Imperial College London

Self-Adaptive Isogeometric Spatial Discretisations of the First and Second-Order Forms of the Neutron Transport Equation with Dual-Weighted Residual Error Measures and Diffusion Acceleration

Academic Supervisors:

Dr Matthew EATON

Dr Michael BLUCK

Industrial Supervisors:

Prof Alan COPESTAKE

Dr Paul WARNER

Author:

Seth G. WILSON

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 2021

The Abstract

As implemented in a new modern-Fortran code, NURBS-based isogeometric analysis (IGA) spatial discretisations and self-adaptive mesh refinement (AMR) algorithms are developed in the application to the first-order and second-order forms of the neutron transport equation (NTE).

These AMR algorithms are shown to be computationally efficient and numerically accurate when compared to standard approaches. IGA methods are very competitive and offer certain unique advantages over standard finite element methods (FEM), not least of all because the numerical analysis is performed over an exact representation of the underlying geometry, which is generally available in some computer-aided design (CAD) software description. Furthermore, mesh refinement can be performed within the analysis program at run-time, without the need to revisit any ancillary mesh generator.

Two error measures are described for the IGA-based AMR algorithms, both of which can be employed in conjunction with energy-dependent meshes. The first heuristically minimises any local contributions to the global discretisation error, as per some appropriate user-prescribed norm. The second employs duality arguments to minimise *important* local contributions to the error as measured in some quantity of interest; this is commonly known as a dual-weighted residual (DWR) error measure and it demands the solution to both the forward (primal) and the adjoint (dual) NTE.

Finally, convergent and stable diffusion acceleration and generalised minimal residual (GMRes) algorithms, compatible with the aforementioned AMR algorithms, are introduced to accelerate the convergence of the within-group self-scattering sources for scattering-dominated problems for the first and second-order forms of the NTE. A variety of verification benchmark problems are analysed to demonstrate the computational performance and efficiency of these acceleration techniques.

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.

Declaration of Copyright

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 on the condition that you credit the author, do not use it for commercial purposes and do not distribute modified versions of the work.

When reusing or sharing this work, ensure that you make the license terms clear to others by naming the licence and linking the licence text.

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 thank my principal academic supervisor, Dr Matthew Eaton, for his support and advice throughout the duration of my research.

I thank Dr Jozsef Kópházi and Dr Alex Owens for helping me to improve my understanding of theoretical and computational neutron transport.

I would like to acknowledge the financial support of the Engineering and Physical Sciences Research Council (EPSRC) under their doctoral training partnership (DTP) scheme (EPSRC grant number: EP/N509486/1). Also, I would like to acknowledge the financial support of Rolls-Royce for the funding of the CASE conversion of the EPSRC DTP PhD.

I also thank my industrial supervisors from Rolls-Royce, Dr Paul Warner and Professor Alan Copestake, for their support and advice throughout my PhD.

Finally, I thank my parents for their love and encouragement.

"As far as the laws of mathematics refer to reality, they are not certain; and as far as they are certain, they do not refer to reality."

Albert Einstein (1879-1955)

Contents

Abstract i

1 The Introduction 1

1.1 The Motivation . . . 1

1.2 Theoretical and Computational Neutronics . . . 4

1.2.1 Useful Quantities and Nuclear Data . . . 6

1.2.2 Verification & Validation . . . 13

1.3 The Scope of the Research . . . 14

1.4 An Outline of the Thesis . . . 16

2 Deterministic Methods 17

2.1 An Introduction . . . 17

2.2 The First-Order Form of the Neutron Transport Equation . . . 18

2.2.1 A Linearised form of the Boltzmann Equation . . . 18

2.2.2 The Steady-State Linearised Form of the Boltzmann Equation . . . 21

2.2.3 The Importance Function . . . 22

2.2.4 Solution Methodologies . . . 24

2.2.4.1 Stochastic Methods . . . 24

2.2.4.2 Deterministic Methods . . . 25

2.3 Second-Order Forms of the Neutron Transport Equation . . . 27

2.3.1 Causality . . . 29

2.3.2 The Self-Adjoint Angular Flux Formulation . . . 31

2.4 A Discretisation in Energy . . . 32

2.4.1 The Multi-Group Approximation . . . 32

2.5 A Discretisation in Angle . . . 35
2.5.1 The Spherical Harmonics Method . . . 35
2.5.2 The Discrete Ordinates Method . . . 40
2.5.2.1 The Ray Effects . . . 44
2.6 Diffusion Theory . . . 47
2.7 A Discretisation in Space . . . 51
2.7.1 The Finite Differences Method . . . 51
2.7.2 The Method of Weighted-Residuals . . . 52
2.7.2.1 The Bubnov-Galerkin Method . . . 54
2.7.3 The Finite Elements Method . . . 55
2.7.3.1 Isoparametric Elements . . . 57
2.7.3.2 The Discontinuous Galerkin Method . . . 58
2.8 The Source-Iteration Procedure . . . 60
2.8.1 The Transport Sweep . . . 62
2.9 A Discussion . . . 64

3 NURBS-Based Isogeometric Analysis **67**

3.1 An Introduction . . . 67
3.2 The Rudiments . . . 70
3.2.1 B-Splines . . . 71
3.2.2 NURBS . . . 73
3.2.3 Gaussian Quadrature . . . 75
3.3 Geometric Design and Mesh Generation . . . 76
3.3.1 Degeneracy . . . 79
3.4 The Global Uniform Refinement Strategy . . . 82
3.5 Energy-Dependent Meshes . . . 83
3.6 The Local Adaptive Refinement Strategies . . . 87
3.6.1 Energy-Driven Error Measures . . . 93
3.6.2 Goal-Based Error Measures . . . 98
3.6.2.1 A Volumetric Detector Response . . . 99

3.6.2.2 The Effective Neutron Multiplication Factor . . . 101
3.7 A Discussion . . . 102

4 Elliptic Systems: Part 1 **105**

4.1 An Introduction . . . 105
4.2 The Multi-Group Neutron Diffusion Equations . . . 108
4.2.1 The Direct (Primal) Equations . . . 108
4.2.2 The Adjoint (Dual) Equations . . . 109
4.3 A Continuous Bubnov-Galerkin Isogeometric Discretisation . . . 111
4.3.1 Enforcing C^0 Continuity on Irregular Meshes . . . 118
4.4 An Interior Penalty Scheme for a Discontinuous Galerkin Isogeometric Discretisation . 124
4.4.1 A Discrete Coercivity Analysis . . . 136
4.4.2 The Trace Inequality Constants . . . 139
4.4.3 Optimised Penalty Parameters . . . 144
4.5 The Numerical Results . . . 145
4.5.1 The Method of Manufactured Solutions . . . 146
4.5.1.1 A Cartesian Grid . . . 147
4.5.1.2 A Pin-Cell Grid . . . 150
4.5.2 The 2D IAEA/ANL BSS-11 (LWR) Benchmark . . . 153
4.5.3 The 2D BIBLIS (PWR) Core Benchmark . . . 157
4.6 A Discussion . . . 159

5 Elliptic Systems: Part 2 **163**

5.1 An Introduction . . . 163
5.2 The Self-Adjoint Angular Flux Form of the Multi-Group Neutron Transport Equations 164
5.2.1 The Direct (Primal) Equations . . . 164
5.2.2 The Adjoint (Dual) Equations . . . 166
5.3 A Continuous-Bubnov Galerkin Isogeometric Discretisation . . . 170
5.4 An Interior Penalty Scheme for a Discontinuous Galerkin Isogeometric Discretisation . 174
5.4.1 A Discrete Coercivity Analysis . . . 183
5.5 The Numerical Results . . . 185

5.5.1 The Method of Manufactured Solutions . . . 185
5.5.1.1 A Cartesian Grid . . . 186
5.5.1.2 A Pincell-like Grid . . . 188
5.5.2 Reed's 1D Problem . . . 191
5.5.3 Ackroyd's 2D Small Reactor Problem . . . 194
5.5.4 Hong and Cho's 2D BWR Problem . . . 198
5.6 A Discussion . . . 200

6 Hyperbolic Systems **205**

6.1 An Introduction . . . 205
6.2 The First-Order Form of the Multi-Group Neutron Transport Equations . . . 207
6.2.1 The Direct (Primal) Equations . . . 207
6.2.2 The Adjoint (Dual) Equations . . . 208
6.3 A Continuous Bubnov-Galerkin Isogeometric Discretisation . . . 210
6.4 An Upwinding Scheme for a Discontinuous Galerkin Isogeometric Discretisation . . . 212
6.4.1 Graph Theory . . . 218
6.4.1.1 Breaking Cycles . . . 219
6.5 The Numerical Results . . . 221
6.5.1 The Method of Manufactured Solutions . . . 222
6.5.1.1 A Cartesian Grid . . . 222
6.5.1.2 A Pincell-like Grid . . . 225
6.5.2 Azmy's 2D Radiation Shielding Problem . . . 227
6.5.3 A 2D CANDU-6 Annular Cell Test Case . . . 230
6.6 A Discussion . . . 233

7 Acceleration Techniques **237**

7.1 An Introduction . . . 237
7.2 Iterative Transport Solvers . . . 240
7.2.1 The Source-Iteration Procedure . . . 241
7.2.1.1 Richardson Iteration . . . 241
7.2.1.2 False Convergence . . . 243

7.2.2 Krylov Subspace Iterations . . . 244
7.2.2.1 The GMRes Method . . . 245
7.3 Diffusion Acceleration . . . 249
7.3.1 Preconditioning the Source-Iteration Procedure . . . 250
7.3.1.1 Consistency . . . 251
7.3.2 Preconditioning the GMRes Method . . . 256
7.4 The Numerical Results . . . 257
7.4.1 A 2D PHI Test Case . . . 258
7.4.2 Ragusa and Wang's 2D Radiation Shielding Problem . . . 262
7.4.3 Wood and William's 2D BWR Problem . . . 267
7.4.4 The 2D IAEA CEA-2 Benchmark . . . 271
7.5 A Discussion . . . 274

8 The Conclusion **279**
8.1 Some Concluding Remarks . . . 279
8.2 Prospective Research . . . 283
8.2.1 Improved Geometry Creation . . . 283
8.2.2 Alternative Bases for Isogeometric Methods . . . 284
8.2.3 Additional Spatial Refinement Strategies . . . 284
8.2.4 Adaptive Angular Refinement Strategies . . . 285
8.2.5 Non-Linear Diffusion Acceleration . . . 285
8.2.6 Parallelisation . . . 285

The Bibliography **287**

List of Figures

2.1 An example of *two* different S_6 angular quadrature sets. 40

2.2 A demonstration of the ray effects for the example of the 2D 1G IAEA CEA-1 benchmark. 45

2.3 A comparison of the typical meshes employed by different spatial discretisation techniques for the example of a 2D pin-cell. 50

2.4 The SI procedure depicted as a flowchart. 66

3.1 The *domain*, as represented in the parametric, the parent and the real physical spaces. 70

3.2 An example of reduced continuity in a basis of quadratic B-splines. 71

3.3 An illustration of a multi-patch configuration for the example of a 2D unit pin-cell. . . 76

3.4 A snippet of the geometry (.geo) input file for the declaration of a single *volume* NURBS patch for the example of a 2D unit pin-cell. 77

3.5 An illustration of the *tiling* of the surface of general spheroids with quartic NURBS patches based upon the topology of a cube. 78

3.6 The Jacobian determinant for the example of a 2D unit circle as represented by a single degenerate bi-quadratic NURBS patch and *five* non-degenerate ones. 80

3.7 The different mechanisms for refinement for the example of a quadratic NURBS curve. 81

3.8 An example of the *implied* parametric super-mesh formed for the conservative interpolation of a scalar field between *two* volume meshes. 83

3.9 A schematic of the interpolation between a more-refined reference mesh and a coarser current mesh; and the projections to prospective subspaces. 93

3.10 The general AMR-iteration procedure depicted as a flowchart. 104

4.1 A comparison between the solutions of a CBG-IGA discretisation and its discontinuous counterpart over the irregular mesh of a toy problem. 119

4.2 An example of regularity enforced over the irregular mesh of a toy problem. 125

4.3 Another comparison between the solutions of an SIP-DG-IGA discretisation and its continuous counterpart over the irregular mesh of a toy problem. 137

4.4 A schematic of a typical pin-cell grid with an annular region for which TICs are calculated over distinct faces; along with a table of those constant values. 140

4.5 A parametric study that compares TIC values for some distinct faces of a 2D unit pin-cell grid as the outer-radius is varied. 141

4.6 A parametric study that compares TIC values for a distinct face of a 2D unit pin-cell grid as the number of QP is varied. 142

4.7 A parametric study of the discretisation error and conditioning of the 1G SIP-DG-IGA NDE as the penalty parameters are varied over a 2D unit pin-cell grid. 144

4.8 Convergence plots for the uniform refinement strategy for the CBG-IGA and SIP-DG-IGA 1G NDE over a 2D Cartesian grid. 148

