi = \Xi_\eta = \{0, 0, 0.25, 0.5, 0.75, 1, 1\}$

over the curve defined by equation (3.10) are local, $N_{i,p}(\xi) = 0$ for $\xi \notin [\xi_i, \xi_{i+p+1})$. Given a knot span $[\xi_j, \xi_{j+1})$ at most $p + 1$ of the B-spline functions are nonzero ('active') in that knot span. Therefore, when one control point is moved it no longer effects the entire curve, as was the issue when using Bézier curves.

Similar to Bézier curves, surfaces and volumes can be generated using a tensor product structure. A p th degree B-spline surface, analogous to equation (3.5), is defined as:

$$\mathbf{S}(\xi, \eta) = \sum_{i=0}^N \sum_{j=0}^M N_{i,p}(\xi) N_{j,p}(\eta) \mathbf{P}_{i,j}, \quad (3.11)$$

where $\xi \in \Xi_\xi$ and $\eta \in \Xi_\eta$ are knot vectors for each parametric direction. An example of two dimensional B-spline basis functions is presented in Figure 3.4. A p th degree B-spline volume, analogous to equation (3.6), is defined as:

$$\mathbf{V}(\xi, \eta, \zeta) = \sum_{i=0}^N \sum_{j=0}^M \sum_{k=0}^L N_{i,p}(\xi) N_{j,p}(\eta) N_{k,p}(\zeta) \mathbf{P}_{i,j,k}. \quad (3.12)$$

When open knot vectors are used the bounding curves or surfaces of a B-spline surface or volume will themselves be B-spline curves or surfaces. This makes the extraction of bounding objects, and the application of boundary conditions, much simpler. However, there is still one more desirable geometric property that the B-spline basis functions do not have, they cannot represent circles and other conic sections [190]. In order to do that, non uniform rational B-splines are developed in the next section.

3.2.3 Non-Uniform Rational B-splines (NURBS)

In order to define Non-Uniform Rational B-splines (NURBS) one more set of parameters must be specified. The set of weights $\{w_i\}_{i=0}^N \in \mathbb{R}$ is chosen such that each control point has a weight associated with it. The NURBS functions of degree p can then be defined as:

$$R_{i,p}(\xi) = \frac{N_{i,p}(\xi)w_i}{\sum_{j=0}^N N_{j,p}(\xi)w_j}, \text{ for } i = 0, 1, 2, \dots, N, \quad (3.13)$$

where $N_{i,p}$ are the B-spline functions defined in equation (3.8). It can be shown that the NURBS are a generalisation of B-splines. If $w_i = 1 \forall i \in [0, N]$ then, due to the partition of unity property of the B-splines, the denominator in equation (3.13) becomes unity and $R_{i,p}(\xi) = N_{i,p}(\xi)$. Furthermore, the requirement that $w_i > 0$ is made to make sure that the basis functions are non-negative and that the convex hull property still holds. Analogous to B-splines, a p th degree NURBS curve is defined as:

$$\mathbf{C}(\xi) = \sum_{i=0}^N R_{i,p}(\xi) \mathbf{P}_i, \quad \text{for } \xi \in \Xi, \quad (3.14)$$

and using a tensor product structure a NURBS surface is defined as:

$$\mathbf{S}(\xi, \eta) = \sum_{i=0}^N \sum_{j=0}^M R_{i,p}(\xi) R_{j,p}(\eta) \mathbf{P}_{i,j}, \quad (3.15)$$

and a NURBS volume by:

$$\mathbf{V}(\xi, \eta, \zeta) = \sum_{i=0}^N \sum_{j=0}^M \sum_{k=0}^L R_{i,p}(\xi) R_{j,p}(\eta) R_{k,p}(\zeta) \mathbf{P}_{i,j,k}. \quad (3.16)$$

NURBS functions can be viewed as B-splines in four dimensional homogeneous coordinates where the weight is the fourth dimension. Dividing the $\mathbf{P}_i$ by their corresponding weight represents a projection from the four dimensional space to the real space $\mathbb{R}^3$. This generalisation of B-splines to four dimensional homogeneous coordinates allows NURBS to represent more complex curves and surfaces. In particular NURBS functions are able to exactly represent conic sections in two dimensions, and quadric surfaces in three dimensions [27].


Similarly to B-spline objects, due to the use of open knot vectors the bounding curves or surfaces of NURBS surfaces or volumes are themselves NURBS curves or surfaces. Since the space of NURBS functions is a superset of both the Bernstein polynomials and B-splines, from now on only NURBS functions will be considered.

3.3 Refinement Strategies


Refinement strategies in IGA present more options for the continuity and order of the basis functions than the traditional hp -refinement found in FE. This is due to the fact that in IGA the p th order basis function over knots are C^{p-m_i} , whereas in FE the basis functions are C^0 over element boundaries. The two mechanisms for refinement in IGA are knot insertion, a generalisation of h -refinement, and order elevation, a generalisation of p -refinement. Knot insertion and order elevation do not commute, so a third refinement strategy is available referred to as k -refinement which has no analogue in FE.

3.3.1 Knot Insertion

Given a knot vector $\Xi = \{\xi_1, \xi_2, \dots, \xi_{n-1}, \xi_n\}$, knot insertion is the process of inserting α new knots to form an extended knot vector $\bar{\Xi} = \{\xi_1, \xi_2, \dots, \xi_{n+\alpha-1}, \xi_{n+\alpha}\}$. Some of the new knots added may be the same as ones that were already present in Ξ , increasing their multiplicity. The knot insertion of α knots has the effect of creating α new basis functions and corresponding control points. A simple




(a) NURBS basis functions for $p = 2$
 $\Xi = \{0, 0, 0, 1, 1, 1\}$.




(b) NURBS basis functions for $p = 2$
 $\Xi = \{0, 0, 0, 0.5, 1, 1, 1\}$.

Figure 3.5: NURBS basis functions for a quadratic open knot vector Ξ before and after insertion of a knot at $\xi = 0.5$.



(a) NURBS basis functions for $p = 1$
 $\Xi = \{0, 0, 0.4, 1, 1\}$



(b) NURBS basis functions for $p = 3$
 $\Xi = \{0, 0, 0, 0, 0.4, 0.4, 0.4, 1, 1, 1, 1\}$

Figure 3.6: NURBS basis functions for a knot vector Ξ where order elevation has been performed twice. Notice all functions at $\xi = 0.4$ maintain C^0 continuity after order elevation.

example of how the NURBS basis functions are changed by knot insertion is displayed in Figure 3.5. To replicate h -refinement in IGA, a newly inserted knot must be inserted p times so that all basis functions would be C^0 at the new knot. Therefore, knot insertion is a generalisation of h -refinement.

3.3.2 Order Elevation

Given a knot vector Ξ , order elevation increases the order of all of the basis functions whilst maintaining the continuity of the basis functions over any extant knots. If Ξ is of order p and we wish to raise it to order q , then in order for C^{p-m_i} continuity to be maintained all extant knots must be inserted $q - p$ times. This includes the 0's and 1's that are present at the beginning and end of open unit knot vectors. In this way the order elevation procedure can be thought of as a series of knot insertions. The order elevated knot vector $\bar{\Xi}$ contains the same knots as Ξ but each knot has one greater multiplicity. An example of order elevation is displayed in Figure 3.6. Order elevation is a generalisation of p -refinement, as it can be applied to a set of basis function with any level of continuity over the element boundaries. Traditional p -refinement is only applicable to basis functions that are C^0 continuous over element boundaries.

3.3.3 k -refinement

The operations of knot insertion and order elevation do not commute. For example, consider the linear knot vector $\Xi = \{0, 0, 0.5, 1, 1\}$. Two cases are considered:

1. The knot $\xi = 0.7$ is inserted (knot insertion), then the degree of the knot vector is raised by one (order elevation) giving the quadratic knot vector:

$$\Xi_1 = \{0, 0, 0, 0.5, 0.5, 0.7, 0.7, 1, 1, 1\}.$$

The basis functions generated by Ξ_1 are displayed in Figure 3.7a. The continuity of the NURBS basis functions at $\xi = 0.7$ is $C^{p-m_3} = C^{2-2} = C^0$.

2. The degree of Ξ is raised by one (order elevation) and then the knot $\xi = 0.7$ is inserted (knot insertion) giving the quadratic knot vector:

$$\Xi_2 = \{0, 0, 0, 0.5, 0.5, 0.7, 1, 1, 1\}.$$

The basis functions generated by Ξ_2 are displayed in Figure 3.7b. The continuity of the NURBS basis functions at $\xi = 0.7$ is $C^{p-m_3} = C^{2-1} = C^1$.


It can be seen that $\Xi_1 \neq \Xi_2$ and that the operations of knot insertion and order elevation do not commute. Furthermore, it can be seen in Figure 3.7 that the sets of NURBS basis functions generated by Ξ_1 and Ξ_2 are not the same. The operation of k -refinement is the second case, that is order elevation followed by knot insertion. It can be thought of as the process that creates knots that allow the basis functions to achieve their maximal possible continuity. That is, any basis function will be C^{p-1} continuous over a knot created by k -refinement. It can also be noted that fewer basis functions are created when k -refinement is performed. There are 7 basis functions in Figure 3.7a and 6 basis functions in Figure 3.7b.

For continuous Bubnov-Galerkin IGA it has been found that having maximal continuity of the basis functions over element boundaries results in the most accuracy per degree of freedom for smooth problems [175]. Therefore, in this work k -refinement is utilised heavily.


Whilst the above refinement methods will create new control points and move existing ones, none of the refinement strategies change the geometry or its parameterisation [191]. Since the geometry is exact from the beginning, it will remain exact for any level of refinement, meaning the geometric error is identically zero. This means that no interface with a CAD or meshing program is needed as refinement can be done in real time with the sole aim of reducing the discretisation error. Refinement does not change the underlying geometry. Therefore, IGA is much more amenable to adaptive mesh refinement (AMR) procedures.

3.4 Geometry Creation

Due to the tensor product nature of higher dimensional NURBS basis functions, see equations (3.15) & (3.16), a d -dimensional patch is topologically equivalent to a hypercube in d -dimensional space. Therefore, in order to analyse problems with geometries that are not topologically equivalent to a hypercube, the concept of a NURBS patch must be introduced. A patch is a collection of knot vectors, control points, and weights. Together these allow NURBS functions to be defined over the



(a) NURBS basis functions generated by $\Xi_1 = \{0, 0, 0, 0.5, 0.5, 0.7, 0.7, 1, 1, 1\}$. Knot insertion followed by order elevation.



(b) NURBS basis functions generated by $\Xi_2 = \{0, 0, 0, 0.5, 0.5, 0.7, 1, 1, 1\}$. Order elevation followed by knot insertion.

Figure 3.7: NURBS basis functions showing that knot insertion and order elevation do not commute.

patch.

It is convenient to think of patches as building blocks that make up the problem domain. The ability to vary the positions of control points in real space and change the value of weights means that a 2D patch, that is a square in the parametric domain, can be mapped to a variety of shapes in real space. This is demonstrated in Figure 3.8, where the square parametric domain is mapped to a circle. Patches can then be placed together to form a consistent geometry in the physical space that is suitable for analysis. Since the basis functions at the start and end of open knot vectors are always interpolatory the patches can be glued together by placing their control point at the same location in the physical space. The basis functions between patches are therefore C^0 continuous over the interface. This does not mean that the geometry must be limited by to this order of continuity. Higher-order geometric continuity can still be generated despite the lower continuity of the NURBS basis functions [192]. A fuel pin that has been constructed from five NURBS patches can be seen on the left side of Figure 3.10.

Within a NURBS patch it is assumed that material coefficients are constant. The main reason for this is that when there is a large discontinuity in macroscopic neutron cross-section spurious numerical artefacts or oscillations can be introduced into the solution. When the basis functions are C^0 over material boundaries, these spurious oscillations tend to be more localised. If the basis functions have higher inter-element continuity then the spurious numerical artefacts or oscillations can spread. This effect is illustrated in Figure 3.9. In this problem a strong central absorber is surrounded by a region containing a strong extraneous (fixed) source. All boundary conditions are prescribed as reflective. For the C^0 and first-order results the domain was made of a four by four grid of $1\text{cm} \times 1\text{cm}$ square patches. For all other results the domain was constructed from one $4\text{cm} \times 4\text{cm}$ patch. The first-order solution uses a discontinuous Bubnov-Galerkin IGA discretisation. All other results use continuous Bubnov-Galerkin IGA discretisations. Figure 3.9 displaying plots of the scalar neutron flux along the line $y = 2\text{cm}$. It can be seen that for higher-order continuity basis functions the spurious numerical artefacts or oscillations propagate to the boundary and also lower the numerical value of the scalar neutron flux there. Furthermore, the higher the order of the basis functions the bigger the discrepancy at the boundary. It would be interesting to see if using a discontinuous interior penalty IGA discretisation [193] for the NDE between patches whilst still maintaining a continuous discretisation within a patch would result in superior qualitative results. Due to these results for the rest of this thesis NURBS patches will be assumed to have constant material properties.

In traditional continuous IGA, NURBS patches that are placed next to one another must have compatible discretisations. This means that control points for neighbouring NURBS patches must be identical [27]. Therefore, control points between NURBS patches that share a side will also get the same global number in the connectivity information. The numbering of the control points is important when minimising the bandwidth of the system of linear equations it produces when discretised and can affect the time taken to invert the system. However, the numbering of the control points can be optimised with respect to bandwidth using the reverse Cuthill-McKee algorithm [194].

Due to the requirement of a compatible discretisation, any refinement performed on a NURBS patch propagates through the domain via neighbouring NURBS patches. This process is illustrated in Figure 3.10. Later in this work the need for neighbouring discretisations to be consistent, compatible, and for refinement to propagate through the domain will be addressed by using a constraint based





Figure 3.8: Mapping from the parametric domain $\hat{V}$ to the physical domain V .

local refinement method.


The majority of the nuclear reactor physics related geometric primitives, such as pincells and areas of moderator, can be decomposed into groupings of these patches. These primitives are then assembled using an ancillary program to create full geometries, typically a CAD program. When modelling more complex geometries generated using commercial CAD software other issues can arise. Geometric continuity between patches is not guaranteed by CAD software, making performing analyses on these geometries difficult [195]. From an analysis point of view these CAD models are unusable for two reasons. First, they contain gaps, geometrically and parametrically, thus causing issues in terms of the continuity of the geometry [196]. Second, a lot of the NURBS patches may be trimmed by the CAD program. A trimmed NURBS patch is a patch that has been cut so that it is no longer topologically equivalent to a hypercube, and it is hard to say how to define the NURBS basis functions over such an geometrical object. This means that the complex commercial CAD software geometries often need to be ‘fixed’ before they are suitable for analysis. Much research has been done into the ‘fixing’ of CAD models to make them geometrically watertight [197], creating analysis software that is able to deal with these non watertight geometries [198, 199], or creating a new generation of CAD software that is able to automatically fix the CAD models [152]. Alternatively, non NURBS based IGA can be employed. For example, T-splines can be formulated over a domain with any number of sides [150].

When generating the 3D geometric objects, commercial CAD software only generates the bounding surface of a volume. There is no native parameterisation for the internal volume of a CAD model. Methods for the parameterisation of these models has been investigated, as well as the generation of complicated 3D NURBS objects that have internal parameterisations [200, 201]. 3D NURBS based IGA problems have been formulated and solved for example, for a simulation of blood flow through a human heart [202]. Nuclear reactor core geometries are relatively simple in the axial direction so 3D volumetric parameterisation may be created by extruding a 2D geometric model.


These disadvantages of commercial CAD software is not surprising as the traditional use of CAD software is to produce an engineering design of the component as opposed to models that are suitable for CAE. Still IGA is beginning to be used in industry, in particular in vehicle crash simulations [203]. In the work presented in this thesis a bespoke NURBS based geometry generation software has been developed. This enables analysis suitable IGA models to be generated for the various nuclear reactor




(a) $p = 2$, 64 elements.




(b) $p = 2$, 256 elements.




(c) $p = 3$, 64 elements.



(d) $p = 3$, 256 elements.



(e) $p = 4$, 64 elements.



(f) $p = 4$, 256 elements.

Figure 3.9: Scalar neutron flux plots along the line $y = 2\text{cm}$. Global high-order continuity of basis functions can cause propagation of spurious oscillations in the solution. Notice for the cases C^1, C^2 , and C^3 with 64 elements that there are oscillations near the domain boundary. In the case of 256 elements the flux near the boundary is lower than the C^0 case. The first-order solution is included for qualitative comparison. All other results are from a neutron diffusion solution. The C^0 results are closer to the neutron transport solution near the boundary in all cases.




Figure 3.10: A five patch pincell. The requirement that patches have compatible, consistent discretisation means that when knots are inserted into one patch, this insertion propagates through the domain. In the second picture a knot has been inserted in each parametric dimension of the southern patch. In the third picture two knots have been inserted in each parametric dimension in the western patch, and one along the horizontal axis of the southern patch.

physics and radiation shielding verification benchmark problems that are presented in this thesis.

3.5 Mathematical Spaces and Numerical Quadrature

When formulating an IGA problem and preparing to perform analysis on it, several different spaces have to be considered: the physical, parameter, and parent spaces.

Physical space: V The physical space is equivalent to the real space. It is the $\mathbb{R}^3$ space in which the geometry resides. There is only one physical space.

Parameter space: $\hat{V}$ Each NURBS patch has its own parameter space. In 2D the parameter space is defined as $[0, 1] \times [0, 1]$ due to the choice of unit knot vectors. By using the curve, surface, or volume definitions from equations (3.14) - (3.16) as a mapping, it is possible to map from the parameter space for a patch into a sub domain of the physical space that corresponds to that patch. More formally:

$$\mathbf{C}(\xi) : \hat{V} \rightarrow V \quad (3.17)$$

where $\mathbf{C}$ is a NURBS curve but could also be a surface or volume.

Parent space: $\tilde{V}$ In 2D the parent space is given by $[-1, 1] \times [-1, 1]$ and is the space in which the quadrature points reside. A portion of a patch's parameter space, defined by a nonzero knot span, is mapped into the parent space so that numerical quadrature can be performed. By using an open quadrature set, such as Gaussian quadrature, this method ensures that quadrature points will never land on knot points, as these are areas of reduced continuity of the basis functions that might effect the accuracy of the quadrature rule. Alternatively the quadrature points can be mapped from $\tilde{V}$ into $\hat{V}$ and then numerical quadrature can be performed.

The above arrangement of spaces allows basis functions to be evaluated in the parameter space. These functions can then be combined so that integrals can be evaluated using a numerical quadrature rule with quadrature points mapped from the parent space into the relevant part of the parameter

space. Finally the integrals can be combined with the Jacobian between the parameter and physical space in order to map them into the physical space.

3.5.1 Quadrature

In this work integration over an element is performed using m point Gaussian quadrature. This quadrature set is well tabulated, and integrates polynomials of degree $2m - 1$ exactly. However, the NURBS functions are rational polynomials so Gaussian quadrature still performs well although not optimally. It has been shown that as the continuity of the basis functions over element boundaries increases, the number of quadrature points required to perform exact numerical integrations of NURBS decreases [191]. Therefore, it is suggested that optimal quadrature sets may span more than one element. This has led to work investigating new strategies utilising weighted quadrature sets for the formation of stiffness matrices by assembling the stiffness matrix row by row as opposed to looping over elements [204]. This assembly methodology shows a decrease in stiffness assembly time of four orders of magnitude for tenth order NURBS functions. However, symmetric matrices assembled using this strategy will be slightly non-symmetric due to the way that the quadrature points are calculated.

In this work we are interested in maintaining the Symmetric Positive Definite (SPD) property of the stiffness matrices formed when spatially discretising second-order forms of the neutron transport equation with Bubnov-Galerkin IGA. Furthermore, in nuclear reactor physics and radiation shielding problems it is often the case that solution time dominates assembly time; this is for two reasons. First, although many stiffness matrices have to be formed, the energy and angle independent dense element matrices can be calculated and stored, then multiplied by the relevant nuclear data to calculate the stiffness, scattering, reflection, and fission matrices for each angle and energy group. Second, the solution times in neutron transport are the dominant factor as there are power iterations, scattering iterations between energy groups and angles, and potentially the need to iterate over reflective boundary conditions when using an S_N angular discretisation. Therefore, the reduction in computational effort that would be afforded by using optimal quadrature and assembly algorithms is assumed negligible and not considered in this work.

3.6 Neutron Diffusion Simulation Results

In order to highlight the impact that high-order continuity of NURBS basis functions and exact geometry have on the accuracy per dof of an IGA solution several test cases are considered. All of the results in this section are from solving neutron diffusion equation with an IGA or FE spatial discretisation.

3.6.1 Two Group Bare Nuclear Fuel Pin

The first test case is a simple circular bare fuel pin with radius $r = 1$. It has two energy groups and zero flux boundary conditions are prescribed on all external boundaries. The macroscopic neutron cross-sections used for this benchmark verification test case are presented in Table 3.1. The K_{eff} 's from solving the NDE discretised with IGA and volume preserved FE are compared. Therefore, the aim of this benchmark verification test case is to highlight the effect that geometric approximations have on the convergence rates of global quantities (such as the K_{eff}). The problem requires a power

Macroscopic neutron cross-section	Value (cm ⁻¹)
σ_t^1	1.0
σ_t^2	2.0
σ_a^1	0.2
σ_a^2	1.9
$\sigma_s^{1 \rightarrow 1}$	0.5
$\sigma_s^{1 \rightarrow 2}$	0.3
$\sigma_s^{2 \rightarrow 2}$	0.1
$\nu\sigma_f^1$	1.0
χ^1	1.0

Table 3.1: Macroscopic neutron cross-sections for the bare pin nuclear reactor physics test case. If a quantity is not specified then it is assumed to be zero.




Figure 3.11: Error in the K_{eff} for the bare nuclear fuel pin problem with radius $r = 1$ and zero flux boundary conditions prescribed on all external boundaries.

iteration algorithm which was halted when the following condition was satisfied:

$$\left| 1 - \frac{K_{\text{eff}}^{(i+1)}}{K_{\text{eff}}^{(i)}} \right| < 1 \times 10^{-11}. \quad (3.18)$$

For the rest of this thesis this will be referred to as the relative change in the K_{eff} . The reference value for this calculation is $K_{\text{eff}} = 4.3148794874600549 \times 10^{-2}$ generated with 5th order NURBS basis functions and 68,121 dof. A graph of the error in the K_{eff} compared to the number of dof for various orders of IGA and quadratic FE basis functions is displayed in Figure 3.11. When a quadratic FE discretisation is used the radius of the nuclear fuel pincell is increased in order to preserve fissile mass. However, there is still a geometric error due to the approximation of the surface area of the nuclear fuel pin.

From Figure 3.11 it is clear that, in terms of accuracy per dof, the quadratic IGA discretisation is superior to the quadratic FE discretisation and this is for two reasons. First, the IGA solution is free from any geometric error. The FE solution has a geometric error which approaches zero as the

mesh is refined. This geometric error is $O(h^2)$ due to the approximation of the surface area of the nuclear fuel pin. This can be seen in Figure 3.11 as the rate of convergence of the FE solution is $O(h^2)$. Indeed, this behaviour would be observed if the order of the FE basis functions was raised higher. If the external sides of the finite elements were mapped to an approximation of the circle using isoparametric FE techniques then this order of convergence would no longer be so limited [130]. However, isoparametric FE methods are still unable to represent the surface area and volume of a circle to machine precision.

Whilst in some cases the preservation of both surface area and volume may not be necessary, in nuclear reactor physics calculations inexact representation of either of these quantities can have a large effect on the solution. The K_{eff} is an important quantity of interest that measures whether the neutron population in a reactor will decrease over time (sub-critical), increase over time (super-critical), or if a steady state solution exists (critical). Therefore, the evaluation of the K_{eff} is key in predicting the behaviour of a nuclear reactor. The K_{eff} can be interpreted physically using the following expression:

$$K_{\text{eff}} = \frac{\text{Neutron production due to fission in the current generation}}{\text{Neutron loss due to leakage and absorption in the previous generation}}. \quad (3.19)$$

Geometric approximation of the volume of fissile regions will lead to errors in the numerator of expression (3.19), whilst errors in the surface area of external boundaries will lead to errors in the denominator of expression (3.19). Therefore, in order for the K_{eff} to be as accurate as possible, exact representation of the volume and surface area of the computational domain is critical.

The second reason for a superior IGA solution is the higher order continuity of the basis functions. Over element boundaries, within a NURBS patch, the NURBS basis functions are C^{p-1} whereas the FE basis functions are C^0 . What is concerning is that the convergence rates of the IGA solution for $p = 3, 4, 5$ are identical. It is suspected that the reason for the limited convergence rates of the IGA discretisation at higher-orders is due to the strong implementation of the zero flux boundary conditions. In order to verify if this is indeed the case the same calculation is performed but with vacuum boundary conditions prescribed on all external boundaries. For this case the reference solution is $K_{\text{eff}} = 0.10148853917922351$ generated with sixth order NURBS basis functions and 68,644 control points.

Figure 3.12 shows the results from this calculation and for $p = 2, 3, 4$ the $O(h^{p+1})$ convergence rate is observed. For $p = 5$ the theoretical $O(h^6)$ convergence rate is not present. In fact for certain refinement levels the $p = 5$ solution is less accurate per dof than the $p = 4$ solution. It is thought that this is because the K_{eff} is only converged until the relative change is less than 1×10^{-11} . Beyond this level the solution is unreliable. For this problem with $p = 5$ basis functions this convergence criteria is satisfied very quickly. The convergence rates for $p = 5$ will be investigated more in the following test case.

Spectral condition number of global stiffness matrix The spectral condition number for a diagonalisable matrix is defined as the ratio of the largest and smallest eigenvalues, that is:

$$\kappa = \frac{\lambda_{\text{max}}}{\lambda_{\text{min}}}. \quad (3.20)$$




Figure 3.12: Error in the K_{eff} for the bare pin problem with radius $r = 1$ and vacuum boundary conditions prescribed on all external boundaries.

The condition number of a system of linear equations $A\mathbf{x} = \mathbf{b}$ is a measure of how much the solution $\mathbf{x}$ will change for a small change in the source $\mathbf{b}$. Therefore, the condition number can also be used as a measure of how many iterations it will take to solve the system $A\mathbf{x} = \mathbf{b}$.

We are not aware if the matrices that result from the spatial discretisation of a given second-order form of the neutron transport equation are diagonalisable in general. Therefore, equation (3.20) is used as an estimate of κ . The eigenvalues are computed using the PETSc GMRES solver with a restart of 10,000. The option `-ksp_monitor_singular_value` is also used. More information can be found on the [PETSc FAQ](#). It should be noted that the use of PETSc's GMRES provides estimates of κ for large matrices. This, combined with the assumption that equation (3.20) is valid, means that the condition numbers here should be used as an illustration of the behaviour of κ rather than providing the correct values.

For the bare nuclear fuel pin the spectral condition number of the global stiffness matrix of the NDE with respect to the number of elements is presented in Table 3.2. It can be seen that in all cases the condition number of the IGA discretised system is less than the condition number of the FE discretised system. It is assumed that this is mostly due to the high order (C^{p-1}) of the NURBS basis functions over element boundaries.

Table 3.3 contains the spectral condition number of the same matrix system preconditioned with a symmetric successive over relaxation (SSOR) preconditioner with $\omega = 1$. The effect of the preconditioner seems to be greater for the FE discretisation than the IGA discretisation in most cases. That is, the reduction in κ is greater. This agrees with the IGA literature that the traditional preconditioners that have been developed for low continuity FE basis functions are not as suited to the high order continuity IGA basis functions [164, 170, 205, 206].

The values of κ for other refinement levels of the unpreconditioned system are displayed in Figure 3.13. It is clear from Figure 3.13 that, when plotted against the number of elements, the IGA discretised system produces a lower κ than the FE discretised system. When the number of dof is considered this is not always the case. However, the spectral condition number of the unpreconditioned system

Number of elements	$p = 2$		$p = 5$	
	IGA	FE	IGA	FE
5	13	14	1664	16070
20	8	23	17160	23730
80	25	81	12820	22180
320	86	351	8673	22980
1280	335	1516	7487	23190
5120	1400	6421	7566	>60000

Table 3.2: Spectral condition numbers of the unpreconditioned global stiffness matrix for the bare nuclear fuel pin problem.

Number of elements	$p = 2$		$p = 5$	
	IGA	FE	IGA	FE
5	2	3	43	46
20	2	3	36	110
80	6	7	125	80
320	14	24	102	98
1280	37	87	86	376
5120	113	336	202	>60000

Table 3.3: Spectral condition numbers of the SSOR preconditioned global stiffness matrix for the bare nuclear fuel pin problem.


seems to grow at the same rate, regardless of spatial discretisation methodology.

3.6.2 Mono-Energetic Lattice Calculation


The mono-energetic lattice calculation has been selected to further investigate the effect of geometric error on convergence rates of high order IGA and FE spatial discretisations. The following benchmark has been taken from an IGA neutron diffusion paper by Welch et al. [175]. It consists of a nuclear fuel pin with reflective boundary conditions prescribed on all boundaries. The geometry for the problem is displayed in Figure 3.14. The external boundaries are straight edges meaning that both IGA and FE capture this aspect of the geometry exactly. However, FE is unable to represent the circular nuclear fuel pin and so there will be some error in the solution due to this approximation. The macroscopic neutron cross-section data for the mono-energetic lattice calculation are presented in Table 3.4.

The quantity of interest (QoI) is the disadvantage factor and a reference solution is given as $\xi = 1.9998237529643$ which is taken from the following paper [175]. The disadvantage factor is the average scalar neutron flux in the moderator region divided by the average scalar neutron flux in the fuel region. A plot of the error in ξ for the NDE discretised with various orders of IGA and FE basis functions can be seen in Figure 3.15. The FE geometry has been volume preserved to maintain fissile mass. Therefore, error in the geometry is due to approximation of the surface area of the nuclear fuel pin. The scalar neutron flux has been converged to $\epsilon = 1 \times 10^{-11}$.

The errors displayed in Figure 3.15 can be used to see if the behaviour seen in the bare nuclear fuel pin test case will be observed in geometries where the misrepresentation of the external boundaries doesn't effect the implementation of the boundary conditions. It can still be observed that the error in



(a) Condition number against number of dof



(b) Condition number against number of elements

Figure 3.13: Condition number of the bare nuclear fuel pin plotted against number of dof and number of elements. It can be seen that κ is lower for the IGA system than the FE system when plotted against number of elements.




Figure 3.14: Pincell geometry for the mono-energetic lattice calculation. Reflective boundary conditions are prescribed on all external boundaries.

Macroscopic neutron cross-section	Value (cm^{-1})
σ_t^F	10.0
σ_t^M	1.0
σ_s^F	9.5
σ_s^M	0.95
Extraneous source strength	Value ($\text{cm}^{-3} \cdot \text{s}^{-1}$)
Q_F	1.0
Q_M	1.0

Table 3.4: Macroscopic neutron cross-sections for the mono-energetic lattice nuclear reactor physics benchmark verification test case. Superscripts F and M denote quantities in the fuel pin and moderator regions respectively




Figure 3.15: Error in the disadvantage factor for the mono-energetic lattice calculation.

the surface area of the nuclear fuel pin causes all of the FE solutions to have $O(h^2)$ convergence despite the nuclear fuel pin not being on an external boundary. Furthermore, it can be seen that with respect to the accuracy per dof, the high-order FE discretisations are inferior to the IGA discretisations. It can also be observed that the rate of convergence of the IGA solution for higher-order basis functions is now $O(h^{p+1})$ as is to be expected, even for $p = 5$, unlike in Figure 3.12 for the bare nuclear fuel pin case. For this problem, it is clear that a high-order IGA discretisation is an effective method for solving the NDE. In order to see how IGA compared to FE for problems with Cartesian geometries the next verification benchmark test case is considered.

3.6.3 ANL Two Dimensional Nuclear Reactor Physics Benchmark

This 2D x - y geometry nuclear reactor physics benchmark is based on a three-dimensional problem from the ANL benchmark problem book [207]. A two-dimensional version is presented in a paper by G. Theler and includes an axial buckling term $B_{z,g}^2 = 0.8 \times 10^{-4}$ [208]. The corrected macroscopic neutron cross-section data used in this thesis are displayed in Table 3.5 and the geometry is displayed in Figure 3.16. If a material property has not been mentioned then it is assumed to be zero. The ANL quarter core benchmark has Cartesian x - y geometry and can be exactly represented by IGA and FE. Therefore, the volume preservation applied to the FE mesh does not change the underlying geometry. The only difference in the solution methodology is the choice of basis functions. The IGA basis functions are C^{p-1} continuous over element boundaries whereas the FE basis functions are C^0 . The Quantity of Interest (QoI) for this problem is the K_{eff} and a reference value of $K_{\text{eff}} = 1.02958867$ was provided from a virtual element (VEM) code using quartic basis functions and with 629,651 dof per group [209]. The error in the K_{eff} generated from solving the NDE spatially discretised with IGA and FE is displayed in Figure 3.17.

It can clearly be seen that for all orders of basis function considered the value of K_{eff} generated from the IGA solution is superior to the value generated by the FE solution. To get a similar level of accuracy from the FE solution as the IGA solution about five times the number of dof are needed. The Cartesian geometry of the ANL quarter core test case is exactly represented by IGA and FE.




Figure 3.16: Geometry for the ANL two dimensional nuclear reactor physics benchmark testcase. The grey square regions correspond to material III.

Region	Macroscopic neutron cross-section (cm^{-1})								χ^1
	σ_t^1	σ_t^2	σ_a^1	σ_a^2	$\sigma_s^{1 \rightarrow 1}$	$\sigma_s^{1 \rightarrow 2}$	$\sigma_s^{2 \rightarrow 2}$	$\nu \sigma_f^2$	
I	$\frac{1}{4.5}$	$\frac{1}{1.2}$	0.01012	0.080032	0.19210222	0.2	0.7533013333	0.135	1.0
II	$\frac{1}{4.5}$	$\frac{1}{1.2}$	0.01012	0.085032	0.19210222	0.2	0.7483013333	0.135	1.0
III	$\frac{1}{4.5}$	$\frac{1}{1.2}$	0.01012	0.130032	0.19210222	0.2	0.7033013333	0.135	1.0
IV	$\frac{1}{6.0}$	$\frac{1}{0.9}$	1.6×10^{-4}	0.010024	0.126506667	0.04	1.101087111	0.0	0.0

Table 3.5: Macroscopic neutron cross-sections for the two dimensional ANL nuclear reactor physics benchmark verification test case. Group 1 represents fast neutrons and group 2 represents thermal neutrons.




Figure 3.17: Error in the K_{eff} for the 2D IAEA quarter core nuclear reactor physics benchmark verification test case. The NDE has been solved for this problem. It can be seen that IGA spatial discretisation is uniformly better than FE. This increase in accuracy per dof is due to the higher inter-element continuity of the IGA basis functions within a NURBS patch.

Therefore, this increase in accuracy is purely due to the increased inter-element continuity of the IGA basis functions within a NURBS patch.

The convergence rates for both the IGA and FE solutions are not of the expected order. Since these convergence rates are theoretical it is assumed that there are too many other factors that effect the convergence rates: material heterogeneity, scattering between groups, and power iterations to name a few potential factors. A similar convergence behaviour was noted by the VEM code used to generate the reference values of the K_{eff} . It is also noted that the convergence rates of the IGA discretisation and the FE discretisation match one another; so neither method has the advantage in terms of convergence rate.

3.6.4 OECD/NEA C5G7 Nuclear Reactor Physics Benchmark

The final problem considered in this section is the OECD/NEA seven group C5G7 quarter core nuclear reactor physics verification benchmark test case. The neutron transport solution of this problem will be considered in later chapters but here the NDE is solved. Several solution sets are tabulated in Table 3.6. Three sets of results are considered. Results generated using the code ICARUS, which is an IGA code capable of solving various second-order forms of the NTE. ICARUS used quadratic NURBS basis functions to spatially discretise the problem. Results from the IGA code PIRANA were taken from an IGA neutron diffusion paper [175] and are considered here as a reference solution. Finally, the results from the FE neutron diffusion code CRONOS2 were taken from the C5G7 benchmark specification [210]. CRONOS2 used quadratic, isoparametric, parabolic triangular elements to discretise the problem domain. All results from all codes are presented in Table 3.6.

Two important points are apparent from Table 3.6. First, the ICARUS result is converging to the PIRANA reference result as additional knots are inserted. This is expected behaviour as both are IGA based codes. Second, the error in the QoIs generated from the ICARUS solution with 65,532 dof


is less than or equal to the error in the CRONOS2 solution with 98,737 dof when compared with the PIRANA solution for most QoIs. The exceptions are the outer UO_2 power and the K_{eff} .