4.9 Convergence plots for the NRG-AMR strategy for the CBG-IGA and SIP-DG-IGA 1G NDE over a 2D Cartesian grid. 150

4.10 Convergence plots for the uniform refinement strategy for the CBG-IGA and SIP-DG-IGA 2G NDE over a 2D pin-cell grid. 151

4.11 Convergence plots for the NRG-AMR strategy for the CBG-IGA and SIP-DG-IGA 2G NDE over a 2D pin-cell grid. 152

4.12 The error effectivity index for the NRG-AMR strategy for the CBG-IGA and SIP-DG-IGA 2G NDE over a 2D pin-cell grid. 152

4.13 The geometry for the 2D 2G IAEA/ANL BSS-11 (LWR) benchmark. 153

4.14 The spatial distribution of the multi-group components of the scalar neutron flux and importance for the 2G NDE for the 2D IAEA/ANL BSS-11 (LWR) benchmark. 154

4.15 The meshes generated at the 9th iteration of the NRG-AMR-h and DWR-AMR-h strategies for the CBG-IGA 2G NDE for the 2D IAEA/ANL BSS-11 (LWR) benchmark. . . 155

4.16 The meshes generated at the 9th iteration of the NRG-AMR-hp and DWR-AMR-hp strategies for the CBG-IGA 2G NDE for the 2D IAEA/ANL BSS-11 (LWR) benchmark. 156

4.17 Convergence plots for the different refinement strategies for the CBG-IGA and SIP-DG-IGA 2G NDE for the 2D IAEA/ANL BSS-11 (LWR) benchmark. 157

4.18 The geometry for the 2D 2G BIBLIS (PWR) benchmark. 157

4.19 The spatial distribution of the multi-group components of the scalar neutron flux and importance for the 2G NDE for the 2D BIBLIS (PWR) core benchmark. 159

4.20 The meshes generated at the 9th iteration of the NRG-AMR-h and DWR-AMR-h strategies for the SIP-DG-IGA 2G NDE for the 2D BIBLIS (PWR) benchmark. 160

4.21 The meshes generated at the 9th iteration of the NRG-AMR-hp and DWR-AMR-hp strategies for the SIP-DG-IGA 2G NDE for the 2D BIBLIS (PWR) benchmark. 161

4.22 Convergence plots for the different refinement strategies for the CBG-IGA and SIP-DG-IGA 2G NDE for the 2D BIBLIS (PWR) benchmark. 162

5.1 Convergence plots for the uniform refinement strategy for the CBG-IGA and SIP-DG-IGA 1G SAAF-S_4 equations over a 2D Cartesian grid. 187

5.2 Convergence plots for the NRG-AMR strategy for the CBG-IGA and SIP-DG-IGA 1G SAAF-S_4 equations over a 2D Cartesian grid. 189

5.3 Convergence plots for the uniform refinement strategy for the CBG-IGA and SIP-DG-IGA 2G SAAF-S_4 equations over a 2D pin-cell grid. 190

5.4 Convergence plots for the NRG-AMR strategy for the CBG-IGA and SIP-DG-IGA 2G SAAF-S_4 equations over a 2D pin-cell grid. 190

5.5 The error effectivity index for the NRG-AMR strategy for the CBG-IGA and SIP-DG-IGA 2G SAAF-S_4 equations over a 2D pin-cell grid. 191

5.6 The geometry for Reed's 1D 1G problem. 191

5.7 The solution to the SIP-DG-IGA 1G SAAF-S_8 equations for Reed's 1D problem. . . . 192

5.8 Convergence plots for the uniform refinement strategy for the CBG-IGA and SIP-DG-IGA 1G SAAF-S_8 equations for Reed's 1D problem. 193

5.9 The geometry for Ackroyd's 2D 1G small reactor problem. 194

5.10 The spatial distribution of the mono-group component of the scalar neutron flux and importance for the 1G SAAF-S_N equations for Ackroyd's 2D small reactor problem. . 195

5.11 The meshes generated at the 9th iteration of the NRG-AMR-*h* and DWR-AMR-*h* strategies for the CBG-IGA 1G SAAF-S_N equations for Ackroyd's 2D small reactor problem. 196

5.12 The meshes generated at the 9th iteration of the NRG-AMR-*hp* and DWR-AMR-*hp* strategies for the CBG-IGA 1G SAAF-S_N equations for Ackroyd's 2D small reactor problem. 197

5.13 Convergence plots for the different refinement strategies for the CBG-IGA and SIP-DG-IGA 1G SAAF-S_N equations for Ackroyd's 2D small reactor problem. 198

5.14 The geometry for Hong and Cho's 2D 2G BWR problem. 199

5.15 The spatial distribution of the multi-group components of the scalar neutron flux and importance for the 2G SAAF-S_6 equations for Hong and Cho's 2D BWR problem. . . 200

5.16 The meshes generated at the 9th iteration of the NRG-AMR-*h* and DWR-AMR-*h* strategies for the SIP-DG-IGA 2G SAAF-S_6 equations for Hong and Cho's 2D BWR problem. 201

5.17 The meshes generated at the 9th iteration of the NRG-AMR-*hp* and DWR-AMR-*hp* strategies for the SIP-DG-IGA 2G SAAF-S_6 equations for Hong and Cho's 2D BWR problem. 202

5.18 Convergence plots for the different refinement strategies for the CBG-IGA and SIP-DG-IGA 2G SAAF-S_6 equations for Hong and Cho's 2D BWR problem. 203

6.1 An example of the spurious oscillations exhibited by the CBG-IGA 1G FONT-S_8 equations for Reed's 1D problem. 211

6.2 A schematic of the cyclic digraph associated to a given direction of neutron flight for the example of a 2D pin-cell constructed from re-entrant NURBS patches. 218

6.3 Convergence plots for the uniform refinement strategy for the sweep-based DG-IGA 1G FONT-S_4 equations over a 2D Cartesian grid. . . . 223

6.4 Convergence plots for the NRG-AMR strategy for the sweep-based DG-IGA 1G FONT-S_4 equations over a 2D Cartesian grid. . . . 224

6.5 Convergence plots for the uniform refinement strategy for the sweep-based DG-IGA 2G FONT-S_4 equations over a 2D pin-cell grid. . . . 225

6.6 Convergence plots for the NRG-AMR strategy for the sweep-based DG-IGA 2G FONT-S_4 equations over a 2D pin-cell grid. . . . 226

6.7 The error effectivity index for the NRG-AMR strategy for the sweep-based DG-IGA 2G FONT-S_4 equations over a 2D pin-cell grid. . . . 226

6.8 The geometry for Azmy's 2D 1G radiation shielding problem. . . . 227

6.9 The spatial distribution of the mono-group component of the scalar neutron flux and importance for the 1G FONT-S_N equations for Azmy's 2D radiation shielding problem. 228

6.10 The meshes generated at the 9th iteration of the NRG-AMR-h and DWR-AMR-h strategies for the sweep-based DG-IGA 1G FONT-S_N equations for Azmy's 2D radiation shielding problem. . . . 229

6.11 The meshes generated at the 9th iteration of the NRG-AMR-hp and DWR-AMR-hp strategies for the sweep-based DG-IGA 1G FONT-S_N equations for Azmy's 2D radiation shielding problem. . . . 230

6.12 Convergence plots for the different refinement strategies for the sweep-based DG-IGA 1G FONT-S_N equations for Azmy's 2D radiation shielding problem. . . . 231

6.13 The geometry for a 2D 2G CANDU-6 annular cell test case. . . . 232

6.14 The spatial distribution of the multi-group components of the scalar neutron flux and importance for the 2G FONT-S_6 equations for a 2D CANDU-6 annular cell test case. 233

6.15 The meshes generated at the 9th iteration of the NRG-AMR-h and DWR-AMR-h strategies for the sweep-based DG-IGA 2G FONT-S_6 equations for a 2D CANDU-6 annular cell test case. . . . 234

6.16 The meshes generated at the 9th iteration of the NRG-AMR-hp and DWR-AMR-hp strategies for the sweep-based DG-IGA 2G FONT-S_6 equations for a 2D CANDU-6 annular cell test case. . . . 235

6.17 Convergence plots for the different refinement strategies for the sweep-based DG-IGA 2G FONT-S_6 equations for a 2D CANDU-6 annular cell test case. . . . 236

7.1 A parametric study of the NSR for the DSA-SI schemes for the CBG-IGA and SIP-DG-IGA SAAF-S_N equations over a 2D pin-cell mesh of an infinite homogeneous medium. 252

7.2 A parametric study of the NSR for the DSA-SI scheme for the sweep-based DG-IGA FONT-S_N equations over a 2D pin-cell mesh of an infinite homogeneous medium. . . . 253

7.3 The geometry for the 2D 1G PHI test case. . . . 258

7.4 The geometry for Ragusa and Wang's 2D 1G radiation shielding problem. . . . 262

7.5 The spatial distribution of the mono-group component of the scalar neutron flux and importance for the 1G FONT-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. 263

7.6 The meshes generated at the 9th iteration of the DWR-AMR-h strategy for each of the 1G SIP-DG-IGA SAAF-S_4 equations and the sweep-based 1G DG-IGA FONT-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. 263

7.7 The meshes generated at the 9th iteration of the DWR-AMR-p strategy for each of the 1G SIP-DG-IGA SAAF-S_4 equations and the sweep-based DG-IGA 1G FONT-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. 265

7.8 The geometry for Wood and William's 2D 1G BWR problem. 267

7.9 The spatial distribution of the mono-group component of the scalar neutron flux and importance for the 1G FONT-S_6 equations for Wood and William's 2D 1G BWR problem. 268

7.10 The meshes generated at the 9th iteration of the DWR-AMR-h strategy for each of the SIP-DG-IGA 1G SAAF-S_6 equations and the sweep-based DG-IGA 1G FONT-S_6 equations for Wood and William's 2D BWR problem. 268

7.11 The meshes generated at the 9th iteration of the DWR-AMR-p strategy for each of the SIP-DG-IGA 1G SAAF-S_6 equations and the sweep-based DG-IGA 1G FONT-S_6 equations for Wood and William's 2D BWR problem. 270

7.12 The geometry for the 2D 1G IAEA CEA-2 benchmark. 272

7.13 The spatial distribution of the mono-group component of the scalar neutron flux and importance for the 1G FONT-S_8 equations for the 2D IAEA CEA-2 benchmark. . . . 273

7.14 The meshes generated at the 9th iteration of the DWR-AMR-h strategy for each of the CBG-IGA 1G SAAF-S_8 equations and the sweep-based DG-IGA 1G FONT-S_8 equations for the 2D IAEA CEA-2 benchmark. 273

7.15 The meshes generated at the 9th iteration of the DWR-AMR-p strategy for each of the CBG-IGA 1G SAAF-S_8 equations and the sweep-based DG-IGA 1G FONT-S_8 equations for the 2D IAEA CEA-2 benchmark. 275

List of Tables

1.1 A recapitulation of a few useful quantities in neutronics. 13

4.1 The nuclear data for the 2D 1G MMS convergence analysis. 147

4.2 Asymptotic orders of accuracy for the uniform refinement strategy for the CBG-IGA and SIP-DG-IGA 1G NDE over a 2D Cartesian grid. 149

4.3 The nuclear data for the 2D 2G MMS convergence analysis. 150

4.4 Asymptotic orders of accuracy for the uniform refinement strategy for the CBG-IGA and SIP-DG-IGA 2G NDE over a 2D pin-cell grid. 151

4.5 The nuclear data for the 2D 2G IAEA/ANL BSS-11 (LWR) benchmark. 153

4.6 The nuclear data for the 2D 2G BIBLIS (PWR) benchmark. 158

5.1 Asymptotic orders of accuracy for the uniform refinement strategy for the CBG-IGA and SIP-DG-IGA 1G SAAF-S_4 equations over a 2D Cartesian grid. 188

5.2 Asymptotic orders of accuracy for the uniform refinement strategy for the CBG-IGA and SIP-DG-IGA 2G SAAF-S_4 equations over a 2D pin-cell grid. 190

5.3 The nuclear data for Reed's 1D 1G problem. 191

5.4 The nuclear data for Ackroyd's 2D 1G small reactor problem. 194

5.5 The nuclear data for Hong and Cho's 2D 2G BWR problem. 199

6.1 Asymptotic orders of accuracy for the uniform refinement strategy for the sweep-based DG-IGA 1G FONT-S_4 equations over a 2D Cartesian grid. 224

6.2 Asymptotic orders of accuracy for the uniform refinement strategy for the sweep-based DG-IGA 2G FONT-S_4 equations over a 2D pin-cell grid. 225

6.3 The nuclear data for Azmy's 2D 1G radiation shielding problem 227

6.4 The nuclear data for a 2D 2G CANDU-6 annular cell test case. 232

7.1 The nuclear data for a 2D 1G PHI test case. 258

7.2 The NSR for a 2D PHI test case for the (CBG-IGA DSA-)SI procedure for the CBG-IGA 1G SAAF-S_4 equations. 258

7.3 The iterations to converge the self-scatter source for a 2D PHI test case for the (CBG-IGA DSA-)SI procedure for the CBG-IGA 1G SAAF-S_4 equations. 259

7.4 The iterations to converge the self-scatter source for a 2D PHI test case for the (CBG-IGA DSA-)GMRes(20) method for the CBG-IGA 1G SAAF-S_4 equations. 259

7.5 The NSR for a 2D PHI test case for the (MIP-DG-IGA DSA-)SI procedure for the SIP-DG-IGA 1G SAAF-S_4 equations. 259

7.6 The iterations to converge the self-scatter source for a 2D PHI test case for the (MIP-DG-IGA DSA-)SI procedure for the SIP-DG-IGA 1G SAAF-S_4 equations. 260

7.7 The iterations to converge the self-scatter source for a 2D PHI test case for the (MIP-DG-IGA DSA-)GMRes(20) method for the SIP-DG-IGA 1G SAAF-S_4 equations. . . . 260

7.8 The NSR for a 2D PHI test case for the (MIP-DG-IGA DSA-)SI procedure for the sweep-based DG-IGA 1G FONT-S_4 equations. 260

7.9 The iterations to converge the self-scatter source for a 2D PHI test case for the (MIP-DG-IGA DSA-)SI procedure for the sweep-based DG-IGA 1G FONT-S_4 equations. . . 261

7.10 The iterations to converge the self-scatter source for a 2D PHI test case for the (MIP-DG-IGA DSA-)GMRes(20) method for the sweep-based DG-IGA 1G FONT-S_4 equations. 261

7.11 The nuclear data for Ragusa and Wang's 2D 1G radiation shielding problem. 262

7.12 A summary of results for the DWR-AMR-h strategy for the (CBG-IGA DSA-)SI and the (CBG-IGA DSA-)GMRes(15) solutions of the 1G CBG-IGA SAAF-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. 264

7.13 A summary of results for the DWR-AMR-h strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the SIP-DG-IGA 1G SAAF-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. 264

7.14 A summary of results for the DWR-AMR-h strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the sweep-based DG-IGA 1G FONT-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. 265

7.15 A summary of results for the DWR-AMR-p strategy for the (CBG-IGA DSA-)SI and the (CBG-IGA DSA-)GMRes(15) solutions of the CBG-IGA 1G SAAF-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. 266

7.16 A summary of results for the DWR-AMR-p strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the SIP-DG-IGA 1G SAAF-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. 266

7.17 A summary of results for the DWR-AMR-p strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the sweep-based DG-IGA 1G FONT-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. 267

7.18 The nuclear data for Wood and William's 2D 1G BWR problem. 268

7.19 A summary of results for the DWR-AMR-h strategy for the (CBG-IGA DSA-)SI and the (CBG-IGA DSA-)GMRes(15) solutions of the CBG-IGA 1G SAAF-S_6 equations for Wood and William's 2D BWR problem. 269

7.20 A summary of results for the DWR-AMR-h strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the SIP-DG-IGA 1G SAAF-S_6 equations for Wood and William's 2D BWR problem. 269

7.21 A summary of results for the DWR-AMR-h strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the sweep-based DG-IGA 1G FONT-S_6 equations for Wood and William's 2D BWR problem. 270

7.22 A summary of results for the DWR-AMR-p strategy for the (CBG-IGA DSA-)SI and the (CBG-IGA DSA-)GMRes(15) solutions of the CBG-IGA 1G SAAF-S_6 equations for Wood and William's 2D BWR problem. 271

7.23 A summary of results for the DWR-AMR-p strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the SIP-DG-IGA 1G SAAF-S_6 equations for Wood and William's 2D BWR problem. 271

7.24 A summary of results for the DWR-AMR-p strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the sweep-based DG-IGA 1G FONT-S_6 equations for Wood and William's 2D BWR problem. 272

7.25 The nuclear data for the 2D 1G IAEA CEA-2 benchmark. 272

7.26 A summary of results for the DWR-AMR-h strategy for the (CBG-IGA DSA-)SI and the (CBG-IGA DSA-)GMRes(15) solutions of the CBG-IGA 1G SAAF-S_8 equations for the 2D IAEA CEA-2 benchmark. 274

7.27 A summary of results for the DWR-AMR-h strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the SIP-DG-IGA 1G IGA-SAAF-S_8 equations for the 2D IAEA CEA-2 benchmark. 274

7.28 A summary of results for the DWR-AMR-h strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the sweep-based DG-IGA 1G FONT-S_8 equations for the 2D IAEA CEA-2 benchmark. 275

7.29 A summary of results for the DWR-AMR-p strategy for the (CBG-IGA DSA-)SI and the (CBG-IGA DSA-)GMRes(15) solutions of the CBG-IGA 1G SAAF-S_8 equations for the 2D IAEA CEA-2 benchmark. 276

7.30 A summary of results for the DWR-AMR-p strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the SIP-DG-IGA 1G SAAF-S_8 equations for the 2D IAEA CEA-2 benchmark. 276

7.31 A summary of results for the DWR-AMR-p strategy for the (SIP-DG-IGA DSA-)SI and the (SIP-DG-IGA DSA-)GMRes(15) solutions of the sweep-based DG-IGA 1G FONT-S_8 equations for the 2D IAEA CEA-2 benchmark. 277

List of Acronyms & Abbreviations

1D One-Dimensional/One Dimensions.
1G One-Group/One-Speed.
2D Two-Dimensional/Two Dimensions.
2G Two-Group.
3D Three-Dimensional/Three Dimensions.

ABWR Advanced Boiling-Water Reactor(s).
AGR Advanced Gas-Cooled Reactor(s).
AMG Algebraic Multi-Grid.
AMR Adaptive Mesh Refinement.
ANL Argonne National Laboratory.

B Basis.
BC Boundary Condition(s).
BEPU Best Estimate with Uncertainty.
BSS Benchmark Source Situation.
BTE Boltzmann Transport Equation.
BVP Boundary Value Problem(s).
BWR Boiling-Water Reactor(s).

CAD Computer-Aided Design.
CANDU Canadian Deuterium Uranium.
CASD Continuous At-Sea Deterrence/Deterrent.
CBG Continuous Bubnov-Galerkin.
CDF Cumulative Distribution Functions(s).
CDS Compressed Diagonal Storage.
CEA Commissariat à l'Energie Atomique et aux Energies Alternatives.
CFD Computational Fluid Dynamics.
CG Conjugate Gradient.
CoM Centre-of-Mass.
CPU Central Processing Unit.
CRS Compressed Row Storage.

DCF Diffusion-Conforming Form.
DD Diamond-Difference.
DE Degree-Elevation.
DE-1 Degree-Elevation ($\times 1$).
DFS Depth-First Search.
DG Discontinuous Galerkin.
DoF Degrees of Freedom.
DP$_\mathbf{N}$ Double-Spherical Harmonics.
DSA Diffusion Synthetic Acceleration.
DWR Dual-Weighted Residual.

EDF Electricité de France.

ENDF (US) Evaluated Nuclear Data File(s).
EOP Even-Parity and Odd-Parity.
EP Even-Parity.
EPR Evolutionary (European) Power Reactor(s).
EVENT Even-Parity Neutral Particle Transport.

FDM Finite Difference(s) Methods.
FEM Finite Element(s) Methods.
FOM Full Orthogonalisation Method.
FONT First-Order Form of the Neutron Transport Equation(s).
Fortran Mathematical Formula Translating System.
FTRAN Finite Element Transport.

GC Gauss-Chebyshev.
GMRes Generalised Minimal Residual.

HPC High-Performance Computing.
HTR High-Temperature Reactor(s).

IAEA International Atomic Energy Agency.
IGA Isogeometric Analysis.
IGES Initial Graphics Exchange Specification.
IIP Incomplete Interior Penalty.
IP Interior Penalty.
IVP Initial Value Problem(s).

JEFF Joint Evaluated Fission and Fussion File.

KBA Koch-Baker-Alcouffe.
KI Knot-Insertion.
KI-h Knot-Insertion ($\dot{m} = p{+}1$).
KI-1 Knot-Insertion ($\dot{m} = 1$).

LAB Laboratory.
LAPACK Linear Algebra Package.
LFR Lead-Cooled Fast Reactor(s).
LR Locally-Refined.
LS Least-Squares.
LS Level-Symmetric.
LU Lower-Upper.
LWR Light-Water Reactor(s).

M4S Modified Four-Step Acceleration.
MAUD Military Application of Uranium Detonation.
MCP Method of Collision Probability.
MFP Mean Free Path(s).
MIP Modified Interior Penalty.
MMS Method of Manufactured Solutions.
MoC Method of Characteristics.
MoD Ministry of Defence.
MOOSE Multi-Physics Object-Oriented Simulation Environment.
MPI Message Passing Interface.
MRA Micro-Reversibility Approximation.
M&S Modelling and Simulation.
MWR Method of Weighted-Residuals.

NDA Non-Linear Diffusion Acceleration.

NDE Neutron Diffusion Equation(s).
NEA Nuclear Energy Agency (OECD).
NIP Non-Symmetric Interior Penalty.
NJOY Nuclear Data Processing System.
NNPP Naval Nuclear Propulsion Programme.
NPP Nuclear Power Plant(s).
NRG Energy-Driven.
NSR Numerical Spectral Radius.
NSRP Nuclear Steam Raising Plant(s).
NTE Neutron Transport Equation(s).
NRTE Naval Reactor Test Establishment.
NURBS Non-Uniform Rational B-Splines.

OECD Organisation for Economic Co-Operation and Development.
OP Odd-Parity.

PANTHER PWR and AGR Neutronics and Thermal Hydraulics Evaluation Route.
PARTISN Parallel Time-Dependent S_N.
PBMR Pebble-Bed Modular Reactor(s).
PCG Preconditioned Conjugate Gradient.
PIDE Partial-Differential Equation(s).
PDF Probability Distribution Function.
PENDF Point-Wise Evaluated Nuclear Data File(s).
PENTRAN Parallel Environment Neutral-Particle Transport.
PETSc Portable Extensible Toolkit for Scientific Computation.
PHI Periodic Horizontal Interface.
PI Power(-)Iteration.
PIDE Partial-Integro-Differential Equation(s).
$\mathbf{P_L}$ L^{th}-Degree Spherical Harmonics.
$\mathbf{P_N}$ N^{th}-Degree Spherical Harmonics.
PWR Pressurised-Water Reactor(s).

QoI Quantity of Interest.
QP Quadrature Point(s).

R&D Research and Development.
RPV Reactor Pressure Vessel.

$\mathbf{S_2}$ Second-Order S_N ($N = 2$).
$\mathbf{S_4}$ Fourth-Order S_N ($N = 4$).
$\mathbf{S_6}$ Sixth-Order S_N ($N = 6$).
$\mathbf{S_8}$ Eighth-Order S_N ($N = 8$).
$\mathbf{S_2SA}$ S_2 synthetic acceleration.
SAAF Self-Adjoint Angular Flux.
SCC Strongly Coupled Component(s).
SI Source(-)Iteration.
SIP Symmetric Interior Penalty.
SMR Small Module Reactor(s).
$\mathbf{S_N}$ N^{th}-Order Discrete Ordinates.
SPD Symmetric Positive Definite.
$\mathbf{SP_N}$ N^{th}-Degree Simplified Spherical Harmonics.
SPSD Symmetric Positive Semi-Definite.
SSBN Submersible Ship Ballistic Missile Nuclear.
SSN Submersible Ship Nuclear.
STEP Spherical Tokamak for Energy Production.

STEP Standard for the Exchange of Product Data.
SUPG Streamline-Upwind Petrov Galerkin.

T Truncated.
TIC Trace Inequality Constant(s).
TRIDENT Triangular Mesh Discrete Ordinates Explicit Neutron Transport.
TRIPLET Triangular Mesh Planar Geometry Explicit Transport.
TSA Transport Synthetic Acceleration.
TWOTRAN Two-Dimensional Transport.

U Unstructured.
UKAEA United Kingdom Atomic Energy Authority.

VBVP Variational Boundary Value Problem(s).
VUKIMS Virtual Uncommon Knot-Inserted Master-Slave.

WIMS Winfrith Improved Multi-Group Scheme.
WLS Weighted Least-Squares.
WLUP WIMS Library Update Project.

Chapter 1

The Introduction

This chapter describes the recent innovations, research and development (R&D) in both civil and defence-related nuclear-fission power technology. The UK is currently enjoying a renaissance of nuclear-fission power, centered upon a new civil nuclear build programme and the replacement of the Continuous At-Sea Deterrence (CASD) programme. A major focus of R&D within the UK's civil and defence nuclear-fission power programmes is to improve the performance, reliability and safety of nuclear power plants (NPP), as well as to reduce the overall associated costs. Best estimate with uncertainty (BEPU) numerical simulation methods play a key role in achieving these aims through the reduction of any pessimism associated with the design of modern NPP. This provides the motivation for the development of those novel high-fidelity and high-performance self-adaptive neutron transport algorithms presented in this thesis. In addition to the recent innovations in nuclear-fission power technology, this chapter introduces the basic physical concepts, principles and mathematics of theoretical and computational neutron transport. Finally, this chapter concludes with the scope for the R&D presented in this thesis; and an outline of the structure of the remaining chapters.