Code name	Number of degrees of freedom (dof)	K_{eff}	Maximum pin power	Minimum pin power	MOX power	Inner UO_2 power	Outer UO_2 power
PIRANA	315320	1.83238	2.5209	0.2406	210.82	495.45	138.90
CRONOS2	98737	1.83230	2.520	0.241	211.7	495.4	138.9
ICARUS	29105	1.83400	2.5157	0.2423	211.04	494.91	138.02
ICARUS	65532	1.83420	2.5213	0.2408	210.79	495.50	138.92
ICARUS	181136	1.83245	2.5209	0.2407	210.82	495.45	138.90


Table 3.6: Various QoI for several C5G7 solves. The results generated by PIRANA are taken from this paper [175] and the results generated by CRONOS2 are taken from the C5G7 benchmark specification [210]

The outer UO_2 value could be slightly more accurate for CRONOS2 for two reasons. First, the errors in the pin powers for all of the nuclear fuel pins in the outer UO_2 assembly are cancelling out and giving a seemingly improved solution. Second, it may be that the rounding that has occurred has made the CRONOS2 result appear improved compared to the ICARUS solution. If only four significant figures were considered for the ICARUS result then it would match exactly both values for the outer UO_2 power given by CRONOS2 and PIRANA.

The reason for the more accurate value of K_{eff} generated by CRONOS is thought to be due to the fact that the CRONOS2 mesh has more dof in the nuclear fuel pins than the corresponding ICARUS mesh. Therefore, when calculating the fission source, CRONOS2 has a slightly higher fidelity discretisation with respect to locations of the dof within the domain. The pincell mesh that was used for the CRONOS2 computation is available and is displayed in Figure 3.18. For comparison the locations of the control points within the pincell of the ICARUS 65,532 dof solution are also displayed. It can be seen that there are many more dof in the CRONOS2 pincell, in particular in the nuclear fuel pin.



(a) Location of nodes in CRONOS2 C5G7 fuel pin. There are 85 nodes.



(b) Location of control points in ICARUS C5G7 fuel pin. There are 52 control points.

Figure 3.18: Layout of dof within a C5G7 fuel pin for the FE code CRONOS2 and the IGA code ICARUS. Both codes use quadratic basis functions. CRONOS2 is using isoparametric parabolic triangular elements. Notice CRONOS2 has higher spatial fidelity than the ICARUS geometry.

3.7 Conclusion

In this chapter the concept of isogeometric analysis (IGA) has been introduced. In particular, the NURBS basis functions have been discussed in detail along with their historical development and a detailed discussion of their properties. Some of the mathematical details of IGA have been explained: mathematical spaces, quadrature, geometry generation algorithms, and refinement of the NURBS basis functions have been covered. In particular, it was stated that the refinement space of NURBS basis functions is much richer than that of standard Lagrangian or Hermitian polynomials that are typically used in FE. Knot insertion and degree refinement are both generalisations of h -refinement and p -refinement, and a new k -refinement strategy, which has no analogue in the classical FE literature was presented.

IGA was applied to the NDE and to several 2D nuclear reactor physics benchmarks test cases. It was shown that IGA can achieve the same theoretical $O(h^{p+1})$ convergence rates as the FE method for the same polynomial order. Furthermore, it was demonstrated that, for both Cartesian and non-Cartesian geometries, IGA is more accurate than the FE method per degree of freedom. This was also demonstrated for the seven group OECD/NEA C5G7 nuclear reactor physics benchmark by comparison with the neutron diffusion code CRONOS2.

Chapter 4

The Self Adjoint Angular Flux Equation (SAAF)

Much of the work in this chapter, including figures and tables, has been adapted from the publication by Latimer et al. [177].

Second-order forms of the NTE, such as the even-parity (EP) and self-adjoint angular flux (SAAF) equation, have been investigated since the mid 1960s [110, 137]. They have a number of numerical advantages over the first-order form of the NTE. One of the key numerical advantages of second-order forms is that they are numerically stable when spatially discretised using a continuous Bubnov-Galerkin discretisation using prescribed test and trial functions. Conversely, the first-order form of the neutron transport equation can produce spurious numerical artefacts, near sharp gradients in the solution, when discretised using the same approach [115]. Therefore, in general, the first-order form has to be discretised using either stabilised FE methods such as streamlined-upwind Petrov-Galerkin (SUPG) or discontinuous Galerkin methods [132, 133]. The use of discontinuous Galerkin methods can lead to significantly more degrees of freedom than the equivalent continuous Bubnov-Galerkin discretisation [134, 135].

A further advantage of second-order forms of the NTE (such as the SAAF and the EP equations) are that they produce symmetric positive definite (SPD) systems of linear equations when their weak form is spatially discretised using a Bubnov-Galerkin approach. SPD linear systems of equations can be solved efficiently with preconditioned conjugate gradient matrix solution algorithms [123]. Amongst the second-order forms of the NTE the most widely used are the EP, self-adjoint angular flux (SAAF), and symmetrised neutron transport (ST) [125] forms. Such forms of the NTE have been used extensively because they are self-adjoint equations and extremum variational principles may be derived for them [137] which allows for upper and lower bounds of the solution to be determined [139]. The EP form of the NTE allows for the efficient computation of the scalar neutron flux. The efficiency of the EP form is due to the number of angular unknowns being approximately half of that required by the SAAF or the first-order form of the NTE. Nevertheless, the additional cost of the SAAF solution provides the full angular neutron flux which, amongst other things, allows for easier implementation of S_N reflective boundary conditions [140].

A disadvantage of the aforementioned second-order forms is that they are unable to model regions containing voids whilst maintaining symmetry of the system and global conservation of neutrons.

The least-squares (LS) form of the NTE [211] developed by Hansen et. al. is symmetric but does not maintain conservation of neutrons except in the limit of the discretisation error vanishing. The hybrid SAAF & conservative LS (SAAF-CLS) formulation by Laboure et. al. [126] maintains conservation of neutrons and can model near-void regions, but is not symmetric. Void treatments for the SAAF form of the NTE have similar issues [124, 140]. Zheng et. al. presented a sub-domain discontinuous least-squares scheme (SDLS) that is globally continuous and forms a symmetric system of linear equations within each sub-domain. However, the sub-domain problems have to be linked together using transport sweeps or solved simultaneously as a non-symmetric global system of equations [212].

The SAAF equation was first derived from a variational principle by Pomraning and Clark [137]. However, this was only for 1D slab geometries where the material cross-sections are spatially constant. Later, Ackroyd derived a more general 3D maximum principle for the SAAF equation from his $\tilde{K}$ functional without any restriction on the spatial variation in the macroscopic neutron cross-sections [119]. Morel & McGhee published the first detailed theoretical and numerical study of the SAAF equation. Their paper presented a source iteration suitable formulation and also considered the relative advantages and disadvantages when compared to the first-order form and the EP & OP forms of the neutron transport equation [140]. In their paper they provided a simple algebraic derivation of the SAAF equation and spatially discretised it using a lumped FE approach. They utilised a discrete ordinate S_N angular discretisation of the SAAF equation and applied it to a 1D problem. They also discussed the issues surrounding purely scattering regions as well as voids for the SAAF equation. The formulation they developed was implemented within a unstructured FE, 3D neutron transport code called DANTE that has options for both S_N and P_N angular representations [213].

Liscum-Powell et. al. applied the discrete ordinates (S_N) form of the SAAF equation to coupled electron-photon transport and derived a consistent DSA scheme [214]. They utilised a continuous Bubnov-Galerkin spatial discretisation with linear FEs and applied this to 1D radiation transport problems. Fan et. al. derived an operator form of the SAAF equation angularly discretised using discrete ordinates (S_N) and with anisotropic scattering [215]. Ackroyd developed another variational derivation by considering relaxed forms of the SAAF equation based on trial functions and vectors for the reaction and leakage rates respectively, spatially discretised using FEs [120]. Morel et. al. presented the SAAF, EP, and OP equations applied to general radiative transfer [123]. The self-adjoint angular intensity (SAAI) equation was applied to two-dimensional (2D) problems that were spatially discretised using bi-linear FEs and a discrete ordinate (S_N) angular discretisation. Cao and Wu compared the spherical harmonics (P_N) and discrete ordinate (S_N) angular discretisation methods when applied to the SAAF equation over an unstructured FE mesh [216]. Wang & Gleicher investigated reflective boundary conditions for the SAAF equation proposing a different method that is symmetric but not positive definite [217]. Wang et. al. developed a consistent DSA scheme for the SAAF equation that included a void treatment [124]. Schunert et. al. investigated the physical and mathematical adjoint. The reason for this is that the SAAF equation is self-adjoint for the mono-energetic equation but in the multi-group case the scattering and fission operators are not self-adjoint. The SAAF- S_N equation was spatially discretised using a continuous Bubnov-Galerkin FE discretisation in the radiation transport code Rattle S_N ake [218]. Rattle S_N ake can use either a discrete ordinate (S_N) or a spherical harmonic (P_N) angular discretisation [219]. The radiation transport code Rattle S_N ake is a physics module within the multi-physics code MOOSE [15].

Contemporary unstructured mesh neutron transport calculations are typically spatially discretised using the FE method [15, 138, 220, 221]. In fact most of the unstructured mesh neutron transport literature is devoted to the development of FE based discretisation schemes. However, more recently high-order accurate spectral element (SE) [222, 223, 224, 225] and isogeometric analysis (IGA) methods have been developed for modelling neutron transport [163, 176, 175, 174, 186, 187]. The IGA method is a spatial discretisation methodology first introduced in 2005 by T. J. R Hughes [148]. In this chapter IGA will be used to discretise the SAAF neutron transport equation.

The advantages of IGA compared to the FE method are two-fold and have been illustrated in Chapter 3 for the NDE. First, there is no error due to approximation of geometry, with IGA, both surface area and volume are represented exactly and this is satisfied even on the coarsest mesh. In a nuclear reactor physics context this allows for the exact representation of volume (fissile mass) and surface area (neutron leakage effects) for geometries that can be represented by NURBS [163]. The preservation of volume and surface area are mutually exclusive when using isoparametric tetrahedral or hexahedral meshes for geometries described by NURBS (e.g. conic sections in 2D and quadric surfaces in 3D) [163, 176, 175]. Second, the NURBS basis functions used with IGA have high-order continuity within a NURBS patch. Conversely, FE basis functions are at most C^0 continuous between FEs. The higher-order continuity within a NURBS patch leads to a higher fidelity solution when representing smooth solutions within a NURBS patch. This is particularly useful when solving elliptic partial differential equations (PDEs) which are often characterised by smooth solutions. For some problems, due to the discontinuous spatial variation in material properties, the solution may exhibit sharp gradients. However, this is usually known ahead of time and the inter-element continuity of the basis functions can be adjusted.

In section 4.1 the SAAF equation is derived in two ways. Firstly the well known source iteration compatible algebraic derivation from Morel & McGhee's paper [140] is repeated for completeness. Then a novel source iteration compatible variational derivation is performed in a manner similar to Ackroyd's work [119]. This derivation gives a theoretical grounding to the incoming boundary condition used in the algebraic derivation. These two derivations are then shown to be equivalent. Following this section 4.2 contains details of the IGA spatial and S_N angular discretisations applied to the SAAF equation. Finally numerical results are given in section 4.3.

4.1 Algebraic Derivation of the Self-Adjoint Angular Flux NTE and its Associated Weak and Variational Forms

The steady-state (time-independent), first-order, energy-dependent, NTE is given by [19]:

$$\boldsymbol{\Omega} \cdot \nabla \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) + \sigma_t(\mathbf{r}, E) \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) = Q(\mathbf{r}, \boldsymbol{\Omega}, E), \quad (4.1)$$

where

$$Q(\mathbf{r}, \boldsymbol{\Omega}, E) = Q_{\text{scatter}}(\mathbf{r}, \boldsymbol{\Omega}, E) + Q_{\text{ext}}(\mathbf{r}, E), \quad (4.2)$$

and

$$Q_{\text{scatter}}(\mathbf{r}, \boldsymbol{\Omega}, E) = \int_0^\infty \int_{4\pi} \sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}, E' \rightarrow E) \psi_o(\mathbf{r}, \boldsymbol{\Omega}', E') d\boldsymbol{\Omega}' dE'$$

is the source of neutrons due to scatter into energy E , $\psi_o(\mathbf{r}, \boldsymbol{\Omega}, E)$ is the angular neutron flux $\sigma_t(\mathbf{r}, E)$ is the total macroscopic neutron cross-section, $\sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}, E' \rightarrow E)$ is the differential scattering macroscopic neutron cross-section, and $Q_{\text{ext}}(\mathbf{r}, E)$ is a source composed of fission sources and fixed (extraneous) sources.

The vacuum and incoming surface source conditions are assumed to be prescribed for the incoming boundary condition:

$$\psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) = 0 \quad \text{for } \boldsymbol{\Omega} \cdot \mathbf{n} < 0 \text{ and } \mathbf{r} \in S_b, \quad (4.3)$$

$$\psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) = T(\mathbf{r}, \boldsymbol{\Omega}, E) \quad \text{for } \boldsymbol{\Omega} \cdot \mathbf{n} < 0 \text{ and } \mathbf{r} \in S_s, \quad (4.4)$$

where $S_b \cup S_s = S$ forms the boundary of the computational domain. It should be noted that a reflective boundary condition is a special case of an incoming source with

$$T(\mathbf{r}, \boldsymbol{\Omega}, E) = \psi_o(\mathbf{r}, \boldsymbol{\Omega}', E) \quad \text{for } \mathbf{r} \in S_r \text{ and } \boldsymbol{\Omega} \cdot \mathbf{n} < 0,$$

where $\boldsymbol{\Omega}'$ is the angle of specular reflection for $\boldsymbol{\Omega}$. Together these three boundary conditions highlight the general form of the boundary conditions used in neutron transport theory [138].

What follows are two derivations of the second-order SAAF equation. Firstly an algebraic source iteration compatible derivation taken from work by Morel & McGhee [140] and secondly a maximum variational principle is derived that is shown to yield the source iteration compatible SAAF equation along with natural boundary conditions for the incoming and outgoing angles.

4.1.1 Algebraic Derivation of the SAAF Form of the NTE

Rearranging equation (4.1) for the angular flux gives:

$$\psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) = \frac{1}{\sigma_t(\mathbf{r}, E)} [Q(\mathbf{r}, \boldsymbol{\Omega}, E) - \boldsymbol{\Omega} \cdot \nabla \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E)], \quad (4.5)$$

which is then substituted back into the first term of the first-order form of the NTE yielding:

$$-\boldsymbol{\Omega} \cdot \nabla \frac{1}{\sigma_t(\mathbf{r}, E)} \boldsymbol{\Omega} \cdot \nabla \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) + \sigma_t(\mathbf{r}, E) \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) = Q(\mathbf{r}, \boldsymbol{\Omega}, E) - \left(\boldsymbol{\Omega} \cdot \nabla \frac{Q(\mathbf{r}, \boldsymbol{\Omega}, E)}{\sigma_t(\mathbf{r}, E)} \right). \quad (4.6)$$

Implicit in the rearrangement step is the requirement that $\sigma_t \neq 0$. This means that in the above form the SAAF equation is not compatible with problems containing void regions. As is noted by Morel & McGhee [140], this formulation of the SAAF equation is suitable for source iteration due to the fact that the operator on the left hand side is self-adjoint in the energy dependent case. This also obviates any issues singularities in regions of pure scattering as there is no $(\sigma_t - \sigma_s)$ term in the denominator as is possible in certain forms of the SAAF equation.

The incoming boundary conditions for SAAF can also be derived algebraically from the first-order boundary conditions:

$$\psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) = \frac{1}{\sigma_t(\mathbf{r}, E)} (Q(\mathbf{r}, \boldsymbol{\Omega}, E) - \boldsymbol{\Omega} \cdot \nabla \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E)) = T(\mathbf{r}, \boldsymbol{\Omega}, E) \quad \text{for } \mathbf{r} \in S_s \quad (4.7)$$

$$\psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) = \frac{1}{\sigma_t(\mathbf{r}, E)} (Q(\mathbf{r}, \boldsymbol{\Omega}, E) - \boldsymbol{\Omega} \cdot \nabla \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E)) = 0 \quad \text{for } \mathbf{r} \in S_b. \quad (4.8)$$

In addition, due to SAAF being second-order, extra boundary conditions for the outgoing flux must be prescribed. Morel & McGhee [140] suggest that the outgoing flux on the boundary should satisfy the first-order form of the NTE, that is:

$$\psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) = \frac{1}{\sigma_t(\mathbf{r}, E)} (Q(\mathbf{r}, \boldsymbol{\Omega}, E) - \boldsymbol{\Omega} \cdot \nabla \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E)) \quad \text{for } \mathbf{r} \in S \quad (4.9)$$

claiming that it is “a natural way to ensure that spurious solutions to the SAAF equation are eliminated.” It will be seen in the next section that this boundary condition can be derived from a maximum variational principle.

4.1.2 The Weak Form of the SAAF NTE

Equation (4.6) is the strong form of the SAAF equation. It is assumed that an approximation to ψ_o must be found. In order to form a system of linear equations that represent equation (4.6) and allow the computation of an approximate solution, a weak form must be generated by using a weighted residual (WR) method.

We consider a weighting function $w(\mathbf{r})$ from the space of admissible functions. A function u is an admissible function if u and $\boldsymbol{\Omega} \cdot \nabla u$ are continuous functions of $\mathbf{r}$ and $\boldsymbol{\Omega}$ almost everywhere within a given patch [82]. No assumptions are made about the boundary conditions that an admissible function must satisfy.

The strong form equation is multiplied by a weighting function and integration is performed over the entire spatial domain:

$$\int_V \left[w(\mathbf{r}) \boldsymbol{\Omega} \cdot \nabla \frac{1}{\sigma_t(\mathbf{r}, E)} \left(-\boldsymbol{\Omega} \cdot \nabla \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) + Q(\mathbf{r}, \boldsymbol{\Omega}, E) \right) + w(\mathbf{r}) \left(\sigma_t(\mathbf{r}, E) \psi_o(\mathbf{r}, \boldsymbol{\Omega}, E) - Q(\mathbf{r}, \boldsymbol{\Omega}, E) \right) \right] d\mathbf{r} = 0. \quad (4.10)$$

Suppressing explicit dependencies for brevity and applying Green’s first identity to the first term in the spatial integral yields:

$$\int_V \left[\left(\frac{\boldsymbol{\Omega} \cdot \nabla w}{\sigma_t} \right) (\boldsymbol{\Omega} \cdot \nabla \psi_o - Q) + w (\sigma_t \psi_o - Q) \right] d\mathbf{r} + \int_S (\boldsymbol{\Omega} \cdot \mathbf{n}) \frac{w}{\sigma_t} (-\boldsymbol{\Omega} \cdot \nabla \psi_o + Q) d\mathbf{r}_s = 0 \quad (4.11)$$

where $\mathbf{r}_s$ refers to points on surface S . The above equation is the weak form of the SAAF equation and the weighting (test) function can be chosen in order to form either a Bubnov-Galerkin or Petrov-Galerkin discretisation of the weak form.

4.1.3 Source Iteration Compatible Variational Formulation of the SAAF NTE

This variational derivation is based on work done by Ackroyd in two papers [119, 82] and the following notation is adopted:

$$\langle u, v \rangle = \int_{4\pi} uv d\boldsymbol{\Omega}. \quad (4.12)$$

It is also noted that:

$$\int_{4\pi} (\mathbf{\Omega} \cdot \mathbf{n}) f(\mathbf{\Omega}) d\mathbf{\Omega} = \int_{\mathbf{\Omega} \cdot \mathbf{n} > 0} |\mathbf{\Omega} \cdot \mathbf{n}| f(\mathbf{\Omega}) d\mathbf{\Omega} - \int_{\mathbf{\Omega} \cdot \mathbf{n} < 0} |\mathbf{\Omega} \cdot \mathbf{n}| f(\mathbf{\Omega}) d\mathbf{\Omega}. \quad (4.13)$$

where $f(\mathbf{\Omega})$ is some function that depends on the solid angle $\mathbf{\Omega}$.

In the original derivation by Ackroyd the following operator is defined:

$$\Sigma f(\mathbf{r}, \mathbf{\Omega}) = \sigma_t(\mathbf{r}) f(\mathbf{r}, \mathbf{\Omega}) - \int_{4\pi} \sigma_s(\mathbf{r}, \mathbf{\Omega}' \cdot \mathbf{\Omega}) f(\mathbf{r}, \mathbf{\Omega}') d\mathbf{\Omega}'. \quad (4.14)$$

and it can be shown that

1. Σ has an inverse,
2. Σ and Σ^{-1} are self-adjoint,
3. Σ and Σ^{-1} are positive definite.

These results are proven in Appendix A of Ackroyd's work [119].

In this derivation, in order to ensure that the variational problem will be compatible with source iteration the scattering kernel is considered as part of the source and a new choice of Σ is made:

$$\Sigma(\mathbf{r}) f(\mathbf{r}, \mathbf{\Omega}) = \sigma_t(\mathbf{r}) f(\mathbf{r}, \mathbf{\Omega}). \quad (4.15)$$

This choice of Σ clearly satisfies the conditions laid out above with the inverse:

$$\Sigma^{-1}(\mathbf{r}) = \frac{1}{\sigma_t(\mathbf{r})},$$

except where $\sigma_t = 0$ as then the inverse is not defined.

Then the following functional is defined:

$$F(u, v) = \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} \mathbf{\Omega} \cdot \nabla v \rangle + \langle u, \Sigma v \rangle] d\mathbf{r} + \int_S \int_{\Omega} |\mathbf{\Omega} \cdot \mathbf{n}| u v d\mathbf{\Omega} dS \quad (4.16)$$

where u, v are admissible functions and the list of dependent variables has been omitted for brevity. The choice of $F(u, v)$ is based on the following observations:

1. Σ^{-1} operates on $\mathbf{\Omega} \cdot \nabla \psi_o$ in equation (4.6),
2. Σ operates on ψ_o in equation (4.6),
3. All terms in $F(u, u)$ are strictly positive for all u (except at $u = 0$ where $F(0, 0) = 0$) as Σ and Σ^{-1} are positive definite.

It is noted that:

$$F(u, \psi_o) = \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} \mathbf{\Omega} \cdot \nabla \psi_o \rangle + \langle u, \Sigma \psi_o \rangle] d\mathbf{r} + \int_S \int_{\Omega} |\mathbf{\Omega} \cdot \mathbf{n}| u \psi_o d\mathbf{\Omega} dS, \quad (4.17)$$

where ψ_o is the exact solution of equation (4.1). The volume integral in equation (4.17) can be

simplified giving:

$$\begin{aligned}
& \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} \mathbf{\Omega} \cdot \nabla \psi_o \rangle + \langle u, \Sigma \psi_o \rangle] d\mathbf{r} \\
&= \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} Q - \psi_o \rangle + \langle u, Q - \mathbf{\Omega} \cdot \nabla \psi_o \rangle] d\mathbf{r} \\
&= \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} Q \rangle - \langle \mathbf{\Omega} \cdot \nabla u, \psi_o \rangle + \langle u, Q \rangle - \langle u, \mathbf{\Omega} \cdot \nabla \psi_o \rangle] d\mathbf{r}.
\end{aligned}$$

The identity $\mathbf{\Omega} \cdot \nabla(u\psi_o) = (\mathbf{\Omega} \cdot \nabla u)\psi_o + u(\mathbf{\Omega} \cdot \nabla \psi_o)$ along with the divergence theorem is used:

$$\begin{aligned}
& \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} Q \rangle + \langle u, Q \rangle] d\mathbf{r} - \int_V \int_{\Omega} \mathbf{\Omega} \cdot \nabla(u\psi_o) d\mathbf{\Omega} d\mathbf{r} \\
&= \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} Q \rangle + \langle u, Q \rangle] d\mathbf{r} - \int_S \int_{\Omega} (\mathbf{\Omega} \cdot \mathbf{n}) u \psi_o d\mathbf{\Omega} dS \\
&= \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} Q \rangle + \langle u, Q \rangle] d\mathbf{r} - \int_S \int_{\mathbf{\Omega} \cdot \mathbf{n} > 0} |\mathbf{\Omega} \cdot \mathbf{n}| u \psi_o d\mathbf{\Omega} dS + \int_S \int_{\mathbf{\Omega} \cdot \mathbf{n} < 0} |\mathbf{\Omega} \cdot \mathbf{n}| u \psi_o d\mathbf{\Omega} dS.
\end{aligned}$$

where the integral over angle has been split into two half range integrals according to equation (4.13).

Therefore, we have that

$$F(u, \psi_o) = \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} Q \rangle + \langle u, Q \rangle] d\mathbf{r} + 2 \int_S \int_{\mathbf{\Omega} \cdot \mathbf{n} < 0} |\mathbf{\Omega} \cdot \mathbf{n}| u \psi_o d\mathbf{\Omega} dS \quad (4.18)$$

is independent of ψ_o once boundary conditions have been prescribed. Here the boundary conditions, given by equations (4.7) & (4.8), for ψ_o can be prescribed yielding the following expression:

$$F(u, \psi_o) = \int_V [\langle \mathbf{\Omega} \cdot \nabla u, \Sigma^{-1} Q \rangle + \langle u, Q \rangle] d\mathbf{r} + 2 \int_{S_s} \int_{\mathbf{\Omega} \cdot \mathbf{n} < 0} |\mathbf{\Omega} \cdot \mathbf{n}| u T d\mathbf{\Omega} dS. \quad (4.19)$$

An admissible function ϕ is considered where ϕ is a variation from the exact solution ψ_o . Applying F to the error gives:

$$F(\psi_o - \lambda\phi, \psi_o - \lambda\phi) = F(\psi_o, \psi_o) - 2\lambda F(\phi, \psi_o) + \lambda^2 F(\phi, \phi) \geq 0, \quad (4.20)$$

which implies

$$2\lambda F(\phi, \psi_o) - \lambda^2 F(\phi, \phi) \leq F(\psi_o, \psi_o). \quad (4.21)$$

This provides a mechanism for generating maximum principles for the NTE, depending on λ , as equality of equation (4.21) is only seen when $(\psi_o - \lambda\phi) \rightarrow 0$. Here $\lambda = 1$ is chosen and the maximum principle $K(\phi) = 2F(\phi, \psi_o) - F(\phi, \phi)$ is defined where:

$$\begin{aligned}
K(\phi) &= \int_V [2\langle \phi, Q \rangle + 2\langle \mathbf{\Omega} \cdot \nabla \phi, \Sigma^{-1} Q \rangle - \langle \mathbf{\Omega} \cdot \nabla \phi, \Sigma^{-1} \mathbf{\Omega} \cdot \nabla \phi \rangle - \langle \phi, \Sigma \phi \rangle] d\mathbf{r} \\
&+ 4 \int_{S_s} \int_{\mathbf{\Omega} \cdot \mathbf{n} < 0} |\mathbf{\Omega} \cdot \mathbf{n}| \phi T d\mathbf{\Omega} dS - \int_S \int_{\Omega} |\mathbf{\Omega} \cdot \mathbf{n}| \phi^2 d\mathbf{\Omega} dS
\end{aligned} \quad (4.22)$$

and $K(\phi)$ is independent of the exact solution ψ_o .

In order to find extrema the first variation of $K(\phi)$ must be considered, in particular where it is

stationary. A variation of the exact solution is considered $\phi = \psi_o + \epsilon\xi$ and $K(\psi_o + \epsilon\xi)$ is evaluated:

$$\begin{aligned} K(\psi_o + \epsilon\xi) &\equiv I(\epsilon) = \int_V [2 \langle \psi_o + \epsilon\xi, Q \rangle + 2 \langle \mathbf{\Omega} \cdot \nabla (\psi_o + \epsilon\xi), \Sigma^{-1} Q \rangle \\ &- \langle \mathbf{\Omega} \cdot \nabla (\psi_o + \epsilon\xi), \Sigma^{-1} \mathbf{\Omega} \cdot \nabla (\psi_o + \epsilon\xi) \rangle - \langle (\psi_o + \epsilon\xi), \Sigma (\psi_o + \epsilon\xi) \rangle] d\mathbf{r} \\ &+ 4 \int_{S_s} \int_{\mathbf{\Omega} \cdot \mathbf{n} < 0} |\mathbf{\Omega} \cdot \mathbf{n}| (\psi_o + \epsilon\xi) T d\mathbf{\Omega} dS - \int_S \int_{\Omega} |\mathbf{\Omega} \cdot \mathbf{n}| (\psi_o + \epsilon\xi)^2 d\mathbf{\Omega} dS. \end{aligned} \quad (4.23)$$

To find stationary points the following equality is considered

$$\left(\frac{dI(\epsilon)}{d\epsilon} \right)_{\epsilon=0} = 0, \quad (4.24)$$

giving

$$\begin{aligned} 0 &= \int_V \left[\int_{\Omega} \left(2\xi Q + 2\mathbf{\Omega} \cdot \nabla \xi (\Sigma^{-1} Q) - (\mathbf{\Omega} \cdot \nabla \psi_o) (\Sigma^{-1} \mathbf{\Omega} \cdot \nabla \xi) - (\mathbf{\Omega} \cdot \nabla \xi) (\Sigma^{-1} \mathbf{\Omega} \cdot \nabla \psi_o) \right. \right. \\ &\quad \left. \left. - \psi_o \Sigma \xi - \xi \Sigma \psi_o \right) d\mathbf{\Omega} \right] d\mathbf{r} + 4 \int_{S_s} \int_{\mathbf{\Omega} \cdot \mathbf{n} < 0} |\mathbf{\Omega} \cdot \mathbf{n}| \xi T d\mathbf{\Omega} dS - 2 \int_S \int_{\Omega} |\mathbf{\Omega} \cdot \mathbf{n}| \psi_o \xi d\mathbf{\Omega} dS. \end{aligned} \quad (4.25)$$

The volume integral terms can be simplified:

$$\begin{aligned} &2 \int_V \left[\int_{\Omega} (\xi Q + \mathbf{\Omega} \cdot \nabla \xi (\Sigma^{-1} Q) - (\mathbf{\Omega} \cdot \nabla \xi) (\Sigma^{-1} \mathbf{\Omega} \cdot \nabla \psi_o) - \xi \Sigma \psi_o) d\mathbf{\Omega} \right] d\mathbf{r} \\ &= 2 \int_V \int_{\Omega} (Q - \Sigma \psi_o) \xi + [\Sigma^{-1} (Q - \mathbf{\Omega} \cdot \nabla \psi_o)] (\mathbf{\Omega} \cdot \nabla \xi) d\mathbf{\Omega} d\mathbf{r}, \end{aligned}$$

with the chain rule and the divergence theorem being used on the second term:

$$\begin{aligned} &\int_V \Sigma^{-1} (Q - \mathbf{\Omega} \cdot \nabla \psi_o) (\mathbf{\Omega} \cdot \nabla \xi) d\mathbf{r} \\ &= \int_V [\nabla \cdot (\mathbf{\Omega} \xi \Sigma^{-1} (Q - \mathbf{\Omega} \cdot \nabla \psi_o)) - \xi \mathbf{\Omega} \cdot \nabla \Sigma^{-1} (Q - \mathbf{\Omega} \cdot \nabla \psi_o)] d\mathbf{r} \\ &= \int_S \int_{\Omega} \mathbf{\Omega} \cdot \mathbf{n} \xi \Sigma^{-1} (Q - \mathbf{\Omega} \cdot \nabla \psi_o) d\mathbf{\Omega} dS - \int_V \xi \mathbf{\Omega} \cdot \nabla \Sigma^{-1} (Q - \mathbf{\Omega} \cdot \nabla \psi_o) d\mathbf{r}. \end{aligned}$$

The surface integral can be written as:

$$\begin{aligned} &4 \int_{S_s} \int_{\Omega} |\mathbf{\Omega} \cdot \mathbf{n}| \xi T d\mathbf{\Omega} dS - 2 \int_S \int_{\Omega} |\mathbf{\Omega} \cdot \mathbf{n}| \psi_o \xi d\mathbf{\Omega} dS \\ &= 2 \int_{S_s} \int_{\Omega} |\mathbf{\Omega} \cdot \mathbf{n}| \xi (2T - \psi_o) d\mathbf{\Omega} dS - 2 \int_{S_b} \int_{\Omega} |\mathbf{\Omega} \cdot \mathbf{n}| \xi \psi_o d\mathbf{\Omega} dS, \end{aligned}$$

where the range of the angular integral over S_S has been increased to all of $\mathbf{\Omega}$ by noting that $T(\mathbf{\Omega}) = 0$ for $\mathbf{\Omega} \cdot \mathbf{n} > 0$. Recombining the volume and surface terms yields:

$$\begin{aligned} 0 &= \int_V \int_{\Omega} \xi [Q - \Sigma \psi_o - \mathbf{\Omega} \cdot \nabla \Sigma^{-1} (Q - \mathbf{\Omega} \cdot \nabla \psi_o)] d\mathbf{\Omega} d\mathbf{r} + \int_{S_S} \int_{\Omega} \xi [|\mathbf{\Omega} \cdot \mathbf{n}| (2T - \psi_o) \\ &+ \mathbf{\Omega} \cdot \mathbf{n} \Sigma^{-1} (Q - \mathbf{\Omega} \cdot \nabla \psi_o)] d\mathbf{\Omega} dS - \int_{S_b} \int_{\Omega} \xi [|\mathbf{\Omega} \cdot \mathbf{n}| \psi_o - \mathbf{\Omega} \cdot \mathbf{n} \Sigma^{-1} (Q - \mathbf{\Omega} \cdot \nabla \psi_o)] d\mathbf{\Omega} dS, \end{aligned}$$

and applying the fundamental lemma of the calculus of variations yields:

$$\begin{aligned}
0 = & \int_V \int_\Omega [Q - \Sigma\psi_o - \boldsymbol{\Omega} \cdot \nabla \Sigma^{-1} (Q - \boldsymbol{\Omega} \cdot \nabla \psi_o)] d\boldsymbol{\Omega} d\mathbf{r} + \int_{S_S} \int_\Omega [|\boldsymbol{\Omega} \cdot \mathbf{n}|(2T - \psi_o) \\
& + \boldsymbol{\Omega} \cdot \mathbf{n} \Sigma^{-1} (Q - \boldsymbol{\Omega} \cdot \nabla \psi_o)] d\boldsymbol{\Omega} dS - \int_{S_b} \int_\Omega [|\boldsymbol{\Omega} \cdot \mathbf{n}| \psi_o - \boldsymbol{\Omega} \cdot \mathbf{n} \Sigma^{-1} (Q - \boldsymbol{\Omega} \cdot \nabla \psi_o)] d\boldsymbol{\Omega} dS.
\end{aligned} \tag{4.26}$$

The SAAF equation along with its boundary conditions can be recovered yielding the following equations:

$$-\boldsymbol{\Omega} \cdot \nabla (\Sigma^{-1} \boldsymbol{\Omega} \cdot \nabla \psi_o) + \Sigma\psi_o = Q - \boldsymbol{\Omega} \cdot \nabla \Sigma^{-1} Q \tag{4.27}$$

$$|\boldsymbol{\Omega} \cdot \mathbf{n}|(2T - \psi_o) + \boldsymbol{\Omega} \cdot \mathbf{n} \Sigma^{-1} (Q - \boldsymbol{\Omega} \cdot \nabla \psi_o) = 0 \quad \text{on } S_s \tag{4.28}$$

$$|\boldsymbol{\Omega} \cdot \mathbf{n}| \psi_o - \boldsymbol{\Omega} \cdot \mathbf{n} \Sigma^{-1} (Q - \boldsymbol{\Omega} \cdot \nabla \psi_o) = 0 \quad \text{on } S_b \tag{4.29}$$

in particular, the boundary conditions can be decomposed into expressions for incoming and outgoing directions to the boundary as follows:

$\boldsymbol{\Omega} \cdot \mathbf{n} < 0$:

$$2T - \psi_o - \Sigma^{-1} (Q - \boldsymbol{\Omega} \cdot \nabla \psi_o) = 0 \Rightarrow \psi_o = T \quad \text{for } \mathbf{r} \in S_s, \tag{4.30}$$

$$\psi_o + \Sigma^{-1} (Q - \boldsymbol{\Omega} \cdot \nabla \psi_o) = 0 \Rightarrow \psi_o = 0 \quad \text{for } \mathbf{r} \in S_b, \tag{4.31}$$

$\boldsymbol{\Omega} \cdot \mathbf{n} > 0$:

$$\psi_o - \Sigma^{-1} (Q - \boldsymbol{\Omega} \cdot \nabla \psi_o) = 0 \Rightarrow \psi_o = \Sigma^{-1} (-\boldsymbol{\Omega} \cdot \nabla \psi_o + Q) \quad \text{for } \mathbf{r} \in S. \tag{4.32}$$

Not only do the incoming boundary conditions match the boundary conditions for the first-order form of the NTE, but the outgoing boundary condition postulated by Morel & McGhee is found as a natural consequence of the maximum principle. This variational derivation will not be valid for void regions due to the choice of the operator Σ . In the case of a void the inverse of the operator Σ will be singular. In addition as this variational principle is not valid for void regions the outgoing boundary condition is also not valid for void regions.

4.1.4 Equivalence of Algebraic and Variational Derivation of the SAAF Form of the NTE

The weak form from the algebraic derivation is given by equation (4.11). Angular integration is performed over the surface and is split into incoming and outgoing angles according to equation (4.13):

$$\begin{aligned}
& \int_{4\pi} \int_V \left[\left(\frac{\boldsymbol{\Omega} \cdot \nabla w}{\sigma_t} \right) (\boldsymbol{\Omega} \cdot \nabla \psi - Q) + w (\sigma_t \psi - Q) \right] d\mathbf{r} d\boldsymbol{\Omega} \\
& + \int_{\boldsymbol{\Omega} \cdot \mathbf{n} > 0} \int_S \left[|\boldsymbol{\Omega} \cdot \mathbf{n}| \frac{w}{\sigma_t} (-\boldsymbol{\Omega} \cdot \nabla \psi + Q) \right] d\mathbf{r}_s d\boldsymbol{\Omega} \\
& - \int_{\boldsymbol{\Omega} \cdot \mathbf{n} < 0} \int_S |\boldsymbol{\Omega} \cdot \mathbf{n}| \frac{w}{\sigma_t} (-\boldsymbol{\Omega} \cdot \nabla \psi + Q) d\mathbf{r}_s d\boldsymbol{\Omega} = 0.
\end{aligned} \tag{4.33}$$

From the variational derivation we consider equation (4.18) substituted into equation (4.21) without applying the boundary conditions that are given for ψ_o . By following through with the rest of the

derivation and considering the zeroing of the first variation (see equation 4.24) an analogous expression to equation (4.25) is derived, that is:

$$0 = \int_V \left[\int_{\Omega} \xi Q + \boldsymbol{\Omega} \cdot \nabla \xi (\Sigma^{-1} Q) - (\boldsymbol{\Omega} \cdot \nabla \psi_o) (\Sigma^{-1} \boldsymbol{\Omega} \cdot \nabla \xi) - \xi \Sigma \psi_o d\boldsymbol{\Omega} \right] d\mathbf{r} \\ + 2 \int_S \int_{\boldsymbol{\Omega} \cdot \mathbf{n} < 0} |\boldsymbol{\Omega} \cdot \mathbf{n}| \xi \psi_o d\boldsymbol{\Omega} dS - \int_S \int_{\Omega} |\boldsymbol{\Omega} \cdot \mathbf{n}| \psi_o \xi d\boldsymbol{\Omega} dS. \quad (4.34)$$

Simplifying the equation and multiplying by -1 we are left with:

$$0 = \int_V \left[\int_{\Omega} (\boldsymbol{\Omega} \cdot \nabla \xi) \Sigma^{-1} (\boldsymbol{\Omega} \cdot \nabla \psi_o - Q) + \xi (\Sigma \psi_o - Q) d\boldsymbol{\Omega} \right] d\mathbf{r} \\ - \int_S \int_{\boldsymbol{\Omega} \cdot \mathbf{n} < 0} |\boldsymbol{\Omega} \cdot \mathbf{n}| \xi \psi_o d\boldsymbol{\Omega} dS + \int_S \int_{\Omega} |\boldsymbol{\Omega} \cdot \mathbf{n}| \psi_o \xi d\boldsymbol{\Omega} dS. \quad (4.35)$$

Using the fact that $\psi_o = \frac{1}{\sigma_t} (-\boldsymbol{\Omega} \cdot \nabla \psi_o + Q)$, replacing ξ with w , and ψ_o with ψ it can be seen that the variational principle yields the same weak form as the algebraic derivation as is expected of symmetric operators [226].

4.2 Spatial and Angular Discretisation of the Weak Form of the SAAF NTE

4.2.1 Bubnov-Galerkin IGA Approximation

Similarly to FE methods the solution is assumed to be a linear combination of trial functions:

$$\psi(\mathbf{r}, \boldsymbol{\Omega}, E) = \sum_{i=1}^{\infty} d_i(\boldsymbol{\Omega}, E) v_i(\mathbf{r}), \quad (4.36)$$

where the $d_i \in \mathbb{R}$ are referred to as the control variables and $v_i \in T_R$ are the trial functions. The infinite dimensional trial function space is defined as:

$$T_R = \{ v \mid v \in H^1(V), v|_{S_D} = g \}. \quad (4.37)$$

Furthermore, w is taken to be a member of an infinite dimensional test function space:

$$T_E = \{ w \mid w \in H^1(V), w|_S = 0 \}, \quad (4.38)$$

where $H^1(V)$ is the Sobolev space $W_2^1(V)$ over the domain V [227]. The condition $v|_{S_D} = g$ comes from a Dirichlet boundary condition that may be given for $\psi(\mathbf{r}) = g$ for $\mathbf{r} \in S_D$.

In order to provide closure to equation (4.36) the expansion must be truncated. To do this finite dimensional analogues of T_R and T_E are defined:

$$T_R^h = \left\{ v_i^h \mid v_i^h \in H^1(V), v_i^h|_{S_D} = g, \text{ for } i \in [1, N] \right\}, \\ T_E^h = \left\{ w_i^h \mid w_i^h \in H^1(V), w_i^h|_S = 0, \text{ for } i \in [1, N] \right\},$$

and the solution is approximated as:

$$\psi(\mathbf{r}, \boldsymbol{\Omega}, E) \approx \psi^h(\mathbf{r}, \boldsymbol{\Omega}, E) = \sum_{i=1}^N d_i(\boldsymbol{\Omega}, E) v_i^h(\mathbf{r}). \quad (4.39)$$

A Bubnov-Galerkin IGA discretisation is then applied by making the choices $v_i(\mathbf{r}) = R_i^p(\mathbf{r}) \in T_R^h$ and $w_j(\mathbf{r}) = R_j^p(\mathbf{r}) \in T_E^h$. This encapsulates the essence of IGA, to use a basis capable of exactly representing the known geometry as a basis for the fields we wish to approximate [27].

4.2.2 Angular ($\boldsymbol{\Omega}$) Discretisation of the SAAF NTE

The method of discrete ordinates (S_N) is chosen to discretise the angular domain. Integration with respect to $\boldsymbol{\Omega}$ is performed via a summation of functions evaluated at certain angles and scaled by certain weights. These angles and weights are referred to as the quadrature set. For example, the scalar neutron flux ϕ can be calculated from the angular neutron flux ψ in the following way:

$$\phi(\mathbf{r}, E) = \int_{4\pi} \psi(\mathbf{r}, \boldsymbol{\Omega}, E) d\boldsymbol{\Omega} \approx \sum_{m=1}^M w_m \psi(\mathbf{r}, \boldsymbol{\Omega}_m, E) = \sum_{m=1}^M w_m \psi_m(\mathbf{r}, E). \quad (4.40)$$

By solving equation (4.6) over the directions $\{\boldsymbol{\Omega}_m\}_{m=1}^M$ the scalar flux can be recovered using the above quadrature rule. Therefore, the weak form of the SAAF equation given by equation (4.11) evaluated at $\{\boldsymbol{\Omega}_m\}_{m=1}^M$ will be considered:

$$\begin{aligned} \int_V \left[\left(\frac{\boldsymbol{\Omega}_m \cdot \nabla w}{\sigma_t} \right) (\boldsymbol{\Omega}_m \cdot \nabla \psi_m - Q_m) + w (\sigma_t \psi_m - Q_m) \right] d\mathbf{r} + \int_S (\boldsymbol{\Omega}_m \cdot \mathbf{n}) \frac{w}{\sigma_t} (-\boldsymbol{\Omega}_m \cdot \nabla \psi_o + Q_m) d\mathbf{r}_s \\ = 0, \quad \text{for } m = 1, \dots, M. \end{aligned} \quad (4.41)$$

The relevant boundary conditions are found by evaluating (4.30)-(4.32) for $\boldsymbol{\Omega} = \boldsymbol{\Omega}_m$. The above equation (4.41) is referred to as the S_N -SAAF equation. As equation (4.40) is a general quadrature formula, a quadrature set that can adequately represent the angular integrals must be chosen. In this work level-symmetric and Legendre-Chebyshev quadrature sets are used, in particular triangular ones where the order of the polar quadrature order is set to two. This yields a quadrature set that gives the same number of directions as the level-symmetric set but is not identical.

4.2.3 IGA Spatial Discretisation of the Weak Form of the SAAF NTE

In order to spatially discretise equation (4.41), equation (4.39), the choice $w(\mathbf{r}) = R_j^q(\mathbf{r})$, and the boundary conditions given in equations (4.30)-(4.32) are substituted into the angularly discretised weak form given in equation (4.41). This yields the following equation set depending on the sign of $\boldsymbol{\Omega}_m \cdot \mathbf{n}$:

$$\underline{\boldsymbol{\Omega}_m \cdot \mathbf{n} < 0:}$$

$$\sum_{i=1}^N d_{i,m} \int_V \left(\frac{1}{\sigma_t} (\boldsymbol{\Omega}_m \cdot \nabla R_i^q) (\boldsymbol{\Omega}_m \cdot \nabla R_j^q) + \sigma_t R_i^q R_j^q \right) d\mathbf{r} = \int_V \left(Q_m R_j^q - \frac{Q_m}{\sigma_t} (\boldsymbol{\Omega}_m \cdot \nabla R_j^q) \right) d\mathbf{r} - \int_S \left(|\boldsymbol{\Omega}_m \cdot \mathbf{n}| Z R_j^q \right) d\mathbf{r}_s, \quad (4.42)$$

$$\underline{\boldsymbol{\Omega}_m \cdot \mathbf{n} > 0:}$$

$$\begin{aligned} \sum_{i=1}^N d_{i,m} \left[\int_V \left(\frac{1}{\sigma_t} (\boldsymbol{\Omega}_m \cdot \nabla R_i^q) (\boldsymbol{\Omega}_m \cdot \nabla R_j^q) + \sigma_t R_i^q R_j^q \right) d\mathbf{r} - \int_S \left(|\boldsymbol{\Omega}_m \cdot \mathbf{n}| R_i^q R_j^q \right) d\mathbf{r}_s \right] \\ = \int_V \left(Q_m R_j^q - \frac{Q_m}{\sigma_t} (\boldsymbol{\Omega}_m \cdot \nabla R_j^q) \right) d\mathbf{r}, \quad \text{for } m = 1, \dots, M \end{aligned} \quad (4.43)$$

where q is the order of the NURBS basis functions, $Z(\mathbf{r}, \boldsymbol{\Omega}, E)$ is a incoming source term that can represent either the vacuum, incoming side source, or reflective boundary conditions given in equation (4.30).

Due to the highly localised nature of the NURBS basis functions the integrals over the domain can be split into a summation of integrals over the knot spans within patches. This leads to an assembly process similar to the one typically seen in the FE method. Each knot span can be visited and a local dense bilinear form computed. The global linear system of equations is then assembled using the geometric connectivity information. The final IGA discretised weak form is given by:

$$\underline{\boldsymbol{\Omega}_m \cdot \mathbf{n} < 0:}$$

$$\begin{aligned} \sum_{i=1}^N d_{i,m} \sum_{p=1}^P \sum_{e=1}^{e_p} \int_{V_e} \left(\frac{1}{\sigma_t} (\boldsymbol{\Omega}_m \cdot \nabla R_i^q) (\boldsymbol{\Omega}_m \cdot \nabla R_j^q) + \sigma_t R_i^q R_j^q \right) d\mathbf{r} = \sum_{p=1}^P \sum_{e=1}^{e_p} \int_{V_e} \left(Q_m R_j^q - \frac{Q_m}{\sigma_t} (\boldsymbol{\Omega}_m \cdot \nabla R_j^q) \right) d\mathbf{r} \\ - \sum_{p=1}^P \sum_{b=1}^{b_p} \int_{S_b} \left(|\boldsymbol{\Omega}_m \cdot \mathbf{n}| Z R_j^q \right) d\mathbf{r}_s, \end{aligned} \quad (4.44)$$

$$\underline{\boldsymbol{\Omega}_m \cdot \mathbf{n} > 0:}$$

$$\begin{aligned} \sum_{i=1}^N d_{i,m} \left[\sum_{p=1}^P \sum_{e=1}^{e_p} \int_{V_e} \left(\frac{1}{\sigma_t} (\boldsymbol{\Omega}_m \cdot \nabla R_i^q) (\boldsymbol{\Omega}_m \cdot \nabla R_j^q) + \sigma_t R_i^q R_j^q \right) d\mathbf{r} - \sum_{p=1}^P \sum_{b=1}^{b_p} \int_{S_b} \left(|\boldsymbol{\Omega}_m \cdot \mathbf{n}| R_i^q R_j^q \right) d\mathbf{r}_s \right] \\ = \sum_{p=1}^P \sum_{e=1}^{e_p} \int_{V_e} \left(Q_m R_j^q - \frac{Q_m}{\sigma_t} (\boldsymbol{\Omega}_m \cdot \nabla R_j^q) \right) d\mathbf{r}, \quad \text{for } m = 1, \dots, M \end{aligned} \quad (4.45)$$

where P is the number of patches in the domain, e_p is the number of knot spans in patch p , b_p is the number of knot spans on the external boundary of patch p , V_e is the volume of knot span e , and S_b is the external surface of knot span b . Equations (4.44) & (4.45) are the IGA-S_N-SAAF equations for the incoming and outgoing angles.

4.3 Numerical Results

In this section the results of several test problems are explored. These include IGA discretisations compared to FE based discretisations and comparing the SAAF equation with the first-order neutron transport equation and the EP formulation.

The continuous IGA SAAF spatial discretisation method developed in this chapter has been implemented with a Fortran code called ICARUS (Isogeometric Continuous self-Adjoint Radiation Using Splines). ICARUS is used to generate the SAAF- S_N results presented in this chapter. The code is also capable of performing FE calculations and is used here to generate the FE results. This ensures that exactly the same solvers and convergence criteria are used for both IGA and FE solutions. Inferno is a first-order neutron transport code with local refinement capabilities [174, 186, 187]. It is a discontinuous IGA discrete ordinate (S_N) code and is used in this chapter to generate reference solutions to a number of nuclear reactor physics verification benchmark test cases.

4.3.1 MMS: Gaussian Modulated Tensor Product of Trigonometric Functions

The method of manufactured solutions (MMS) [228] is used here to investigate the rates of convergence of the IGA discretised SAAF equation for different orders of basis function. The function used to generate the source terms for the SAAF equation is a Gaussian function that is modulated in each direction using a high frequency trigonometric function:




Figure 4.1: Example MMS function, $C = 10, m = 20, n = 10$

$$f(x, y) = Cx(x-1)y(y-1) \exp\left(-(x-0.5)^2 - (y-0.5)^2\right) \sin(nx) \sin(my) \quad (4.46)$$

where each term is chosen to ensure that the solution has the following properties:

- $x(x-1)y(y-1)$: enforces homogeneous Dirichlet boundaries on the unit square.
- $C \exp(-(x-0.5)^2 - (y-0.5)^2)$: to make the solution smooth and infinitely differentiable.

- $\sin(nx)\sin(my)$: adds oscillations in both directions. This ensures that higher-order basis functions cannot trivially represent the solution.




Figure 4.2: Spatial convergence plots of MMS verification test case

An example of this function can be seen in Figure 4.1. The function $f(x, y)$ is substituted into the LHS of equation (4.6) and the remaining terms are taken as a source term $S(x, y, \Omega)$. This source term is then used on the RHS of the SAAF equation which is discretised using IGA with NURBS basis functions.

The convergence plot of the L_2 -error in the solution computed using S_{20} IGA SAAF is shown in Figure 4.2. The solution of the linear system of equations formed from the spatial and angular discretisation of the SAAF equation was performed using the incomplete Cholesky factorisation preconditioned conjugate gradient (ICC-PCG) algorithm implemented within PETSc using a iterative convergence tolerance of 10^{-13} for each angle. The solution algorithm was terminated when the magnitude of the residual in the scalar neutron flux was less than 10^{-10} . Basis functions of order one through six have been used and the theoretical order of convergence for degree p basis functions is $O(p + 1)$.

The order of convergence of all basis functions is limited initially until a sufficient spatial resolution has been achieved to allow the solution to capture the oscillations in expression (4.46). Once the level of discretisation is sufficient to capture the oscillatory behaviour, all solutions achieve an order of convergence close to the theoretical one. One exception to this is the highly refined $p = 6$ solution where the rate of convergence begins to decrease as the error in the solution begins to approach the convergence tolerance of the scalar neutron flux.

4.3.2 The IAEA Swimming Pool Nuclear Reactor Physics Verification Benchmark Test Case

The IAEA swimming pool nuclear reactor physics verification test case was first defined in a IAEA report by J. Stepanek [229]. The problem is relatively insensitive to approximation in angle beyond

a certain point [230]. The insensitivity to angular resolution can be seen in Table 4.1 where the K_{eff}




Figure 4.3: Geometry and region numbers for the IAEA swimming pool nuclear reactor physics verification test case. Vacuum boundary conditions are prescribed on all boundaries of the solution domain.

has been calculated using a discontinuous Galerkin IGA discrete ordinate (S_N) code, called Inferno, for different orders of angular quadrature with similar levels of spatial refinement. For this reason the IAEA swimming pool problem has been selected to explore the spatial convergence of the IGA discretised SAAF equation when compared to a FE discretisation. Furthermore, the problem itself is numerically challenging due to some regions being optically thick and the difference in order of magnitude between volume averaged scalar neutron fluxes in neighbouring regions. In this chapter both an extraneous (fixed) and a fission source version of the IAEA swimming pool problem are considered. The geometry and macroscopic neutron cross-section data are presented in Figure 4.3 and Tables 4.2 & 4.3. The reference results for each problem were calculated using the Inferno code with roughly 1,200,000 quadratic elements per calculation and using S_{10} Legendre-Chebyshev quadrature. The reference quantity of interest (QoI) for the criticality problem is $K_{\text{eff}} = 1.008875$ and for the extraneous (fixed) source problem the volume averaged scalar neutron flux in each region is presented in Table 4.4.

S_N	Number of elements	K_{eff}
8	1,199,569	1.0088657332
10	1,199,554	1.0088746303
20	1,199,560	1.0088852053
40	1,199,569	1.0088878513

Table 4.1: K_{eff} for various discrete ordinate (S_N) orders as calculated by Inferno for the IAEA swimming pool nuclear reactor physics benchmark.

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

Region	1	2	3	4	5
Neutron source strength (cm ⁻³ .s ⁻¹)	1.0	0.0	1.0	0.0	0.0

Table 4.3: Extraneous (fixed) neutron source strengths for the fixed source IAEA swimming pool nuclear reactor physics verification test case.

Region	Volume averaged scalar neutron flux (cm ⁻² .s ⁻¹)
1	11.9569693948
2	0.54057779903
3	19.1960438341
4	0.83475246098
5	1.52653294652

Table 4.4: Reference volume averaged scalar neutron flux solutions for the extraneous (fixed) source IAEA swimming pool nuclear reactor physics verification test case. These solutions were generated using the Inferno code.

Figure 4.4 shows the error in the K_{eff} for S_{10} quadratic IGA and FE solutions when compared to S_{20} Inferno for the fission source problem. The IGA solution always preserves the geometry (volume and surface area) exactly for geometries that can be represented by NURBS basis functions. However, due to the Cartesian geometry nature of the problem the FE method also preserves geometry (volume and surface area) of all regions.

It is seen that for the same number of dof the K_{eff} calculated using the IGA method is more accurate than the K_{eff} calculated using FEs. This is due to the high-order continuity of the NURBS basis over element boundaries within a NURBS patch. In this case they are C^1 continuous over element boundaries within a NURBS patch, whereas the FE basis functions are C^0 over element boundaries.

Plots of the error in the region averages for the extraneous (fixed) source problem are shown in Figure 4.5. The same general behaviour can be observed, that is fewer dof are needed to reduce the error by a prescribed amount when using quadratic IGA compared to quadratic FEs. In particular, the error in the FE solution is about half an order of magnitude larger than the error in the IGA solution.





Figure 4.4: Error in the K_{eff} for the eigenvalue IAEA swimming pool reactor physics benchmark. Both the IGA and FE use quadratic basis functions and a S_{10} angular quadrature set. Both solutions were computed using ICARUS where a triangular Legendre-Chebyshev discrete ordinate angular quadrature has been used. The results are compared to the reference solution provided by an Inferno S_{20} solution.


4.3.3 Seven-Group, Bare UO_2 Fuel Pin Nuclear Reactor Physics Verification Benchmark Test Case

This seven-group, eigenvalue (K_{eff}), nuclear reactor physics verification benchmark test case consists of a two-dimensional UO_2 bare fuel pin of radius 0.54 cm with prescribed vacuum boundary conditions. The seven-group macroscopic neutron cross-section data for this verification test case are taken from the OECD/NEA C5G7 nuclear reactor physics benchmark for the UO_2 fuel-clad material [210]. The geometry for this verification test case is a modified version of a pincell from the two-dimensional OECD/NEA C5G7 quarter core test case. The fuel pin has the same dimensions as the original pincell but the moderating material has been removed. The modified bare fuel pin geometry is shown in Figure 4.6. This benchmark has been specifically designed to illustrate the effects of geometrical discretisation errors on the numerical accuracy and rate of convergence of the solution. The problem is solved using S_2 triangular Legendre-Chebyshev quadrature. The reference solution used is a highly refined sixth order IGA solution with $K_{\text{eff}} = 7.30533835\text{E}^{-003}$. The convergence plots of the IGA and volume preserved FEs are shown in Figure 4.7.


It can be seen that for both sets of FE calculations the rate of convergence is limited by the rate of decrease of geometric error. In the non-volume preserved FE solution the errors in surface area, volume, and the total geometric error have been plotted and it can be observed that the rate of convergence of all solutions is limited to the rate of decrease in geometric error. Similar behaviour can be seen in the volume preserved FE solutions. Here the error in the surface area is slightly higher, however the error in volume is zero. This is due to the fact that the nodes on the surface are being pushed out by different amounts for different numbers of elements in order to maintain the volume of the geometry. Therefore, the surface polygon circumscribes the true circular geometry. In the non-volume preserved case the polygon inscribes the geometry.




(a) Region 1




(b) Region 2



(c) Region 3



(d) Region 4



(e) Region 5

Figure 4.5: Error in the volume averaged scalar neutron flux for the extraneous (fixed) source IAEA swimming pool reactor physics benchmark in each region computed by ICARUS using quadratic IGA and quadratic FEs with an S_{10} angular quadrature. The reference solutions were produced using the discontinuous Galerkin IGA S_N code Inferno.




Figure 4.6: Geometry for the seven-group, bare UO_2 fuel pin eigenvalue (K_{eff}) nuclear reactor physics verification benchmark test case. Vacuum boundary conditions have been prescribed on the boundary of the bare UO_2 fuel pin.

The errors in the non-volume preserved and volume preserved FE solutions differ by roughly two orders of magnitude for higher-order basis functions. Therefore, the larger error in the surface area incurred by preserving volume would appear to be worthwhile. The difference between the volume preserved FE and IGA results are even greater. This is mainly due to the fact that the rate of convergence of the IGA solution is not limited by the geometric approximation as the geometry is represented exactly for all levels of refinement. Further accuracy in the IGA solution is obtained due to the high inter-element continuity of the NURBS basis functions within a NURBS patch, as was demonstrated in section 4.3.2.


A more in depth study of the effects of higher-order isoparametric FE basis functions on the surface and volume errors has been performed [129, 231]. Furthermore, the comparison of these methods to IGA has also been extensively studied as well [175].

4.3.4 OECD/NEA, Seven-Group, Two-Dimensional (2D) C5G7 Nuclear Reactor Physics Verification Benchmark Test Case


The OECD/NEA C5G7 verification test case is a 2D, seven-group, heterogeneous, quarter core nuclear reactor physics benchmark used by the OECD/NEA and code developers as a means of verifying the implementation and assessing the accuracy of deterministic neutron transport codes. The material composition can be seen in Figure 4.8 and the macroscopic neutron cross-section data for the benchmark can be found in Appendix A of the benchmark specification [210].

Figures 4.9a, 4.9b, & 4.9c show the neutron scalar flux profiles for the C5G7 quarter core verification test case. It can be seen that the fast and epi-thermal neutrons are born in the fuel regions. As they are transported through the domain down scattering causes a drop in the fast and epi-thermal scalar neutron fluxes as the neutron energy is attenuated. As neutrons enter the reflector region from the fuel region there is a large spike in the thermal scalar neutron flux caused by the large scattering cross-section causing any fast neutrons to become thermalised.


Several results are presented for comparison. First are the solutions of the benchmark from the first-order neutron transport IGA discrete ordinate (S_N) code Inferno. The results presented are from a reference solution with 2,164,032 quadratic elements for the whole geometry and a S_8 angular



(a) Non-volume preserved FEs



(b) Volume preserved FEs



(c) IGA

Figure 4.7: Error in the K_{eff} for the bare fuel pin problem for varying levels of spatial and polynomial refinement within a NURBS patch. All results calculated using ICARUS .







Figure 4.8: Geometry specification for the 2D OECD/NEA C5G7 benchmark.



(a) Group 1 (fast scalar neutron flux with units $\text{cm}^{-2} \cdot \text{s}^{-1}$).



(b) Group 4 (epithermal scalar neutron flux with units $\text{cm}^{-2} \cdot \text{s}^{-1}$).



(c) Group 7 (thermal scalar neutron flux with units $\text{cm}^{-2} \cdot \text{s}^{-1}$).

Figure 4.9: Scalar neutron flux profiles for the 2D OECD/NEA C5G7 quarter core nuclear reactor physics verification test case for several energy groups. The scalar neutron flux is normalised according to the C5G7 benchmark specification [210].

discretisation. The eigenvalue was converged up to 10^{-11} and can be found in Owens' paper [174]. Second are the results from the even-parity S_N code CRONOS2 [232]. These results were discretised using linear triangular FEs that give 24,804 degrees of freedom for the entire geometry. An S_8 angular discretisation was used and these results can be found in the OECD/NEA C5G7 benchmark book [210]. Finally, results from ICARUS using quadratic basis functions are presented. Tables 4.5, 4.6, 4.7 and 4.8 present the errors in several quantities of interest (QoI) for CRONOS2 and ICARUS against Inferno. It is assumed that the CRONOS2 FE mesh is mass preserving as ICARUS ran a similar FE problem with no mass preservation and the error in the K_{eff} was significantly greater.

Number of degrees of freedom (dof)	K_{eff} (pcm)	Maximum pin power (%)	Minimum pin power (%)	Inner UO ₂ assembly (%)	Outer UO ₂ assembly (%)	MOX assembly (%)
24804	58	0.0049	-0.041	0.0031	-0.0041	-0.0023

Table 4.5: Error in CRONOS2 solution compared to the Inferno C5G7 reference result. The K_{eff} is converged to 10^{-6} . The particular S_8 angular quadrature set is unknown.

Number of degrees of freedom (dof)	K_{eff} (pcm)	Maximum pin power (%)	Minimum pin power (%)	Inner UO ₂ assembly (%)	Outer UO ₂ assembly (%)	MOX assembly (%)
29105	52	0.0030	-0.0177	0.0016	-0.0014	-0.0014

Table 4.6: Error in ICARUS solution compared to Inferno C5G7 reference result. The K_{eff} has been converged to 10^{-6} and a triangular Legendre-Chebyshev S_8 angular quadrature set has been used.

Number of degrees of freedom (dof)	K_{eff} (pcm)	Maximum pin power (%)	Minimum pin power (%)	Inner UO ₂ assembly (%)	Outer UO ₂ assembly (%)	MOX assembly (%)
29105	57	0.0027	-0.0179	0.0014	-0.0014	-0.0012

Table 4.7: Error in ICARUS solution compared to Inferno C5G7 reference result. The K_{eff} has been converged to 10^{-6} and a level-symmetric S_8 angular quadrature set has been used.

Number of degrees of freedom (dof)	K_{eff} (pcm)	Maximum pin power (%)	Minimum pin power (%)	Inner UO ₂ assembly (%)	Outer UO ₂ assembly (%)	MOX assembly (%)
29105	311	-0.0051	-0.0154	-0.0039	0.0012	0.0040
65332	302	-0.0089	-0.0014	-0.0060	0.0025	0.0062
181136	296	-0.0083	0.00053	-0.0057	0.0026	0.0057

Table 4.8: Error in ICARUS solution compared to the Monte-Carlo C5G7 reference solution found in the benchmark specification [210]. The K_{eff} has been converged to 10^{-7} and a S_8 triangular Legendre-Chebyshev angular quadrature set has been used.

Tables 4.5, 4.6, and 4.7 show the errors in various quantities of interest (QoI). The solution algorithm has been terminated when the eigenvalue for successive iterations differs by less than 10^{-6} . It

can be observed that the IGA solution is superior to FE solution. This is due, in part, to the quadratic basis functions used and also to the exact representation of the surface area and volume of each pin, whereas the FE solution only preserves fuel pin volume. While the coarsest ICARUS solve has more dof than the CRONOS2 solve the number of dof per pincell is the same. The ICARUS IGA mesh has slightly more dof in the water near the boundaries than the CRONOS2 FE mesh. Therefore, the extra dof in the IGA mesh should affect the pin powers minimally. Furthermore, it can be seen from Figure 4.9 that the flux profiles in the water near the boundaries are relatively flat so fewer dof are required to represent the solution here.

Whilst the pin power results for the ICARUS solution in Table 4.7 are clearly more accurate than the CRONOS2 solution, the difference in the two K_{eff} values is only 1 pcm. The angular quadrature set used by CRONOS2 is not known to the authors so, in an effort to investigate how much of a difference angular quadrature sets might make, tables 4.6 and 4.7 show the errors in ICARUS calculations where the only difference is the use of triangular Legendre-Chebyshev (LCT) and level-symmetric set (LSS) quadrature sets. Both quadrature sets result in the same number of angular unknowns. It can be seen that there is a difference of 5pcm in the K_{eff} whilst the pin power errors change minimally. Therefore, we conclude that the small difference seen between the K_{eff} for the CRONOS2 and ICARUS results could be caused by the use of different angular quadrature sets.

Table 4.8 shows the errors in various QoI calculated from IGA-SAAF solutions of the 2D OECD/NEA C5G7 verification test case using ICARUS. The results have been compared against the reference Monte-Carlo solution from the 2D OECD/NEA C5G7 benchmark report [210]. Three levels of discretisation are presented. The coarsest mesh is identical to the one used in Tables 4.6 & 4.7 and the two finer meshes are the result of refining each patch once or twice respectively in each parametric direction. The tolerance for the relative change in K_{eff} is 10^{-7} to make sure that errors in the QoI were not caused by a lack of convergence.

It can be seen that most QoI are converging slowly to the reference solution, whilst others appear to diverge. It is postulated that the reason for the slow convergence of some QoI, the divergence of other QoI, and the large pcm difference in the K_{eff} is caused by the low angular quadrature set used. Furthermore, there will be an error introduced into the S_8 -IGA solution due to the ray effect. These errors are not present in the reference Monte-Carlo solution and cause a discrepancy. Indeed, the pcm error between the K_{eff} from the finest solution in Table 4.8 and the highly refined S_8 first-order transport reference solution generated using Inferno is only 37pcm.

4.4 Conclusion

In this chapter a maximum principle compatible with source iteration based upon the SAAF form of the neutron transport equation was presented. It was shown to be mathematically equivalent to a Bubnov-Galerkin discretisation of the weak form of the SAAF NTE. This variational derivation provides a rigorous mathematical justification for the use of the first-order form of the NTE as a boundary condition for the outgoing angular flux on the boundary of the solution domain. The SAAF equation was spatially discretised using isogeometric analysis (IGA) with non-uniform rational B-spline (NURBS) test and trial functions. A method of manufactured solutions (MMS) verification test case demonstrates that this spatial discretisation yields close to the theoretical order of convergence for several orders of basis functions. Several extraneous (fixed) source and eigenvalue (k_{eff}) nuclear reactor physics verification benchmark test cases comparing FE and IGA spatial discretisations of the SAAF equation were performed. It was demonstrated that in the case where the geometry could be represented exactly by both IGA and FE basis functions that the IGA spatial discretisation was more accurate per dof. This improvement in spatial accuracy, using the IGA method, was due to the high inter-element continuity of the NURBS basis functions within a NURBS patch. Furthermore, in the case where the FE method is unable to exactly represent the geometry it was demonstrated that this geometric error limits the rate of convergence of the overall solution.

The 2D OECD/NEA, seven-group, C5G7 quarter core nuclear reactor physics verification benchmark problem was solved using a discrete ordinate (S_N) IGA spatial discretisation of the SAAF equation. The IGA SAAF solutions were compared against a second-order even-parity (EP) solution from the CRONOS2 FE code. This comparison was performed on meshes with similar numbers of dof and shows that the IGA solution was more accurate spatially than a similar FE solution. Furthermore, the effect of discrete ordinate (S_N) angular quadrature sets on the solution was also investigated.

Chapter 5

Constraint Based Locally Refined IGA Applied to the SAAF Equation

Much of the work in this chapter, including figures and tables, has been submitted for publication in the Journal of Computational Physics.

The self-adjoint angular flux (SAAF) equation is a second-order form of the NTE that is derived from the first-order form of the NTE as was shown in Chapter 3. It is a symmetric, elliptic partial integro-differential equation (PIDE). When discretised using FE or IGA methods the resulting system of linear equations is typically solved using preconditioned conjugate gradient (PCG) algorithms [136]. Whilst these require more storage than the sweep based solvers they have the potential to be more efficient due to the use of preconditioners to reduce the range of the eigenspectrum [233, 234]. Furthermore, the SAAF equation is stable when discretised using continuous Bubnov-Galerkin methods. Therefore, the same fidelity of solution can be achieved with fewer degrees of freedom (dof) than for discontinuous Bubnov-Galerkin discretisations. However, local refinement for continuous Bubnov-Galerkin methods is more challenging due to the requirement of a continuous basis function space.

Due to the tensor product nature of the NURBS basis functions, and the requirement for a consistent discretisation, the application of continuous NURBS based IGA does not allow for hanging node formulations natively. Refinements propagate through the domain making true local refinement impossible. In this chapter the development and application of constraint based local refinement for continuous Bubnov-Galerkin IGA is described. In section 5.1 the constraint based local refinement method is explained and applied to IGA. In section 5.2 the physical adjoint of the SAAF equation is derived so that it can be used to create a goal-based error measure. In section 5.3, two error indicators are derived. A heuristic error indicator based upon the local error in the solution, and a dual weighted residual (DWR) or goal based error indicator. The two error indicators are then used to drive an adaptive mesh refinement (AMR) procedure. Finally, in section 5.4 the magnitude of errors and rate of convergence of locally refined problems is investigated. Furthermore, an AMR algorithm driven by the two derived error indicators is applied to a nuclear reactor physics and a radiation shielding problem.




Figure 5.1: The geometry for the two patch case. The coarse patch is on the left and its control points are red. The fine 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 $\Xi^C = \{0, 0, 0.5, 1, 1\}$ and $\Xi^F = \{0, 0, 0.25, 0.5, 0.75, 1, 1\}$ in both parametric directions.

5.1 Constraint Based Local Refinement

Constraint based local refinement is an idea that was originally used in FE to restore the continuity of the FE space when hanging nodes were present [235]. The methodology has been extended to the more general IGA framework [27] in order to restore continuity of the NURBS basis function space for arbitrary order basis functions between patches when hanging nodes are present.

Two cases are considered in this section. The first case is two square patches side by side (see Figures 5.1 and 5.2). This case will illustrate how a constraint is applied and also how the global stiffness matrix is formed. The second case is three square patches in a corner configuration (see Figure 5.3). This case details how to deal with a point being constrained multiple times. These two cases can be combined to deal with most geometries.

5.1.1 The Two Patch Case

The case of two square NURBS patches that share a side is considered. The geometry with an example mesh superimposed onto it is presented in Figure 5.1. 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 patches form a compatible discretisation. It is also assumed that the control points on the shared interface of the fine patch are the result of knot insertion of the control points from the coarse patch. However, this 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 patches are then constrained to match this 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 the coarse patch. That is:

$$\mathbf{B}^F = T\mathbf{B}^C. \quad (5.1)$$

where the matrix operator T represents knot insertion, and $\mathbf{B}^F$ and $\mathbf{B}^C$ are the control points in the fine and coarse patches respectively. Following from this it is apparent that the control variable associated with each control point in each patch must satisfy a similar relationship, that is:

$$\mathbf{d}^F = T\mathbf{d}^C, \quad (5.2)$$

where $\mathbf{d}$ is a vector representing the control variables used to form the solution function as given in equation (4.39). The control variables in each patch are now partitioned into two sets: control variables that are on an interface with another patch, denoted by a f subscript, and control variables that are not, denoted by a n subscript. Therefore, from equation (5.2) it can be written that:

$$\begin{pmatrix} \mathbf{d}_n^F \\ \mathbf{d}_f^F \end{pmatrix} = \begin{pmatrix} T_n & 0 \\ 0 & T_f \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^C \\ \mathbf{d}_f^C \end{pmatrix}, \quad (5.3)$$

and in particular that:

$$\mathbf{d}_f^F = T_f \mathbf{d}_f^C. \quad (5.4)$$

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

$$\begin{pmatrix} K_{nn}^F & K_{nf}^F \\ K_{fn}^F & K_{ff}^F \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^F \\ \mathbf{d}_f^F \end{pmatrix} = \begin{pmatrix} S_n^F \\ S_f^F \end{pmatrix}, \quad (5.5)$$

and on the coarse NURBS patch:

$$\begin{pmatrix} K_{nn}^C & K_{nf}^C \\ K_{fn}^C & K_{ff}^C \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^C \\ \mathbf{d}_f^C \end{pmatrix} = \begin{pmatrix} S_n^C \\ S_f^C \end{pmatrix}, \quad (5.6)$$

where K^F and K^C are the stiffness matrices for the fine and coarse patches respectively over the sets of control variables implied by the subscripts. Expression (5.4) is then substituted into equation (5.5) giving:

$$\begin{pmatrix} K_{nn}^F & K_{nf}^F \\ K_{fn}^F & K_{ff}^F \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^F \\ T_f \mathbf{d}_f^C \end{pmatrix} = \begin{pmatrix} S_n^F \\ S_f^F \end{pmatrix}. \quad (5.7)$$

In order to assemble the patch wise systems of linear equations in equations (5.6) and (5.7) into one global system that represents the problem over the entire domain, it is necessary for both equations to operate on the same set of control variables on the shared interface. Therefore, the unknowns on the shared interface of the fine patch system must be constrained to the same set as the coarse patch system. This is achieved by operating on the control variables in the fine patch that are on the interface with T_f^T . This gives:

$$\begin{pmatrix} K_{nn}^F & K_{nf}^F T_f \\ T_f^T K_{fn}^F & T_f^T K_{ff}^F T_f \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^F \\ \mathbf{d}_f^C \end{pmatrix} = \begin{pmatrix} S_n^F \\ T_f^T S_f^F \end{pmatrix}. \quad (5.8)$$

This is an expression for the stiffness matrix over the fine patch that operates on the same set of control variables on the interface as the coarse patch. In other words, there is now one equation for each control variable in $\mathbf{d}_f^C$ in equation (5.8). Now the two linear systems of equations, (5.6) and (5.8),




Figure 5.2: The layout of the control points for quadratic basis functions. The coarse patch is on the left and its control points are red. The fine 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.

can be assembled into a global system of linear equations. In this particular case yielding:

$$\begin{pmatrix} K_{nn}^C & K_{nf}^C & 0 \\ K_{fn}^C & K_{ff}^C + T_f^T K_{ff}^F T_f & T_f^T K_{fn}^F \\ 0 & K_{nf}^F T_f & K_{nn}^F \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^C \\ \mathbf{d}_f^C \\ \mathbf{d}_n^F \end{pmatrix} = \begin{pmatrix} S_n^C \\ S_f^C + T_f^T S_f^F \\ S_n^F \end{pmatrix}. \quad (5.9)$$

This system can then be solved using a CG solver as symmetry of the system is maintained.