1.1 The Motivation

Civil: World-wide energy demands continue to increase whilst conventional fossil fuel reserves are steadily depleted. This, compounded with international commitments to global climate governance, such as the 1997 Kyoto protocol and the 2015 Paris Agreement, provide the motivation to develop low-carbon power generation technologies; e.g. the exploitation of solar, wind and tidal power [1, 2].

The UK government has set itself the target of becoming a net-zero carbon economy by 2050; cf. the climate change act of 2008 [3]. This ambitious target presents many inter-related technical challenges associated with both power generation and energy distribution, as well as transport infrastructure and

use within the UK. One of the major challenges is associated with power generation, in particular, base-load power generation. The base-load on a power distribution grid is the minimum level of demand on the grid over a prescribed period of time. In the consideration of an optimal energy-mix of low cost, availability and reliability, this power demand on the grid is usually met by power plants that provide continuous and constant power, dispatchable power generation or by a collection of smaller intermittent energy sources. The remainder of demand, which varies throughout the day, is met by dispatchable generation, such as load-following power plants, peaking power plants or energy storage. Whilst it is important for the UK to maintain a diverse, low-carbon, energy generation capacity through the use of *green* technologies, they are often rife with problems of intermittency and distribution [4]. These types of power sources can give rise to both intentional or unintentional drops in voltage, brown-outs, and power cuts or outages, black-outs. As such, solar, wind and tidal power technologies can be unreliable providers of a constant base-load of electricity, even for the use of smart-grid and energy storage technologies. The national grid has very little storage capacity and a balance between supply and demand must be maintained at all times. This is facilitated by a variable source of power; unfortunately, however, this comes primarily in the form of gas-fired power stations, which are one of the most important sources of power in the UK for providing base-load power to supplement the variable power. That being said, the UK does rely upon a fleet of 15 operating NPP, including 14 advanced gas-cooled reactors (AGR) and *one* pressurised water reactor (PWR), that provide $\sim 20\%$ of the net generation of electricity [5]. A similar figure is recorded for the US and OECD Europe, although their nuclear cultures are very different [6].

The first half of the twentieth century was a period of profound development for the physical sciences; the *two* most significant scientific advances being the special theory of relativity and quantum theory [7]. These theoretical developments in the physical sciences were made concurrently with experimental discoveries in atomic and nuclear physics; including the discovery of the neutron in 1932, by Chadwick [8]. In the same decade, this was followed by the discovery of nuclear fission by Germans Hahn and Strassmann at the suggestion of Austrian-Swedish physicists, Meitner and Frisch [9, 10]. The famous equation, $E = mc^2$, first proposed by Poincaré and then made famous by Einstein's special theory of relativity, provides the fundamental mass-energy equivalence [11, 12]. In 1940, this led expatriate German physicists, Frisch and Peierls to write the Frisch-Peierls memorandum, which was the first technical exposition of a practical nuclear weapon. Subsequently, this persuaded the UK government to form the Military Application of Uranium Detonation (MAUD) committee, which produced *two* reports: *Use of Uranium for a Bomb* and *Use of Uranium as a Source of Power*. The latter laid the foundations for the civilian use of nuclear power and the development of NPP.

The most widespread economic ramification of Einstein's famous mass-energy equivalence relationship is that embodied by NPP; where, on average, the thermal fission of a single Uranium-235 isotope

releases ~ 200.0 MeV; cf. the combustion of *one* carbon atom releases a mere ~ 4.08 eV. In other words, 1 kg of natural Uranium yields about $\times 20,000$ as much energy as the same amount of coal, which makes the former a very portable and tradable commodity [13]. The initial cost of a NPP does tend to be significantly more than that of a conventional coal-fired plant, or gas-fired one, due to the relatively long period of construction; and the need for special materials and sophisticated safety features. Once built, however, the operating costs become minimal, which is a reflection of the low fuel costs. The US Nuclear Energy Institute suggests that this may be as much as 78% of the total costs for a typical coal-fired plant; and 87% for a gas-fired plant. In the procurement of Uranium for an NPP, this may be as low as 14%. More realistically, this figure rises to 34% when accounting for the front-end costs associated with the enrichment and fabrication of the fuel pellets; and the back-end costs associated with the management of radioactive waste and the reprocessing of spent fuel [14].

All of the UK AGR are scheduled to be decommissioned by 2030. The current programme for new build reactors focusses on conventional generation III, such as the evolutionary power reactor (EPR) by Framatome, previously Areva NP; *two* of which are currently under construction at the new Hinkley Point C site [15]. This is the first new nuclear power station for a generation, which has the intended capacity to meet $\sim 7\%$ of the UK's electrical power requirements. Also under consideration are the generation III+ light-water reactors (LWR); such as Hitachi-GE's UK advanced boiling-water reactor (ABWR) [16]. Unfortunately, however, the Horizon Nuclear project plans for the Wylfa Newydd site have been recently terminated [17, 18]. Turning towards more innovative technologies, the UK is investing in advanced nuclear fission and fusion reactor designs; which include Westinghouse's generation IV lead-cooled fast reactor (LFR) technology and the UK's Atomic Energy Authority (UKAEA) Spherical Tokamak for Energy Production (STEP) programme [19, 20]. In addition, Rolls-Royce is leading a consortium of UK companies in the R&D of their first-of-kind generation III small modular reactor (SMR) technology, the first of which is due in the early 2030s [21].

Defence: Since 1969, the UK's Ministry of Defence (MoD) has provided a CASD in the form of ballistic missile submarines, or submersible ship ballistic missile nuclear (SSBN). The new Dreadnought-class of SSBN will replace the current Vanguard-class by the 2030s and each new submarine is fitted with Rolls-Royce's most advanced reactor cores, the PWR3 nuclear steam raising plant (NSRP), which allows for more than 20 years of continued operation without refuelling. Also within the MoD's naval nuclear propulsion programme (NNPP) is the requirement to replace the current fleet of Trafalgar-class attack submarines, or submersible ship nuclear (SSN), by the new Astute-class; the last of which will be in service by 2024. These submarines have been fitted with the latest core H design of the PWR2 NSRP by Rolls-Royce; originally developed for the Vanguard-class of SSBN.

Within the MoD's NNPP, the development of modern, high-fidelity, BEPU modelling and simulation (M&S) methods and the associated advances in high-performance computing (HPC) technology has reduced the need for the experimental validation of new nuclear reactor designs. The advent of modern high-fidelity, as-built, digital twin models will enable numerical simulations not only to inform the design of the NNPP, but it will also provide assurance for the long-term behaviour of the NNPP, without recourse to additional integral experiments, such as on-shore prototypes and test facilities. Indeed, this is one of the reasons for decomissioning the Vulcan naval reactor test establishment (NRTE) for PWR2 by 2033, along with the Dounreay site for PWR1 [22]. The closure of Vulcan NRTE is driving innovation in the NNPP through the development of high-fidelity numerical codes that are less reliant upon experimental data adjustment and assimilation; and ones that can test particular design features outside of the in-service limits of an NPP through sub-system test rigs, individual component tests and various other models and simulations; viz. in the application to advanced reactor physics and radiation shielding, computational fluid dynamics (CFD) and fuel performance [23].

A better understanding of advanced nuclear reactors is crucial to improving their overall reliability, performance, safety and availability; all of which can be achieved through improvements to high-fidelity BEPU M&S methods. Certainly, there is the motivation to improve the underlying numerical methods to allow for full three-dimensional (2D) heterogeneous nuclear reactor core calculations [24]. A nuclear reactor is governed by the complex interaction of many physical processes, including: neutron and gamma-ray transport; material composition and behaviour; fluid dynamics and thermal-hydraulic effects. In an ideal and holistic approach, the physical processes involved would be modelled in high fidelity and the coupling between these processes would be modelled explicitly within integrated physics or multi-physics models. However, in practice, this may not be feasible, nor tractable, due to the physical constraints of memory and performance of modern-day computers. Practically, each of the physics models can be resolved separately and any interdependence accounted for by code coupling; cf. the open-source multi-physics object-oriented simulation environment (MOOSE) framework [25, 26].

1.2 Theoretical and Computational Neutronics

The migration of rarefied neutral particles through a host medium can be modelled by the neutron transport equation (NTE), which is a linearised form of the non-linear Boltzmann transport equation (BTE) and is central to a variety of academically and industrially important areas of engineering, including: reactor physics, radiation shielding, nuclear criticality, medical imaging, nuclear security and safeguards, cloud radiation physics, photon transport in plasmas and solar atmospheres, as well as quantum transport in nano-scale electronic devices [27–31]. The BTE was first derived by the Austrian physicist, Boltzmann, in 1872; whereas the NTE was first derived by Dutch physicists, Ornstein and

Uhlenbeck, in 1937 [32, 33]. There exist many forms of the NTE; q.v. Chapter 2. The solution of any one of these is challenging, not least of all due to the seven-dimensional phase space that sufficiently defines a neutron population.

Additional complexities are introduced when considering the multi-scale nature of most engineering problems; viz. ranging from the meso-scale, including small air gaps in the cladding (mm) and the fuel pins (cm); to the macro-scale, including assemblies and the entire reactor core (m). In the consideration of reactor shielding, the computational domain extends beyond the reactor core where it must be possible to model any radiation that has escaped. For civil NPP, the primary and secondary reactor shields are typically very large in comparison to the size of the reactor core; whereas the opposite is true for NSRP [34]. This is important from a design perspective in terms of both ensuring the health and safety of those working in proximity to the NPP as well as being able to determine accurately any possible irradiation damage to the primary and secondary systems [35]. Generally, the lifetime of an NPP is determined by the integrity of its reactor pressure vessel (RPV), which ages and becomes increasingly more brittle with neutron fluence [36, 37].

In general, most nuclear reactor core and shield designs exhibit a high-degree of spatial heterogeneity in terms of both the material properties and the geometric configuration, which may also incur significant uncertainties. These complexities may be challenging to represent with a high degree of accuracy. A further complication arises in the behaviour of the NTE, which is often not a smooth function with respect to its independent variables of spatial position, energy, angle and time. Indeed, any discontinuities in the solution to the first-order form of the neutron transport (FONT) equation may propagate along its characteristics, which thus, exhibits the hyperbolic nature of this partial-integral-differential equation (PIDE) [38]. Conversely, solutions to the second-order forms of the NTE exhibit more diffusive-like properties, as one would expect from elliptic or parabolic PIDE [39, 40]. That being said, these second-order forms of the NTE and even the neutron diffusion approximation can still exhibit singularities in certain situations [41, 42].

Generally, analytical solutions to the NTE exist only over (semi-)infinite media for singular eigenfunction expansions and Weiner-Hopf methods; and where analytical solutions do not exist, numerical approximations are employed to obtain physically useful and computationally tractable solutions [43, 44]. A mathematical model serves as an idealisation of the true system as it pertains to the real physical situation; yet, described by significantly less variables, or degrees of freedom (DoF), than that required to define the true system in all of its generality [45]. A well-posed initial value problem (IVP) with a unique solution demands the prescription of initial conditions; and similarly, a boundary value problem (BVP) demands the prescription of appropriate boundary conditions (BC).

The mathematical character of the underlying equations will often provide a guide to the appropriate

numerical scheme to employ. That is, the numerical scheme should reflect the manner in which the relevant information propagates throughout the domain. There exist, however, many trade-offs; e.g. between the complexity of the model and the accuracy of the approximation; and the simplicity of the data structures and the complexity of the algorithms [46]. To exploit the full potential of HPC hardware architectures for solving very large, linear systems of algebraic equations, the independent variables must first be discretised and any derivatives and integrals replaced by discrete representations to yield a set of algebraic relations of M-number of linear equations in N-number of *unknowns*, for $i = 1, \ldots, M$:

$$\sum_{j=1}^{N} a_{ij}\varphi_j = b_i; \tag{1.1}$$

with the system coefficients, a_{ij}; the known *sources*, b_i; and the unknowns, φ_j, to solve for as any N-uple of values that satisfy the above, which can be posed as a general matrix equation [47]:

$$\underline{\underline{A}} \cdot \underline{\varphi} = \underline{b}; \tag{1.2a}$$

$$\implies \underline{\varphi} = [\underline{\underline{A}}]^{-1} \cdot \underline{b}. \tag{1.2b}$$

To comment on the notation above, an N-dimensional vector is denoted by an underlined lower-case letter and an $(N \times M)$-dimensional matrix is denoted by a double-underlined upper-case letter; e.g. $\underline{b}$ and $\underline{\underline{A}}$, respectively. When necessary, the entries (i) of a vector or (i, j) of a matrix are denoted by $b_i = (\underline{b})_i \in \mathbb{R}^N$ and $a_{ij} = (\underline{\underline{A}})_{ij} \in \mathbb{R}^{N \times M}$, respectively.

Such high-fidelity numerical methods as those employed extensively in current practice are fast approaching the point where the major uncertainties in reactor theory are due mostly to those inaccuracies in the fundamental parametric data rather than any approximations due to the discretisation methods themselves; q.v. Chapter 2.