In Figure 5.1, due to the linear basis functions along the interface, the control points that belong to the coarse patch are a proper subset of the control points that belong to the interface patch. This will not be the case for non-linear basis functions. The case of quadratic basis functions is illustrated in Figure 5.2. It can be seen that there is a control point on the interface that belongs only to the coarse patch. The constraint based methodology detailed above deals with this case by design as all of the refinement information is contained in T_f .

5.1.2 The Three Patch Case

In order to fully describe the constraint based local refinement methodology the case of three 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. The geometry for this case is displayed in Figure 5.3 and the point that will have more than one constraint applied to it is coloured black.

The methodology begins as in the two 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 followed until three patch wise stiffness matrices are formed, two for the coarse patches which are unconstrained:

$$\begin{pmatrix} K_{nn}^{C_w} & K_{nf_w}^{C_w} \\ K_{f_w n}^{C_w} & K_{f_w f_w}^{C_w} \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^{C_w} \\ \mathbf{d}_{f_w}^{C_w} \end{pmatrix} = \begin{pmatrix} S_n^{C_w} \\ S_{f_w}^{C_w} \end{pmatrix}, \quad (5.10)$$




Figure 5.3: The geometry for the three patch case. The coarse patch to the west is referred to as C_w and its control points are red. The coarse patch to the south is referred to as C_s and its control points are yellow. The fine 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, $\Xi^{C_e} = \Xi^{C_s} = \{0, 0, 0.5, 1, 1\}$ and $\Xi^F = \{0, 0, 0.25, 0.5, 0.75, 1, 1\}$ in both parametric directions.

$$\begin{pmatrix} K_{nn}^{C_s} & K_{nf_s}^{C_s} \\ K_{f_s n}^{C_s} & K_{f_s f_s}^{C_s} \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^{C_s} \\ \mathbf{d}_{f_s}^{C_s} \end{pmatrix} = \begin{pmatrix} S_n^{C_s} \\ S_{f_s}^{C_s} \end{pmatrix}, \quad (5.11)$$

and one for the fine patch which has had two sets of constraints applied to it:

$$\begin{pmatrix} K_{nn}^F & K_{nf_w}^F T_{f_w} & K_{nf_s}^F T_{f_s} \\ T_{f_w}^T K_{f_w n}^F & T_{f_w}^T K_{f_w f_w}^F T_{f_w} & T_{f_w}^T K_{f_w f_s}^F T_{f_s} \\ T_{f_s}^T K_{f_s n}^F & T_{f_s}^T K_{f_s f_w}^F T_{f_w} & T_{f_s}^T K_{f_s f_s}^F T_{f_s} \end{pmatrix} \begin{pmatrix} \mathbf{d}_n^F \\ \mathbf{d}_{f_w}^{C_w} \\ \mathbf{d}_{f_s}^{C_s} \end{pmatrix} = \begin{pmatrix} S_n^F \\ T_{f_w}^T S_{f_w}^F \\ T_{f_s}^T S_{f_s}^F \end{pmatrix}. \quad (5.12)$$

The subscripts f_w and f_s refer to the interface between the central patch and western patch and the interface between the central patch and the southern 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 5.3 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 patch. We refer to this point as d_l . Since d_l has been constrained twice that means that it appears twice in equation (5.12), so there is some abuse of notation going on.

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 $d_l = d_l$. This means that d_l is never actually constrained. So when equations (5.10), (5.11), and (5.12) are assembled into a global system the row that relates all of the control variables to d_l will be added twice. Therefore, to get the correct answer the coefficients in the row corresponding to d_l must be divided by two.

In three dimensions this repeated assembly of the row corresponding to d_l could happen three or four times, depending on the geometry. Therefore, in three dimensions the number of times that d_l is constrained must be monitored so that the row corresponding to d_l can be 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.

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

5.2 The Continuous (Physical) Adjoint of the SAAF Equation

The adjoint of the neutron transport equation has several uses in nuclear reactor physics and radiation shielding calculations. The solution of the adjoint neutron transport equation can be used as a variance reduction technique for Monte Carlo methods [41] and is also useful for calculating the effect of local perturbations in material coefficients on the reactivity of a system [236]. The solution of the adjoint neutron transport equation can be thought of as describing the importance of a neutron in the neutron flux [78]. This means that the adjoint neutron flux can be used as part of a goal-based error measure and will be discussed in this chapter as part of an spatially adaptive mesh refinement (AMR) scheme for specified quantities of interest (QoI).

The adjoint of an operator H is defined as follows:

$$\langle Hu, v \rangle = \langle u, H^\dagger v \rangle, \quad (5.13)$$

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 (4.6), $\mathbf{\Omega} \cdot \nabla \frac{Q(\mathbf{r}, \mathbf{\Omega}, E)}{\sigma_t} \equiv 0$ which is satisfied for mono-energetic, homogeneous problems with no scattering or fission source. 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 $(\mathbf{r}, \mathbf{\Omega}, E)$ and taking the adjoint according to expression (5.13) do not, in general, commute. These adjoints are commonly referred to as the continuous adjoint or the discrete adjoint, or sometimes the physical adjoint and the mathematical adjoint in the nuclear community [218], depending on the order these operations are performed [237]. However, for the SAAF equation there are several possibilities for the discrete adjoint. Following from the notation used in the paper by Schunert et al. [218] a vector is used to signify the order in which discretisation and the action of taking the adjoint are applied. For example, the vector $(E, \mathbf{\Omega}, \dagger, \mathbf{r})$ 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: $(\dagger, E, \mathbf{\Omega}, \mathbf{r})$
2. Continuous adjoint: $(E, \dagger, \mathbf{\Omega}, \mathbf{r}) \equiv (E, \mathbf{\Omega}, \dagger, \mathbf{r})$
3. Discrete adjoint: $(E, \mathbf{\Omega}, \mathbf{r}, \dagger)$

In this work the continuous adjoint is used for the following reasons. First, if the true adjoint is used then adjoint flux weighted neutron cross-sections are necessary. That is, equation (2.21) would become weighted by the adjoint neutron angular flux, $\psi \rightarrow \psi^\dagger$. 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 forward (primal) scalar neutron flux [20], see equation (2.21). Since adjoint weighted neutron cross-section data is not generally available, the true adjoint is not considered in this work.

The comparison of the continuous and discrete adjoints is more subtle. In the case of the first-order neutron transport equation with a linear spatial discretisation the continuous and discrete adjoints are the same so the comparison between the two adjoint formulations is not often considered [218]. When

the discrete adjoint is used the discretisation errors in the forward (primal) and adjoint calculations are the same. Therefore, the eigenvalues of the primal (forward) equation and discrete adjoint equation agree exactly as they both include the same discretisation error [218]. This is not the case when the continuous adjoint is used [238]. However, in the asymptotic limit as the discretisation error goes to zero the eigenvalues of the primal (forward) equation, discrete adjoint equation, and the continuous adjoint equation are identical.

It was mentioned earlier that the adjoint can be thought of as a measure of a neutron's importance to the system [78] which makes the adjoint solution suited to a goal-based error measure. This physical interpretation is the case for the continuous (physical) adjoint. It is unclear what the physical significance of the discrete adjoint is outside of the asymptotic limit [239], but it is assumed that, since the solutions of the continuous and discrete adjoints differ only by discretisation error that the discrete adjoint would also be suitable. However, despite the literature suggesting that the discrete adjoint is superior to the mathematical adjoint in terms of accuracy and computational time [218, 238, 239] the mathematical adjoint is used here, mainly due to the fact that it is consistent with the forward SAAF equation [218].

5.2.1 Derivation of the Continuous (Physical) Adjoint

The start point in the derivation of the physical adjoint of the SAAF equation is the multigroup adjoint first-order neutron transport equation. This is given by:

$$-\mathbf{\Omega} \cdot \nabla \psi_g^\dagger + \sigma_{t,g} \psi_g^\dagger = Q_g^\dagger, \text{ for } g = 1, \dots, G, \quad (5.14)$$

with

$$Q_g^\dagger = Q_{\text{ext},g}^\dagger(\mathbf{r}, \mathbf{\Omega}) + \nu \sigma_{f,g}(\mathbf{r}) \sum_{g'=1}^G \chi_{g'} \phi_{g'}^\dagger(\mathbf{r}) + \sum_{g'=1}^G \int_{4\pi} \sigma_{s,g \rightarrow g'}(\mathbf{r}, \mathbf{\Omega} \cdot \mathbf{\Omega}') \psi_{g'}^\dagger(\mathbf{r}, \mathbf{\Omega}') d\mathbf{\Omega}'. \quad (5.15)$$

with the boundary condition:

$$\psi_g^\dagger(\mathbf{r}) = F_g^\dagger(\mathbf{r}, \mathbf{\Omega}), \text{ for } \mathbf{\Omega} \cdot \mathbf{n} > 0, \quad (5.16)$$

where the function $F_g^\dagger(\mathbf{r}, \mathbf{\Omega})$ can represent vacuum boundary conditions ($F_g^\dagger = 0$), reflective boundary conditions ($F_g^\dagger(\mathbf{r}, \mathbf{\Omega}) = \psi^\dagger(\mathbf{r}, \mathbf{\Omega}_\mathbf{r})$, where $\mathbf{\Omega}_\mathbf{r}$ is the angle of specular reflection), or an incoming boundary source. In accordance with the last section the multigroup approximation has been applied before the adjoint has been taken so that the neutron cross sections are not weighted by $\psi^\dagger(\mathbf{r}, \mathbf{\Omega}, E)$.

The derivation of the physical adjoint of the SAAF equation then proceeds in the same fashion as the algebraic derivation of the SAAF equation [140, 177]. Equation (5.14) is rearranged to form the following expression:

$$\psi_g^\dagger = \frac{1}{\sigma_{t,g}} (Q_g^\dagger + \mathbf{\Omega} \cdot \nabla \psi_g^\dagger), \text{ for } g = 1, \dots, G. \quad (5.17)$$

Expression (5.17) is then substituted into the streaming term of equation (5.14) giving:

$$-\mathbf{\Omega} \cdot \nabla \frac{1}{\sigma_{t,g}} \mathbf{\Omega} \cdot \nabla \psi_g^\dagger + \sigma_{t,g} \psi_g^\dagger = Q_g^\dagger + \mathbf{\Omega} \cdot \nabla \frac{Q_g^\dagger}{\sigma_{t,g}}, \text{ for } g = 1, \dots, G. \quad (5.18)$$

The outgoing boundary conditions for equation (5.18) can be derived from equation (5.16) as:

$$\frac{1}{\sigma_{t,g}}(Q_g^\dagger + \mathbf{\Omega} \cdot \nabla \psi_g^\dagger) = F_g^\dagger(\mathbf{r}, \mathbf{\Omega}), \text{ for } \mathbf{\Omega} \cdot \mathbf{n} > 0. \quad (5.19)$$

For the incoming boundary condition it is assumed that the adjoint angular neutron flux should satisfy equation (5.14) yielding:

$$\frac{1}{\sigma_{t,g}}(Q_g^\dagger + \mathbf{\Omega} \cdot \nabla \psi_g^\dagger) = \psi_g^\dagger, \text{ for } \mathbf{\Omega} \cdot \mathbf{n} < 0. \quad (5.20)$$

Equations (5.18), (5.19), and (5.20) form the physical adjoint of the SAAF equation. This derivation is equivalent to the transform suggested by Schunert et al. [218]. 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 neutron per fission and the macroscopic fission neutron cross-section and the fission spectrum are switched. This is equivalent to: $\nu \sigma_{f,g'}(\mathbf{r}) \rightarrow \nu \sigma_{f,g}(\mathbf{r})$ and $\chi'_{g'} \rightarrow \chi_g$.
- The forward (primal) scalar and angular neutron fluxes are replaced with the adjoint (dual) scalar and angular neutron fluxes. $\psi_g(\mathbf{r}, \mathbf{\Omega}) \rightarrow \psi_g^\dagger(\mathbf{r}, \mathbf{\Omega})$, $\phi_g(\mathbf{r}) \rightarrow \phi_g^\dagger(\mathbf{r})$.
- Reinterpret solutions in the opposite direction. This is equivalent to replacing every instance of $\mathbf{\Omega}$ with $-\mathbf{\Omega}$.

The physical adjoint of the SAAF equation can now be used to form a goal-based error measure.

5.3 Adaptive Mesh Refinement (AMR) Methods

Adaptive mesh refinement (AMR) algorithms refine specific areas of the problem domain based upon some driving mechanism. AMR algorithms can refine spatially, [240, 241], polynomially [242], or in a hybrid of the two [222, 243, 244, 245]. In this work only spatial refinement strategies are investigated. The constraint based mechanism by which local refinement is achieved in NURBS patches is detailed in section 5.1. 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 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 [245, 246, 247]. As the mechanism for achieving local refinement has already been discussed, in this section the strategy for flagging areas for refinement is detailed.

The strategies for flagging areas for refinement can broadly be split into two groups. Heuristic methods and goal-based error measures. Heuristic methods refine based on some generic 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 [186]. Heuristic methodologies can be applied to any solution, that is they are black box techniques.

Goal-based error measures are tailored towards the specific problem being solved. In this work this is achieved by solving the adjoint neutron transport equation. As mentioned earlier, the solution

of the adjoint equation can broadly be considered to give a measure of a neutron's importance to the scalar neutron flux [78]. By correctly setting the source terms of the adjoint solution it can be further refined to give a measure of a neutron's importance to some quantity of interest (QoI). An example of a QoI is a reaction rate, or a dose that a given area receives.

Both a heuristic driven AMR strategy and a goal-based error measure AMR strategy are considered for the SAAF equation in the following sections. Both of these methods are based up the error between a reference solution on a fine mesh and an approximate solution over a coarse mesh. The approximate solution can either be acquired by solving the SAAF equation over the coarse mesh or by projecting the reference solution from the fine mesh to the coarse mesh. Due to the fact that refinement of NURBS patches can be performed within the program, no interfacing with an ancillary mesh generator is necessary.

5.3.1 The Forward Error Indicator (FEI)

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

$$\|\psi - \psi_h\|_N, \quad (5.21)$$

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:

$$\|f(\mathbf{r})\|_{L^2} = \int_V f(\mathbf{r})^2 d\mathbf{r}, \quad (5.22)$$

or the H^1 norm:

$$\|f(\mathbf{r})\|_{H^1} = \|f(\mathbf{r})\|_{L^2} + \|\nabla f(\mathbf{r})\|_{L^2} = \int_V f(\mathbf{r})^2 + \nabla f(\mathbf{r}) \cdot \nabla f(\mathbf{r}) d\mathbf{r}. \quad (5.23)$$

The H^1 seminorm can also be considered:

$$|f(\mathbf{r})|_{H^1} = \|\nabla f(\mathbf{r})\|_{L^2} = \int_V \nabla f(\mathbf{r}) \cdot \nabla f(\mathbf{r}) d\mathbf{r}, \quad (5.24)$$

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 we consider the H^1 norm as this accounts for not only 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:

$$e^F = \|\psi - \psi_h\|_{H^1}. \quad (5.25)$$

The FEI e^F is calculated in each patch yielding e_p^F . All the e_p^F 's are then formed into an ascending list and a method for selecting patches to flag for refinement is chosen. Two options are considered here. First, a patch is flagged for refinement if the error in that patch is within a certain threshold of the maximum error:

$$e_p^F > \epsilon \max\{e_p^F\}_{p=1}^N, \quad (5.26)$$

where $0 \leq \epsilon \leq 1$ and N is the number of patches. This is referred to as the threshold method. The

second method 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 patches NURBS patches are indicated for refinement.

The second strategy will ensure that the number of patches flagged for refinement in each AMR iteration will be the same. The first strategy should produce a more suitable mesh for a given problem but could require more AMR iterations. This will be investigated in section 5.4.

In practice the error indicator from expression (5.25) is not suited to all problems as it contains the exact solution ψ . The exact solution is only available ahead of time for a small subset of neutron transport problems that can be solved analytically [34]. 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 mesh where every patch has been uniformly refined to a higher level of spatial resolution than the 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 mesh with higher order basis functions [248]. A more flexible FEI is then given by:

$$\mu^F = \|\psi_{h+1} - \psi_h\|_{H^1}. \quad (5.27)$$

The process of calculating the FEI μ^F now involves solving two neutron transport problems: one over a coarse mesh to calculate ψ_h and one over a fine 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 grid. This choice is justified by the following lemma [245].

Lemma 1. *Let V_{h+1} and V_h be two finite real Hilbert spaces such that $V_h \subset V_{h+1}$ and let V_{h+1} have the norm $\|\cdot\|$. Let $\phi_{h+1} \in V_{h+1}$ and $\phi_h \in V_h$. Let $a : V_{h+1} \times V_{h+1} \rightarrow \mathbb{R}$ be a bilinear form that is:*

- *continuous:* $|a(u, v)| \leq \gamma \|u\| \|v\| \quad \forall u, v \in V_{h+1} \text{ with } \gamma > 0,$
- *coercive:* $a(v, v) \geq \alpha \|v\|^2 \quad \forall v \in V_{h+1} \text{ with } \alpha > 0.$

Then

$$\|\phi_{h+1} - \phi_h\| \leq \frac{\gamma}{\alpha} \|\phi_{h+1} - \Pi_h^{h+1}\phi_{h+1}\|$$

$\forall \phi_{h+1} \in V_{h+1}, \phi_h \in V_h$ and for any $\Pi_h^{h+1} : V_{h+1} \rightarrow V_h$.

Proof. By Céa's lemma:

$$\|\phi_{h+1} - \phi_h\| \leq \frac{\gamma}{\alpha} \|\phi_{h+1} - \varphi_h\| \quad \forall \varphi_h \in V_h.$$

Therefore:

$$\|\phi_{h+1} - \phi_h\| \leq \frac{\gamma}{\alpha} \min_{\varphi_h \in V_h} \|\phi_{h+1} - \varphi_h\|,$$

and a weaker condition is:

$$\|\phi_{h+1} - \phi_h\| \leq \frac{\gamma}{\alpha} \|\phi_{h+1} - \Pi_h^{h+1}\phi_{h+1}\|.$$

□

A flexible, relatively cheap, heuristic FEI is given by:

$$\eta^F = \|\psi_{h+1} - \Pi_h^{h+1}\psi_{h+1}\|_{H^1}. \quad (5.28)$$

Furthermore, Lemma 1 means that minimising η^F will also minimise μ^F regardless of the choice of the projection Π_h^{h+1} .

5.3.2 The Weighted Error Indicator (WEI)

The weighted error indicator (WEI) is a goal-based error measure that builds upon the FEI η_F given in expression (5.28). This is done by weighting η^F with the error in the adjoint neutron transport equation. This is expressed mathematically as:

$$e^W = \eta^F \|\psi^\dagger - \psi_h^\dagger\|_{H^1} = \|\psi_{h+1} - \Pi_h^{h+1}\psi_{h+1}\|_{H^1} \|\psi^\dagger - \psi_h^\dagger\|_{H^1}. \quad (5.29)$$

This computation of the WEI e^W requires the solution of three neutron transport equations: one primal and two adjoint. An analogous series of approximations to the ones performed in section 5.3.1 can be made to arrive at a WEI that requires the solution of one primal neutron transport equation and one adjoint neutron transport equation. A flexible, relatively cheap, goal-based WEI is given by:

$$\eta^W = \eta^F \|\psi_{h+1}^\dagger - \Pi_h^{h+1}\psi_{h+1}^\dagger\|_{H^1} = \|\psi_{h+1} - \Pi_h^{h+1}\psi_{h+1}\|_{H^1} \|\psi_{h+1}^\dagger - \Pi_h^{h+1}\psi_{h+1}^\dagger\|_{H^1}. \quad (5.30)$$

Furthermore, the computation cost of this procedure could be halved by solving a special neutron transport formulation that yields the primal scalar neutron flux and the adjoint scalar neutron flux as its solution in a manner similar to the even/odd parity neutron flux [125].

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 $Q_{\text{ext}}^\dagger$ in expression (5.15). To set the goal as some neutron reaction rate in the system, the choice:

$$Q_{\text{ext},g}^\dagger(\mathbf{r}) = \sigma_{d,g}(\mathbf{r})\phi_g(\mathbf{r}), \quad (5.31)$$

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 at the QoI then the choice $\sigma_{d,g} = \nu\sigma_{f,g}$ is made. This is based upon the fact that the K_{eff} is calculated by:

$$K_{\text{eff}}^{(i)} = \frac{\int_V \int_0^\infty \nu\sigma_f(\mathbf{r}, E)\phi^{(i)}(\mathbf{r}, E)dEd\mathbf{r}}{\int_V \int_0^\infty \nu\sigma_f(\mathbf{r}, E)\phi^{(i-1)}(\mathbf{r}, E)dEd\mathbf{r}} K_{\text{eff}}^{(i-1)} \quad (5.32)$$

where $\phi^{(i)}$ is the neutron scalar flux from the i^{th} power iteration. 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.

5.4 Results

In this section the results of several test cases are explored. These include an MMS test case to investigate the rate of convergence of locally refined solutions, two test cases to investigate how errors




Figure 5.4: Geometry for the MMS verification test case.

caused by the constraint based local refinement propagate, and an AMR scheme applied to an IAEA nuclear reactor physics test case.

The continuous, locally refined, S_N IGA SAAF discretisation method developed in this paper has been implemented within the modern Fortran code ICARUS (Isogeometric Continuous self-Adjoint Radiation Using Splines). Some reference results have been generated with the discontinuous S_N IGA first-order transport code Inferno [174, 186, 187].

5.4.1 Method of Manufactured Solutions Test Case

The method of manufactured solutions (MMS) is a popular technique for the verification of numerical computer codes [228, 249]. It is used here to investigate the effect that the 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:

$$\psi_M(x, y) = x(x - 10)y(y - 10) \left(\cos\left(\frac{4\pi x}{5}\right) \cos\left(\frac{4\pi y}{5}\right) + 2 \right). \quad (5.33)$$


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 patches are one level more refined than the bottom right and top left 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 5.4. Figure 5.5 shows an illustration of the errors between the manufactured solution and the discrete solution for locally refined, and uniformly refined geometries.

Plotted in Figure 5.6 are the errors in the MMS test in the fine patches, in the coarse patches, and in the case of uniform refinement. The rate of convergence of a solution is indicated by the gradient of the line joining the solution points.

The first thing of note is that the rate of convergence and level of error in the solution in the coarse



(a) Error in uniformly refined MMS solution. The knot vectors in each patch have been subdivided four times.



(b) Error in locally refined MMS solution. The knot vectors in the fine patches have been subdivided five times and the knot vectors in the coarse patches have been subdivided four times.

Figure 5.5: Plot of the error in the MMS verification test case. The colour scale in both plots is the same.

patches is identical to the rate of convergence and level of error in the solution for a uniformly refined geometry. This is to be expected as the local constraints apply to the dof in the fine patch not in the coarse patch. This also explains why, for the same element size, the error in the fine patches is greater than in the coarse patches. The dof in the fine patches are being constrained to match the dof in the coarse patches along the shared boundary, introducing an error into the fine patch solution.

Secondly it can be observed that, whilst the error in the fine patches is greater than in the coarse patches for the same level of refinement, the rate of convergence of the solution in the fine patches is greater than the rate of convergence in the coarse patches. Furthermore, the error in the fine patches is tending towards the error in the coarse patches.

It is expected that the level of error in the fine patches will decrease toward the level of error in the coarse patches as the problem is refined. In the fine 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 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 inline with the theoretical rate of $O(h^{p+1})$, in fact they seem to be $O(h^p)$. This is assumed to be due to the strong imposition of the zero flux boundary conditions caused by the big-spring method that is effecting the rate of convergence of the solution as seen in section 3.6.1.

5.4.2 Pincell Lattice Calculation

The aim of the next test case is to investigate how errors introduced by constraints and areas of low spatial resolution propagate through the domain. The geometry for the test cases consists of a




Figure 5.6: Spatial convergence plots for the MMS verification test case. The rates of convergence for the coarse patches and the uniformly refined case match exactly for all orders of basis function.




Figure 5.7: Geometry of the pincell row problem. All boundaries have reflective boundary conditions applied to them.

5×1 array of pincells with reflective boundary conditions. All pincells have the same neutron cross-sections and are geometrically identical. Therefore, this test case is similar to an infinite pincell lattice calculation. The geometry is shown in Figure 5.7. A specific refinement configuration is chosen shown in Figure 5.8. All pincells except the far left one are formed of 5 biquadratic NURBS patches, each consisting of 256 elements. The far left pincell is formed of 5 biquadratic NURBS patches, each one consisting of one element. The problem is discretised with quadratic NURBS basis functions and angularly discretised with S_{12} discrete ordinates.

The QoI considered for this test case is the disadvantage factor in each 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:

$$\zeta = \frac{\pi}{16 - \pi} \frac{\int_{\text{moderator}} \phi(\mathbf{r}) d\mathbf{r}}{\int_{\text{fuel}} \phi(\mathbf{r}) d\mathbf{r}}. \quad (5.34)$$

Six different sets of material properties are considered and are given in Table 5.1. In each set of material properties the scattering ratio $c = \frac{\sigma_s}{\sigma_t} = 0.95$ is kept constant whilst σ_t is varied. The ratio of




Figure 5.8: Discretised geometry of the pincell row problem.

Case	Macroscopic neutron cross-section data (cm^{-1})				Source data (neutrons. cm^{-3})	
	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 5.1: Macroscopic neutron cross-section data for the pincell row test cases. A F superscript refers to the fuel and a M superscript refers to the moderator.

the 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 5.9 displays the errors in the cell disadvantage factors for each pincell. The error is calculated by comparing the locally refined solution with an uniformly refined 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 pincell one and by the constraint between pincells one and two no longer effects the value of ζ in pincell three due to a small λ . In the cases where $\sigma_t^F < 10$ there is still the error can be seen to propagate throughout the domain due to a large λ . This manifests as an error greater than 1×10^{-11} in ζ in pincell five.

Another trend apparent from Figure 5.9 is that the error in ζ for 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 pincell one will lead to greater error in ζ as this gradient gets steeper.

5.4.3 Adaptive Mesh Refinement Applied to the IAEA Swimming Pool Problem

To illustrate how constraint based local refinement would be used in practice an adaptive mesh refinement (AMR) algorithm has been implemented. Here it is applied to the IAEA swimming pool nuclear reactor physics verification test case [229]. The geometry and neutron cross-section data for the test case can be seen in Figure 5.10 and Table 5.2. The IAEA swimming pool problem has been chosen as a benchmark due to the difference in the order of magnitude of the neutron scalar flux in each region.




Figure 5.9: Error in ζ for the pincell row problem.

The region averaged neutron scalar fluxes vary up to three orders of magnitude between regions [250]. Therefore, the problem will have strong gradients that can be efficiently resolved by local refinement.

The AMR algorithm uses the WEI η^W detailed in equation (5.30). The operation Π_h^{h+1} is performed using the energy dependent mesh interpolation methods developed by Welch [176] and Owens [187] 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 patch, a patch is flagged for refinement if $\eta_p^W > \epsilon \max\{\eta_p^W\}_{p=1}^N$ with $\epsilon = 0.3$. The algorithm is then repeated. Due to fact that in IGA knot insertion is performed at run time, no interfacing with an ancillary mesh generator is necessary.

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

As the aim of this problem is investigate the effectiveness of constraint based local refinement on a more realistic nuclear reactor physics problem the angular domain is discretised using a S_4 triangular Legendre-Chebyshev quadrature set. The quantity of interest (QoI) chosen is the K_{eff} . The goal for the WEI is set for the K_{eff} as described in section 5.3.2. A reference value of $K_{\text{eff}} = 1.0087598988$ was generated using the discontinuous IGA, first-order neutron transport code Inferno [174, 186, 187] using fourth order NURBS basis functions with 468,571 elements.

The results of the AMR scheme applied to the IAEA swimming pool problem are displayed in




Figure 5.10: Geometry of the IAEA swimming pool nuclear reactor physics test case.




Figure 5.11: Error in the K_{eff} calculated from the S_4 -SAAF equation solved for the IAEA swimming pool nuclear reactor physics test case. An AMR scheme driven by the WEI is compared to uniform refinement.

Figure 5.11. 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 scheme is uniformly better than the uniformly refined solution. However, this is to be expected as the primal (forward) and adjoint SAAF equations have been solved for the WEI case. For a given number of degrees of freedom the WEI refined solution takes roughly twice as long as 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 [125].

Region	Macroscopic neutron cross-section data (cm ⁻¹)			Source data (neutrons.cm ⁻³)
	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 5.3: Macroscopic neutron cross-section data for shielding problem I.

5.4.4 Lathouwers Radiation Shielding Problem

The following test cases is a radiation shielding problem from work by Lathouwers [241] and is referred to as shielding problem I. The geometry is displayed in Figure 5.12. Region I contains a unit isotropic extraneous source of neutrons and the QoI is a volumetric detector response in region III. The neutron cross section data for the problem is given in Table 5.3. 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 5.3. This is used to defined the adjoint source in expression (5.31). The problem has been solved with a S_4 discrete ordinate quadrature set to match the calculation performed by Lathouwers [241].

Figure 5.13 shows the error in the FEI and WEI AMR solutions for two quantities. Both the WEI and FEI AMR algorithms had patches selected for refinement using the error threshold method defined in expression (5.26). Figure 5.13a shows the error in the quantity $\int_V \phi(\mathbf{r})d\mathbf{r}$. It can be seen that both the FEI and the WEI lead to an improvement in the quantity $\int_V \phi(\mathbf{r})d\mathbf{r}$. Both AMR solutions require about five times fewer dof to reach the same error tolerance as the uniformly refined solution. The WEI error is also slightly less than the FEI error for large numbers of dof. This is to be expected as the WEI requires the solution of both a primal (forward) and an adjoint neutron transport equation.

Figure 5.13b shows the error in the quantity $\int_{III} \phi(\mathbf{r})d\mathbf{r}$. It is clear that the FEI refinement strategy led to a mesh that was unable to calculate the QoI accurately. The mesh generated from the WEI strategy seems to be better as the error in the QoI is one order of magnitude less than the FEI error. Indeed, the FEI error is greater than the uniformly refined solution. This implies that the FEI refined mesh failed to sufficiently resolve area III. The FEI mesh can be seen in Figure 5.14. It can indeed be seen that region III is relatively unrefined. The WEI mesh is shown in Figure 5.15. The extra computational effort of solving the adjoint equation can be seen as the interface between region II and III has been heavily refined.

In order to check that the threshold method was not biased towards the WEI method the same calculation was run but this time the list method was used to select patches for refinement. Patches that were within the top 10% of the list were flagged for refinement and Figure 5.16 displays the results. Similar trends to the ones in the threshold case can be observed. One strange behaviour in the FEI error is that at one refinement step the error in $\int_V \phi(\mathbf{r})d\mathbf{r}$ 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 algorithm. The error of the WEI solution is monotonically decreasing in Figures 5.13 and 5.16 and implies the method is more predictable and stable.






Figure 5.12: Geometry specification for shielding problem I.



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



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

Figure 5.13: Error flux integral quantities. The WEI η^W and the FEI η^F are used to drive AMR. The decision to refine a 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.




Figure 5.14: Mesh produced from the FEI strategy with 8469 control points.






Figure 5.15: Mesh produced from the WEI strategy with 7581 control points.



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



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

Figure 5.16: Error flux integral quantities. The WEI η^W and the FEI η^F are used to drive AMR. 10% of the patches that form the domain are flagged for refinement in each AMR step. Therefore, the total number of refinements is equal for FEI and WEI. This does not imply that the number of control points is the same.

5.5 Conclusions

In this chapter a constraint based method for the local refinement of NURBS based isogeometric analysis (IGA) has been described and applied to the self-adjoint angular flux (SAAF) neutron transport equation. Furthermore, two error indicators have been derived, a heuristic error indicator referred to as FEI, and a dual-weighted goal-based error indicator referred to as WEI. In addition, the physical adjoint of the SAAF equation has been derived so that the WEI can be applied to the SAAF equation.

The efficacy of the constraint based local refinement method has been investigated using a method of manufactured solutions (MMS) test case and a pincell lattice calculation with several material configurations. It was shown that the errors introduced by the constraint based local refinement method affect neither the convergence rates, nor the error in the solution significantly for the MMS test case. It can be observed that as the problem is further refined the rate of convergence of the locally refined solutions approaches the rate of convergence of the uniformly refined solution. In the lattice 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 illustrated. It was noted that the propagation of this error depends upon the mean free path of neutrons in the system.

An adaptive mesh refinement (AMR) algorithm driven by heuristic (FEI) and goal-based (WEI) error measures was applied to a nuclear reactor physics and a radiation shielding test case. It was demonstrated that both error indicators performed well for global quantities of interest (QoI). However, if a local QoI is considered, then WEI, with a suitable goal functional, was superior to FEI.

Further research is required to find a way to remove the requirement that the coarsest geometric representation must form a consistent discretisation for NURBS based IGA. This would simplify the initial creation of geometries and allow for easier processing of ‘dirty’ CAD geometries. The effect of the neutron mean free path as well as other material properties should be taken into account when deciding where the system should be refined so that large errors are not introduced into the solution.

Chapter 6

The Weighted Least Squares Equation

Much of the work in this chapter, including figures and tables, has been adapted from the publication by Latimer et al. [251].

Neutron transport problems that contain void regions are typically solved using the first-order form of the neutron transport equation or Monte Carlo simulations [252]. Both of these methods can be combined in order to reduce the variance in the Monte Carlo solution. The adjoint first-order neutron transport equation is solved and the solution is used as an importance map for the Monte-Carlo algorithm [41, 104].

There are certain desirable numerical properties of the second-order forms of the neutron transport equation that are advantageous compared to the first-order form. In particular, the elliptic property of most second-order forms of the neutron transport equation means that the discretised equations yield symmetric positive definite (SPD) matrices. Such SPD matrix systems of linear equations can be solved using computationally efficient preconditioned conjugate gradient (PCG) solution algorithms [125, 136]. However, traditional second-order forms of the neutron transport equation, such as the even-parity (EP) [110] and odd-parity (OP) equations, as well as the self-adjoint angular flux (SAAF) equation derived in Chapter 4 cannot be solved in void regions. This is due to the reciprocal of the total macroscopic neutron cross-section that appears in these formulations. Solving neutron transport problems that contain voids is important in two cases. First, in a boiling water reactor (BWR) the coolant boils and forms regions of steam that are modelled as voids. Second, in radiation shielding problems, often a large portion of the domain will be outside of the reactor core. For example, in order to calculate the dosage the operator might receive the neutron flux in the space between the reactor core and the operator must be modelled. Even though this space is filled with air, it is seen as a void by the neutrons travelling through it. Furthermore, some second-order forms of the neutron transport equation can exhibit a lack of causality which can cause non-physical behaviour of the neutron flux [127].

There are two standard approaches to solving second-order forms of the neutron transport equation in void regions. The first approach involves using a very low absorption macroscopic neutron cross-section within the void region [147]. This approach can lead to an increase in the condition number of the resultant system of linear equations and a potentially erroneous solution [253]. In particular, a flat neutron flux is often observed across the void region [117]. Further research has been performed into improving the accuracy of this approach. These more complex methods involve solving the second-

order neutron transport equation several times whilst varying the value of the macroscopic absorption neutron cross-section used in the void region. For example, the iteration and extrapolation methods proposed by Ackroyd [147, 146].

The other approach is to interface different forms of the neutron transport equation at the boundaries between void and non-void regions. Research has been performed on interfacing with both first-order [124] and other second-order forms [126, 212] of the neutron transport equations. These attempts often lead to a loss of symmetry of the resulting system of linear equations which precludes the possibility of using PCG solvers.

A second-order form of the neutron transport equation recently proposed by Hansen et al. is a source iteration compatible self-adjoint least squares (LS) method that is compatible with void regions [211]. The source iteration compatible LS equation removes the reciprocal of the total macroscopic neutron cross-section and therefore, avoids the increase in condition number that is typically seen by second-order forms in optically transparent or fully void regions. However, this approach means that conservation of neutrons is only guaranteed in the fine-mesh limit. This issue of conservation only in the fine mesh limit is seen in other second-order forms of the neutron transport equation that are utilised for modelling neutron transport in void regions [140].

Least squares formulations have been used in many areas of research such as: quantum electrodynamics [254], fluid mechanics [255], and biomolecular systems [256]. In particular, LS formulations of the neutron transport equation, where the scattering kernel is taken to be part of the transport operator, have been considered before within the literature. LS based maximum principles for the first-order neutron transport equation were derived by Ackroyd and shown to be equivalent to the even and odd parity maximum principles [119]. Schofield expanded upon this work with a generalised LS form of the first-order neutron transport equation that is valid in void regions. Second-order maximum variational principles were derived for both the steady state [257, 258] and time dependent neutron transport equation [259]. These maximum variational principles were applied to neutron transport involving void problems and shown to compare reasonably well to point kernel calculations, as well as the aforementioned interpolation and extrapolation methods [147].

Manteuffel & Ressel presented proofs for boundedness and coercivity of the LS formulation. This means that the existence of a unique solution of the formulation can be proven using the Lax-Milgram theorem [122]. Error bounds, independent of neutron cross-section, that hold in the diffusion limit were also derived by applying a scaling to the neutron transport equation prior to the discretisation. The LS equation was solved in slab geometry discretised spatially and angularly using linear finite elements and spherical harmonics (P_N) respectively. Varin & Samba performed a similar analysis but compared the LS- P_N equation to the diamond differenced first-order neutron transport equation with discrete ordinate (S_N) angular discretisation [260]. Both of the previous references [122, 260] only addressed isotropic scattering and Austin & Manteuffel produced a similar analysis but allowed for anisotropic scattering [261].

Chaudhry et al. investigated adding an extra term to the LS formulation to provide additional accuracy in the evaluation of prescribed quantities of interest (QoI) [262]. This was found to provide more accuracy in the evaluation of the QoI but produce a larger L_2 error for the solution as a whole for both diffusion and convection-diffusion equations. Drumm showed a simpler method of scaling the LS equation that allowed it to be solved accurately in thick diffusive media. This method was applied

to one and two-dimensional void problems using finite elements with P_N angular discretisation [221]. Ketelsen et al. derived the LS equation from the first-order neutron transport equation formulated in one-dimensional (1D) spherical coordinates [68]. This mitigates the ray effects from the S_N discretised solution at the cost of more strongly coupling all of the angular unknowns. However, all of the problems that they analysed in their paper were purely absorbing as opposed to problems involving scattering.

Laboure et al. presented the source iteration compatible LS equation angularly discretised with both P_N and S_N [263]. The effect that the lack of conservation of neutrons has on the source iteration compatible LS equation's ability to solve criticality problems when compared to the SAAF equation for void free problems was analysed. Furthermore, Laboure et al. presented a hybrid SAAF & conservative least squares (SAAF-CLS) formulation of the neutron transport equation that is valid in void and near void regions, but whose variational form is not symmetric [126]. This methodology was shown to perform much better than the non-conservative LS formulation, and performed similarly to other void treated second-order forms of the neutron transport equation. Hammer et al. derived a source iteration compatible weighted LS formulation by including a weighting operator [127]. This weighting operator was chosen so that in non-void regions the WLS formulation was identical to the SAAF formulation. In void, or near-void, regions the formulation becomes more like Hansen's LS form [211]. Therefore, conservation of neutrons is not maintained in void regions. However, the weighting term helps to restore causality to the formulation [127].

In this chapter two issues are investigated. First, the magnitude and effect of the error induced by the lack of conservation of neutrons will be investigated. Second, the benefits of the IGA discretisation methodology applied to the WLS equation, as well as its ability to solve neutron transport problems containing voids, will be studied. In section 6.1 the source iteration compatible WLS form of the neutron transport equation is derived. Furthermore, a procedure for determining the weighting operator is discussed. Section 6.2 describes how the weak form of the WLS equation is spatially and angularly discretised using IGA and S_N respectively. In section 6.3 the lack of causality of the WLS in void regions is illustrated. Finally the results of the WLS and LS equation applied to several nuclear reactor physics verification benchmark test cases are presented in section 6.4.

6.1 Derivation of the Weighted Least Squares (WLS) Equation

The steady-state (time-independent), first-order, energy-dependent, neutron transport equation is given by:

$$\boldsymbol{\Omega} \cdot \nabla \psi(\mathbf{r}, \boldsymbol{\Omega}, E) + \sigma_t(\mathbf{r}, E) \psi(\mathbf{r}, \boldsymbol{\Omega}, E) = Q_{\text{scatter}}(\mathbf{r}, \boldsymbol{\Omega}, E) + Q_{\text{ext}}(\mathbf{r}, \boldsymbol{\Omega}, E), \quad (6.1)$$

where

$$Q_{\text{scatter}}(\mathbf{r}, \boldsymbol{\Omega}, E) = \int_0^\infty \int_{4\pi} \sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}, E' \rightarrow E) \psi(\mathbf{r}, \boldsymbol{\Omega}', E') d\boldsymbol{\Omega}' dE', \quad (6.2)$$

and $Q_{\text{ext}}(\mathbf{r}, E)$ consists of neutron fission and extraneous (fixed) sources. Furthermore, the total source is defined as:

$$Q(\mathbf{r}, \boldsymbol{\Omega}, E) = Q_{\text{scatter}}(\mathbf{r}, \boldsymbol{\Omega}, E) + Q_{\text{ext}}(\mathbf{r}, \boldsymbol{\Omega}, E).$$

The transport operator is defined as $\mathcal{L} = \boldsymbol{\Omega} \cdot \nabla + \sigma_t(\mathbf{r}, E)$ such that equation (6.1) can be written as:

$$\mathcal{L}\psi(\mathbf{r}, \boldsymbol{\Omega}, E) = Q(\mathbf{r}, \boldsymbol{\Omega}, E). \quad (6.3)$$

By considering the inner product $\int_V \int_{4\pi} \int_0^\infty dE d\mathbf{\Omega} d\mathbf{r}$ an expression for the adjoint transport operator $\mathcal{L}^\dagger$ can be derived by solving the equation:

$$\int_V \int_{4\pi} \int_0^\infty u(\mathbf{r}, \mathbf{\Omega}, E) \mathcal{L}v(\mathbf{r}, \mathbf{\Omega}, E) dE d\mathbf{\Omega} d\mathbf{r} = \int_V \int_{4\pi} \int_0^\infty \mathcal{L}^\dagger u(\mathbf{r}, \mathbf{\Omega}, E) v(\mathbf{r}, \mathbf{\Omega}, E) dE d\mathbf{\Omega} d\mathbf{r}. \quad (6.4)$$

This yields the following expression for the adjoint transport operator:

$$\mathcal{L}^\dagger = -\mathbf{\Omega} \cdot \nabla + \sigma_t(\mathbf{r}, E). \quad (6.5)$$

For this derivation the following notation for the inner product over the spatial domain is introduced:

$$\begin{aligned} \left(u, v \right)_V &= \int_V uv d\mathbf{r}, \\ \left\langle u, v \right\rangle_\Gamma &= \int_\Gamma uv d\mathbf{r}_s, \end{aligned}$$

where V is the physical domain, Γ is the boundary of V , and $\mathbf{r} \in V$ and $\mathbf{r}_s \in \Gamma$. Then the first-order neutron transport equation in operator form, see equation (6.3), is operated on by $\mathcal{L}^\dagger \mathcal{W}$ yielding:

$$\mathcal{L}^\dagger \mathcal{W} \mathcal{L} \psi(\mathbf{r}, \mathbf{\Omega}, E) = \mathcal{L}^\dagger \mathcal{W} Q(\mathbf{r}, \mathbf{\Omega}, E), \quad (6.6)$$

where $\mathcal{W}$, known as the weighting operator, is to be chosen later on. This operator equation can be written algebraically. The left hand side (LHS) and right hand side (RHS) of equation (6.6) are considered separately yielding:

LHS:

$$\begin{aligned} \mathcal{L}^\dagger \mathcal{W} \mathcal{L} \psi &= (-\mathbf{\Omega} \cdot \nabla + \sigma_t) \mathcal{W} (\mathbf{\Omega} \cdot \nabla \psi + \sigma_t \psi) \\ &= (-\mathbf{\Omega} \cdot \nabla + \sigma_t) (\mathcal{W} \mathbf{\Omega} \cdot \nabla \psi + \mathcal{W} \sigma_t \psi) \\ &= -\mathbf{\Omega} \cdot \nabla (\mathcal{W} \mathbf{\Omega} \cdot \nabla \psi) - \mathbf{\Omega} \cdot \nabla (\mathcal{W} \sigma_t \psi) + \sigma_t \mathcal{W} \mathbf{\Omega} \cdot \nabla \psi + \mathcal{W} \sigma_t^2 \psi, \end{aligned} \quad (6.7)$$

RHS

$$\mathcal{L}^\dagger \mathcal{W} Q = (-\mathbf{\Omega} \cdot \nabla + \sigma_t) \mathcal{W} Q = -\mathbf{\Omega} \cdot \nabla \mathcal{W} Q + \sigma_t \mathcal{W} Q. \quad (6.8)$$

It is noted that the application of the product rule to the LHS would result in the combination of the second and third terms in equation (6.7). However, this requires that both functions be differentiable. Since $\mathcal{W}$ has not been chosen yet nothing can be said about its differentiability. Furthermore, σ_t is allowed to be discontinuous at patch boundaries. Since both functions are not differentiable over the domain the product rule is not applicable here. Despite not being allowed, its application would yield the same weak form. This derivation continues without any application of the product rule. Equating sides gives:

$$-\mathbf{\Omega} \cdot \nabla (\mathcal{W} \mathbf{\Omega} \cdot \nabla \psi) - \mathbf{\Omega} \cdot \nabla [(\mathcal{W} \sigma_t) \psi] + \sigma_t \mathcal{W} \mathbf{\Omega} \cdot \nabla \psi + \mathcal{W} \sigma_t^2 \psi = (\sigma_t - \mathbf{\Omega} \cdot \nabla) \mathcal{W} Q, \quad (6.9)$$

which is then multiplied by a test function v and integrated over the entire spatial domain to give the weak form:

$$\begin{aligned} & - \left(\mathbf{\Omega} \cdot \nabla (\mathcal{W} \mathbf{\Omega} \cdot \nabla \psi), v \right)_V - \left(\mathbf{\Omega} \cdot \nabla (\mathcal{W} \sigma_t \psi), v \right)_V + \left(\sigma_t \mathcal{W} \mathbf{\Omega} \cdot \nabla \psi, v \right)_V + \left(\mathcal{W} \sigma_t^2 \psi, v \right)_V \\ & = \left(\mathcal{W} Q, \sigma_t v \right)_V - \left(\mathbf{\Omega} \cdot \nabla \mathcal{W} Q, v \right)_V. \end{aligned} \quad (6.10)$$

Performing integration by parts on the first and second terms on the left hand side and the second term on the right hand side gives:

$$\begin{aligned} & \left(\mathcal{W} \mathbf{\Omega} \cdot \nabla \psi, \mathbf{\Omega} \cdot \nabla v \right)_V + \left(\mathcal{W} \mathbf{\Omega} \cdot \nabla \psi, \sigma_t v \right)_V + \left(\mathcal{W} \sigma_t \psi, \mathbf{\Omega} \cdot \nabla v \right)_V + \left(\mathcal{W} \sigma_t \psi, \sigma_t v \right)_V \\ & - \left\langle \mathcal{W} \mathbf{\Omega} \cdot \nabla \psi, (\mathbf{\Omega} \cdot \mathbf{n}) v \right\rangle_\Gamma - \left\langle \mathcal{W} \sigma_t \psi, (\mathbf{\Omega} \cdot \mathbf{n}) v \right\rangle_\Gamma = \\ & \left(\mathcal{W} Q, \sigma_t v \right)_V + \left(\mathcal{W} Q, \mathbf{\Omega} \cdot \nabla v \right)_V - \left\langle \mathcal{W} Q, (\mathbf{\Omega} \cdot \mathbf{n}) v \right\rangle_\Gamma, \end{aligned} \quad (6.11)$$

which can be simplified to:

$$\begin{aligned} & \left(\mathcal{W} \mathbf{\Omega} \cdot \nabla \psi, \mathbf{\Omega} \cdot \nabla v + \sigma_t v \right)_V + \left(\mathcal{W} \sigma_t \psi, \mathbf{\Omega} \cdot \nabla v + \sigma_t v \right)_V - \left\langle \mathcal{W} \mathbf{\Omega} \cdot \nabla \psi + \mathcal{W} \sigma_t \psi, (\mathbf{\Omega} \cdot \mathbf{n}) v \right\rangle_\Gamma \\ & = \left(\mathcal{W} Q, \mathbf{\Omega} \cdot \nabla v + \sigma_t v \right)_V - \left\langle \mathcal{W} Q, (\mathbf{\Omega} \cdot \mathbf{n}) v \right\rangle_\Gamma. \end{aligned} \quad (6.12)$$

The boundary terms can be removed as they satisfy the first-order neutron transport equation given in equation (6.1), leaving:

$$\left(\mathcal{W} \mathbf{\Omega} \cdot \nabla \psi, \mathbf{\Omega} \cdot \nabla v + \sigma_t v \right)_V + \left(\mathcal{W} \sigma_t \psi, \mathbf{\Omega} \cdot \nabla v + \sigma_t v \right)_V = \left(\mathcal{W} Q, \mathbf{\Omega} \cdot \nabla v + \sigma_t v \right)_V. \quad (6.13)$$

The removal of the boundary integral terms leaves no way to satisfy natural boundary conditions. Therefore, these are implemented weakly by introducing an additional term with a penalty parameter, c into equation (6.13). The final weak form is then:

$$\begin{aligned} & \left(\mathcal{W} \mathbf{\Omega} \cdot \nabla \psi, \mathbf{\Omega} \cdot \nabla v + \sigma_t v \right)_V + \left(\mathcal{W} \sigma_t \psi, \mathbf{\Omega} \cdot \nabla v + \sigma_t v \right)_V + c \int_{\Gamma^-} \mathcal{W} v |\mathbf{\Omega} \cdot \mathbf{n}| (\psi - F) d\mathbf{r}_s \\ & = \left(\mathcal{W} Q, \mathbf{\Omega} \cdot \nabla v + \sigma_t v \right)_V, \end{aligned} \quad (6.14)$$

where $\Gamma^- = \{\Gamma \mid \mathbf{\Omega} \cdot \mathbf{n} < 0\}$ and $F(\mathbf{r}_s, \mathbf{\Omega}, E)$, defined for $\mathbf{r}_s \in \Gamma^-$, is a function that represents the boundary conditions. For example:

- $F = 0 \Rightarrow$ vacuum boundary
- $F = \psi(\mathbf{r}_s, \mathbf{\Omega}_R, E) \Rightarrow$ reflective boundary, where $\mathbf{\Omega}_R$ is the angle of specular reflection of $\mathbf{\Omega}$
- $F = \zeta(\mathbf{r}_s, \mathbf{\Omega}, E) \Rightarrow$ incoming boundary source.

The penalty parameter c is taken to be:

$$c = \max \left(\sigma_t, \frac{1}{h} \right), \quad (6.15)$$

where h is a measurement of element size, taken here to be the largest distance between successive knots in a patch. This is suggested by Laboure in order to ensure that the LS formulation is consistent with the SAAF formulation, that is they differ only by their discretisation error when $c = \sigma_t$ [264].

Equation (6.14) with $\mathcal{W} = 1$ is the source iteration compatible LS form presented by Hansen et. al [211]. A non-unit weighting factor gives the source iteration compatible weighted LS (WLS) form given by Hammer et al. [127].

The weighting factor suggested by Hammer et al. [127] is:

$$\mathcal{W} = \min \left(\frac{1}{\sigma_t}, w_{\max} \right), \quad (6.16)$$

where $w_{\max} \in \mathbb{R}$ is the maximum weight. In non-void regions $\mathcal{W} = \frac{1}{\sigma_t}$ implies that the WLS equation will be equivalent to the SAAF equation, up to boundary conditions. This means that the WLS equation conserves neutrons in non-void regions, unlike Hansen's LS equation. However, in near-void and void regions $\frac{1}{\sigma_t} \rightarrow \infty$, meaning equation (6.14) will tend towards a singularity. When this would happen the weighting factor becomes $w_{\max}$, the singularity is avoided but at the cost of losing conservation of neutrons in near-void and void regions.

6.1.1 Choice of the Weighting Operator $\mathcal{W}$

In the original WLS paper by Hammer et al. [127] the dependence of the solution of the WLS equation on the value $w_{\max}$ was analysed. It was suggested that the optimal value of $w_{\max}$ is independent of mesh size. Furthermore, it is suggested that $w_{\max}$ should be angularly dependent to account for the fact that neutrons streaming in different directions may travel different distances in a given void region. A strategy, based on the assumption that $w_{\max}$ is independent of the mesh size, for determining an appropriate value of $w_{\max}$ is presented here. The algorithm, which requires the solution of the first-order neutron transport equation, is presented in Figure 6.1 and is also explained below.

First a quantity of interest (QoI) is chosen, for example: the K_{eff} , a reaction rate over part of the domain, or a detector response at a given location. Once this has been selected the first-order neutron transport equation can be solved for the given problem and the QoI recorded, denoted here by Q_f . In this chapter, the WLS equation is spatially and angularly discretised using IGA and S_N respectively. Therefore, a first-order neutron transport IGA- S_N code is used to ensure that no errors caused by geometric approximation or differing angular discretisation enter the problem here.

An interval $\mathcal{I} \subset \mathbb{R}$ is selected and a set of weights $w_i \in \mathcal{I}$ is chosen. This set of weights is denoted $\mathcal{P}$. The WLS equation is then solved for each $w_i \in \mathcal{P}$ over a coarse mesh and the QoI recorded, denoted here by Q_{w_i} . If there exists a Q_{w_i} such that $|Q_f - Q_{w_i}| < \epsilon$, where ϵ is some tolerance, then $w_{\max} = w_i$ is chosen. If such a w_i is not found then a new $\mathcal{I}$ and $\mathcal{P}$ are chosen until a suitable weighting factor is found. In accordance with the assumption that $w_{\max}$ should be almost independent of mesh size, this procedure should give a good value of $w_{\max}$ for a highly spatially refined solution of the WLS neutron transport equation with the same level of angular discretisation.

One potential issue with this algorithm is related to the error in the QoI calculated from the first-order neutron transport solution, Q_f . Since Q_f is calculated over a coarse mesh there will be some error in Q_f relative to the exact answer, Q_f^e . It would be ideal to minimise the error in the QoI using $|Q_f^e - Q_{w_i}| < \epsilon$. However, the exact solution Q_f^e is not available so the expression $|Q_f - Q_{w_i}|$ is used.

Therefore, it would be beneficial to minimise $|Q_f^e - Q_f|$ to ensure that Q_{w_i} is being matched with a value of Q_f that is sufficiently close to the exact solution Q_f^e . Minimising $|Q_f^e - Q_f|$ implies that the mesh for the first-order neutron transport calculation should be highly refined. This causes two issues with respect to the execution time of the algorithm proposed in this chapter. First, the more refined the ‘coarse’ first-order neutron transport calculation, the more refined the ‘coarse’ WLS calculation must be in order to find the value of $w_{\max}$ that produces a matching QoI. Therefore, as a consequence, the more refined the ‘coarse’ mesh is, the longer the algorithm takes to execute.

Second, if it is possible to obtain fine mesh solutions of the first-order neutron transport equation in a reasonable time for void problems then this algorithm, along with the WLS equation, would be unnecessary. If it is not possible to obtain fine mesh solutions of the first-order neutron transport equation in a reasonable time then the computational cost of applying this algorithm for a fine ‘coarse’ mesh would be too high. Therefore, a coarse ‘coarse’ mesh should be used. The effect of the refinement of the ‘coarse’ mesh will be investigated in section 6.4.

6.2 Spatial and Angular Discretisation of the Weak Form of the WLS Neutron Transport Equation

6.2.1 Discrete Ordinate (S_N) Discretisation of the WLS Neutron Transport Equation

In this chapter the angular variable is discretised using the discrete ordinates (S_N) method. The S_N method is a quadrature rule that can be used to evaluate angular integrals as follows:

$$\int_{4\pi} f(\boldsymbol{\Omega}) d\boldsymbol{\Omega} \approx \sum_{m=1}^N w_m f(\boldsymbol{\Omega}_m). \quad (6.17)$$

A quadrature set defining the quadrature weights and directions $\{w_m, \boldsymbol{\Omega}_m\}_{m=1}^N$ must be chosen. Some examples of suitable angular quadrature sets are level-symmetric quadrature sets, and Legendre-Chebyshev quadrature sets.

In order to find the scalar neutron flux, the weak form is evaluated and solved for in each direction $\boldsymbol{\Omega}_n \in \{\boldsymbol{\Omega}_m\}_{m=1}^N$. Then the quadrature rule can be used to recover the scalar neutron flux from the angular neutron fluxes using equation (6.17):

$$\phi(\mathbf{r}, E) = \int_{4\pi} \psi(\mathbf{r}, \boldsymbol{\Omega}, E) d\boldsymbol{\Omega} \approx \sum_{m=1}^N w_m \psi(\mathbf{r}, \boldsymbol{\Omega}_m, E) = \sum_{m=1}^N w_m \psi_m(\mathbf{r}, E), \quad (6.18)$$

where $\psi_m(\mathbf{r}, E) := \psi(\mathbf{r}, \boldsymbol{\Omega}_m, E)$. This means that in order to use expression (6.18), equation (6.14) must be solved for $\forall \boldsymbol{\Omega}_m$. That is:

$$\langle \mathcal{W} \mathcal{L}_m \psi_m, \mathcal{L}_m v \rangle + \int_{\Gamma_m^-} c \mathcal{W} v (\boldsymbol{\Omega}_m \cdot \mathbf{n}) (\psi_m - F(\mathbf{r}, \boldsymbol{\Omega}_m)) d\mathbf{r}_s = \langle \mathcal{W} Q_m, \mathcal{L}_m v \rangle. \quad (6.19)$$

where $\mathcal{L}_m := \boldsymbol{\Omega}_m \cdot \nabla + \sigma_t(\mathbf{r}, E)$ and $\Gamma_m^- = \{\Gamma \mid \boldsymbol{\Omega}_m \cdot \mathbf{n} < 0\}$.

In Morel & McGhees paper [140] it is noted that the SAAF- S_N formulation can be derived either directly, by evaluating the final weak form in a direction $\boldsymbol{\Omega}_m$, or indirectly by evaluating the strong




Figure 6.1: The algorithm for determining a value of $w_{\max}$ based upon a solution of the first-order neutron transport equation.

form in a direction $\boldsymbol{\Omega}_m$ and then proceeding to derive the weak form. It is stated that these methods lead to the same set of equations. This is also the case for the WLS-S_N formulation. This can be seen by defining $\mathcal{L}_m = \boldsymbol{\Omega}_m \cdot \nabla + \sigma_t(\mathbf{r}, E)$ and $\mathcal{L}_m^\dagger = -\boldsymbol{\Omega}_m \cdot \nabla + \sigma_t(\mathbf{r}, E)$. Then it can be noted that equation (6.6) evaluated at $\boldsymbol{\Omega} = \boldsymbol{\Omega}_m$ gives

$$\mathcal{L}_m^\dagger \mathcal{W} \mathcal{L}_m \psi_m(\mathbf{r}, E) = \mathcal{L}_m^\dagger \mathcal{W} Q_m(\mathbf{r}, E).$$

Following through with the same weak form procedure specified in section 6.1 yields a weak form that is equivalent to equation (6.19).

6.2.2 IGA Spatial Discretisation of the Weak Form of the WLS Neutron Transport Equation

The IGA discretisation of the weak form of the WLS begins by expanding the spatial dependence of the solution as a linear combination of trial functions:

$$\psi_m(\mathbf{r}, E) = \sum_{i=1}^{\infty} d_{i,m}(E) v_i(\mathbf{r}), \quad (6.20)$$

where the $d_{i,m}(E)$ are referred to as the control variables. The $v_i \in T_R$ are referred to as the trial functions and belong to the infinite dimensional trial function space:

$$T_R = \{ v \mid v \in H^1(V), v|_{\Gamma_D} = g \}, \quad (6.21)$$

where $H^1(V)$ is the Sobolev space $W_2^1(V)$ over the domain V [227]. The condition $v|_{S_D} = g$ allows for the implementation of essential boundary conditions, such as a Dirichlet boundary condition: $\psi(\mathbf{r}) = g$ for $\mathbf{r} \in \Gamma_D$ [27]. A test function w is chosen to belong to another infinite dimensional space:

$$T_E = \{ w \mid w \in H^1(V), w|_{\Gamma} = 0 \}, \quad (6.22)$$

which is known as the test function space. The condition $w|_{\Gamma} = 0$ is enforced so that the test functions do not interfere with the essential boundary conditions.

In order to provide closure to expression (6.20) the expansion must be truncated. This is achieved by defining finite dimensional analogues of the trial and test function spaces:

$$\begin{aligned} T_R^h &= \left\{ v_i^h \mid v_i^h \in H^1(V), v_i^h|_{\Gamma_D} = g, \text{ for } i \in [1, N] \right\}, \\ T_E^h &= \left\{ w_i^h \mid w_i^h \in H^1(V), w_i^h|_{\Gamma} = 0, \text{ for } i \in [1, N] \right\}, \end{aligned}$$

and the solution can then be approximated by a truncated expansion:

$$\psi_m(\mathbf{r}, E) \approx \psi_m^h(\mathbf{r}, E) = \sum_{i=1}^N d_{i,m}(E) v_i^h(\mathbf{r}). \quad (6.23)$$

The Bubnov-Galerkin IGA discretisation is then formulated by making the choice $v_i(\mathbf{r}) = R_i^q(\mathbf{r}) \in T_R^h$ and $w_j(\mathbf{r}) = R_j^q(\mathbf{r}) \in T_E^h$. The truncated expansion of the angular flux given by expression (6.23) is




Figure 6.2: Geometry for the void wall problem.

substituted into the weak form given in equation (6.19). Furthermore, $w_j = R_j^q$ is substituted into equation (6.19) in place of v yielding:

$$\begin{aligned} \sum_{i=1}^N d_{i,m} \left[\int_V (\mathcal{W} \boldsymbol{\Omega}_m \cdot \nabla R_i + \mathcal{W} \sigma_t R_i) (\boldsymbol{\Omega}_m \cdot \nabla R_j + \sigma_t R_j) d\mathbf{r} + c \int_{\Gamma_m^-} |\boldsymbol{\Omega}_m \cdot \mathbf{n}| \mathcal{W} R_j (R_i - F) d\mathbf{r}_s \right] \\ = \int_V \mathcal{W} Q_m (\boldsymbol{\Omega}_m \cdot \nabla R_j + \sigma_t R_j) d\mathbf{r}. \end{aligned} \quad (6.24)$$

This global bilinear form can be split into the summation of local bilinear forms computed over knot spans as follows:

$$\begin{aligned} \sum_{i=1}^N d_{i,m} \sum_{p=1}^P \sum_{e=1}^{e_p} \int_{V_e} (\mathcal{W} \boldsymbol{\Omega}_m \cdot \nabla R_i + \mathcal{W} \sigma_t R_i) (\boldsymbol{\Omega}_m \cdot \nabla R_j + \sigma_t R_j) d\mathbf{r} \\ + \sum_{i=1}^N d_{i,m} \sum_{p=1}^P \sum_{b=1}^{b_p} c \int_{\Gamma_{m,b}^-} |\boldsymbol{\Omega}_m \cdot \mathbf{n}| \mathcal{W} R_j (R_i - F) d\mathbf{r}_s \\ = \sum_{p=1}^P \sum_{e=1}^{e_p} \int_{V_e} \mathcal{W} Q_m (\boldsymbol{\Omega}_m \cdot \nabla R_j + \sigma_t R_j) d\mathbf{r}. \end{aligned} \quad (6.25)$$

where P is the number of patches in the domain, e_p is the number of knot spans in patch p , b_p is the number of knot spans on the external boundary of patch p , V_e is the volume of knot span e , and $\Gamma_{m,b}^-$ is the external surface of knot span b such that $\boldsymbol{\Omega}_m \cdot \mathbf{n} < 0$.

6.3 Causality in Second-Order Forms of the Neutron Transport Equation

Causality is a fundamental property of neutrons travelling through a medium. It states that information only travels downstream along the direction of neutron travel [127]. The first-order form of the

Region	Macroscopic neutron cross-section (cm^{-1})		Extraneous source strength ($\text{cm}^{-3}.\text{s}^{-1}$)
	σ_t	σ_a	Q
Source	0.0	0.0	1.0
Void	0.0	0.0	0.0
Absorber	10.0	10.0	0.0
Component	1.0	1.0	0.0


Table 6.1: Macroscopic neutron cross-sections and source strengths for the void wall extraneous (fixed) source radiation shielding test case.

neutron transport equation is a hyperbolic PDE and therefore, it obeys causality as information travels along the characteristic curves of the equation [42]. Second-order forms of the neutron transport equation are elliptic PDEs. They have no real characteristic curves and therefore, information doesn't propagate in any one direction. A disturbance in the domain is propagated in all directions at once [265]. The problem specified in Figure 6.2 has been designed to illustrate the qualitative effects that lack of causality can have on the scalar neutron flux and is referred to as the void wall problem. The macroscopic neutron cross-sections and source strengths for the problem are specified in Table 6.1.


Figure 6.3 displays scalar neutron flux profiles for the void wall problem plotted along the line $y = 1.5\text{cm}$. For the given problem the scalar neutron flux before 3cm along the chosen arc length should be constant. This is obvious when it is noted that all neutrons that leave the source region will be travelling through a void for $x < 3\text{cm}$. Given the reflective boundaries and absence of neutron absorption, there is no way for a neutron to leave the system before it reaches $x = 3$. In all of the WLS scalar flux profiles in Figure 6.3 it can be observed that there is a dip in the scalar neutron flux before 3cm. For the given problem, this non-physical dip in the scalar neutron flux can be explained by the lack of causality of the WLS equation. It can be also be observed that this dip in the scalar neutron flux becomes smaller as the problem is refined. This illustrates that causality is a property of the discretised WLS equation only in the fine mesh limit.

One other non-physical feature of the WLS scalar neutron flux solution is the oscillation at about 3.1cm in Figure 6.3a. It is suspected that this spurious oscillation is caused by using a continuous IGA discretisation to approximate a strongly discontinuous solution. Figure 6.3c shows that these spurious oscillations have been resolved with further refinement reinforcing this hypothesis. An area for further research could be the application of patch discontinuous IGA discretisations to second-order forms of the neutron transport equation. With this methodology, discontinuities in the scalar neutron flux can be properly resolved at material interfaces without adding too many extra dof to the problem.


The final feature of Figure 6.3 is the qualitative lack of ray effects in the WLS solution compared to the first-order solution in the region $4\text{cm} < x < 5\text{cm}$. Characteristic curves are the only places that solutions to PDEs are allowed to have discontinuous derivatives. Due to the aforementioned fact that elliptic PDEs have no real characteristic curves, their solutions can not have discontinuous derivatives anywhere. This means that the spurious oscillations seen when the first-order neutron transport equation is angularly discretised with the S_N method are mitigated to some extent when the same problem is solved using S_N discretised second-order forms of the neutron transport equation. However, this does not necessarily mean that the solution of the WLS equation is entirely free from



(a) 1152 elements



(b) 4608 elements



(c) 18432 elements

Figure 6.3: First-order and WLS solutions of the void wall problem. S_{10} angular quadrature has been used and $w_{\max} = 100.0$ has been chosen.

ray effects, merely that they are not qualitatively present. That is, the ray effects have been smoothed out to some extent by numerical diffusion. This could be advantageous for certain applications.

6.4 Numerical Results

The nuclear reactor physics and radiation shielding verification benchmark test cases in this section have been chosen to investigate two aspects of the numerical discretisation method presented in this chapter. First, the solutions of the FE and IGA discretisation methods are compared for several nuclear reactor physics extraneous (fixed) source and fission source problems. These problems will illustrate the magnitude of errors that are incurred when the two-dimensional geometries are approximated by quadrilaterals, as is often the case in traditional FE discretisations. Whilst volume or surface area preservation techniques can be applied to FE meshes, it is impossible to accurately preserve both the fissile mass and the surface area of conic sections in 2D and quadric surfaces in 3D using standard higher order FE basis functions. This is due to the fact that these basis functions are not able to represent conic sections and quadric surfaces exactly [27]. However, the basis functions associated with the IGA method are chosen such that they are able to represent both the volume and surface area of these geometries exactly.

Using curvilinear high order isoparametric FEs to approximate pincell geometries has been investigated for the first-order neutron transport equation [129]. Furthermore, comparisons between curvilinear high order isoparametric FEs and IGA have been performed for the first-order neutron transport equation [186] and the multigroup diffusion equation [175]. It was seen that, even with a similar number of elements per pincell, the error in detector response for FE solutions could be over one order of magnitude greater than similarly discretised IGA solutions. Since isoparametric FEs preserve neither volume or surface area exactly for a pincell, in this chapter volume preserved higher order isoparametric finite elements with straight sides are investigated. This means that the error in approximating the surface area of the geometry is easily determined and fissile mass can be preserved to machine precision.

The second issue investigated in this section is the lack of conservation of neutrons over the entire domain inherent in the LS method and, over void regions in the WLS method. Several of the nuclear reactor physics and radiation shielding verification benchmark test cases will highlight the errors associated with this property. Furthermore, the effect of the magnitude of $w_{\max} \in \mathbb{R}$ on the solution of the WLS equation will be investigated for void and non-void problems. There may also be some errors due to the lack of causality, as was demonstrated in section 6.3. However, these errors are difficult to isolate from discretisation and conservation errors, so shall not be investigated here.

The continuous IGA-WLS spatial discretisation method developed in this chapter has been implemented with a modified version of the Fortran code ICARUS (Isogeometric Continuous self-Adjoint Radiation Using Splines). ICARUS was originally developed to solve the SAAF form of the neutron transport equation [177]. ICARUS is used to generate the WLS-S_N, LS-S_N, and the SAAF-S_N results in this chapter for both IGA and FE discretisations. Inferno is a first-order neutron transport code with local refinement capabilities [174, 186, 187]. It is a discontinuous IGA-S_N code and is used to generate reference solutions to both void and non-void nuclear reactor physics verification benchmark test cases.

In all results in this section the number of degrees of freedom (dof) are considered along the x-axis

Region	Macroscopic neutron cross-section (cm^{-1})		Extraneous source strength ($\text{cm}^{-3}.\text{s}^{-1}$)
	σ_t	σ_a	Q
Fuel	0.764	0.367	1.0

Table 6.2: Macroscopic neutron cross-sections and source strengths for the bare fuel pin extraneous (fixed) source nuclear reactor physics verification test case.





Figure 6.4: Geometry for the bare fuel pin extraneous (fixed) source nuclear reactor physics verification benchmark test case. A vacuum boundary condition is prescribed on the outer boundary of the bare fuel pin.

as opposed to the number of elements. In this chapter the degree of the basis functions will remain fixed and an increase in the number of dof refers to mesh refinement. This is done so that the accuracy of a given degree of basis functions can be fairly compared on a per dof basis. This leads to a more realistic comparison between linear and higher order basis functions than per element.


6.4.1 Bare Fuel Pin Extraneous (Fixed) Source Nuclear Reactor Physics Verification Benchmark Test Case

This nuclear reactor physics test case is taken from work by Hall et al. [163]. The verification test case comprises of a bare fuel pin with an extraneous (fixed) source. The macroscopic cross-section data and the extraneous (fixed) source strength are presented in Table 6.2 and the geometry in Figure 6.4. It has been noted that the LS equation lacks conservation of neutrons except in the fine mesh limit. However, it is difficult to determine where the fine mesh limit will be for particular problems. In an attempt to quantify the magnitude of the error that lack of conservation of neutrons may introduce in the solution, the bare pin test case has been run with several different discretisation techniques. The LS and SAAF equations have been spatially discretised using FEs, volume-preserved FEs, and IGA and angularly discretised using an S_{10} quadrature set. The numerical results for the bare fuel pin verification benchmark test case are presented in Figure 6.5.


The quantity of interest (QoI) considered is the integral of the scalar flux over one quarter of the fuel pin, that is: $\frac{\int_V \phi(\mathbf{r}) d\mathbf{r}}{4}$. The reference solution is generated using Inferno with an S_{10} discrete ordinate discretisation and is given as 0.73596725097 with 2,359,296 dof using quadratic NURBS basis functions.