1.2.1 Useful Quantities and Nuclear Data

The migration of neutrons within a host medium is often assimilated to that of rarefied gas molecules; Boltzmann modelled such systems using statistical methods [32]. A similar approach was adopted by Maxwell for the derivation of the Maxwell-Boltzmann distribution in 1860 [48, 49]. Employing the concept of a statistical ensemble of a given system, this was then formalised by Gibbs in 1884 [50]. Due to their sheer number at any given moment, a population of neutrons can be approximated as a probability distribution function (PDF) for the state of the system, as characterised by the momenta and the positions of the particles. The solution to the BTE yields this PDF. In most situations, the density of neutrons is such that the population can be modelled using averaged values. However, in certain other situations, such as during nuclear reactor start up, this would not be appropriate and

instead, one should consider the individual neutron fission chains using the Pal-Bell equation [51, 52].

The most detailed description of a neutron population is encapsulated in a quantity known as the angular neutron density, $n(\vec{r}, E, \vec{\Omega}, t)$; such that $n(\vec{r}, E, \vec{\Omega}, t)dVdEd^2\vec{\Omega}$ is the number of neutrons at time t contained within the differential volume dV about $\vec{r}$, a point represented by a triplet of coordinates (x, y, z) in a Cartesian coordinate system, travelling in direction $\vec{\Omega}$ contained within the differential solid angle element $d^2\vec{\Omega}$, with energy E contained within the differential energy bin dE.

For an orthogonal geometric coordinate system, the components of the angular direction, $\vec{\Omega}$, may be identified as the *three* direction cosines, μ_x, μ_y and μ_z that align locally with the corresponding unitary vectors, e_x, e_y and e_z, respectively, on the *three* axes:

$$\vec{\Omega} = \Omega_x\vec{e}_x + \Omega_y\vec{e}_y + \Omega_z\vec{e}_z, \tag{1.3a}$$

$$= \mu_x\vec{e}_x + \mu_y\vec{e}_y + \mu_z\vec{e}_z. \tag{1.3b}$$

As $\vec{\Omega}$ represents a vector on the unit sphere, its components are constrained such that:

$$\vec{\Omega}\cdot\vec{\Omega} = (\Omega_x)^2 + (\Omega_y)^2 + (\Omega_z)^2 = 1. \tag{1.4}$$

The differential solid angle may also be expressed in spherical coordinates as:

$$d^2\vec{\Omega} = \sin\theta d\theta d\psi, \tag{1.5a}$$

$$= d\mu d\psi; \tag{1.5b}$$

where θ is the co-latitude, or polar angle formed with some directional axis; ψ is the longitude, or azimuthal angle specifying the angle of rotation about that same axis; and letting $\mu = \cos\theta$:

$$\int_{4\pi} d^2\vec{\Omega} = \int_0^{2\pi}\int_{-1}^{1} d\mu d\psi = 4\pi. \tag{1.6}$$

Even under steady state conditions, the exact number of neutrons per infinitesimal phase space volume varies continuously with time. The angular neutron density therefore describes the mean statistical behaviour of a large population of neutrons, which oscillates about an average expected value that is related to the solution of the NTE.

More useful quantities for neutron transport are the angular neutron flux, φ, which is defined as the product of the neutron speed, v, and the angular neutron density, n:

$$\varphi(\vec{r}, E, \vec{\Omega}, t) = v(E)n(\vec{r}, E, \vec{\Omega}, t)\left[\frac{1}{\text{cm}^2.\text{eV.sr.s}}\right]; \tag{1.7}$$

and the angular neutron current, $\vec{j}$, which represents the number of neutrons per unit volume, energy,

solid angle and time traversing a differential surface, dS, with a unit-normal vector, $\vec{n}$, located at $\vec{r}$:

$$\vec{\jmath}(\vec{r}, E, \vec{\Omega}, t) = \vec{\Omega}\varphi(\vec{r}, E, \vec{\Omega}, t) \left[\frac{1}{\text{cm}^2.\text{eV.sr.s}}\right]. \tag{1.8}$$

Performing a distribution reduction on $\vec{\Omega}$, one obtains the neutron scalar flux, ϕ:

$$\phi(\vec{r}, E, t) = \int_{4\pi} \varphi(\vec{r}, E, \vec{\Omega}, t)\, d^2\vec{\Omega} \left[\frac{1}{\text{cm}^2.\text{eV.s}}\right]; \tag{1.9}$$

and the net neutron current, $\vec{J}$, which represents the number of neutrons per unit volume, energy and time crossing a differential surface, dS, with the unit-normal vector, $\vec{n}$, located at $\vec{r}$:

$$\vec{J}(\vec{r}, E, t) = \int_{4\pi} \vec{\Omega}\varphi(\vec{r}, E, \vec{\Omega}, t)\, d^2\vec{\Omega} \left[\frac{1}{\text{cm}^2.\text{eV.s}}\right]. \tag{1.10}$$

The angular current, $\vec{\jmath}$, is positive if the neutrons cross the differential surface, dS, in the direction of $\vec{n}$; and negative if they do not. As such, it is possible to define an *outgoing* partial current, J^+, as well as an *incoming* one, J^-:

$$J^{\pm}(\vec{r}, E, t) = \int_{\vec{\Omega}\cdot\vec{n}\gtrless 0} \vec{\Omega}\varphi(\vec{r}, E, \vec{\Omega}, t) \cdot \vec{n}(\vec{r})\, d^2\vec{\Omega}; \tag{1.11}$$

such that:

$$\vec{J}(\vec{r}, E, t) \cdot \vec{n}(\vec{r}) = J^+(\vec{r}, E, t) + J^-(\vec{r}, E, t). \tag{1.12}$$

The Cross-Sections: A sufficient description of a neutron population requires knowledge of *how* the neutral particles migrate in phase space and interact with the host medium. At the core of such interactions is the nuclear *cross-section*; that is, the cross-section area of a (target) nucleus as it appears to oncoming (projectile) neutrons. The basic physical data required by computational neutron transport codes are the microscopic cross-sections, σ_i (cm^2), which are highly dependent upon the energy of the incident neutron, as well as the relative probabilities of collision; and characterise either absorption or scattering events, as denoted by the index i. Although no time-dependence has been assumed, the material properties do vary as a function of state-point parameters; e.g. temperature, burn-up and poisoning during reactor operation. Generally, the microscopic cross-sections are very small and are more commonly expressed in units of barn (b), which is equivalent to 10^{-24} cm^2.

Should a neutron be absorbed by a target nucleus it may induce a fission event, as denoted by the subscript f, which is accompanied by the release of a few prompt neutrons. Otherwise, the neutron has been captured, as denoted by the subscript c, and no other neutrons are released by the compound nucleus, which itself may de-excite by γ-ray emission, (n, γ), or α-particle emission, (n, α).

The nuclear cross-sections are expressed in terms of the relative speeds of the incident neutron and the target nucleus. The latter may be assumed to be stationary with respect to the former because

the kinetic energy due to the thermal motion of the background is generally so much less than that possessed by the incoming neutron. Clearly, however, the velocity distribution of the target nuclei is temperature-dependent; and therefore so is the rate of interaction. For those cross-sections that are a relatively smooth function of energy, typically below 1.0 eV, the effect of varying temperature may not be important. Yet, it becomes so for those cross-sections that are sharply peaked, typically in the epi-thermal energy range, $\sim$ 1.0 eV to 0.1 MeV, which is characteristic of the (resolved) resonance structure of the target nucleus. The temperature dependence of resonance absorption is important to the stability of NPP as it allows for a negative feedback effect; whereby the rate of absorption increases with temperature as the resonances become effectively wider and their peaks become smaller, this is known as the Doppler broadening effect [53].

Scattering is characterised by those interactions through which a neutron changes in energy and angular direction as it collides with a target nucleus. Across most of the range of energies, potential elastic scatter is the dominant mode of scatter in LWR, where the slowing down, or down-scatter, of neutrons occurs primarily in the moderator on Hydrogen atoms. Neither the scalar nor the angular neutron fluxes are generally affected by the presence of any associated resonances, since they are *narrow* in comparison to the maximum energy loss by collision [54]. Conversely, should a neutron be travelling slowly enough that it finds itself in thermal equilibrium with its surroundings, it is as likely to gain in kinetic energy as it is to lose it when colliding with a target nucleus. That is, the thermal motion of the nucleus is transferred to the incoming neutron in an up-scattering event. In the current context, it is sufficient to consider only potential elastic scattering; i.e. (n, n); and it will not be necessary to consider higher-order scattering components; e.g. $(n, 2n\prime)$. Assuming that any quantum mechanical effects are unimportant, one can propose a treatment by classical mechanics in the consideration of neutrons as point particles, or billiard balls [55].

For those neutron-induced nuclear reactions that emit neutrons, it is fitting to define a differential microscopic cross-section, $\sigma_i(\vec{r}, E\prime \rightarrow E, \vec{\Omega}\prime \rightarrow \vec{\Omega})$, which represents the *area* of interaction per target nucleus for a neutron at position $\vec{r}$, with energy $E\prime$, and travelling in direction $\vec{\Omega}\prime$, to scatter into a differential energy element dE about E and into a differential cone of direction $d^2\vec{\Omega}$ about $\vec{\Omega}$.

For a given region of the host medium, the target nuclei may equally be characterised by a number density, N; by which it is assumed that the material of the region is homogeneous and static, such that N is constant. Interpreted as the probability of collision per unit path length, the macroscopic cross-sections, Σ_i (cm^{-1}), are related to the microscopic ones, σ_i, by:

$$\Sigma_i(\vec{r}, E) = N(\vec{r})\sigma_i(\vec{r}, E). \tag{1.13}$$

Similarly, one may define a differential macroscopic cross-section, $\Sigma_i(\vec{r}, E\prime \rightarrow E, \vec{\Omega}\prime \rightarrow \vec{\Omega})$, which repre-

sents the probability per unit path length that a neutron at position $\vec{r}$, with energy $E\prime$, and travelling in direction $\vec{\Omega}\prime$, will scatter into a differential energy element dE about E and into a differential solid angle of direction $d^2\vec{\Omega}$ about $\vec{\Omega}$; where:

$$\Sigma_i(\vec{r}, E\prime) = \int_0^\infty \int_{4\pi} \Sigma_i(\vec{r}, E\prime \to E, \vec{\Omega}\prime \to \vec{\Omega})\ d^2\vec{\Omega} dE. \tag{1.14}$$

Collecting the different types of interactions, one defines a macroscopic cross-section of total collision:

$$\Sigma_t(\vec{r}, E) = \Sigma_a(\vec{r}, E) + \Sigma_s(\vec{r}, E), \tag{1.15a}$$

$$= \Sigma_c(\vec{r}, E) + \Sigma_f(\vec{r}, E) + \Sigma_s(\vec{r}, E). \tag{1.15b}$$

Continuous energy nuclear cross-section data, as it has been presented thus far, are available for many different isotopes in general-purpose point-wise evaluated nuclear data files (PENDF) such as the Joint Evaluated Fission and Fusion (JEFF) Library of Nuclear Energy Agency (NEA) countries and the (US) Evaluated Nuclear Data File (ENDF)/-VII [56, 57]. Nuclear data processing codes, such as NJOY, employ various methods for the treatment of resonance effects and use approximate energy spectra to condense the point-wise data into a group-wise format, which can then be edited to form nuclear data libraries; e.g. the original 69-group WIMS library [58].

In the consideration of thermal reactors, the energy dependence of the neutron flux is dominated by *three* principal phenomena in *three* distinct energy sub-ranges. The emission of prompt fission neutrons at high energies, 1.0E−03 eV $< E <$ 1.0E+07 eV, may be approximated by the Watt spectrum, $\chi(E)$; where most neutrons are born with an average energy of 2.0 MeV [59]. The fraction of fast fission neutrons born with energy between E and $E+dE$ is given by $\chi(E)dE$; whereby the fission spectrum is normalised as $\int_0^\infty \chi(E)dE = 1$. The slowing down of epi-thermal neutrons at intermediate energies, 1.0E+00 eV $< E <$ 1.0E+05 eV, varies as $1/E$; and those *slow-moving* neutrons at low energies, 1.0E−03 eV $< E <$ 1.0E+00 eV, that are in thermal equilibrium with their surroundings, may be approximated by a Maxwell-Boltzmann distribution [60, 61].

The mention of group-wise, or energy-averaged, cross-section data begs the introduction of the multi-group approximation; q.v. Section 2.4.1. Whilst this is formally reserved for Chapter 2, one mentions briefly that the practical range of neutron energies, $(E_G, E_0]$, is partitioned into distinct values, $E_0 > E_1 > \cdots > E_G$, forming G energy groups, where the gth-group spans the half-closed interval: $(E_g, E_{g\text{-}1}]$; and one emphasises the importance of preserving the reaction rates of the form:

$$\int_{E_g}^{E_{g-1}} \Sigma_i(\vec{r}, E)\varphi(\vec{r}, E, \Omega)\ dE. \tag{1.16}$$

Such integral quantities as Equation 1.16 appear in the NTE and any discretisation with respect to energy should not alter the neutron balance; as such, a preservation of the reaction rates also ensures

preservation *criticality* of the system, as quantified by the k_{eff}. Furthermore, the reaction rates are a measurable quantity, whereas the scalar and the angular neutron fluxes are not.

To preserve any isotropy of the medium, the scalar neutron flux is employed as a weighting function in the definition of the energy-averaged macroscopic cross-sections, for $g = 1, \ldots, G$:

$$\Sigma_i^g(\vec{r}) = \frac{\int_{E_g}^{E_{g-1}} \Sigma_i(\vec{r}, E)\phi(\vec{r}, E)\ dE}{\int_{E_g}^{E_{g-1}} \phi(\vec{r}, E)\ dE}; \tag{1.17}$$

and similarly for those differential macroscopic cross-sections, $g\prime, g = 1, \ldots, G$:

$$\Sigma_i^{g\prime \to g}(\vec{r}, \vec{\Omega}\prime \to \vec{\Omega}) = \frac{\int_{E_g}^{E_{g-1}} \int_{E_{g\prime}}^{E_{g\prime-1}} \Sigma_i(\vec{r}, E\prime \to E, \vec{\Omega}\prime \cdot \vec{\Omega})\hat{\phi}(\vec{r}, E\prime)\ dE\prime dE}{\int_{E_{g\prime}}^{E_{g\prime-1}} \phi(\vec{r}, E\prime)\ dE\prime}. \tag{1.18}$$

In a *Catch-22* situation, the energy-averaged macroscopic cross-sections cannot be evaluated before the scalar neutron flux is known; and, *vice versa*, the multi-group components of the scalar neutron flux cannot be evaluated before the energy-averaged macroscopic cross-sections are known. Practically, a good-enough approximation of the true scalar neutron flux is employed as a weighting function, such that $\phi(\vec{r}, E) \approx \hat{\phi}(\vec{r}, E)$; where $\hat{\phi}$ approximates the energy-dependence of the scalar neutron flux, ϕ.

Due to constraints in available computing resources, it is not possible to model the day-to-day steady-state operation of an entire reactor in all of its generality. Provided an appropriate nuclear data library, a *lattice physics code* can perform a very detailed neutronics analysis over a portion of the reactor at different core conditions. For most modern-day lattice physics codes, it would be too time consuming to solve the NTE in the same fine-group energy structure of the nuclear data library; and therefore, it is practical to condense the cross-section data into a smaller, more manageable many-group energy structure. Once the neutron flux distributions in space and energy are calculated over the heterogenous problem, it is then possible to perform a spatial homogenisation, or *smearing*, to define a simpler model over which *effective* cross-sections may be evaluated in a further-condensed few-group energy structure. These homogenised few-group cross-sections are then tabulated against the different state-point parameters and passed to a *nodal code* for full-core analyses. For example, the WIMS code can employ the method of collision probability, q.v. Section 2.2.4.2, to generate approximate fine-group flux spectra to condense the 172-group XMAS library data into *nine* to 20 groups, typically, to perform detailed many-group lattice calculations, which, subject to smearing and a further condensation, are used to generate *effective* two-group cross-section data for PANTHER, a nodal code by EDF [62, 63]. Where a good approximation of the spatial and energy dependence of the neutron flux is provided by a detailed *fine-mesh* calculation, upon convergence, the reaction rates must be conserved in the *coarse-mesh* one.

With an appreciation of how the basic nuclear data are obtained and processed, it is enough now to

assume that effective multi-group macroscopic cross-sections will simply be supplied; along with all other relevant problem specifications. As the coupling to thermal-hydraulics models is beyond the scope of this work, it can be assumed that any temperature dependence has already been accounted for in the data provided; and finally, it is assumed that any damage due to the irradiation of the host medium is negligible within the time-scales considered.

That being said, the group-wise differential macroscopic scattering cross-sections demand further attention. In addition to their dependence upon energy, they also depend upon both the neutron's incoming and outgoing angular directions. Rather than working in the laboratory (LAB) reference system, it is much simpler to resolve the kinematics of elastic scattering in the centre-of-mass (CoM) frame; where the moduli of the velocity vectors of each of the projectile and the target do not change, remaining collinear after the collision [54]. Therefore, for isotropic media, the differential cross-section depends only upon the angle of deviation formed by the incoming and outgoing directions in the CoM frame, $\vec{\Omega}\prime$ and $\vec{\Omega}$, respectively; or rather its cosine, $\mu_C = \vec{\Omega}\prime \cdot \vec{\Omega}$, such that, for $g\prime, g = 1, \ldots, G$:

$$\Sigma_i^{g\prime \to g}(\vec{r}, \vec{\Omega}\prime \to \vec{\Omega}) \equiv \Sigma_s^{g\prime \to g}(\vec{r}, \vec{\Omega}\prime \cdot \vec{\Omega}). \tag{1.19}$$

Due to rotational invariance in the azimuthal angle, ψ, in either of the *two* reference systems, the probability of a neutron being emitted along any angle ψ is $1/2\pi$; and for isotropic elastic scattering in the CoM frame, the probability of a neutron being emitted in any angular direction, $\vec{\Omega}$, is $1/4\pi$. However, isotropy in the CoM frame does not imply isotropy in the LAB frame; where it may be demonstrated that the average cosine of the angle of deviation is approximately [53]:

$$\bar{\mu}_L \approx \frac{2}{3A}; \tag{1.20}$$

where A is the atomic mass number of the target nuclide relative to the incident neutron.

As per the persistence of velocity, the above reflects a bias for scattering in the LAB frame to be forward-peaked; whereby it may only become reasonable to assume isotropic scatter on target nuclides of infinitely large mass number; i.e. $\bar{\mu}_L \to 0$ as $A \to \infty$. In the consideration of fission events, it can be assumed that all prompt neutrons are emitted isotropically in both reference systems; and therefore, no further treatment is required.

Those most useful quantities and terms mentioned above are recapitulated in Table 1.1; for they will be referred to throughout the remainder of this thesis.

Symbol	Units	Description
$\vec{r}$	-	A position in the real physical space, V, as represented by a vector in a Cartesian coordinate system; i.e. (x, y, z).
E	eV	Neutron energy.
$\vec{\Omega}$	-	An angular direction, as defined by a polar angle, θ, and an azimuthal angle, ψ, and represented by a vector on the unit sphere; i.e. $(\Omega_x, \Omega_y, \Omega_z)$.
t	s	Time.
$v(E)$	cms^{-1}	Neutron speed; viz. the velocity magnitude, $v = \lvert\vec{v}\rvert$.
$N(\vec{r})$	cm^{-3}	Density of target nuclei in the host medium.
$n(\vec{r}, E, \vec{\Omega}, t)$	$\mathrm{cm}^{-3}\mathrm{eV}^{-1}\mathrm{sr}^{-1}$	Angular neutron density.
$\varphi(\vec{r}, E, \vec{\Omega}, t)$	$\mathrm{cm}^{-2}\mathrm{eV}^{-1}\mathrm{sr}^{-1}\mathrm{s}^{-1}$	Angular neutron flux.
$\vec{\jmath}(\vec{r}, E, \vec{\Omega}, t)$	$\mathrm{cm}^{-2}\mathrm{eV}^{-1}\mathrm{sr}^{-1}\mathrm{s}^{-1}$	Angular neutron current.
$\phi(\vec{r}, E, t)$	$\mathrm{cm}^{-2}\mathrm{eV}^{-1}\mathrm{s}^{-1}$	Scalar neutron flux.
$\vec{J}(\vec{r}, E, t)$	$\mathrm{cm}^{-2}\mathrm{eV}^{-1}\mathrm{s}^{-1}$	Net neutron current.
$\sigma_i(\vec{r}, E)$	cm^2 or b	Point-wise microscopic cross-section for i-type reactions.
$\sigma_s(\vec{r}, E\prime \to E, \vec{\Omega}\prime \cdot \vec{\Omega})$	cm^2 or b	Point-wise differential microscopic scatter cross-section.
$\Sigma_i^g(\vec{r})$	cm^{-1}	Group-wise macroscopic cross-section for i-type reactions.
$\Sigma_s^{g\prime \to g}(\vec{r}, \vec{\Omega}\prime \cdot \vec{\Omega})$	cm^{-1}	Group-wise differential macroscopic scatter cross-section.
$\chi(\vec{r}, E)$	-	Prompt fission spectrum.

Table 1.1: A recapitulation of a few useful quantities in neutronics.

1.2.2 Verification & Validation

In the development of computational methods, one must ensure that a numerical code has met its design specifications by verifying that it *correctly* solves the governing field equations. Importantly, code verification should not be confused with code validation, which itself may be more interested in determining whether or not the correct underlying equations were chosen in the first place, by comparing numerical results to experimental data [64–66]. Within the confines of the current research, an emphasis is placed on code verification.

Throughout this thesis, to give an indication of the quality of the numerical approximation, one will employ the subscripts h and p to denote, respectively, the characteristic mesh spacing and the polynomial degree of the basis functions used to approximate the true solution. Although the specific techniques have not yet been formally introduced, it is sufficient to appreciate that Galerkin discretisations of partial-differential equations (PIDE) asymptotically provide for consistently better approximations under refinement; i.e. as $h \to 0$ and as $p \to \infty$. In fact, the theoretical rate at which the error due to discretisation becomes vanishing may be known *a priori*. One may therefore verify the numerical code by comparing theoretical orders-of-accuracy with observed ones. Clearly, this demands knowledge of some *true* solution with which to compare against; whether that be in the form of analytical solutions to simple benchmarks or reference solutions provided by *Monte Carlo* codes. Instead, however, one will rely primarily upon the method of manufactured solutions (MMS), which provides

a formal procedure for code verification as presented by Salari [67]. The current implementation of MMS follows the modified procedure as presented by Warsa [68].

In addition to tractability, one has the advantage of selecting the desired smoothness of the manufactured solution, such that one may test the maximum achievable convergence rates with an infinitely differentiable solution. The few available analytical solutions to the NTE generally contain discontinuities, which tend degrade the accuracy of the discretisation [69].

When employed, one should recall that the manufactured solution will be chosen to assess the convergence properties of the spatial discretisation and not any other discretisation of the other independent variables; and each part of the numerical code will be verified, including: the consistency of the interior equations with prescribed BC; the conservative interpolation of a scalar field between volume and surface meshes; and an ability to enforce continuity across irregular meshes.

1.3 The Scope of the Research

This thesis serves as an attempt by the author to collect all knowledge and appreciation of those theoretical and computational aspects of neutronics assimilated throughout the period of research in an effort to improve the spatial discretisation techniques and strategies thereof.

Due to the multi-scale and multi-physics challenges that nuclear reactors pose, it may only be feasible, at a reasonable computational cost, to employ deterministic methods in operational situations. With that in mind, the aim of this project is to develop both continuous and discontinuous self-adaptive Galerkin IGA methods for the spatial discretisation of first-order and second-order forms of the NTE; namely, the steady-state multi-group discrete-ordinates FONT(-S_N) equations; the steady-state multi-group discrete-ordinates self-adjoint angular flux (SAAF-S_N) formulation; and the steady-state multi-group neutron diffusion equations (NDE). These semi-discrete equations assume a preliminary discretisation by the multi-group approximation and the discrete ordinates method with respect to energy and angle, respectively; and will thus, provide a point of departure for the ensuing chapters.

In the application to reactor physics and radiation shielding, the current *état de l'art* of spatial discretisations are the standard FEM; however, the emerging (NURBS-based) IGA technology is very competitive, offering certain unique advantages [70–72]:

- The integration of design and analysis.
- The analysis is performed on an exact representation of the underlying geometry; as described in a basis of NURBS by some CAD software.

- The exploitation of higher-order NURBS bases is facilitated by simple recursion formulae.
- The higher-order continuity exhibited by NURBS basis may offer more accuracy per DoF.
- The refinement of the mesh is performed within the analysis program at run-time, without the need to (re-)communicate with any ancillary mesh generator.
- The integration of a cell-based adaptive refinement strategy is relatively straightforward.

As comparisons between IGA and FEM have been made aplenty, they will not be repeated here.

Within a new computational transport code, this work derives interpolation-based energy-driven (NRG) and dual-weighted residual (DWR) *a posteriori* error measures to drive an AMR strategy based upon the hierarchical nesting of meshes from previous iterations, whose *modus operandi* is analogous to those options of h-refinement and p-refinement available in standard FEM, with an additional capacity for energy-dependent meshes. A global uniform refinement strategy will rarely allow for an *optimal* use of computing resources, as defined by the greatest attainable accuracy per the fewest possible DoF. As such, the NRG-AMR strategy aims to minimise the local contributions to the global discretisation error, as measured in an appropriate norm. However, this may not be the ultimate *goal* of the analysis; and again, such a strategy may not be optimal. Making use of duality arguments, the more sophisticated DWR-AMR strategy minimises instead those *important* local contributions to the error as measured in some specified quantity of interest (QoI).