(a) FE no volume preservation



(b) FE volume preservation



(c) IGA

Figure 6.5: Errors in the integral of the scalar neutron flux over the domain for the bare fuel pin extraneous (fixed) source nuclear reactor physics verification test case. S_{10} angular quadrature has been used.




Figure 6.6: Geometry for the seven group, C5G7 UOX pincell nuclear reactor physics verification benchmark test case. All of the boundary conditions are prescribed to be reflective.

Figure 6.5a shows the SAAF and LS equations discretised using quadratic, non volume-preserved FEs. It can be observed that the SAAF and LS solutions seem to agree, that is the error caused by the lack of conservation of neutrons is small compared to the error caused by approximating the volume and surface area of the geometry. Figure 6.5b shows the SAAF and LS equations discretised using quadratic, volume-preserved FEs. The error between the two formulations can now be observed demonstrating that the ratio of the magnitude of the conservation error relative to the magnitude of the surface area approximation is greater. Figure 6.5c shows the SAAF and LS formulations discretised using IGA. Now the difference in the two formulations can be clearly observed and gives a clearer picture of the magnitude of the conservation error. This leads to the conclusion that, at least for some problems, the error induced by the lack of conservation of neutrons is small when compared to the errors caused by the approximation of geometry.

A second point to note is that the error in the numerical solutions for the non volume preserving FE compared to the IGA differs by about three orders of magnitude for the SAAF solution. The same comparison between volume preserved FE and IGA shows a difference of about 1-2 orders of magnitude. Whilst some of this difference in accuracy is due to the C^1 inter-element continuity of the NURBS basis functions [175, 176], the main reason for this difference is the ability of IGA to exactly represent the problem geometry.

6.4.2 Seven Group, C5G7 UOX Pincell Lattice Fission Source Nuclear Reactor Physics Verification Benchmark Test Case

The seven-group, C5G7 UOX pincell lattice nuclear reactor physics test case has been selected to further explore the relative magnitude of errors due to the lack of conservation of neutrons and geometric approximation. Lack of conservation of neutrons has been seen to more adversely affect nuclear criticality and reactor physics problems, so the C5G7 UOX pincell problem could be considered




Figure 6.7: Error in the K_{eff} for the UOX C5G7 pincell fission source nuclear reactor physics verification test case angularly discretised with S_8 discrete ordinate angular quadrature.

more challenging for the LS method [127]. The geometry of the C5G7 UOX pincell is displayed in Figure 6.6. The macroscopic neutron cross-sections for the fuel pin and moderator are taken from the C5G7 nuclear reactor physics benchmark specification for materials ‘UO₂ fuel-clad’ and ‘moderator’ respectively [210]. The macroscopic transport neutron cross-section has been used in place of the macroscopic total neutron cross-section to account for anisotropic scattering. All of the domain boundaries are prescribed reflective boundary conditions in order to model an infinite lattice of pincells. This C5G7 UOX pincell test case is a nuclear reactor physics problem, so the K_{eff} will be used to investigate the magnitude of errors that arise from different discretisation methodologies and different second-order forms of the neutron transport equation.

The errors in the K_{eff} using various discretisations and second order formulations for this problem can be seen in Figure 6.7. All problems were angularly discretised using an S_8 quadrature set. The same general trend as was observed for the bare fuel pin test case in section 6.4.1 can be seen. That is, the error in the IGA solution is less than both of the FE solutions for both the LS and SAAF formulations. However, what is more apparent here is the error caused by lack of conservation of neutrons. The K_{eff} values from the LS and SAAF formulations differ by at least one order of magnitude for matching levels of spatial discretisation. The errors in the K_{eff} values for the LS-FE results are of a similar order of magnitude as LS-FE results that have been presented in other works [126, 127]. However, the difference in the K_{eff} values for the SAAF-IGA and LS-IGA schemes for similar refinement levels is much larger than the results presented in section 6.4.1. This error is caused by lack of conservation of neutrons and is suspected to be larger for the following reason.

The difference in the scalar flux and the fission source for the SAAF and LS formulations at the i^{th} iteration is denoted $\phi_E^{(i)}$ and $F_E^{(i)}$ respectively. After one outer iteration, $\phi_E^{(1)}$ is small and, since the




Figure 6.8: Geometry of the BWR supercell extraneous (fixed) source nuclear reactor physics verification benchmark test case.

Region	Macroscopic neutron cross-section (cm^{-1})		Extraneous source strength ($\text{cm}^{-3}.\text{s}^{-1}$)
	σ_t	σ_s	Q
Fuel	0.5145	0.3432	0.0
Moderator	0.7557	0.7504	1.0
Burnable poison circle	14.270	0.3400	0.0
Burnable poison annulus	21.60	0.3414	0.0

Table 6.3: Macroscopic neutron cross-sections and source strengths for the BWR supercell extraneous (fixed) source nuclear reactor physics verification test case.

flux is used to generate the fission source, $F_E^{(1)}$ will also be small. However, this difference in fission sources causes a larger difference in fluxes for the second outer iteration. Therefore, $\phi_E^{(2)} > \phi_E^{(1)}$ and similarly, since $\phi_E^{(2)}$ is used to generate the fission source, $F_E^{(2)} > F_E^{(1)}$. This carries on until both solutions have converged and the final error is greater than it would be in the fixed source case.

6.4.3 BWR Supercell Extraneous (Fixed) Source Nuclear Reactor Physics Verification Benchmark Test Case

The boiling water reactor (BWR) supercell nuclear reactor physics test case is based on a mono-energetic, linearly anisotropic scattering problem presented by Wood & Willams [250] and adapted for isotropic scattering by Owens [186]. It contains eight moderated fuel pins that surround a central pin simulating a mixture of fuel and burnable poison. The description of the geometry of the domain is presented in Figure 6.8. All boundaries of the domain are prescribed with reflective boundary conditions to simulate an infinite lattice. The macroscopic neutron cross-section data for the extraneous (fixed) source problem with isotropic scattering are tabulated in Table 6.3. The BWR supercell nu-




Figure 6.9: Error in the QoI for the BWR supercell extraneous (fixed) source nuclear reactor physics verification benchmark test case discretised using an S_6 LCT angular quadrature set.

clear reactor physics test case was selected in order to investigate how the choice of $w_{\max}$ will effect the WLS equation when applied to a realistic problem with strongly varying coefficients. The accuracy of the solution shall be investigated for the SAAF equation, the LS equation, and the WLS equation with several values of $w_{\max}$. The quantity of interest (QoI) for the BWR supercell problem is the total flux in the middle side fuel pin. The reference solution was generated using Inferno with quartic basis functions and 189,064 elements per angle and using a S_6 triangular Legendre-Chebyshev (LCT) discrete ordinate quadrature set.

The comparisons of the error in the QoI for the SAAF, LS, and WLS second-order forms of the neutron transport equation are show in Figure 6.9. All results were run using ICARUS with quadratic basis functions and with S_6 triangular Legendre-Chebyshev angular quadrature. The WLS formulation has been run with three different weightings, $w_{\max} = 0.06, 1.0, 1.5$. These values were chosen as they cause the weighting operator to become active over different regions of the BWR supercell geometry. When $w_{\max} = 1.0$ the weighting operator will be active in the fuel and moderator regions, $w_{\max} = 1.5$ it will be active in only the moderator region, and $w_{\max} = 0.06$ it will be active in fuel, moderator, and the interior of the central pin.

In regions where the weighting operator is not active, that is $\mathcal{W} \neq w_{\max}$, the WLS formulation is identical to the SAAF formulation, up to the weak enforcement of boundary conditions. Given this, it would seem reasonable to predict that the value of $w_{\max}$ that produces a weighting operator that is active over the *smallest* part of the domain would produce results closest to the SAAF solution. Inversely, the value of $w_{\max}$ that produces a weighting operator that is active over the *largest* part of the domain would produce results furthest from the SAAF solution.

This trend can be observed in Figure 6.9. It can also be seen that WLS formulation with $w_{\max} = 1.5$ has a lower error than the SAAF formulation at all points. This is most likely due to the error caused by the lack of conservation of neutrons in certain areas of the problem cancelling with the numerical

Region	Macroscopic neutron cross-section (cm ⁻¹)		Extraneous source strength (cm ⁻³ .s ⁻¹)
	σ_t	σ_a	Q
Absorber	0.5	0.5	0.0
Void	0.0	0.0	0.0
Source	0.0	0.0	1.0

Table 6.4: Macroscopic neutron cross-sections and source strengths for the dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.

discretisation error. The advantageous cancellation of errors may seem like a benefit. However, the inability to predict the cancellation of errors introduces uncertainty into the numerical solution. It is only clear that errors are cancelling out in the case of $w_{\max} = 1.5$ as a ‘reference’ SAAF solution is available. One final point that is evident from Figure 6.9 is that the WLS solutions are superior to the LS solution for all chosen values of $w_{\max}$.

6.4.4 Dog-Leg Duct Extraneous (Fixed) Source Radiation Shielding Verification Benchmark Test Case

The mono-energetic dog-leg duct radiation shielding test case is taken from work done by Ackroyd [138, 147]. The problem consists of an extraneous (fixed) source in a void duct surrounded by a pure absorber. Ackroyd originally used the problem to investigate the efficacy of the iteration and extrapolation methods for solving second-order forms of the neutron transport equation in void regions [147]. In this chapter the dog-leg duct test case is chosen to further investigate the effect of the value of $w_{\max}$ on solutions produced by the WLS equation, in particular for void regions. The geometry of the dog-leg duct test case can be seen in Figure 6.10. The macroscopic neutron cross-sections and extraneous (fixed) source strengths are displayed in Table 6.4. All problems are run with S₁₄ triangular Legendre-Chebyshev angular quadrature. The reference results were produced using Inferno, a first-order neutron transport IGA-S_N code

Figure 6.11 shows the errors for three different quantities of interest (QoI) defined here as:

$$I_n = \int_{V_n} \phi(\mathbf{r}) d\mathbf{r} \quad \text{for } n = 1, 2, 3, \quad (6.26)$$

where V_n is the n^{th} region shown in Figure 6.10. The value of $w_{\max}$ is varied and for the neutron cross-section data in this problem $w_{\max} = 1$ corresponds to the source iteration compatible LS equation originally presented by Hansen [211].

In Figure 6.11 it can be seen that increasing the value of $w_{\max}$ doesn’t necessarily yield a better solution or QoI as was suggested by Figure 6.9. A reason for this loss of accuracy in the QoI can be observed in Figure 6.12a. It can be seen that the solution for $w_{\max} = 100$ exaggerates the peaks and troughs in the computed scalar flux. Similarly, in Figure 6.12b the solution is over estimated over most of the line AB when $w_{\max} = 100$. Furthermore, it is noted that the sign of the error in I_n for $w_{\max} = 1$ and $w_{\max} = 10$ is negative, whilst the sign of the error for $w_{\max} = 100$ is positive. This implies that there is a locally optimal value of $w_{\max}$ located at some crossover point.

To investigate this crossover point further, the algorithm proposed in section 6.1.1 (Figure 6.1) is







Figure 6.10: Geometry of the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.



(a) Error in I_1 .




(b) Error in I_2 .




(c) Error in I_3 .

Figure 6.11: Error in various QoI for the two-dimensional S_{14} dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.



(a) Scalar neutron flux profile along the line OX



(b) Scalar neutron flux profile along the line AB

Figure 6.12: Scalar neutron flux profiles for the S_{14} solution of the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.




Figure 6.13: Error in I_n with $w_{\max} = 16.8$ and an S_{14} angular quadrature set for the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.


Number of elements	Value of $w_{\max}$
480	15.2
1920	16.8
7680	19.0

Table 6.5: Values of $w_{\max}$ produced by the algorithm in Figure 6.1 for different coarse meshes applied to the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case.


applied. I_1 is chosen as the QoI and the first-order neutron transport equation is solved with 1920 elements per angle. The Inferno code is used to determine the solution of the first-order neutron transport equation. Then a series of WLS calculations with the same number of elements per angle are run with varying weights until the error in I_1 is as small as possible. Using this method a value of $w_{\max} = 16.8$ is decided upon to one decimal place. The WLS equation is then solved over a highly refined spatial grid. All WLS solutions were calculated using the code ICARUS. The errors in all I_n are displayed in Figure 6.13.

It can be observed that for the finest mesh, with 123585 dof per angle, the errors in I_n for $\forall n$ are lower than the I_n calculated from other values of $w_{\max}$ with the same resolution presented in Figure 6.11. Furthermore, it can be seen that the error in I_1 is clearly the lowest. This is to be expected as I_1 is the QoI that was used to pick the value of $w_{\max}$. The sudden dip in I_1 that can be observed is located at the mesh refinement level that was used to generate the reference value of I_1 .


In order to investigate the efficacy of the algorithm presented in Figure 6.1 two more ‘coarse’ meshes are chosen for the first-order calculation. The two chosen levels of refinement yield 480 elements and 7680 elements per angle for each coarse mesh. The values of $w_{\max}$ produced for each mesh are tabulated in Table 6.5. The errors in each QoI for the values of $w_{\max}$ tabulated in Table 6.5 are presented in Figure 6.14. It can be seen that the value of $w_{\max}$ calculated from the finest ‘coarse’ mesh causes errors in all QoI to be the smallest. This is not unexpected as the value of Q_f used in the algorithm



(a) Error in I_1 .



(b) Error in I_2 .



(c) Error in I_3 .

Figure 6.14: Error in various QoI for the two-dimensional dog-leg duct extraneous (fixed) source radiation shielding verification benchmark test case for values of $w_{\max}$ selected using the algorithm in Figure 6.1. An S_{14} angular quadrature set was used for all calculations.

shown in Figure 6.1 will become more converged as the ‘coarse’ mesh is refined. This means that the matching performed in the algorithm is against a more converged QoI. However, this does not necessarily imply that the optimal value of $w_{\max}$ changes with mesh refinement in two-dimensions. Further investigation is necessary to determine whether this is the case.

It can be seen that the methodology presented in section 6.1.1 is useful for determining a value of $w_{\max}$ for use in the WLS equation when solving problems involving void regions, particularly where a given QoI is desired. The computational cost of applying this algorithm can be kept low by choosing a sufficiently coarse mesh for the solution of the first-order neutron transport equation. However, Figure 6.14 suggests that using a finer coarse mesh for the first-order transport solution leads to a value of $w_{\max}$ that is closer to optimal.

6.5 Conclusion

In this chapter the weighted least squares (WLS) neutron transport equation was derived and spatially discretised using non-uniform rational B-spline (NURBS) based isogeometric analysis (IGA). The angular domain was discretised using the discrete ordinate (S_N) method. IGA is a spatial discretisation technique that is able to exactly represent the volume and surface area of geometries that can be created using commercial CAD software. This includes conic sections in two-dimensions (2D) and quadric surfaces in three-dimensions (3D).

The WLS formulation is an elliptic second-order form of the neutron transport equation. It can be solved over void regions and is compatible with the source iteration method. The WLS equation involves two weighting parameters. The first is used as a penalty parameter to weakly impose natural boundary conditions, whilst the second is used to weight regions of the problem domain that contain voids. For the correct choices of these weighting parameters the WLS equation is equivalent to the self-adjoint angular flux (SAAF) form of the neutron transport equation for non-void problems. Therefore, the WLS equation conserves neutrons in non-void regions of the domain. However, in void regions the choice of $w_{\max}$ can significantly impact the solution due to lack of conservation of neutrons and the impact on the causality of the WLS formulation [127].

In an attempt to ameliorate the error in void regions, a methodology for determining a value of $w_{\max}$ that involves solving the first-order neutron transport equation over a coarse mesh was presented. This algorithm is computationally efficient compared to solving the WLS on a fine mesh. Furthermore, it was demonstrated that the WLS does not maintain causality of neutrons. This is an intrinsic property of elliptic operators such as second-order forms of the neutron transport equation.

Several nuclear reactor physics and radiation shielding verification benchmark test cases were selected to evaluate the impact of two factors on the error in the solution: the effect of using a NURBS based IGA discretisation, and the effect the choice of $w_{\max}$ can have on two-dimensional problems. A single-group, extraneous (fixed) source bare fuel pin problem and an OECD/NEA seven-group, two-dimension C5G7 UOX pincell fission source problem were discretised using IGA and the FE method. The source iteration compatible least square (LS) equation and the SAAF equation were solved for both of these discretisations. The LS equation is a particular case of the WLS equation where $\mathcal{W} = 1$. It was shown that the error due to the geometric approximation inherent in the FE discretisation was of a larger magnitude than the error due to lack of conservation of neutrons for the bare fuel pin problem. For the fission source problem the two errors were of roughly equal magnitudes.

The effect of the value of $w_{\max}$ on a selected QoI was investigated using a non-void two-dimensional (2D) boiling water reactor (BWR) supercell nuclear reactor physics test case. It was found that the greater the area of the domain where $\mathcal{W} = w_{\max}$, the closer the solution of the WLS equation was to the solution of the LS equation. Inversely, the greater the area of the domain where $\mathcal{W} = \frac{1}{\sigma_t}$, the closer the solution of the WLS equation was to the solution of the SAAF equation. Furthermore, a two-dimensional dog-leg duct radiation shielding test case that included a void duct region was used to investigate the efficacy of the algorithm presented to determine a value of $w_{\max}$. The algorithm requires solving the first-order transport equation over a coarse mesh. It was shown that the more refined the coarse solution was, the more suitable the value of $w_{\max}$ that was calculated.

Chapter 7

Constraint Based Locally Refined IGA Applied to the WLS equation

Much of the work in this chapter, including figures and tables, has been submitted for publication in the Journal of Computational Physics.

In Chapter 6 it was demonstrated that the portion of the domain over which the weighting operator was active dictated the difference between the solution of the weighted least squares (WLS) and self-adjoint angular flux (SAAF) neutron transport equations. Furthermore, it was shown in Chapter 5 that the constraint based local refinement methodology introduces errors into the solution through two mechanisms: by the constraints that are applied along patch interfaces, and by the under resolution of certain areas. Both of these error sources cause perturbations in the solution locally and globally. It was shown that the extent of these perturbations was dependent upon the neutron mean free path.

In this chapter, the interaction between the WLS equation and the constraint based local refinement method will be investigated. The heuristic and goal-based error indicators derived in Chapter 5, referred to as FEI and WEI respectively, will also be used to drive an adaptive mesh refinement (AMR) procedure for the WLS equation

7.1 The Physical (Continuous) Adjoint of the WLS Equation

The adjoint of the WLS equation is derived here in a manner similar to the derivation of the adjoint of the SAAF equation in Section 5.2. For the same reasons given in that section, the physical (continuous) adjoint of the WLS equation is derived as opposed to either the true or mathematical (discrete) adjoint. The physical (continuous) adjoint of the WLS equation gives a measure of a neutrons importance to the system [78] and is derived so that it may be used to drive a goal-based AMR algorithm. In order to avoid the issues mentioned with cross-sections requiring weighting by the adjoint of the neutron flux, this derivation begins with the multigroup approximation of the first-order neutron transport equation. The adjoint of the multigroup first-order neutron transport equation is then taken:

$$-\mathbf{\Omega} \cdot \nabla \psi_g^\dagger(\mathbf{r}, \mathbf{\Omega}) + \sigma_{t,g}(\mathbf{r}) \psi_g^\dagger(\mathbf{r}, \mathbf{\Omega}) = Q_g^\dagger(\mathbf{r}, \mathbf{\Omega}), \quad (7.1)$$

with boundary condition:

$$\psi_g(\mathbf{r}, \boldsymbol{\Omega}) = F_g(\mathbf{r}, \boldsymbol{\Omega}), \quad \text{for } \boldsymbol{\Omega} \cdot \mathbf{n} > 0, \quad (7.2)$$

where F_g represents the incoming boundary source term. In operator form equation (7.1) can be written as:

$$\mathcal{L}_g^\dagger \psi_g^\dagger(\mathbf{r}, \boldsymbol{\Omega}) = Q_g^\dagger(\mathbf{r}, \boldsymbol{\Omega}). \quad (7.3)$$

Equation (7.3) is then operated on by a within group weighting operator $\mathcal{W}_g$ followed by the within group transport operator $\mathcal{L}_g$ giving:

$$\mathcal{L}_g \mathcal{W}_g \mathcal{L}_g^\dagger \psi_g^\dagger(\mathbf{r}, \boldsymbol{\Omega}) = \mathcal{L}_g \mathcal{W}_g Q_g^\dagger(\mathbf{r}, \boldsymbol{\Omega}), \quad (7.4)$$

where the within group transport operators are given by:

$$\begin{aligned} \mathcal{L}_g &= \boldsymbol{\Omega} \cdot \nabla + \sigma_{t,g}(\mathbf{r}), \\ \mathcal{L}_g^\dagger &= -\boldsymbol{\Omega} \cdot \nabla + \sigma_{t,g}(\mathbf{r}). \end{aligned} \quad (7.5)$$

in analogy of expression (6.5). Expanding the operators in equation (7.4) yields the strong form of the physical (continuous) adjoint of the WLS equation:

$$\begin{aligned} & -\boldsymbol{\Omega} \cdot \nabla \mathcal{W}_g \boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger - \sigma_{t,g} \mathcal{W}_g \boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger + \boldsymbol{\Omega} \cdot \nabla (\mathcal{W}_g \sigma_{t,g} \psi_g^\dagger) + \sigma_{t,g}^2 \mathcal{W}_g \psi_g^\dagger \\ & = (\boldsymbol{\Omega} \cdot \nabla + \sigma_{t,g}) \mathcal{W}_g Q_g^\dagger. \end{aligned} \quad (7.6)$$

To find the weak form of the physical (continuous) adjoint of the WLS equation, equation (7.6) is multiplied by a test function v and integration is performed over the problem domain. This yields:

$$\begin{aligned} & \int_V -v \boldsymbol{\Omega} \cdot \nabla \mathcal{W}_g \boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger - v \sigma_{t,g} \mathcal{W}_g \boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger + v \boldsymbol{\Omega} \cdot \nabla (\mathcal{W}_g \sigma_{t,g} \psi_g^\dagger) \\ & + v \sigma_{t,g}^2 \mathcal{W}_g \psi_g^\dagger - v (\boldsymbol{\Omega} \cdot \nabla + \sigma_{t,g}) \mathcal{W}_g Q_g^\dagger d\mathbf{r} = 0. \end{aligned} \quad (7.7)$$

Integration by parts is performed on the first, third, and fifth terms, which adds a surface integral to the equation:

$$\begin{aligned} & \int_V (\boldsymbol{\Omega} \cdot \nabla v) \mathcal{W}_g (\boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger) - (\boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger) v \sigma_{t,g} \mathcal{W}_g - (\mathcal{W}_g \sigma_{t,g} \psi_g^\dagger) (\boldsymbol{\Omega} \cdot \nabla v) \\ & + \sigma_{t,g}^2 v \mathcal{W}_g \psi_g^\dagger + (\boldsymbol{\Omega} \cdot \nabla v) \mathcal{W}_g Q_g^\dagger - \sigma_{t,g} v \mathcal{W}_g Q_g^\dagger d\mathbf{r} \\ & + \int_\Gamma (\boldsymbol{\Omega} \cdot \mathbf{n}) v \mathcal{W}_g [-\boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger + \sigma_{t,g} \psi_g^\dagger - Q_g^\dagger] d\mathbf{r}_s = 0. \end{aligned} \quad (7.8)$$

It can then be noted that all of the terms under the surface integral ($\int_\Gamma d\mathbf{r}_s$) satisfy the adjoint first-order neutron transport equation given in equation (7.1). This leaves the following equation:

$$\begin{aligned} & \int_V (\boldsymbol{\Omega} \cdot \nabla v) \mathcal{W}_g (\boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger) - (\boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger) v \sigma_{t,g} \mathcal{W}_g - (\mathcal{W}_g \sigma_{t,g} \psi_g^\dagger) (\boldsymbol{\Omega} \cdot \nabla v) \\ & + \sigma_{t,g}^2 v \mathcal{W}_g \psi_g^\dagger + (\boldsymbol{\Omega} \cdot \nabla v) \mathcal{W}_g Q_g^\dagger - \sigma_{t,g} v \mathcal{W}_g Q_g^\dagger d\mathbf{r} = 0. \end{aligned} \quad (7.9)$$

In order to be able to weakly satisfy the natural boundary conditions that are given in expression (7.2) for the adjoint first-order neutron transport equation, a boundary term is added and enforced by

the penalty parameter $c_g \in \mathbb{R}$. The final weak form of the physical (continuous) adjoint of the WLS equation is then given by:

$$\begin{aligned} & \int_V (\boldsymbol{\Omega} \cdot \nabla v) \mathcal{W}_g (\boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger) - (\boldsymbol{\Omega} \cdot \nabla \psi_g^\dagger) v \sigma_{t,g} \mathcal{W}_g - (\mathcal{W}_g \sigma_{t,g} \psi_g^\dagger) (\boldsymbol{\Omega} \cdot \nabla v) \\ & + \sigma_{t,g}^2 v \mathcal{W}_g \psi_g^\dagger + (\boldsymbol{\Omega} \cdot \nabla v) \mathcal{W}_g Q_g^\dagger - \sigma_{t,g} v \mathcal{W}_g Q_g^\dagger d\mathbf{r} \\ & + c_g \int_{\Gamma^+} H(\boldsymbol{\Omega} \cdot \mathbf{n})(\boldsymbol{\Omega} \cdot \mathbf{n}) \mathcal{W}_g v (\psi_g^\dagger - F_g) d\mathbf{r}_s = 0, \end{aligned} \quad (7.10)$$

where $\Gamma^+ = \{\Gamma \mid \boldsymbol{\Omega} \cdot \mathbf{n} > 0\}$. The values of c_g and $\mathcal{W}_g$ are chosen as the multigroup analogies of expressions (6.16) and (6.15). That is:

$$c_g = \max(\sigma_{t,g}, \frac{1}{h}), \quad \mathcal{W}_g = \min(\frac{1}{\sigma_{t,g}}, w_{\max}). \quad (7.11)$$

7.2 Generation of the Constraint Operator Using Bézier Extraction

Typically the constraint operator is generated using a form of the Cox de Boor recursion algorithm. Given the knot vector Ξ , and an extended knot vector $\bar{\Xi}$, such that $\Xi \subset \bar{\Xi}$, the one dimensional knot insertion matrix T is generated by:

$$\begin{aligned} T_{ij}^0 &= \begin{cases} 1 & \bar{\xi}_i \in [\xi_j, \xi_{j+1}), \\ 0 & \text{otherwise,} \end{cases} \\ T_{ij}^{q+1} &= \frac{\bar{\xi}_{i+q} - \xi_j}{\xi_{j+q} - \xi_j} T_{ij}^q + \frac{\xi_{j+q+1} - \bar{\xi}_{i+q}}{\xi_{j+q+1} - \xi_{j+1}} T_{ij+1}^q \quad \text{for } q = 0, 1, 2, \dots, p-1, \end{aligned} \quad (7.12)$$

where $\xi_j \in \Xi$, $\bar{\xi}_j \in \bar{\Xi}$ and T_{ij}^n is the entry in the i^{th} row and j^{th} column of the matrix T_f^n . The two dimensional knot insertion matrix T^{2d} is generated by applying equation (7.12) to the knot vector in each parametric dimension and the resulting matrices are combined using a tensor product structure. The entries relevant to T_f can then be extracted from T^{2d} . However, since open knot vectors are considered in this work, the surfaces of d -dimensional NURBS objects are themselves $(d-1)$ -dimensional NURBS objects. Therefore, in two dimensions the T_f matrices can be formed directly by applying equation (7.12) to the relevant sides of matching patches. This saves computational time as the unnecessary components do not have to be formed and no tensor product operation is required.

In this section a different methodology is investigated for the generation of the T_f operator. The Bézier extraction framework was presented by Thomas et al. [178] and can be used to easily turn FE codes into IGA codes without being too intrusive. Bézier extraction can be used as an alternative to the Cox de Boor algorithm in equation (7.12) to generate T_f . This method based on Bézier extraction is explained below. However the method and its nomenclature corresponds to algorithm 4.1 in the paper by Thomas et al. [178], which is recommended if more information is required.

7.2.1 Methodology

Two elements are considered, e and $\bar{e}$, with $\bar{e}$ a sub element of e . That is $\bar{e} \subset e$. For the purpose of generating T_f the element $\bar{e}$ is considered to be in the fine patch and the element e is considered to




Figure 7.1: The elements e and $\bar{e}$ used for the generation of T_f^e using Bézier extraction. As the fine patch is assumed to be a refinement of the coarse patch both elements are illustrated on the same patch to signify that $\bar{e} \subset e$.

be in the coarse patch. This is illustrated in Figure 7.1. The aim is to generate the operator T_f^e such that when it is applied to the control points relevant to element e on the coarse patch, it yields the control points relevant to element $\bar{e}$ on the fine patch, that is:

$$\mathbf{B}_{\bar{e}} = T_f^e \mathbf{B}_e. \quad (7.13)$$

The operator T_f^e is given as a matrix by:

$$T_f^e = (R^{\bar{e}})^T A_{e \rightarrow \bar{e}} (C^e)^T. \quad (7.14)$$

The matrices in the above prescription are:

- C^e is the matrix representing the transformation from the NURBS basis functions over e to the Bernstein basis over e . It is referred to as the *Bézier element extraction operator* and the procedure to generate C^e can be found in Borden et al. [157].
- $A_{e \rightarrow \bar{e}}$ is the matrix representing the transformation between the Bernstein basis over e and the Bernstein basis over $\bar{e}$. Information concerning the generation of $A_{e \rightarrow \bar{e}}$ can be found in Farouki & Neff [189]. It is referred to here as the *Bernstein transformation matrix*.
- $R^{\bar{e}}$ is the matrix representing the transformation from the Bernstein basis over $\bar{e}$ to the NURBS basis functions over $\bar{e}$. It is defined as such that

$$R^{\bar{e}} = (C^{\bar{e}})^{-1},$$

and is referred to as the *spline element reconstruction operator*.

For a two-dimensional NURBS patch T_f^e represents the knot insertion matrix for element e on the one-dimensional B-spline surface. By assumption $\bar{e} \subset e$, and therefore the Bézier projection will be exact so no smoothing or weighting algorithms are necessary [157]. When the T_f^e matrices are generated for

all elements in a patch and assembled using the global connectivity information the resulting matrix operator will be identical the one generated using the Cox de Boor algorithm.

7.3 Results

7.3.1 Timing Comparison of the Cox DeBoor and Bézier Extraction Methodologies


As mentioned above the operator T_f will be the same whether it is generated using the Cox de Boor (CDB) algorithm or Bézier extraction. Therefore, in order to evaluate the efficacy of the Bézier extraction (BE) methodology proposed in section 7.2 a timing comparison will be made. The algorithm used for the CBD methodology is a slightly adapted version of algorithm 5.1 taken from the 2nd edition of The NURBS Book by L. Piegl and W. Tiller [154]. The case of two square patches next to each other is considered. One patch gets steadily more refined and the other patch is left unrefined. Since the patches only share one interface only one T_f matrix is needed. Figure 7.2 displays the wall timings for the generation of the T_f matrix using CDB and BE run 1000 times to help remove some uncertainty in the timings. Linear through quartic basis functions are considered. For the generation of T_f using BE the number of element matrices is equal to the number of elements.

It can be seen that the CDB method is uniformly faster than the BE method. Indeed, in at least half of the cases the inversion of the Bézier element extraction operator needed for the BE method took longer than the CDB method. The Bézier element extraction operator is a $p + 1 \times p + 1$ matrix. For $p \leq 3$ well known closed form analytical expressions for the matrix inverse were used. When $p \geq 4$ direct inversion was implemented using LAPACK. As the polynomial order increases the time to invert the matrix required for the Bézier element extraction operator will also increase. Therefore, it is expected that as the order of the basis functions increases further, the time required to invert the matrix will begin to dominate the BE process.


However, this doesn't rule out the BE method for two reasons. First, the BE method allows for the formation of the element wise knot insertion matrices T_f^e , whereas the CDB method forms the patch wise T_f matrix directly. If one wanted to insert the dense element stiffness matrices into the global stiffness matrix directly rather than form the intermediary sparse patch wise stiffness matrices then the BE method would be suitable. Skipping the formation of the patch wise stiffness matrices could save more time than the time lost by using the BE methodology. Second, if non NURBS basis functions are used in the IGA methodology then the BE method might be the fastest available algorithm. Certainly the BE methodology is more flexible than the CDB. For example, the BE method can be used to refine T-splines, LR B-splines, or hierarchical B-splines by choosing the correct components listed in section 7.2.1 [178].

7.3.2 KAIST Pincell


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 pincell problem is considered. The geometry of the pincell is displayed in Figure 7.3. The neutron-cross section data for all 7 groups can be found in [266]. This nuclear reactor physics pincell benchmark is more realistic when compared to previous pincells considered in this thesis. A section of cladding and a small airgap




(a) Linear basis functions, $p = 1$.



(b) Quadratic basis functions, $p = 2$.



(c) Cubic basis functions, $p = 3$.



(d) Quartic basis functions, $p = 4$.

Figure 7.2: Timings for the generation of the T_f matrix run 1000 times. Four different orders of basis function are displayed. For each order the timing of the Cox-DeBoor algorithm (CDB), Bézier Extraction algorithm (BE), and the time taken to invert the Bézier element extraction operator C^e required in the BE algorithm are plotted.




Figure 7.3: Geometry of the KAIST pincell nuclear reactor physics verification benchmark test case. Reflective boundary conditions are prescribed on all surfaces. The pincell consists of fuel surrounded by a moderator. There is a layer of cladding between the fuel and moderator and a small air gap is present. This air gap is included in real reactors to allow fuel pellets to swell without causing cracking in the cladding.

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 airgap can cause the even/odd parity and SAAF equations 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 due to the weighting operator $\mathcal{W}$.

The disadvantage factor for each group is considered. For this particular geometry the disadvantage factor for each group is given by the following expression:

$$\zeta_g = \frac{\int_{\text{moderator}} \phi_g(\mathbf{r}) d\mathbf{r}}{0.878778} \frac{0.5268814}{\int_{\text{fuel}} \phi_g(\mathbf{r}) d\mathbf{r}}. \quad (7.15)$$

The error in ζ_g for all groups are plotted in Figure 7.5 and have been generated by solving the WLS equation spatially discretised with quadratic NURBS based IGA, angularly discretised with S_8 LCT discrete ordinates and with $w_{\max} = 1.0$. Two refinement schemes have been used. For uniform refinement all areas have been refined equally. In the local case, everywhere but the void has been refined as is displayed in Figure 7.4.

From Figures 7.5a and 7.5b it can be seen 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 a shallower gradient. However, this increase in error and drop in convergence rate does not show the entire picture. Table 7.1 displays the condition number of a SSOR preconditioned spatial system matrix for a given angle and group in both the uniform and locally refined cases. It can clearly be seen that κ is between one and three orders of magnitude less for the locally refined





Figure 7.4: Element boundaries of the KAIST pincell nuclear reactor physics verification benchmark test case as defined by the knot spans. This is the refinement referred to in the local case. The thin void region is unrefined whilst the neighbouring fuel pin and cladding are highly refined.

case. As mentioned previously in section 3.6, a GMRES solver inside of PETSc is used to estimate κ . Furthermore, κ is not a direct measure of how long it will take to solve a system of linear equations but is an indication of the time required to invert a matrix.


For example, in the second row of Table 7.1 the ratio $\frac{\kappa_{\text{uniform}}}{\kappa_{\text{local}}} = 211.5$, but solution time of the the uniformly refined case was only 2.36 times longer than for the locally refined case. However, the increase in κ and runtime versus the potential increase in error is worth considering when flagging regions for refinement. For example, the error introduced in ζ by not refining the gas region is quite severe. However, ζ is a global quantity in this problem. If the QoI was a quantity that is 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.

7.3.3 Reed Cell Problem

The Reed cell problem is a one dimensional fixed source nuclear radiation shielding benchmark that contains a void region [267]. The geometry for the problem is displayed in figure 7.6 and the material properties are specified in Table 7.2. The problem is solved using the WLS equation spatially discretised with IGA and angularly discretised with a S_8 discrete ordinate 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 neutron scalar flux. The flux profiles for the first-order solution, analytical solution, and several values of w_{max} are displayed in Figure 7.7. It can be seen that the higher the value of w_{max} the more qualitatively accurate the solution, particularly at material interfaces. However, there



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



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

Figure 7.5: Error in the disadvantage factor ζ for two refinement schemes. ‘Uniform’ refines all areas uniformly, and ‘local’ refines everywhere except the void. The WLS equation spatially discretised with quadratic IGA, angularly discretised with S_8 LCT discrete ordinates and with $w_{\max} = 1.0$ was used to generate results.

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



Figure 7.6: Geometry and boundary conditions for the one dimensional Reed cell radiation shielding benchmark test case.

Region	Macroscopic neutron cross-section data (cm ⁻¹)		Source data (neutrons.cm ⁻³)
	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 7.2: Macroscopic neutron cross-section data for the one dimensional Reed cell radiation shielding benchmark test case.

is still a quantitative difference between the $w_{\max} = 100$ solution and the analytical solution. The S_8 analytical solution of the Reed cell problem is available in the following reference [60].