The manner in which continuity is enforced across an irregular mesh is dependent upon the spatial discretisation, which itself is dictated by the nature of the governing equations. For a continuous Bubnov-Galerkin (CBG-)IGA discretisation of elliptic systems, continuity is enforced across irregular meshes via constraint equations derived from master/slave relationships based upon a Galerkin projection, which is not dissimilar to that employed when evaluating the interpolation-based error measures or the group-to-group sources on energy-dependent meshes. Separately, an interior penalty scheme is applied to a discontinuous Galerkin (IP-DG-)IGA discretisation of the same elliptic systems. Changing tack to a *matrix-free* approach, a marching scheme is applied to a DG-IGA discretisation of hyperbolic systems; whereby continuity is weakly enforced through a Galerkin projection of upwinded fluxes that follows the direction of neutron flight, as defined by some predetermined sweeping order. Each of these discretisation schemes have been verified in terms of consistency, accuracy and efficiency for a variety of test cases.

In the interest of developing robust and rapidly convergent iterative schemes that are compatible with the aforementioned AMR strategies, an *accelerated* means of solving the discretised NTE is formulated in the guise of the generalised minimal residual (GMRes) method, implemented as a *wrapper* around the source-iteration (SI) procedure; in addition to diffusion synthetic acceleration (DSA) schemes, that may benefit either of the transport solvers. In tandem, these techniques allow for a stable and

marked acceleration of the convergence of the self-scattering sources.

To recapitulate, the novelties of this work are:

- A CBG-IGA discretisation of the multi-group NDE, with projection-based constraint equations.
- An IP-DG-IGA discretisation of the multi-group NDE, with optimised penalty parameters.
- A CBG-IGA discretisation of the multi-group SAAF-S_N equations, with projection-based constraint equations.
- An IP-DG-IGA discretisation of the multi-group SAAF-S_N equations, with optimised penalty parameters
- A sweep-based DG-IGA discretisation of the multi-group FONT-S_N equations.
- The capacity for energy-dependent meshes in each of the above discretisations.
- A wholly-local NRG-AMR-$h(p)$ strategy that may benefit each of the above discretisations.
- A wholly-local DWR-AMR-$h(p)$ strategy that may benefit each of the above discretisations.
- Verification of the above discretisations and AMR strategies over rational geometries.
- Stable and compatible DSA schemes that may benefit either of the SI procedure or the GMRes method for each of the above AMR-$h(p)$ strategies for each of the discretised forms of the NTE.

Based upon the object-oriented architecture of a standard FEM code, all of the above have been developed, implemented and investigated in a new numerical code, written in Modern Fortran [73].

1.4 An Outline of the Thesis

The rest of the thesis has been organised into the following: Chapter 2 introduces the NTE and presents some of the numerical techniques used to solve it. Chapter 3 acquaints the lector with IGA as the spatial discretisation technique of choice; and importantly, explains the process of geometry creation and conservation interpolation between different meshes; and sees the development of the AMR strategies. Separate chapters have been dedicated to the discretisation of each of the multi-group NDE; the multi-group SAAF-S_N equations; and the multi-group FONT-S_N equations; qq.v. Chapter 4, Chapter 5 and Chapter 6, respectively. Chapter 7 applies acceleration techniques for the discretised NTE that are compatible with the AMR strategies; namely, an accelerated solver in the guise of GMRes; and stable DSA schemes. Finally, Chapter 8 summarises the findings of the current thesis, upon which is based the proposal for some possible future works.

Chapter 2

Deterministic Methods

This chapter introduces the neutron transport equation (NTE) in several of its most common partial-integral-differential forms; including, the first-order form of the NTE (FONT), a linearised form of the Boltzmann transport equation (BTE); along with its self-adjoint angular flux (SAAF) formulation and the neutron diffusion approximation. In the theme of deterministic methods, a broad overview of the legacy techniques are presented for their discretisation with respect to each of the independent variables; yet, details of the actual spatial discretisation are reserved for later chapters.

2.1 An Introduction

Central to a complete understanding of nuclear reactor physics, radiation shielding and criticality safety assessment is the NTE; various forms of which arise through the different assumptions and approximations made in their derivation and manipulation. This in turn affects the mathematical properties of the underlying model, which may be better suited to particular numerical methods.

The rest of the chapter is organised as follows: The FONT equation is introduced, namely a steady-state linearised form of the BTE, along with its adjoint. This is followed by the second-order forms; namely, the SAAF formulation. Whilst either *first-order* or *second-order* partial differential with respect to space, the integral-differential forms of the neutron transport equation are integral with respect to energy and angle. In turn, each of the independent variables is subject to discretisation by some numerical approximation and such deterministic methods are discussed; and a solution methodology for the discretised system of algebraic equations is presented in the guise of the SI procedure.

2.2 The First-Order Form of the Neutron Transport Equation

From the principle of particle conservation, the first-order form of the NTE was first established by Ornstein and Uhlenbeck; who presented a systematic method for modelling the slowing down of a neutron through paraffin; or simply put [33]:

$$Accumulation = Gains + (Losses), \tag{2.1a}$$

$$= Sources - Leakage - Total\ Collision. \tag{2.1b}$$

Essentially, in the above one may consider a backward balance of neutrons, amongst those that have been scattered or emitted by fission events in the immediate past, the *gains*; and amongst those that are currently escaping or interacting with target nuclei, the *losses*.

2.2.1 A Linearised form of the Boltzmann Equation

The non-linear Boltzmann equation describes the average statistical behaviour of a large population of particles; where even under steady-state conditions, the number density of neutrons, n, varies continuously in time, oscillating about an expected value [74]. By definition, neutrons are neutral particles unaffected by any charges of the background; it may therefore be assumed that they travel in rectilinear paths until they either escape or collide with something else. The population of neutrons is relatively sparse, 10^9 cm^{-3}, in comparison to that of the target nuclei, 10^{22} cm^{-3}, that make up the surrounding host medium; with neutron-neutron interactions being so rare, one may consider the linearised form:

$$\frac{1}{v}\frac{\partial}{\partial t}\varphi(\vec{r}, E, \vec{\Omega}, t) = -\frac{d}{ds}\varphi(\vec{r}, E, \vec{\Omega}, t) - \Sigma_t(\vec{r}, E)\varphi(\vec{r}, E, \vec{\Omega}, t) + q(\vec{r}, E, \vec{\Omega}, t). \tag{2.2}$$

The above states simply that the time-dependent accumulation of neutrons per unit energy and solid angle within a differential volume, dV, is due to the imbalance between the gains and losses of neutrons within dV per unit energy and solid angle. The first term on the right-hand side of Equation 2.2 is that of the streaming of neutrons, which actually represents the rate of change of the angular neutron flux along the streaming path, s, in the direction of flight, $\vec{\Omega}(\mu, \psi)$. It is more convenient to express the directional derivative in terms of a rectangular Cartesian coordinate system; as per Equation 1.3 [75]:

$$\frac{d}{ds} = \frac{dx}{ds}\frac{\partial}{\partial x} + \frac{dy}{ds}\frac{\partial}{\partial y} + \frac{dz}{ds}\frac{\partial}{\partial z} + \cancelto{0}{\frac{d\mu}{ds}\frac{\partial}{\partial \mu}} + \cancelto{0}{\frac{d\psi}{ds}\frac{\partial}{\partial \psi}}, \tag{2.3a}$$

$$= \mu_x \frac{\partial}{\partial x} + \mu_y \frac{\partial}{\partial y} + \mu_z \frac{\partial}{\partial z}, \tag{2.3b}$$

$$= \vec{\Omega} \cdot \vec{\nabla}; \tag{2.3c}$$

where $\vec{\nabla}$ retains its usual definition as a differential operator with respect to space; and as the angular direction, $\vec{\Omega}$, has no spatial dependence, it may be written:

$$\vec{\Omega} \cdot \vec{\nabla}\varphi(\vec{r}, E, \vec{\Omega}) = \vec{\nabla} \cdot [\vec{\Omega}\varphi(\vec{r}, E, \vec{\Omega})], \tag{2.4a}$$

$$= \vec{\nabla} \cdot \vec{\jmath}(\vec{r}, E, \vec{\Omega}). \tag{2.4b}$$

The streaming of neutrons may also represent the *leakage* of neutrons that leave, or enter, the differential volume, dV, through its bounding surface, dS; where the transfer of neutrons through dS per unit energy and solid angle in the time interval dt can be expressed via the divergence theorem:

$$\int_V \vec{\nabla} \cdot \vec{\Omega}\varphi(\vec{r}, E, \vec{\Omega}, t)\, dV = \int_{\partial V} \vec{\Omega}\varphi(\vec{r}, E, \vec{\Omega}, t) \cdot \vec{n}(\vec{r})\, dS; \tag{2.5}$$

where the outward pointing unit-normal vector on the surface is denoted by $\vec{n}$.

Any interaction between a neutron and the host medium is described in terms of the probability of such an interaction, or rather the macroscopic cross-section, Σ_i, of i-type interaction, as presented in Chapter 1; and the cross-section of total collision may be defined, as per Equation 1.15.

One considers that the emission source density, q, comprises of *three* principal contributions to the *gains* of a neutron population; respectively, those neutrons that are emitted at position $\vec{r}$ and time t with energy E and angular direction $\vec{\Omega}$ from scatter and fission events that are induced by neutrons of any initial energy $E\prime$ and angular direction $\vec{\Omega}\prime$; as well as from extraneous or fixed neutron sources that do not depend upon the angular neutron flux distribution:

$$\begin{aligned} q(\vec{r}, E, \vec{\Omega}, t) = &\int_0^\infty \int_{4\pi} \Sigma_s(\vec{r}, E\prime \to E, \vec{\Omega}\prime \cdot \vec{\Omega})\varphi(\vec{r}, E\prime, \vec{\Omega}\prime, t)\, d^2\vec{\Omega}\prime dE\prime \\ &+ \frac{1}{4\pi}\chi(\vec{r}, E) \int_0^\infty \nu(\vec{r}, E\prime)\Sigma_f(\vec{r}, E\prime)\phi(\vec{r}, E\prime, t)\, dE\prime \\ &+ q_x(\vec{r}, E, \vec{\Omega}, t). \end{aligned} \tag{2.6}$$

For simplicity, it has been assumed in the above that the material properties are time-independent and isotropic, such that any effects of temperature, burn-up and poisoning have been ignored during reactor operation; and any interaction with a target nucleus is independent of the incoming direction of the projectile neutron, as per Equation 1.19.

With regard to the fission events, it is assumed that on average, ν prompt neutrons are always emitted isotropically and instantaneously with a distribution in energy as that defined by the fission spectrum, χ. One notes that ν is only slightly dependent upon the energy of the incident neutron, increasing with $E\prime$; and often simply the product $\nu\Sigma_f$ will be specified. Conversely, it has been assumed that χ is independent of those neutrons that induce the fission events. As per Equation 1.9, the fission sources may be expressed in terms of the scalar neutron flux, ϕ, instead.

There has been no consideration made for the emission of delayed neutrons accompanying the β-decay of those certain fission fragments that have been left in an excited state. Although they typically account for a very small proportion of the overall population, $\sim 1\%$, they play a significant role in the kinetic behaviour of a nuclear reactor; and thereby in its safe operation under transient conditions [76]. Despite the very different time-scales associated with the mean generation time of prompt neutrons, $\sim 10^{-14}$ s, and the half-life of β-decay, $\lesssim 1$ min, it may be reasonable, under steady-state conditions, to ignore any time-dependence and thus, any distinction between a prompt and a delayed neutron [61].

Finally, to reiterate, by extraneous fixed sources, q_x, it is understood that there may be some volumetric distribution of source neutrons that exists independently of the neutron flux distribution; viz. those reactions (γ, n) or (α, n) that emit neutrons in a Beryllium-9 source assembly [74]. The strength of this source is therefore known *a priori*; whereas those neutron-induced scatter and fission sources in Equation 2.6 are known *a posteriori*.

Well-Posedness: In order for the mathematical quantity termed the angular neutron flux, φ, to serve as a *sensical* and realistic representation of the true variations of a neutron distribution, it must be real and non-negative; and with the exception of localised point sources, it should also be bounded. Since Equation 2.2 contains derivatives in both time and space, the existence and the uniqueness of the solution demand the prescription of appropriate initial conditions and BC [47]. Moreover, non-negative solutions may be guaranteed by non-negative sources [77]. Accordingly, Equation 2.2 constitutes a well-posed mathematical problem.

The introduction of specific conditions at the physical boundaries are reserved for those chapters that are dedicated to specific forms of the neutron transport equation. Yet, it is worth considering now a different condition that must be applied to the interior of the problem; namely, the interface condition. One elucidates this in the example of a problem of *two* different regions, V_1 and V_2, that are separated by the interface, ∂V_{12}. Even though the material properties of each region may be different, the number of neutrons travelling in a given direction, $\vec{\Omega}$ will not change by simply traversing the interface. The conservation of neutrons is thus underpinned by the assumption that the the angular

neutron flux is a continuous distribution in $\vec{r}$ across ∂V_{12}, such that [75]:

$$\varphi(\vec{r}, E, \vec{\Omega}, t)|_{V_1} = \varphi(\vec{r}, E, \vec{\Omega}, t)|_{V_2}, \quad \forall \vec{r} \in \partial V_{12} \quad \text{and} \quad \forall \vec{\Omega} \in 4\pi. \tag{2.7}$$

The above continuity condition holds only in those directions that run parallel to the neutron's trajectory. As nothing has been demanded of the spatial derivatives of φ, one is wary that any lack of smoothness, or loss of regularity, as exhibited by the analytical solution will have adverse affects on the accuracy as well as the rates of convergence of the numerical approximation [69].