Two AMR schemes have been applied to the one dimensional Reed cell problem to illustrate the potential of local refinement when applied to neutron transport problems that contain voids. A heuristic error indicator (FEI) and a goal based error indicator (WEI) have been used to drive the AMR algorithm. Both of these error indicators have been detailed in Chapter 5 and the physical (continuous) adjoint of the WLS equation, derived in section 7.1, will be solved in order to calculate the WEI in each patch. During each AMR iteration one patch is selected for refinement. Each patch represents a material, so there are five patches. The L_2 error in the scalar neutron flux for the locally refined cases and uniformly refined case are displayed in Figure 7.8.

Figure 7.8a 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 levels of refinement whilst they are better in other cases. 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 neighbouring regions 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. Furthermore, neither the FEI nor the WEI has been tuned to reduce the functional $\int_0^8 |\phi_{\text{ref}}(x) - \phi(x)|^2 dx$ which represents the global L_2 error plotted in Figure 7.8a. The heuristic FEI scheme refines based on the error between the solution and that same solution projected to a coarse grid within a patch. This is given by :

$$\|\phi^{h+1}(x) - \Pi_h^{h+1}\phi^{h+1}(x)\|_{H_1},$$





Figure 7.7: Plot of the scalar neutron flux for the reed cell problem with an S_8 discrete ordinate quadrature. All problems have been uniformly refined. A larger value of $w_{\max}$ results in a more qualitatively accurate solution.


as detailed in section 5.3.1. Therefore, the FEI scheme could be considered to minimise the patch local H_1 error rather than the global L_2 error. The goal for the WEI scheme has been set as the quantity $\int_3^5 \phi(x)dx$.

The values of the L_2 error in the void region ($\int_3^5 |\phi_{\text{ref}}(x) - \phi(x)|^2 dx$) have been plotted in Figure 7.8b. It can be seen that the WEI and FEI solutions are superior to the uniformly 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 local H_1 error in all patches simultaneously. This means that neither the H_1 or L_2 error in any one patch is necessarily minimised as the solutions in neighbouring patches affect one another. Despite the fact that the goal for the WEI solution is the integral quantity $\int_3^5 \phi(x)dx$ 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 7.9 shows the error in the integral of the scalar neutron flux. Figure 7.9b shows the error in $\int_3^5 \phi(x)dx$ 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 7.7 that the $w_{\max} = 100$ solution is very similar to the analytical solution for $3 \leq x \leq 5$. Figure 7.9a shows the error in $\int_0^8 \phi(x)dx$. The behaviour of the uniformly refined solution implies that the integral quantity is not converging to the analytical value. From Figure 7.7 it can be seen that the WLS solution over and under estimates the analytical solution in various places. These errors do not cancel out and result in the convergence of the integral of the flux to the wrong value and a constant error level. The convergence of the integral of the flux to the wrong value means that not much can be concluded from the results of the FEI and WEI refinement schemes other than that, initially, both local refinement



(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 7.8: L_2 error in the neutron scalar flux for several refinement schemes. Two AMR schemes are used, FEI which is a heuristic based scheme and WEI which is a goal-based error measure. The goal for the WEI has been set as $\int_3^5 \phi(x) dx$. The L_2 error is calculated using an analytical S_8 solution.

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 7.3: Macroscopic neutron cross-section data for the two dimensional ANL nuclear reactor physics benchmark test case.

schemes have a higher rate of convergence than the uniformly refined solution.

Further to the results in section 7.3.2, Figure 7.10 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 [164] and in cases where low values of σ_t have been used to fill void or near void regions an increase in solve time, and therefore condition number, can be observed [253]. Therefore, 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 7.3.2.


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 patches. That is, the more dof that have to be constrained between patches, the greater the increase in κ . However, in the case where region 3 is unrefined, the increase in κ caused by having to apply constraints between neighbouring patches is much smaller in magnitude than the decrease in κ caused 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 system where possible.


7.3.4 Cartesian ANL Quarter Core Nuclear Reactor Physics Benchmark Test Case

A one group two dimensional nuclear reactor physics benchmark has been selected to compare the effectiveness of the FEI and WEI schemes for a monoenergetic void problem. The geometry and neutron cross section data are specified in Figure 7.11 and Table 7.3 respectively. The problem is taken from a paper by D. Lathouwers [240] and is based on the “Multi-dimensional (x - y - z) LWR Model” benchmark from Argonne National Laboratory [207].

The problem has been spatially discretised with linear NURBS and angularly discretised with S_8



(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 7.9: L_2 error in the neutron scalar flux for several refinement schemes. Two AMR schemes are used, FEI which is a heuristic based scheme and WEI which is a goal-based error measure. The goal for the WEI has been set as $\int_3^5 \phi(x) dx$.




Figure 7.10: Condition number of the global stiffness matrix for various refinement cases.




Figure 7.11: Geometry for the IAEA problem. Quadrilateral areas coloured grey correspond to region III

LCT discrete ordinates. A reference solution has been taken from A. Owens' Thesis [29]. Results are presented for the same refinement methodologies as in section 7.3.3 that is: uniform refinement, FEI driven AMR, and WEI driven AMR. The goal for the WEI is the K_{eff} . Both AMR schemes use the list method where 10% of the patches that make up the domain are flagged for refinement. The error in the K_{eff} for four cases is displayed in Figure 7.12. There are two cases for uniform refinement, one solving the problem over the geometry shown in Figure 7.11, and one over the same problem but with region V not modelled which is pictured in Figure 3.16. This means the effect of the weighting operator on the solution and its convergence rate can be investigated. When the void region is not modelled the WLS equation becomes equivalent to the SAAF equation. Therefore, the case '*Uniform $p = 1$, void*' is equivalent to solving the problem using SAAF.

It can be seen that the modelling of the void region does not effect the convergence rate of the K_{eff} . However, the case where the void region is not modelled produces a slightly smaller error in the K_{eff} . The downside of not modelling the external void region is that a QoI inside the external void region can not be calculated, for example, if the dosage to the operator was desired, it would be unavailable in the case where the void is not modelled.

Figure 7.12 also shows that the error in the K_{eff} is almost identical in the FEI and WEI cases. This is not unexpected as the FEI method is trying to reduce the error in the scalar neutron flux in the H_1 sense. Since the K_{eff} is a global quantity that depends on the accuracy in the scalar neutron flux between power iterations, it makes sense that the FEI method would accurately calculate it. The goal for the WEI method was set as the K_{eff} however, it requires roughly twice the work as the FEI method as it solves the forward and adjoint WLS equation. The meshes for the third AMR iteration for the FEI method and the WEI method are displayed in Figures 7.13 and 7.14 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 is superior to the WEI. This is because it gives roughly the same answer for roughly half of the computational work.




Figure 7.12: Error in the K_{eff} for several refinement schemes. The WLS equation with $w_{\text{max}} = 1$ has been spatially discretised with linear IGA and angularly discretised with S_8 discrete ordinates. The goal for the WEI method has been set as the K_{eff} . The no void description refers to the ANL problem with the void in region V not modelled.




Figure 7.13: Mesh for the ANL problem for the third iteration of the FEI driven AMR scheme.




Figure 7.14: Mesh for the ANL problem for the third iteration of the WEI driven AMR scheme. The goal functional has been set as the K_{eff} .

7.4 Conclusion

In this chapter the application of constraint based local refinement to the Weighted Least Squares (WLS) neutron transport equation was considered. In order to apply the goal-based error indicator, the physical (continuous) adjoint of the WLS equation was derived from the adjoint first-order neutron transport equation. The Bézier Extraction (BE) methodology was described and used to generate the element-wise constraint operators. These element operators can be combined to form the same patch-wise constraint operators generated by the Cox de Boor (CDB) algorithm. It was demonstrated that the generation of the patch-wise constraint operators using BE takes longer than the generation of the same operator using CDB. However, it is posited that the BE methodology could provide savings further down the line as the element-wise constraint operators can be used to form global, locally refined stiffness matrices directly, instead of having to form the intermediate patch-wise stiffness matrices as is necessary when CBD is used.

The KAIST pincell nuclear reactor physics benchmark test case demonstrated that noticeable errors are introduced into the global solution by under refining geometrically small regions in the domain. In this case the void gas gap was left unrefined. However, it was also shown that the spectral condition number, and therefore the time taken to solve the system of linear equations, increased greatly as the void gas gap was refined. The Reed cell problem further demonstrated that the refinement of void regions causes a large increase in the spectral condition number of the global matrix system. As the WLS equation is equivalent to the SAAF equation in regions where the weighting operator is not active the analysis of the constraint based local refinement method applied to the SAAF method in Chapter 5 is also valid for the WLS equation.

Adaptive mesh refinement (AMR) was applied to the Reed cell problem. It was shown that for local quantities of interest (QoI) the rate of convergence can stagnate when a heuristic error indicator, the FEI, is used. Using a goal-based error indicator, the WEI, solved this issue. It was also demonstrated that FEI and WEI methods produced similar levels of error in global QoI. The ANL quarter core nuclear reactor physics benchmark further reinforced this finding. It was shown that both FEI and WEI produced similar meshes when the goal functional of the WEI was global in nature. This was reflected in the similar level of error in the K_{eff} generated by the FEI and WEI schemes. It is concluded that if the QoI is global then the FEI methodology is suitable. However, if a local QoI is desired then the WEI methodology with a suitable goal functions should be employed.

Chapter 8

Conclusions

8.1 Conclusion

Isogeometric analysis (IGA) was detailed in Chapter 3. It was applied to the neutron diffusion equation and used to demonstrate that the IGA discretised system can deliver the same theoretical convergence rates as the finite element (FE) method. It was shown that, for a non-Cartesian geometry, the convergence rate of the FE discretised equation was limited by geometric error. In both Cartesian and non-Cartesian geometries the IGA method outperformed the FE method with respect to accuracy per degree of freedom (dof) for all 2D benchmark test cases considered.

In Chapter 4, a variational, source iteration compatible formulation of the self-adjoint angular flux (SAAF) equation was derived. It was shown to be equivalent to the algebraic derivation when both formulations were converted into their Galerkin weak forms. IGA was applied to the discrete ordinate SAAF equation. A method of manufactured solutions (MMS) test case demonstrated theoretical convergence rates of $O(h^{p+1})$. The IGA-SAAF equation was compared to the FE-SAAF equation and it was shown that the IGA-SAAF equation was superior, with respect to accuracy per degree of freedom, for several nuclear reactor physics benchmarks including the IAEA swimming pool test case and the OECD/NEA C5G7 quarter core physics test case. Furthermore, it was explicitly shown that the convergence rate of the FE-SAAF equation was limited by geometric errors.

A constraint based local refinement methodology was derived in Chapter 5. Furthermore, two adaptive mesh refinement (AMR) procedures were described: one based upon a heuristic error indicator (FEI), and the other using an adjoint-weighted goal-based error measure (WEI). In order to utilise the WEI the physical adjoint of the SAAF equation was derived. An MMS test case demonstrated that the convergence rates of the L_2 error of locally refined geometries tends towards the rate of convergence of the L_2 error of uniformly refined geometries. Furthermore, it was shown that the propagation of errors introduced by constraint based local refinement and the under resolution of certain areas depends upon the neutron mean free path. The IAEA swimming pool nuclear reactor physics test case was revisited and demonstrated that the FEI AMR scheme produces a superior K_{eff} when compared to a uniformly refined system. Finally the Lathouwers radiation shielding problem demonstrated that both the FEI and WEI schemes lead to accurate global quantities of interest (QoI). However, the adjoint-weighting of the WEI scheme allows it to produce bespoke meshes for the accurate evaluation of local QoIs.

In Chapter 6 the weighted least squares (WLS) equation was derived. The WLS equation is a

second-order neutron transport equation that can be solved in void regions. A weighting operator is chosen for void regions which influences the behaviour of the solution in, and near, these regions. This is because neutrons are not conserved in void regions. However, in non-void regions the WLS equation is equivalent to the SAAF equation, up to the implementation of boundary conditions. An algorithm for choosing the strength of the weighting operator was presented. It utilises a coarse first-order neutron transport solution and picks the strength based upon a specified QoI. The WLS equation was spatially discretised with NURBS based IGA and angularly discretised using the discrete ordinate method. The unweighted least squares (LS) neutron transport equation does not conserve neutrons in any region of the domain. The solution of the FE-LS, IGA-LS, FE-SAAF, IGA-SAAF equations were compared to demonstrate that errors introduced in the solution by inexact geometry were greater in magnitude than errors introduced by the lack of conservation of neutrons. However, it was also shown that the error caused by the lack of conservation of neutrons was greater for criticality problems than fixed source problems. The BWR supercell benchmark was used to demonstrate that the less of the domain the weighting operator is active over the closer the solution of the IGA-WLS equation is to the IGA-SAAF equation. Finally the efficacy of the algorithm for choosing the value of the weighting operator was demonstrated using the dog-leg duct radiation shielding problem. Furthermore, the qualitative differences between the IGA-first-order and IGA-WLS solutions was illustrated for the dog-leg duct problem.

In Chapter 7 the physical adjoint of the WLS equation was derived so that the WEI scheme can be used to drive AMR of the WLS equation. The use of Bézier extraction (BE) to generate the constraint operator T_f was investigated. It was shown that the BE methodology is slower than the widely used Cox de Boor (CDB) algorithm and that the BE method is limited by dense matrix inversion. However, the BE method is much more flexible than the CDB algorithm and could provide performance savings later in a program due to requiring fewer memory allocations. The locally refined IGA-WLS equation was applied to the KAIST pincell. The magnitude of the error introduced in the solution by not refining the gas gap was shown to be sizeable. However, it was also shown that the refinement of the gas gap led to a large increase in the spectral condition number of the global system of linear equations. The FEI and WEI driven AMR methods from Chapter 5 were applied to the one dimensional Reed cell problem. It was shown that the FEI method can lead to stagnation of convergence in the L_2 errors in certain regions. Using the WEI method with the correct goal functional can prevent this stagnation. It was also demonstrated that the qualitative difference in the first-order reference solution and the WLS solution can lead to the incorrect convergence of certain quantities. Both AMR algorithms were applied to the ANL quarter core nuclear reactor physics benchmark. It was demonstrated that if the QoI is global, then the FEI and WEI methodologies perform similarly. However, the WEI requires the solution of a forward and adjoint transport equation so should potentially only be used if local quantities are of interest.

8.2 Future Work

8.2.1 Trigonometric and Exponential Splines

The B-spline and NURBS basis functions that were considered in this thesis are piecewise polynomial in nature. This leads to the exact representation of many algebraic curves such as conic sections.

However, the B-spline formulation can be extended by considering Generalised B-splines (GB-splines). By including sinusoidal or exponential functions into the piecewise nature of B-splines, trigonometric and exponential splines, which are types of GB-splines, can be derived. These allow the exact representation of the same curves as B-splines and NURBS but can also exactly represent transcendental curves such as helices and cycloids, as well as curves that can be characterised by trigonometric or hyperbolic functions. [268].

The application of trigonometric and exponential splines to Uncertainty Quantification (UQ) is of interest in the future. They would enable the exact representation of surface roughness and errors in components due to manufacturing tolerances. By modelling these imperfections exactly, fewer unknowns should be necessary to quantify uncertainties.

8.2.2 Geometric Spatial Multigrid with Energy Dependent Meshes and Adaptivity

Multigrid solvers have been developed for IGA in order to provide scalable solution techniques with respect to the polynomial order of the basis functions, as well as the size of an element. These multigrid solvers can be used as preconditioners for the conjugate gradient algorithm. Due to the structured nature of elements within a NURBS patch, geometric multigrid (GMG) methods are often used. On the other hand, algebraic multigrid (AMG) methods are used for FEs due to the unstructured nature of FE meshes. Whilst AMG solvers are black box in nature due to the geometry agnostic methods they use to coarsen the system, GMG solvers are typically faster for the same problem, assuming a method for coarsening the problem based on the geometry can be found.

Whilst GMG solvers have been well developed for single patch systems [165, 269], the extension to multipatch geometries is typically achieved using domain decomposition techniques. For example, the Balancing Neumann-Neumann (BNN) or Balancing Domain Decomposition by Constraints (BDDC) methods [169, 170, 270]. By combining the adaptive mesh refinement algorithms detailed in Chapter 5, energy dependent meshes, and scalable domain decomposition GMG solvers, a scalable solution algorithm that minimises the number of unknowns required to evaluate a QoI to a given tolerance could be developed.

8.2.3 Spatial Domain Decomposition on Shared and Distributed Memory Systems

In this thesis, the Modern Fortran code ICARUS was run entirely in serial. A parallel implementation was developed for distributed systems using MPI and PETSc. It was based upon a spatial domain decomposition and significantly improved assembly time for the systems of linear equations. However, it did not offer good scaling for the solution of the discrete ordinate system.

The first-order neutron transport IGA reference code, Inferno [29], allowed for the simultaneous solution of several angular neutron fluxes via a sweep based algorithm. This was achieved for shared memory architectures using OpenMP. However, the geometry creation, refinement algorithms, and the sweep ordering algorithm were not parallelised. The OpenMP methodology is not scalable for two reasons. First, large modern HPC systems are built upon many interconnected shared memory nodes that form a distributed memory system [271]. Therefore, using OpenMP, only one node is able to be utilised on a modern system. Second, in Inferno, only one thread was assigned to each ordinate

direction. Since there will rarely be more than 1000 directions in which to solve, if there are more than 1000 threads available then not all threads will be utilised by the program simultaneously.

For an optimal parallel implementation, both spatial domain decomposition and the simultaneous solution in several directions should be combined. By having several threads solving in each direction the downsides of the current parallel implementations of ICARUS and Inferno can be removed. This could be achieved with a variety of parallel APIs that could be combined in a hybrid method. For CPU based systems: OpenMP, MPI, OpenACC, or Modern Fortran Coarrays [272] could be used. If a GPU based system is where the code will be run: CUDA, DirectCompute, or OpenCL could be utilised.

8.2.4 Spherical NURBS for Angular IGA

Both methods mentioned in Chapter 2 have drawbacks when used to discretise the angular dependence of the neutron transport equation. The discrete ordinate method suffers from ray effects which are never entirely absent, regardless of how many directions the transport equation is solved in [66]. On the other hand, the spherical harmonic method is unable to satisfy certain boundary conditions exactly and is harder to derive and implement [273]. Angular FE has been used to angularly discretise the neutron transport and radiative transport equations [274, 275] and treads a line between the two as it has a simpler implementation and derivation. However, because the interpolatory basis functions used to represent that solid angle are typically not rotationally invariant, the ray effect will still be present [276].

Bernstein-Bézier polynomials have been used to form triangulations over the unit sphere. They can be formulated such that they have local support over triangular elements and also remove spurious oscillations in interpolated solutions [277, 278]. Spherical splines are formed from a piecewise combination of Bernstein-Bézier polynomials over the sphere [279]. Using spherical splines to represent the unit sphere in an angular IGA formulation could produce an angular discretisation that further reduces or is free from ray effects

8.2.5 Boundary Element IGA

For a given equation, the Boundary Element Method (BEM) involves using the fundamental Green's solution of the equation to form a Boundary Integral Equation (BIE). This BIE only utilises unknowns on the boundary of the domain, effectively reducing the physical dimension of the problem by one [280]. BEM has been applied in many fields successfully but in neutronics the situation is slightly more complex. Scattering and fission integrals must be converted from volume integrals into surface integrals [281]. Furthermore, heterogeneous domains must be decomposed into regions of constant material cross-section and then coupled together, as BEM can only be applied to homogeneous domains [282].

BEMIGA represents a significant improvement to both BEM and IGA for two reasons. First, IGA will allow the external surfaces of the domain to be modelled exactly. It was demonstrated in Chapter 3 that surface approximations in FE cause a large error, as well as restricting the order of convergence of the solution. By pairing BEM with IGA, a significant increase in accuracy should be observed for the same number of degrees of freedom. A convergence rate unaffected by geometric error will also allow higher order basis functions to achieve their full potential. Second, as mentioned,

the CAD programs used to create the geometry can not natively generate fully parameterised three dimensional models. Instead, they generate boundary representations (b-reps) using a collection of surface patches. By applying BEM to a three dimensional problem that has been formulated with IGA, the dimension of the problem can be reduced by one. Then the b-rep model generated by the CAD program can be used without necessitating an internal parameterisation [283].

8.2.6 T-splines

T-splines have been mentioned several times in this work. They overcome many of the deficiencies of NURBS based IGA, whilst still being able to represent the same geometries exactly. The two major examples of this are that first, T-splines allow for patches to differ topologically from a hypercube. This means that more general domains can be modelled in a single, water tight patch. Furthermore, holes can be included in the domain. However, these are not holes in a parametric sense. Second, hanging nodes, and therefore local refinement, is natively built into the T-spline function space. There are several rules that must be followed when deciding where to refine in order for the T-spline space to be analysis suitable [195]. This allows for local refinement and adaptive mesh refinement algorithms to be implemented without the added complexity of constraint based local refinement.

There are two major downside of a T-spline implementation. First, it is significantly more complex than implements NURBS based IGA. Second, the T-spline technology was acquired by Autodesk [159]. This means that tooling and CAD support for T-splines is not as available as it is for NURBS. Assuming the tools for a T-spline based IGA implementation could be developed it would prove to be much more flexible than NURBS based IGA.

8.2.7 Multiscale IGA

The industrial workflow of calculating the scalar neutron flux for nuclear reactor physics and radiation shielding problems uses many methods for capturing the effect of small-scale features, such as pincells, on the the large-scale without necessitating the explicit modelling of the small-scale features. For example, small-scale lattice calculations are used to generate flux weighted neutron cross sections for nodal diffusion codes in order to determine the large-scale scalar neutron flux. This calculation can be performed for a fraction of the cost of a full three dimensional transport solve [20]. The small-scale features can then be recovered using both the small-scale and large-scale solutions [23]. These are referred to as flux synthesis methods.

It is desirable to develop methods that capture the effects of the small-scale layout of the reactor on the global neutron scalar flux without requiring the explicit modelling of the small-scale features [284]. Representative Volume Element IGA, where the basis functions are adapted to the problem by solving a small-scale problem, applied to neutronics problems could represent a way to capture the effect of small-scale effects more efficiently, and with a simplified work flow, than the extant methods used in industry [285].

8.2.8 Reduced Order Models (ROM)

If the rapid calculation of global quantities of interest is of concern, then reduced order modelling could be applied. This can be achieved through two methods. Domain Mode Decomposition (DMD)

considers the action of the transport operator in order to find subspaces that can be used to approximate the solution [286]. Proper Orthogonal Decomposition (POD) is a similar technique that can generate a basis for the modes of the solution from a series of approximate solution vectors [287]. Reduced order models have already been applied to IGA and used to demonstrate drastic reductions in run time [288]. Similar behaviour has been demonstrated for time-dependent neutron transport calculations and the calculation of the α eigenvalue [286, 289]. Both of these reduced order models could be applied to many aspects of neutron transport. One such example from this thesis could be using a reduced order model to drive the AMR procedure until a suitably refined mesh is achieved. Then a full order neutron transport solve could be performed.

Bibliography

- [1] C. Böhringer. The Kyoto Protocol: A Review and Perspectives. Technical Report 03-61, Centre for European Economic Research, 2003.
- [2] Q. Chao, Y. Zhang, and X. Gao. Paris Agreement: a new start for global climate governance. *Advances in Climate Change Research*, 12:61–67, 2016.
- [3] M. Lockwood. The political sustainability of climate policy: The case of the UK Climate Change Act. *Global Environmental Change*, 23:1339–1348, 2013.
- [4] The UK’s Nuclear Future: Nuclear Industrial Strategy. Technical Report BIS/13/67, HM Government, 2013.
- [5] W. E. Cummins and R. Matzie. Design evolution of PWRs: Shippingport to generation III+. *Progress in Nuclear Energy*, 102:9–37, 2018.
- [6] L. Fennern. Design evolution of BWRs: Dresden to generation III+. *Progress in Nuclear Energy*, 102:38–57, 2018.
- [7] V. Maurel. EPR, spearhead of the European nuclear renaissance. *Nuclear Engineering and Design*, 236:1447–1454, 2006.
- [8] M. Marshall. Hinkley Point reactors herald new nuclear age for UK. *New Scientist*, 220:6–7, 2013.
- [9] J. Gates. *Astute Class Nuclear Submarine: 2010 to Date (Owners’ Workshop Manual)*. J. H. Haynes & Co Ltd, 2018.
- [10] House of Commons. Ministry of Defence nuclear programme: Sixty-First Report of Session 2017-2019. Technical Report HC 1028, Ministry of Defence, 2018.
- [11] M. K. Rowinski, T. J. White, and J. Zhao. Small and Medium sized Reactors (SMR): A review of technology. *Renewable and Sustainable Energy Reviews*, 44:643–656, 2015.
- [12] G. Locatelli, M. Mancini, and N. Todeschini. Generation IV nuclear reactors: Current status and future projects. *Energy Policy*, 61:1503–1520, 2013.
- [13] B. Lindley, D. Allen, J. Lillington, A. Smethurst, P. Smith, D. Bowman, L. Dwyer, K. Lai, A. Levers, K. Vikhorev, D. Litskevich, B. Merk, E. Patelli, E. Patterson, M. Bankhead, A. Peakman, S. de Haas, C. Jackson, J. Draup, E. Galenne, G. Parks, E. Shwageraus, and M. D. Eaton. Modelling and Simulation Activities in Support of the UK Nuclear R&D Programme on Digital

- Reactor Design. In *Proceedings of the 2018 26th International Conference on Nuclear Engineering ICONE26*. London, UK, 2018.
- [14] E. A. Patterson, R. J. Taylor, and M. Bankhead. A framework for an integrated nuclear digital environment. *Progress in Nuclear Energy*, 87:97–103, 2016.
 - [15] D. Gaston, C. Newman, G. Hansen, and D. Lebrun-Grandié. MOOSE: A parallel computational framework for coupled systems of nonlinear equations. *Nuclear Engineering and Design*, 239:1768–1778, 2009.
 - [16] C. S. Brown and H. Zhang. Uncertainty quantification and sensitivity analysis with CASL Core Simulator VERA-CS. *Annals of Nuclear Energy*, 95:188–201, 2016.
 - [17] E. D. Walker, B. Collins, and J. C. Gehin. Low-order multiphysics coupling techniques for nuclear reactor applications. *Annals of Nuclear Energy*, 132:327–338, 2019.
 - [18] D. R. Gaston, C. J. Permann, J. W. Peterson, A. E. Slaughter, D. Andrs, Y. Wang, M. P. Short, D. M. Perez, M. R. Tonks, J. Ortensi, L. Zou, and R. C. Martineau. Physics-based multiscale coupling for full core nuclear reactor simulation. *Annals of Nuclear Energy*, 84:45–54, 2015.
 - [19] E. E. Lewis and W. F. Miller. *Computational Methods of Neutron Transport*. American Nuclear Society, Inc., La Grange Park, Illinois USA, 1993.
 - [20] Y. Oka, editor. *Nuclear Reactor Design, An Advanced Course in Nuclear Engineering*. Springer-Verlag, 2010.
 - [21] D. G. Cacuci, editor. *Handbook of Nuclear Engineering*. Springer-Verlag, 2010.
 - [22] R. D. Lawrence. Progress in Nodal Methods for the Solution of the Neutron Diffusion and Transport Equations. *Progress in Nuclear Energy*, 17(3):271–301, 1986.
 - [23] K. S. Smith. Assembly Homogenization Techniques for Light Water Reactor Analysis. *Progress in Nuclear Energy*, 17(3):303–335, 1986.
 - [24] R. Sanchez. Assembly homogenization techniques for core calculations. *Progress in Nuclear Energy*, 51:14–31, 2009.
 - [25] J. M. Risner, D. Wiarda, M. E. Dunn, T. M. Miller, D. E. Peplow, and B. W. Patton. Production and Testing of the VITAMIN-B7 Fine-Group and BUGLE-B7 Broad-Group Coupled Neutron/Gamma Cross-Section Libraries Derived from ENDF/B-VII.0 Nuclear Data. Technical Report NUREG/CR-7045, United States Nuclear Regulatory Commission, 2011.
 - [26] K. Almenas and R. Lee. *Nuclear Engineering: An Introduction*. Springer-Verlag, 1992.
 - [27] T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs. *Isogeometric Analysis: Toward Integration of CAD and FEA*. John Wiley and Sons, 2009.
 - [28] J. Welch. *Isogeometric Analysis for Second-Order forms of the Neutron Transport Equation with Applications to Reactor Physics*. PhD thesis, Imperial College London, 2017.

- [29] 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.
- [30] E. E. Lewis. *Fundamentals of Nuclear Reactor Physics*. Academic Press, 2008.
- [31] L. S. Ornstein and G. E. Uhlenbeck. Some kinetic problems regarding the motion of neutrons through paraffine. *Physica*, 4:478–486, 1937.
- [32] K. H. Beckurts and K. Wirtz. *Neutron Physics*. Springer-Verlag, 1964.
- [33] E. Amaldi, U. Fano, L. V. Spencer, and M. J. Berger. *Neutrons and Related Gamma Ray Problems*. Springer-Verlag, 1959.
- [34] B. D. Ganapol. *Analytical Benchmarks for Nuclear Engineering Applications - Case Studies in Neutron Transport Theory*. Nuclear Energy Agency, 2008.
- [35] M. Clark and K. F. Hansen. *Numerical Methods of Reactor Analysis*. Academic Press, 1964.
- [36] A. R. Foster and R. L. Wright. *Basic Nuclear Engineering 3rd edition*. Allyn and Bacon, Inc, 1978.
- [37] K. M. Case, G. Placzek, and F. de Hoffmann. Introduction to the Theory of Neutron Diffusion. Technical Report LAMD-1273, Los Alamos National Laboratory, 1953.
- [38] S. Dababneh, K. Khasawneh, and Z. Odibat. An alternative solution to the neutron diffusion equations in cylindrical symmetry. *Annals of Nuclear Energy*, 38:1140–1143, 2011.
- [39] A. G. Gibbs. Analytical solutions of the neutron transport equation in arbitrary convex geometry. *Journal of Mathematical Physics*, 10(5):875–890, 1969.
- [40] S. Marguet. *The Physics of Nuclear Reactors*. Springer-Verlag, 2017.
- [41] 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.
- [42] M. B. Abbott. *An Introduction to the Method of Characteristics*. Thames and Hudson London, 1966.
- [43] M. van Noort, I. Hubeny, and T. Lanz. Multidimensional Non-LTE Radiative Transfer. I. A Universal Two-Dimensional Short-Characteristics Scheme for Cartesian, Spherical, and Cylindrical Coordinate Systems. *The Astrophysical Journal*, 568:1066–1094, 2002.
- [44] G. Marleau and A. Hébert. An Integral Transport Method for Treading CANDU and GCR Clusters. *Progress in Nuclear Energy*, 18(1):197–205, 1986.
- [45] A. Hébert. A Collision Probability Analysis of the Double-Heterogeneity Problem. *Nuclear Science and Engineering*, 115:177–184, 1993.
- [46] I. Suslov. Improvements in the Long Characteristics Method and Their Efficiency for Deep Penetration Calculations. *Progress in Nuclear Energy*, 39(2):223–242, 2001.

- [47] S. W. Davis, J. M. Stone, and Y. Jiang. A Radiation Transfer Solver for Athena Using Short Characteristics. *The Astrophysical Journal Supplement Series*, 1999(9):1–19, 2012.
- [48] M. Eklund, M. Alamaniotis, H. Hernandez, and T. Jevremovic. Method of characteristics - A review with applications to science and nuclear engineering computation. *Progress in Nuclear Energy*, 85:548–567, 2015.
- [49] A. Yamamoto. Reduction of the Spatial Discretization Error in the Method of Characteristics using the Diamond-difference Scheme. *Journal of Nuclear Science and Technology*, 43(11):1327–1335, 2006.
- [50] A. Hébert. *Applied Reactor Physics, Second Edition*. Presses Internationales Polytechnique, 2016.
- [51] F. Févotte, S. Santandrea, and R. Sanchez. Advanced Transverse Integration for the Method of Characteristics. In *Joint International Topical Meeting on Mathematics & Computing and Supercomputing in Nuclear Applications*, Monterey, California, USA, April 2007.
- [52] M. Tabuchi, A. Yamamoto, T. Endo, and N. Sugimura. Preservation of transmission probabilities in the method of characteristics. *Journal of Nuclear Science and Technology*, 50(8):837–843, 2013.
- [53] G. S. Lee and N. Z. Cho. 2D/1D fusion method solutions of the three-dimensional transport OECD benchmark problem C5G7 MOX. *Progress in Nuclear Energy*, 48:410–423, 2006.
- [54] Z. Liu, C. Zhao, L. Cao, H. Wu, and L. Cao. The material-region-based 2D/1D transport method. *Annals of Nuclear Energy*, 128:1–11, 2019.
- [55] R. M. Ferrer and J. D. Rhodes III. A Linear Source Approximation Scheme for the Method of Characteristics. *Nuclear Science and Engineering*, 182:151–165, 2016.
- [56] P. T. Petkov and T. Takeda. Transport Calculations of MOX and UO₂ Pin Cells by the Method of Characteristics. *Journal of Nuclear Science and Technology*, 35(12):874–885, 1998.
- [57] T. Jevremovic, J. Vujic, and K. Tsuda. ANEMONA - a neutron transport code for general geometry reactor assemblies based on the method of characteristics and R-function solid modeler. *Annals of Nuclear Energy*, 28:125–152, 2001.
- [58] Q. Chen, H. Wu, and L. Cao. Auto MOC - A 2D neutron transport code for arbitrary geometry based on the method of characteristics and customization of AutoCAD. *Nuclear Engineering and Design*, 238:2828–2833, 2008.
- [59] G. I. Reyes, S. V. Escamilla, A. M. G. Torres, and E. V. Gallegos. Gemma: Development of a two-dimensional neutron transport code with the method of characteristics. *Annals of Nuclear Energy*, 115:367–376, 2018.
- [60] J. S. Warsa. Analytical S_N solutions in heterogeneous slabs using symbolic algebra Computer programs. *Annals of Nuclear Energy*, 29:851–874, 2002.

- [61] S. Chandrasekhar. The transfer of radiation in stellar atmospheres. *Bulletin of the American Mathematical Society*, pages 641–711, 1947.
- [62] P. J. Coelho. The role of ray effects and false scattering on the accuracy of the standard and modified discrete ordinates methods. *Journal of Quantitative Spectroscopy & Radiative Transfer*, 73:231–238, 2002.
- [63] P. S. Brantley and E. W. Larsen. The Simplified P_3 Approximation. *Nuclear Science and Engineering*, 134:1–21, 2000.
- [64] E. W. Larsen, J. E. Morel, and J. M. McGhee. Asymptotic Derivation of the Simplified P_N Equations. In *Proceedings of Joint International Conference on Mathematical Methods and Supercomputing in Nuclear Applications*, 1993.
- [65] E. W. Larsen, J. E. Morel, and W. F. Miller. Asymptotic solutions of numerical transport problems in optically thick, diffusive regimes. *Journals of Computational Physics*, 69:283–324, 1987.
- [66] J. C. Chai, H. S. Lee, and S. V. Patankar. Ray effect and false scattering in the discrete ordinates method. *Numerical Heat Transfer*, 24:373–389, 1993.
- [67] F. Févotte. Piecewise diffusion synthetic acceleration scheme for neutron transport simulations in optically thick systems. *Annals of Nuclear Energy*, 118:71–80, 2018.
- [68] C. Ketelsen, T. Manteuffel, and J. B. Schroder. Least-squares finite element discretization of the neutron transport equation in spherical geometry. *SIAM Journal of Scientific Computing*, 37(5):S71–S89, 2015.
- [69] K. D. Lathrop. Ray effects in discrete ordinates equations. *Nuclear Science and Engineering*, 32:357–369, 1968.
- [70] J. E. Morel, T. A. Wareing, R. G. Lowrie, and D. K. Parsons. Analysis of ray-effect mitigation techniques. *Nuclear Science and Engineering*, 144:1–22, 2003.
- [71] A. Gelb. The Resolution of the Gibbs Phenomenon for Spherical Harmonics. *Mathematics of Computation*, 66(218):699–717, 1997.
- [72] R. Gast. On the Equivalence of the Spherical Harmonics Method and the Discrete Ordinate Method using Gauss Quadrature for the Boltzmann Equation. Technical Report WAPD-TM-118, Westinghouse Electric Corp., 1958.
- [73] J. H. Jeans. The equations of radiative transfer of energy. *Monthly Notices, Royal Astronomical Society*, 78:28, 1917.
- [74] J. C. Mark. The spherical harmonics method, part II. Technical Report MT 97, National Research Council of Canada, 1944.
- [75] J. C. Mark. The spherical harmonics method, part I. Technical Report MT 92, National Research Council of Canada, 1944.

- [76] G. C. Pomraning. An Asymptotically Correct Approximation to the Multidimensional Transport Equation. *Nuclear Science and Engineering*, 22:328–338, 1965.
- [77] R. G. McClarren. Theoretical Aspects of the Simplified P_N Equations. *Transport Theory and Statistical Physics*, 39:73–109, 2011.
- [78] G. I. Bell and S. Glasstone. *Nuclear Reactor Theory*. Van Nostrand Reinhold Inc., 1970.
- [79] B. Sigg. The Extended P_L Approximation of the Multidimensional Neutron Transport Equation. *Nuclear Science and Engineering*, 57(4):277–291, 1975.
- [80] R. D. M. Garcia and C. E. Siewert. A Stable Shifted-Legendre Projection Scheme for Generating P_N Boundary Conditions. *Annals of Nuclear Energy*, 23:321–332, 1996.
- [81] S. P. Hamilton and T. M. Evans. Efficient solution of the simplified P_N equations. *Journal of Computational Physics*, 284:155–170, 2015.
- [82] R. T. Ackroyd. A Finite Element Method for Neutron Transport I. Some Theoretical Considerations. *Annals of Nuclear Energy*, 5:75–94, 1978.
- [83] E. M. Gelbard. Applications of spherical harmonics method to reactor problems. Technical Report WAPD-BT-20, Bettis Atomic Power Laboratory, 1960.
- [84] E. M. Gelbard. Simplified spherical harmonics equations and their use in shielding problems. Technical Report WAPD-T-1182, Bettis Atomic Power Laboratory, 1961.
- [85] E. M. Gelbard. Applications of the simplified spherical harmonics equations in spherical geometry. Technical Report WAPD-TM-294, Westinghouse Electric Corp. Bettis Atomic Power Lab, 1962.
- [86] M. F. Modest, J. Cai, W. Ge, and E. Lee. Elliptic formulation of the simplified spherical harmonics method in radiative heat transfer. *International Journal of Heat and Mass Transfer*, 76:459–466, 2014.
- [87] G. C. Pomraning. Asymptotic and Variational Derivation of the Simplified P_N Equations. *Annals of Nuclear Energy*, 20(9):623–637, 1993.
- [88] E. W. Larsen, J. E. Morel, and J. M. McGhee. Asymptotic Derivation of the Multigroup P_1 and Simplified P_N Equations with Anisotropic Scattering. *Nuclear Science and Engineering*, 123:328–342, 1996.
- [89] J. E. Morel, J. M. McGhee, and E. W. Larsen. A Three-Dimensional Time-Dependent Unstructured Tetrahedral-Mesh SP_N Method. *Nuclear Science and Engineering*, 123:319–327, 1996.
- [90] R. Ciolini, G. G. M. Coppa, B. Montagnini, and P. Ravetto. Simplified P_N and A_N Methods in Neutron Transport. *Progress in Nuclear Energy*, 40(2):237–264, 2002.
- [91] R. Rinaldi, L. Liang, and H. Schober, editors. *Neutron Applications in Earth, Energy and Environmental Sciences*. Springer-Verlag, 2008.

- [92] J. J. Duderstadt and W. R. Martin. *Transport Theory*. John Wiley & Sons, 1979.
- [93] H. Hurwitz Jr and R. Ehrlich. Multigroup Methods for Neutron Diffusion Problems. Technical Report HECD-3595, Knolls Atomic Power Laboratory, 1953.
- [94] C. Demazière. *Modelling of Nuclear Reactor Multi-physics: From Local Balance Equations to Macroscopic Models in Neutronics and Thermal-Hydraulics*. Academic Press, 2019.
- [95] B. E. Watt. Energy Spectrum of Neutrons from Thermal Fission of U^{235} . *Physical Review*, 87(6):1037–1041, 1952.
- [96] R. E. MacFarlane, D. W. Muir, R. M. Boicourt, and A. C. Kahler. The NJOY Nuclear Data Processing System, Version 2012. Technical Report LA-UR-12-27079, Los Alamos National Laboratory, 2012.
- [97] M. B. Chadwick, M. Herman, P. Oblozinsky, M. E. Dunn, Y. Danon, A. C. Kahler, D. L. Smith, B. Pritychenko, G. Arbanas, R. Arcilla, R. Brewer, D. A. Brown, R. Capote, A. D. Carlson, Y. S. Cho, H. Derrien, K. Guber, G. M. Hale, S. Hoblit, S. Holloway, T. D. Johnson, T. Kawano, B. C. Kiedrowski, H. Kim, S. Kunieda, N. M. Larson, L. Leal, J. P. Lestone, R. C. Little, E. A. McCutchan, R. E. MacFarlane, M. MacInnes, C. M. Mattoon, R. D. McKnight, S. F. Mughabghab, G. P. A. Nobre, G. Palmiotti, A. Palumbo, M. T. Pigni, V. G. Pronyaev, R. O. Sayer, A. A. Sonzogni, N. C. Summers, P. Talou, I. J. Thompson, A. Trkov, R. L. Vogt, S. C. van der Marck, A. Wallner, M. C. White, D. Wiarda, and P. G. Young. ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data. *Nuclear Data Sheets*, 112:2887–2996, 2011.
- [98] B. A. Lindley, J. G. Hosking, P. J. Smith, D. J. Powney, B. S. Tollit, T. D. Newton, R. Perry, T. C. Ware, and P. N. Smith. Current status of reactor physics code WIMS and recent developments. *Annals of Nuclear Energy*, 102:148–157, 2017.
- [99] M. Knight, P. Bryce, and S. Hall. WIMS/PANTHER Analysis of UO_2 /MOX Cores Using Embedded Supercells. *Nuclear Technology*, 183:398–408, 2012.
- [100] B. G. Carlson. Numerical formulation and solution of neutron transport problems. Technical Report LA-2996, Los Alamos National Laboratory, 1964.
- [101] K. D. Lathrop. Spatial Differencing of the Transport Equation: Positivity vs Accuracy. *Journal of Computational Physics*, 4:475–498, 1969.
- [102] W. A. Rhoades and W. W. Engle Jr. New weighted-difference formulation for discrete ordinates calculations. Technical Report CONF-771109-53, Oak Ridge National Laboratory, 1977.
- [103] S. Hamilton, M. Benzi, and J. Warsa. Negative Flux Fixups in Discontinuous Finite Element S_N Transport. In *International Conference on Mathematics, Computational Methods & Reactor Physics*, Saratoga Spring, New York, USA, May 2009.
- [104] T. M. Evans, A. S. Stafford, R. N. Slaybaugh, and K. T. Clarno. Denovo: A New Three-Dimensional Parallel Discrete Ordinates Code in SCALE. *Nuclear Technology*, 171:171–200, 2010.

- [105] M. J. Turner, R. W. Clough, H. C. Martin, and L. J. Topp. Stiffness and deflection analysis of complex structures. *Journal of the Aeronautical Sciences*, 23(9), 1956.
- [106] W. Ritz. Theorie der Transversalschwingungen einer quadratische Platte mit freien Randern. *Annalen der Physik*, 28:737–786, 1909.
- [107] A. W. Leissa. A historical bases of the Rayleigh and Ritz methods. *Journal of Sound and Vibration*, 287:961–978, 2005.
- [108] M. F. N. Mohsen. Some details on the Galerkin finite element method. *Applied Mathematical Modelling*, 6:165–170, 1982.
- [109] J. Jin. *The Finite Element Method in Electromagnetics*. John Wiley & Sons, 2nd edition, 2002.
- [110] V. S. Vladimirov. Mathematical problems in the one-velocity theory of particle transport. Technical Report AECL-1661, Atomic Energy of Canada Ltd., Chalk River, Ontario, January 1963.
- [111] H. G. Kaper, G. K. Leaf, and A. J. Lindeman. A Timing Comparison Study for Some High Order Finite Element Approximation Procedures and a Low Order Finite Difference Approximation Procedure for the Numerical Solution of the Multigroup Neutron Diffusion Equation. *Nuclear Science and Engineering*, 49:27–48, 1972.
- [112] C. M. Kang and K. F. Hansen. Finite element methods for reactor analysis. *Nuclear Science and Engineering*, 51:456–495, 1973.
- [113] W. F. Miller, E. E. Lewis, and E. C. Rossow. The Application of Phase-Space Finite Elements to the One-Dimensional Neutron Transport Equation. *Nuclear Science and Engineering*, 51:148–156, 1973.
- [114] W. F. Miller, E. E. Lewis, and E. C. Rossow. The Application of Phase-Space Finite Elements to the Two-Dimensional Neutron Transport Equation in X - Y Geometry. *Nuclear Science and Engineering*, 52:12–22, 1973.
- [115] W. R. Martin and J. J. Duderstadt. Finite Element Solutions of the Neutron Transport Equation with Applications to Strong Heterogeneities. *Nuclear Science and Engineering*, 62:371–390, 1976.
- [116] W. R. Martin. *The Application of the Finite Element Method to the Neutron Transport Equation*. PhD thesis, University of Michigan, 1976.
- [117] R. T. Ackroyd and D. T. Grenfell. A Finite Element Method for Neutron Transport III. Two-Dimensional One-Group Test Problems. *Annals of Nuclear Energy*, 6:563–577, 1979.
- [118] R. T. Ackroyd. Completely Boundary-Free Minimum and Maximum Principles for Neutron Transport and their Least-Squares and Galerkin equivalents. *Annals in Nuclear Energy*, 9:95–124, 1982.
- [119] R. T. Ackroyd. Least-Squares Derivation of Extremum and Weighted-Residual Methods for Equations of Reactor Physics - I. The First-Order Boltzmann Equation and a First-Order Initial-Value Equation. *Annals of Nuclear Energy*, 10:65–99, 1983.

- [120] R. T. Ackroyd. The self-adjoint angular flux equation for neutron transport: I - the maximum principle, and its potential for self-adaptive synthesis of solutions. *Annals of Nuclear Energy*, 33:1271–1296, 2006.
- [121] C. R. E. de Oliveira. An arbitrary geometry finite element method for multigroup neutron transport with anisotropic scattering. *Progress in Nuclear Energy*, 18(1/2):227–236, 1986.
- [122] T. Manteuffel and K. Ressel. Least-squares finite-element solution of the neutron transport equation in diffusive regimes. *Society for Industrial and Applied Mathematics*, 35(2):806–835, April 1998.
- [123] J. E. Morel, B. T. Adams, T. Noh, J. M. McGhee, T. M. Evans, and T. J. Urbatsch. Spatial discretizations for self-adjoint forms of the radiative transfer equations. *Journal of Computational Physics*, 214:12–40, 2006.
- [124] 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.
- [125] 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.
- [126] 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.
- [127] 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.
- [128] G. Birkhoff and R. E. Lynch. *Numerical Solution of Elliptic Problems*. Society for Industrial and Applied Mathematics, 1984.
- [129] J. Y. Moller and J. J. Lautard. MINARET, a Deterministic Neutron Transport Solver for Nuclear Core Calculations. In *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, Rio de Janiro, Brazil, May 2011.
- [130] C. M. Anderson and R. J. Y. McLeod. Integration Techniques for Isoparametric and higher order bases on finite elements with a curved side. *Computers & Mathematics with Applications*, 5:285–295, 1979.
- [131] R. W. Clough and J. L. Tocher. Finite Element Stiffness Matrices for Analysis of Plate Bending. Technical Report AFFDL-TR-66-80, Air Force Flight Dynamics Laboratory, 1966.
- [132] C. C. Pain, M. D. Eaton, R. P. Smedley-Stevenson, A. J. H. Goddard, M. D. Piggot, and C. R. E. de Oliveira. Space-time streamline upwind Petrov-Galerkin methods for the Boltzmann transport equation. *Computer Methods in Applied Mechanics and Engineering*, 195:4334–4357, 2006.

- [133] C. C. Pain, M. D. Eaton, R. P. Smedley-Stevenson, A. J. H. Goddard, M. D. Piggott, and C. R. E. de Oliveira. Steamline upwind Petrov-Galerkin methods for the steady-state Boltzmann transport equation. *Computer Methods in Applied Mechanics and Engineering*, 195:4448–4472, 2006.
- [134] W. H. Reed. Triangular Mesh Difference Schemes for the Transport Equation. Technical Report LA-4769, Los Alamos Scientific Lab, 1971.
- [135] W. H. Reed and T. R. Hill. Triangular mesh methods for the neutron transport equation. Technical Report LA-UR-73-479, Los Alamos National Laboratory, 1973.
- [136] H. A. Van Der Vorst and K. Dekker. Conjugate gradient type methods and preconditioning. *Journal of Computational and Applied Mathematics*, 24:73–87, 1988.
- [137] G. C. Pomraning and M. Clark. The Variational Method Applied to the Monoenergetic Boltzmann Equation, Part II. *Nuclear Science and Engineering*, 16:155–164, 1963.
- [138] R. T. Ackroyd. *Finite Element Methods for Particle Transport: Applications to Reactor and Radiation Physics*. John Wiley & Sons, 1997.
- [139] J. N. Reddy. *Energy Principles and Variational Methods in Applied Mechanics*. John Wiley & Sons, 2017.
- [140] J. E. Morel and J. M. McGhee. A self-adjoint angular flux equation. *Nuclear Science and Engineering*, 132:312–325, 1999.
- [141] C. R. E. de Oliveira. *Finite Element Techniques for Multigroup Neutron Transport Calculations with Anisotropic Scattering*. PhD thesis, Imperial College London, 1989.
- [142] Y. Zhang. Even-Parity S_N Adjoint Method Including SP_N Model Error and Iterative Efficiency. Master’s thesis, Texas A&M University, 2014.
- [143] Y. Saad. *Iterative Methods for Sparse Linear Systems*. 2000.
- [144] 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.
- [145] H. Lof. *Parallelizing the Method of Conjugate Gradients for Shared Memory Architectures*. PhD thesis, 2004.
- [146] R. T. Ackroyd, J. G. Issa, and N. S. Riyait. Treatment of voids in finite element transport methods. *Progress in Nuclear Energy*, 18:85–89, 1986.
- [147] R. T. Ackroyd and N. S. Riyait. Iteration and extrapolation methods for the approximate solution of the even-parity transport equation for systems with voids. *Annals of Nuclear Energy*, 16(1):1–32, 1989.

- [148] 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.
- [149] 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.
- [150] Y. Bazilevs, V. M. Calo, J. A. Cottrell, J. E. Evans, T. J. R. Hughes, S. Lipton, M. A. Scott, and T. W. Sederberg. Isogeometric analysis using T-splines. *Computer Methods in Applied Mechanics and Engineering*, 199:229–263, 2010.
- [151] C. Giannelli, B. Juttler, and H. Speleers. THB-splines: The Truncated Basis for Hierarchical Splines. *Computer Aided Geometric Design*, 29:485–498, 2012.
- [152] D. C. Thomas, L. Engvall, S. K. Schmidt, K. Tew, and M. A. Scott. U-splines: splines over unstructured meshes, 2018.
- [153] L. Piegl. On NURBS: A Survey. *IEEE Computer Graphics and Applications*, 11:55–71, 1991.
- [154] L. Piegl and W. Tiller. *The NURBS Book 2nd edition*. Springer-Verlag, 1997.
- [155] L. Piegl and W. Tiller. Curve and surface constructions using rational B-splines. *Computer-Aided Design*, 19:485–498, 1987.
- [156] W. Tiller. Knot-removal algorithms for NURBS curves and surfaces. *Computer Aided Design*, 24:445–453, 1992.
- [157] M. J. Borden, M. A. Scott, J. A. Evans, and T. J. R. Hughes. Isogeometric finite element data structures based on Bézier extraction of NURBS. *International Journal for Numerical Methods in Engineering*, 87:15–47, 2010.
- [158] R. K. C. Cheng. *Inside Rhinoceros 5*. Cengage Learning, 2014.
- [159] E. Moss. *Autodesk AutoCAD 2019 Fundamentals*. SDC Publications, 2018.
- [160] Application Protocol For Product Life Cycle Support (PLCS). Technical Report AP 239, National Institute of Standards and Technology (NIST), 2015.
- [161] R. N. Nagel, W. W. Braithwaite, and P. R. Kennicott. Initial Graphics Exchange Specification IGES Version 1.0. Technical Report NBSIR 80-1978, National Bureau of Standards, 1980.
- [162] Parasolid XT Format Reference. Technical report, Siemens, 2008.
- [163] 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.
- [164] 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.

- [165] C. Hofreither and S. Takacs. Robust multigrid for isogeometric analysis based on stable splittings of spline spaces. *SIAM Journal of Numerical Analysis*, 55(4):2004–2024, 2017.
- [166] J. Sogn and S. Takacs. Robust multigrid solvers for the biharmonic problem in isogeometric analysis. *Computers and Mathematics with Applications*, 77:105–124, 2019.
- [167] N. Collier, D. Pardo, L. Dalcin, M. Paszynski, and V. M. Calo. The cost of continuity: A study of the performance of isogeometric finite elements using direct solvers. *Computer Methods in Applied Mechanics and Engineering*, 213-216:353–361, 2012.
- [168] D. Garcia, D. Pardo, L. Dalcin, M. Paszynski, N. Collier, and V. M. Calo. The value of continuity: Refined isogeometric analysis and fast direct solvers. *Computer Methods in Applied Mechanics and Engineering*, pages 586–605, 2017.
- [169] L. Beirao da Veiga, D. Cho, L. F. Pavarino, and S. Scacchi. BDDC preconditioners for isogeometric analysis. *Mathematical Models and Methods in Applied Sciences*, 23(6):1099–1142, 2013.
- [170] 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.
- [171] T. J. R. Hughes, A. Reali, and G. Sangalli. Duality and unified analysis of discrete approximations in structural dynamics and wave propagation: Comparison of p-method finite elements with k-method NURBS. *Computer Methods in Applied Mechanics and Engineering*, 197:4104–4124, 2008.
- [172] I. Akkerman, Y. Bazilevs, V. M. Calo, T. J. R. Hughes, and S. Hulshoff. The role of continuity in residual-based variational multiscale modelling of turbulence. *Computational Mechanics*, 41:371–378, 2008.
- [173] 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.
- [174] 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.
- [175] 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.
- [176] 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.

- [177] 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, 2020.
- [178] D. C. Thomas, M. A. Scott, J. A. Evans, K. Tew, and E. J. Evans. Bézier projection: A unified approach for local projection and quadrature-free refinement and coarsening of NURBS and T-splines with particular application to isogeometric design and analysis. *Computer Methods in Applied Mechanics and Engineering*, 284:55–105, 2015.
- [179] S. Shahrokhbabadi, T. D. Cao, and F. Vahedifard. Isogeometric analysis through Bézier extraction for thermal-hydro-mechanical modeling of saturated porous media. *Computers and Geotechnics*, 107:176–188, 2019.
- [180] F. L. Sun, C. Y. Dong, and H. S. Yang. Isogeometric boundary element method for crack propagation based on Bézier extraction of NURBS. *Engineering Analysis with Boundary Elements*, 99:76–88, 2019.
- [181] M. S. Alanaes, J. Hake, R. C. Kirby, H. P. Langtangen, A. Logg, and G. N. Wells. *The FEniCS Manual: Excerpts from the FEniCS Book*. Springer-Verlag, 2011.
- [182] D. Kamensky and Y. Bazilevs. tIGAr: Automating isogeometric analysis with FEniCS. *Computer Methods in Applied Mechanics and Engineering*, 344:477–498, 2019.
- [183] P. Phung-Van, C. H. Thai, H. Nguyen-Xuan, and M. A. Wahab. Porosity-dependent nonlinear transient responses to functionally graded nanoplates using isogeometric analysis. *Composites Part B*, 164:215–225, 2019.
- [184] 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.
- [185] F. Auricchio, L. Beirao Da Veiga, T. J. R. Hughes, A. Reali, and G. Sangalli. Isogeometric collocation methods. *Mathematical Models and Methods in Applied Sciences*, 20(11):2075–2107, 2010.
- [186] 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.
- [187] 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.
- [188] M. Hazewinkel. *Encyclopaedia of Mathematics: Supplement Volume 1*, page 119. Springer-Verlag, 1997.
- [189] R. T. Farouki and C. A. Neff. On the Numerical Condition of Bernstein-Bézier Subdivision Processes. *Mathematics of Computation*, 55(192):637–647, October 1990.

- [190] L. A. Piegl and W. Tiller. Circle approximation using integral B-splines. *Computer-Aided Design*, 35:601–607, 2003.
- [191] T. J. R. Hughes, A. Reali, and G. Sangalli. Efficient quadrature for NURBS-based isogeometric analysis. *Computer Methods in Applied Mechanics and Engineering*, 199:301–313, 2010.
- [192] X. Zhou. G^2 continuity algorithms between adjacent NURBS patches along common cubic boundary curve. *Chinese Journal of Aeronautics*, 16(4):241–246, 2003.
- [193] A. R. Owens, J. Kópházi, and M. D. Eaton. Optimal trace inequality constants for interior penalty discontinuous Galerkin discretisations of elliptic operators using arbitrary elements with non-constant Jacobians. *Journal of Computational Physics*, 350:847–870, 2017.
- [194] E. Cuthill and J. McKee. Reducing the bandwidth of sparse symmetric matrices. In *Proceedings of the 1969 24th national conference*, New York, August 1969.
- [195] 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.
- [196] T. W. Sederberg, J. Zheng, A. Bakenov, and A. Nasri. T-splines and T-NURCCs. *ACM Transactions on Graphics*, 22:477–484, 2003.
- [197] B. Urick, B. Marussig, E. Cohen, R. H. Crawford, T. J. R. Hughes, and R. F. Riesenfeld. Watertight Boolean operations: A framework for creating CAD-compatible gap-free editable solid models. *Computer-Aided Design*, 115:147–160, 2019.
- [198] B. Marussig, J. Zechner, G. Beer, and T. Fries. Stable isogeometric analysis of trimmed geometries. *Computational Methods in Applied Mechanical Engineering*, 316:497–521, 2017.
- [199] F. Scholz and B. Jüttler. A. Mantzaflaris. *First Order Error Correction for Trimmed Quadrature in Isogeometric Analysis*, chapter 15, pages 297–321. 2017.
- [200] Y. J. Zhang. *Geometric Modeling and Mesh Generation from Scanned Images*. Chapman and Hall, 2016.
- [201] 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.
- [202] Y. Zhang, Y. Bazilevs, S. Goswami, C. L. Bajaj, and T. J. R. Hughes. Patient-specific vascular NURBS modeling for isogeometric analysis of blood flow. *Computer Methods in Applied Mechanics and Engineering*, 196:2943–2959, 2007.
- [203] L. Rorris, I. Chalkidis, and A. Vafeidis. From CAD to IGA ready geometry and model build, for vehicle crash simulation. In *IGA 2019*, Munich, Germany, 2019.
- [204] F. Calabrò, G. Sangalli, and M. Tani. Fast Formation of Isogeometric Galerkin Matrices by Weighted Quadrature. *Computer Methods in Applied Mechanics and Engineering*, 316:606–622, 2017.

- [205] G. Sangalli and M. Tani. Isogeometric preconditioners based on fast solvers for the sylvester equation. *Journal of Scientific Computing*, 38(6):A3644–A3671, 2016.
- [206] M. Montardini, M. Negri, G. Sangalli, and M. Tani. Space-time least-squares isogeometric method and efficient solver for parabolic problems. *arXiv*, 2019.
- [207] ANL. Argonne Code Center: Benchmark Problem Book. Technical Report ANL-7416 Supplement 2, Argonne National Laboratory, 1977.
- [208] G. Theler. Unstructured Grids and the Multigroup Neutron Diffusion Equation. *Science and Technology of Nuclear Installations*, 2013, 2013.
- [209] J. Ferguson, J. Kópházi, and M. D. Eaton. Virtual Element Methods for the Spatial Discretisation of the Multigroup Neutron Diffusion Equation in Polygonal Meshes with applications to Nuclear Reactor Physics. *Submitted to Computer Methods in Applied Mechanics and Engineering*, 2019.
- [210] E. E. Lewis, G. Palmiotti, T. A. Taiwo, M. A. Smith, and N. Tsoulfanidis. Benchmark on Deterministic Transport Calculations Without Spatial Homogenisation - A 2-D/3-D MOX Fuel Assembly Benchmark. *Nuclear Science*, 2003.
- [211] J. Hansen, J. Peterson, J. Morel, J. Ragusa, and Y. Wang. A Least-Squares Transport Equation Compatible with Voids. *Journal of Computational and Theoretical Transport*, 43:374–401, 2014.
- [212] 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.
- [213] J. M. McGhee, R. M. Roberts, and J. E. Morel. The DANTE Boltzmann transport solver: An unstructured mesh, 3-D, spherical harmonics algorithm compatible with parallel computer architectures. Technical Report LA-UR-97-1031, Los Alamos National Lab, 1997.
- [214] J. L. Liscum-Powell, A. K. Prinja, J. E. Morel, and L. J. Lorence. Finite element solution of the self-adjoint angular flux equation for coupled electron-photon transport. *Nuclear Science and Engineering*, 142:270–291, 2002.
- [215] W. C. Fan, C. F. Drumm, and J. L. Powell. Discrete ordinates approximations to the first and second-order radiation transport equations. Technical Report SAND2002-1880, Sandia National Laboratory, 2002.
- [216] L. Cao and H. Wu. A spherical harmonics - finite element discretization of the self-adjoint angular flux neturon transport equation. *Nuclear Engineering and Design*, 237:2232–2239, 2007.
- [217] Y. Wang and F. N. Gleicher. Revisit boundary conditions for the self-adjoint angular flux formulation. In *PHYSOR 2014 - The Role of Reactor Physics toward a Sustainable Future*, Kyoto, Tokyo, September 2014.
- [218] 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.

- [219] Y. Wang, M. D. DeHart, D. R. Gaston, F. N. Gleicher, R. C. Martineau, J. Ortensi, J. W. Peterson, and S. Schunert. Convergence study of rattlesnake solutions for the two-dimensional C5G7 MOX benchmark. In *ANS M&C2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method*. Nashville, TN, USA, April 2015.
- [220] E. E. Lewis. Second order neutron transport methods. In *Nuclear Computational Science: A Century in Review*. Springer, 2010.
- [221] C. Drumm. Spherical Harmonics (P_N) Methods in the SCEPTRE Radiation Transport Code. In *ANS M&C2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method*. Nashville, TN, USA, April 2015.
- [222] 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.
- [223] E. H. Mund. Spectral element solutions for the P_N neutron transport equations. *Computers & Fluids*, 43:102–106, 2011.
- [224] A. Barbarino, S. Dulla, E. H. Mund, and P. Ravetto. The spectral element method for static neutron transport in A_N approximation. Part I. *Annals of Nuclear Energy*, 53:372–380, 2013.
- [225] A. Barbarino, S. Dulla, E. H. Mund, and P. Ravetto. The spectral element method for static neutron transport in A_N approximation. Part II. *Annals of Nuclear Energy*, 86:108–115, 2015.
- [226] O. C. Zienkiewicz and K. Morgan. *Finite Elements & Approximation*. Dover Publications, 2006.
- [227] J. T. Oden and J. N. Reddy. *An Introduction to the Mathematical Theory of Finite Elements*. Dover Publications, 1976.
- [228] P. J. Roache. Code verification by the method of manufactured solutions. *Journal of Fluids Engineering*, 124:4–10, 2002.
- [229] 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.
- [230] J. Stepanek, T. Auerbach, and W. Haig. Calculation of Four Thermal Reactor Benchmark Problems in X-Y Geometry. Technical Report 464, Federal Institute of Technology Zurich, 1982.
- [231] J. Y. Moller. *Eléments finis courbes et accélération pour le transport de neutrons*. PhD thesis, Université Henri Poincaré, 2012.
- [232] F. Moreau, R. Sanchez, S. Santandrea, I. Zmijarevic, and E. Masiello. CRONOS2 and APOLLO2 results for the NEA C5G7 MOX benchmark. *Progress in Nuclear Energy*, 45(2-4):179–200, 2004.

- [233] C. W. Oosterlee and T. Washio. An evaluation of parallel multigrid as a solver and a preconditioner for singularly perturbed problems. *Journal of Scientific Computing*, 19(1):87–110, January 1998.
- [234] C. W. Oosterlee and T. Washio. On the use of multigrid as a preconditioner. In *Proceedings of the 9th International Conference on Domain Decomposition Methods*, Bergen, Norway, 1998.
- [235] 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.
- [236] J. R. Lamarsh. *Introduction to Nuclear Reactor Theory*. American Nuclear Society Inc., La Grange Park, Illinois, USA, 2002.
- [237] J. E. Hicken and D. W. Zingg. Dual consistency and functional accuracy: a finite-difference perspective. *Journal of Computational Physics*, 256:161–182, 2014.
- [238] 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.
- [239] E. Müller. On the non-uniqueness of the nodal mathematical adjoint. *Annals of Nuclear Energy*, 64:333–342, 2014.
- [240] D. Lathouwers. Spatially adaptive eigenvalue estimation for the S_N equations on unstructured triangular meshes. *Annals of Nuclear Energy*, 38:1867–1876, 2011.
- [241] D. Lathouwers. Goal-oriented spatial adaptivity for the S_N equations on unstructured triangular meshes. *Annals of Nuclear Energy*, 38:1373–1381, 2011.
- [242] J. S. Warsa and A. K. Prinja. p -adaptive numerical methods for particle transport. *Transport Theory and Statistical Physics*, 28(3):229–270, 1999.
- [243] Y. Wang. HP-mesh adaptation for 1-D multigroup neutron diffusion problems. Master’s thesis, Texas A&M, 2006.
- [244] L. Demkowicz. *Computing with hp-adaptive finite elements - volume 1: one and two dimensional elliptic and Maxwell problems*. Chapman and Hall, 2007.
- [245] Y. Wang and J. Ragusa. Application of hp Adaptivity to the Multigroup Diffusion Equations. *Nuclear Science and Engineering*, 161:22–48, 2009.
- [246] 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.
- [247] B. Guo and I. Babuška. The $h - p$ version of the finite element method: Part 2: General results and applications. *Computational Mechanics*, 1:203–220, 1986.

- [248] 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.
- [249] K. Salari and P. Knupp. Code verification by the method of manufactured solutions. Technical Report SAND2000-1444, Sandia National Laboratory, 2000.
- [250] 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.
- [251] C. Latimer, M. D. Eaton J. Kópházi, and R. G. McClarren. A Geometry Conforming, Iso-geometric, Weighted Least Squares (WLS) Method for the Neutron Transport Equation with Discrete Ordinate (S_N) Angular Discretisation. *Progress in Nuclear Energy*, 121, 2020.
- [252] J. Spainer and E. M. Gelbard. *Monte Carlo Principles and Neutron Transport Problems*. Dover Publications, 2nd edition, 2008.
- [253] 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.
- [254] J. Brannick, C. Ketelsen, T. Manteuffel, and S. McCormick. Least-squares finite element methods for quantum electrodynamics. *SIAM Journal of Scientific Computing*, 32(1):398–417, 2010.
- [255] J. J. Heys, E. Lee, T. A. Manteuffel, and S. F. McCormick. An alternative least-squares formulation of the Navier-Stokes equations with improved mass conservation. *Journal of Computational Physics*, 226:994–1006, 2007.
- [256] S. D. Bond, J. H. Chaudhry, E. C. Cyr, and L. N. Olsen. A First-Order System Least-Squares Finite Element Method for the Poisson-Boltzmann Equation. *Journal of Computational Chemistry*, 31:1625–1635, 2009.
- [257] S. L. Schofield. A Maximum Principle for the First-Order Boltzmann Equation, Incorporating a Potential Treatment of Voids. *Annals of Nuclear Energy*, 15(12):543–551, 1988.
- [258] S. L. Schofield and M. Pourzand. Single-pass treatment of voids in F.E. transport theory. *Progress in Nuclear Energy*, 25(2-3):217–229, 1991.
- [259] S. L. Schofield and R. T. Ackroyd. A Maximum Principle for Time Dependent Transport in Systems with Voids. *Annals of Nuclear Energy*, 23(7):567–574, 1996.
- [260] E. Varin and G. Samba. Spherical harmonics finite element transport equation solution using a least-squares approach. *Nuclear Science and Engineering*, 151:167–183, 2005.
- [261] T. M. Austin and T. A. Manteuffel. A least-squares finite element method for the linear Boltzmann equation with anisotropic scattering. *SIAM Journal of Numerical Analysis*, 44(2):540–560, 2006.

- [262] J. H. Chaudhry, E. C. Cyr, K. Liu, T. A. Manteuffel, L. N. Olson, and L. Tang. Enhancing least-squares finite element methods through a quantity-of-interest. *SIAM Journal of Numerical Analysis*, 52(6):3085–3105, 2014.
- [263] V. M. Laboure, Y. Wang, and M. D. DeHart. Least-Squares P_N Formulation of the Transport Equation Using Self-Adjoint-Angular-Flux Consistent Boundary Conditions. In *PHYSOR 2016 - Unifying Theory and Experiments in the 21st Century*, INL/CON-15-36402, Sun Valley, Idaho, USA, May 2016.
- [264] V. M. Laboure. *Improved Fully-Implicit Spherical Harmonics Methods for First and Second Order Forms of the Transport Equation using Galerkin Finite Element*. PhD thesis, Texas A & M University, August 2016.
- [265] E. Zauderer. *Partial Differential Equations of Applied Mathematics*. John Wiley & Sons, 1989.
- [266] N. Z. Cho. Benchmark Problem 2B: MOX Fuel-Loaded Small PWR Core. Technical report, Korean Advanced Institute of Science and Technology (KAIST), 2000.
- [267] W. H. Reed. New Difference Schemes for the Neutron Transport Equation. *Nuclear Science and Engineering*, 46:309–314, 1971.
- [268] C. Manni, A. Reali, and H. Speleers. Isogeometric collocation methods with generalized B-splines. *Computers and Mathematics with Applications*, 70:1659–1675, 2015.
- [269] C. Hofreither and W. Zulehmer. Mass smoothers in geometric multigrid for isogeometric analysis. In *Curves and Surfaces 2014*, number 8, pages 272–279, Paris, France, June 2015. Springer-Verlag.
- [270] C. Latimer, J. Kópházi, and M. D. Eaton. Isogeometric multi-level iterative solution algorithms with applications in nuclear reactor physics. In *Proceedings of the 2018 26th International Conference on Nuclear Engineering*, volume 3, London, UK, 2018.
- [271] B. Schmidt, J. G. Domínguez, C. Hundt, and M. Schlarb. *Parallel Programming Concepts and Practice*. Morgan Kaufmann, 2018.
- [272] R. W. Numrich. *Parallel Programming with Co-Arrays*. CRC Press, 2019.
- [273] R. Becker, R. Koch, M. F. Modest, and H. J. Bauer. An Angular Finite Element Method for the Solution of the Even-Parity Equation. In *Proceedings of ASME 2009 Heat Transfer Summer Conference HT2009*, 2009.
- [274] S. Ukai. Solution of Multi-Dimensional Neutron Transport Equation by Finite Element Method. *Nuclear Science and Technology*, 9(6):366–373, 1972.
- [275] R. O. Castro and J. P. Trelles. Spatial and angular finite element method for radiative transfer in participating media. *Journal of Quantitative Spectroscopy & Radiative Transfer*, 157:81–105, 2015.
- [276] W. F. Miller and W. H. Reed. Ray-Effect Mitigation Methods for Two-Dimensional Neutron Transport Theory. *Nuclear Science and Engineering*, 1977.

- [277] P. Alfred, M. Neamtu, and L. L. Schumaker. Bernstein-Bézier polynomials on spheres and sphere-like surfaces. *Computer Aided Geometric Design*, 13:333–349s, 1996.
- [278] M. Schreiner. Locally Supported Kernals for Spherical Spline Interpolation. *Journal of Approximation Theory*, 89:172–194, 1997.
- [279] J. Hoschek and G. Seemann. Spherical splines. *Modélisation mathématique et analyse numérique*, 26(1):1–22, 1992.
- [280] S. Cavdar and H. A. Ozgener. A finite element/boundary element hybrid method for 2-D neutron diffusion calculations. *Annals of Nuclear Energy*, 31:1555–1582, 2004.
- [281] S. Han, S. Dulla, and P. Ravetto. Computational Methods for Multidimensional Neutron Diffusion Problems. *Science and Technology of Nuclear Installations*, 2009:1–11, 2009.
- [282] M. Maiani and B. Montagnini. A Boundary Element - Response Matrix Method for the Multi-group Neutron Diffusion Equations. *Annals of Nuclear Energy*, 26:1341–1369, 1999.
- [283] R. N. Simpson, S. P. A. Bordas, J. Trevelyan, and T. Rabczuk. A two-dimensional Isogeometric Boundary Element Method for elastostatic analysis. *Computational Methods in Applied Mechanics and Engineering*, 209-212:87–100, 2012.
- [284] Y. Efendiev and T. Y. Hou. *Multiscale Finite Element Methods: Theory and Applications*, volume 4. Springer-Verlag, 2009.
- [285] R. Alberdi, G. Zhang, and K. Khandelwal. A framework for implementation of RVE-based multiscale models in computational homogenization using isogeometric analysis. *International Journal for Numerical Methods in Engineerings*, 114:1018–1051, 2018.
- [286] R. G. McClarren. Calculating Time Eigenvalues of the Neutron Transport Equation with Dynamic Mode Decomposition. *Nuclear Science and Engineering*, 193:854–867, 2019.
- [287] G. Berkooz, P. Holmes, and J. L. Lumley. The Proper Orthogonal Decomposition in the Analysis of Turbulent Flows. *Annual Review of Fluid Mechanics*, 25:539–575, 1993.
- [288] R. Li, Q. Wu, and S. Zhu. Proper orthogonal decomposition with SUPG-stabilized isogeometric analysis for reduced order modelling of unsteady convection-dominated convection-diffusion-reaction problems. *Journal of Computational Physics*, 387:280–302, 2019.
- [289] Y. Choi, W. Arrighi, P. Brown, and R. Anderson. Space-time reduced order model for Boltzmann transport equation. In *26th International Conference on Transport Theory*, Paris, France, 2019.