2.2.2 The Steady-State Linearised Form of the Boltzmann Equation

It is well understood that the distribution of neutrons is sensitive to perturbations in temperature, power, density, material irradiation, the core configuration and other general operating conditions, which all inevitably vary over time. In theory, it should be relatively straightforward to discretise Equation 2.2 with respect to time. However, the mean *life-span* of a neutron in a thermal reactor is of the order $\sim 10^{-8}$ to 10^{-4} s and therefore, in practice, would require extremely small time intervals to capture any transient behaviours, which would become computationally prohibitive [61].

Within the confines of this work, one is content in the limit of normal reactor operation under steady-state conditions. Although this may be an idealised concept, some assumptions have already been made that would support this. It is generally acceptable to assume a steady-state regime when modelling for a reactor core long after its divergence, or start-up; whereby it is not expected, nor intended, that the neutron flux profile should change rapidly during normal operation [47]. Under such steady-state conditions, one assumes constant temperatures and the cross-sections may be evaluated accordingly.

Moreover, the spatial discretisation techniques developed within this thesis for simplified time-independent models are equally applicable to more complex time-dependent ones; and so, the steady-state linearised form of the Boltzmann equation for neutral particles is presented as:

$$\vec{\Omega} \cdot \vec{\nabla} \varphi(\vec{r}, E, \vec{\Omega}) + \Sigma_t(\vec{r}, E) \varphi(\vec{r}, E, \vec{\Omega}) = q(\vec{r}, E, \vec{\Omega}). \tag{2.8}$$

When modelling for neutrons, the above is the FONT equation, and the appropriate cross-sections of interaction are those presented in Section 1.2.1 [33]. Where, now, q is the steady-state emission source density; defined similarly to Equation 2.6, with the exclusion of time-dependence.

In the absence of any extraneous fixed sources of neutrons, $q_x = 0$, a self-sustaining multiplying system is said to be *critical* if every fission event induces exactly *one* other fission event; that is to say that under steady-state conditions, a non-negative asymptotic distribution of neutrons will develop as the rate of *gains* is balanced by the rate of *losses*.

In most situations, however, a system will not be exactly critical until fine adjustments have been made to its configuration. In the interest of criticality calculations, Equation 2.8 may be fictitiously balanced by the introduction of a numerical artefact, the effective neutron multiplication factor, k_{eff}. In doing so, it may be assumed that instead ν/k_{eff} neutrons are emitted on average with every fission event; and Equation 2.8 may be solved as an eigenvalue problem for the eigenvalue $1/k_{eff}$, with arbitrary normalisation of the flux profile [76]. The eigenvector associated with the largest absolute eigenvalue, $|1/k_{eff}|$, or the smallest positive value of k_{eff} is the fundamental mode, which corresponds to the only non-trivial solution that has any physical meaning and the only one that persists over time.

The k_{eff} is the ratio of *gains* to *losses* and thus, gives a measure of the increase or decrease in the size of the neutron population with each new *generation*; according to:

$$k_{eff} \begin{cases} > 1, & \text{if supercritical;} \\ = 1, & \text{if critical;} \\ < 1, & \text{if subcritical.} \end{cases} \tag{2.9}$$

For sub-critical systems, where an extraneous fixed source, $q_x \neq 0$, must sustain the fission reactions, the normalisation will not be arbitrary and the flux shape is determined by the magnitude of q_x.

Equation 2.8 finds its use in a wide range of static problems; e.g. in the design of reactor shields. Equally, when modelling a reactor core during its service, Equation 2.8 may be coupled with the Bateman equations of depletion to update the material compositions and properties; in essence, both sets of equations can be solved in turn, iteratively, at successive snapshots in time [78].

2.2.3 The Importance Function

In addition to the *primal* problem, as posed by Equation 2.8 or by any of its other subsequent forms, to be presented, an adjoint equation may be formulated, thus, posing a *dual* problem.

The adjoint function, $\varphi^\dagger$, is a useful mathematical construct whose physical interpretation as the neutron *importance* can only be established once it relates to some functional of the neutron flux [45]. More specifically, this is the importance of a neutron as it pertains to some observable physical QoI. Equally, one may consider the adjoint function as a flow of pseudo-particles that streams in a direction contrary to that of the neutron current. In the process of down-scatter, these pseudo-particles gain in energy and, *vice versa*, which makes the adjoint equation non-conservative [79].

Where the direct equation was based upon a backward balance of neutrons, the adjoint equation is based upon a forward balance of importance, amongst that of those neutrons currently escaping or

interacting with target nuclei, the *losses*; and that amongst those neutrons that will be scattered or emitted by fission events in the immediate future, the *gains*.

The primal problem, Equation 2.8, may be expressed in operator notation as:

$$L\varphi = q; \tag{2.10}$$

where L is the operator that represents the left-hand side of Equation 2.8; and the dual problem may be posed as:

$$L^\dagger\varphi^\dagger = q^\dagger; \tag{2.11}$$

where the adjoint operator, $L^\dagger$, is *formally* defined by the following identity:

$$(L\varphi, \varphi^\dagger) = (\varphi, L^\dagger\varphi^\dagger); \tag{2.12}$$

where the inner product, $(\cdot, \cdot)$, implies an integration over all independent variables; i.e. in the case of Equation 2.8, for some functions, u and v:

$$(u, v) \coloneqq \int_V \int_0^\infty \int_{4\pi} v(\vec{r}, E, \vec{\Omega}) u(\vec{r}, E, \vec{\Omega})\, d^2\vec{\Omega}\, dE dV; \tag{2.13}$$

where only real functions are considered since the cross-section data, the sources and the neutron flux, as well as the operators that act upon them, are all assumed to be real and well defined. Rather informally, the adjoint of the steady-state NTE may be written explicitly by following some general rules; viz. apply a change of sign to any odd-parity differential operators and interchange the arguments of the kernels of any integral operators [74]:

$$-\vec{\Omega} \cdot \vec{\nabla}\varphi^\dagger(\vec{r}, E, \vec{\Omega}) + \Sigma_t(\vec{r}, E, \vec{\Omega})\varphi^\dagger(\vec{r}, E, \vec{\Omega}) = q^\dagger(\vec{r}, E, \vec{\Omega}); \tag{2.14}$$

with an adjoint source, $q^\dagger$, defined in a similar manner to Equation 2.12, that includes for rotationally symmetric scattering, fission and other non-homogeneous fixed sources, $q_x^\dagger$:

$$\begin{aligned} q^\dagger(\vec{r}, E, \vec{\Omega}) = &\int_0^\infty \int_{4\pi} \Sigma_s(\vec{r}, E \to E\prime, \vec{\Omega} \cdot \vec{\Omega}\prime)\varphi^\dagger(\vec{r}, E\prime, \vec{\Omega}\prime)\, d^2\vec{\Omega}\prime dE\prime \\ &+ \frac{1}{4\pi}\nu(\vec{r}, E)\Sigma_f(\vec{r}, E) \int_0^\infty \chi(\vec{r}, E\prime)\phi^\dagger(\vec{r}, E\prime)\, dE\prime \\ &+ q_x^\dagger(\vec{r}, E, \vec{\Omega}); \end{aligned} \tag{2.15}$$

where the time-independent scalar neutron importance, $\phi^\dagger$, is defined as:

$$\phi^\dagger(\vec{r}, E) = \int_{4\pi} \varphi^\dagger(\vec{r}, E, \vec{\Omega})\, d^2\vec{\Omega}. \tag{2.16}$$

The solution to Equation 2.14 serves a variety of purposes; such as expediting certain calculations in the case of fixed-source problems [80]. It also finds extensive use in perturbation theory estimates for determining changes in the reactivity due to small changes in the cross-section data [81].

In the context of the current work, the adjoint equation will be used to derive the dual-weighted residual error measures to drive an adaptive spatial refinement strategy, presented in Chapter 3; in which case, the fixed adjoint source, $q^\dagger$, is arbitrary and user-dependent.

2.2.4 Solution Methodologies

The NTE, as presented in Equation 2.8, has been introduced in its partial-integral-differential form; viz. in a rectangular Cartesian coordinate system, it is a first-order PIDE with respect to space and an integral one with respect to energy and angle. It may also take on an equivalent forms; viz. its integral form and its characteristic form. The former is formulated by integrating the angular neutron flux, φ, along its characteristic; viz. along the straight line of angular direction $\vec{\Omega}$ that corresponds to the neutron trajectory. Similarly, the latter is formulated by integrating the streaming operator, $\vec{\Omega}\cdot\vec{\nabla}\varphi$, in Equation 2.8 along its characteristic. It is important to note that although these forms may be mathematically equivalent, they will differ in the numerical approach for their solution and thus, by any approximations that are employed in their spatial discretisation.

Broadly, there exist *two* different schools of numerical methods for solving the NTE, in either of its forms; these are the stochastic methods and the deterministic methods. The former simulate the underlying physical processes via Monte Carlo sampling methods; whereas, the latter rely upon numerical approximations to discretise the underlying variables and operators of the NTE.

2.2.4.1 Stochastic Methods

A *Monte Carlo* method may be employed to solve any integral equation; whereby the expected value of some QoI is estimated, along with its standard deviation. In the application to neutron transport, the history of a single neutron is tracked from its point of emission up until the point where it escapes or is captured, all of which can be modelled as a continuous line of straight free paths connected by a number of points of successive collisions. This entire process is assumed to be Markovian, in the sense that the laws of transport, as they apply to the present, do not depend upon anything that may have occurred in the past.

In the words of Lux and Koblinger, the *random walk* of a single neutron constitutes the *game*, which must be played, or simulated, many times and independently of one another. The *score*, X, is the value attributed to the QoI at the end of a game; e.g. the response in some detector or the k_{eff}. By virtue of the law of large numbers, for an increasingly large number of games played, H, the arithmetic average of the scores tends ever closer to the expected value, $E(X)$, which ultimately becomes the

Monte Carlo estimate for the QoI; i.e. $E(X) = QoI$ as $H \to \infty$. Furthermore, the central limit theorem states that for an increasingly large number of histories simulated, the distribution of the QoI converges to the normal distribution of mean $E(X)$ and of standard deviation $\sigma(X)/\sqrt{H}$ [82].

In practice, the histories are often grouped into batches of sufficiently large and equal size, typically of the order $\sim 10^3$ to 10^6, such that the QoI may be evaluated using batch-averaged tallies. It may be trivial to parallelise such an approach by launching separate batches on separate processors and collecting the results at the end. However, methods that sample by batch are often subject to more statistical variability than those that sample individual histories [83].

The series of random physical events experienced by a neutron throughout its history are determined by the random sampling of continuous and discrete random variables, x; e.g. the mean free path and the angle of deviation due to scattering; and the type of interaction and the number of neutrons emitted after fission, respectively. The probability densities, $p(x)$, or the integral of, the cumulative distribution functions (CDF), $P(x)$, that describe the frequency and outcome of different interactions can be determined from the collision laws and the cross-section data. The actual sampling itself necessitates a random number generator; whence the stochastic nature of the Monte Carlo method. Where possible, the simplest of the sampling rules takes a random number, r, uniformly distributed on $[0, 1]$, and has it equal to the CDF of the random variable, $P(x)$, such that $x = P^{-1}(r)$ [76].

A minimum number of approximations are made in a Monte Carlo simulation since no discretisation techniques are employed; although some codes can read in group-wise cross-section data [84]. Consequently, the method has the potential to achieve extreme accuracy and arbitrarily complex geometries can be modelled exactly without incurring any penalties to performance. This makes the Monte Carlo method well suited to solving difficult non-standard problems; and to the verification and benchmarking of deterministic codes. However, as the statistical variance decreases as the inverse of the number of games played, the statistical error in the QoI becomes increasingly slow to converge for an increasingly large H. This obviously becomes very computer intensive for large problems. Where the most basic simulations are described as *analog*, such that they emulate as closely as possible the real physical events, more advanced and efficient ones employ variance reduction schemes that introduce some form of useful bias; e.g. MCBEND, a 3D code used in radiation shielding and dosimetry calculations, employs a splitting and Russian roulette strategy underpinned by an importance map, as made available by the adjoint problem [85]. Other notable Monte Carlo codes include MONK, a 3D nuclear criticality and safety code; and TRIPOLI-4, a 3D code used for radiation shielding, reactor physics, criticality and nuclear instrumentation calculations [84, 86].

2.2.4.2 Deterministic Methods

A variety of deterministic methods exist for solving the neutron transport equation in any of its forms. As the rest of this chapter is dedicated to the discretisation of its partial-integral-differential form, Equation 2.8, the most popular methods for each of the integral and the characteristic forms are presented here, respectively and briefly.

The Method of Collision Probability: The method of collision probabilities (MCP) solves the integral form of the NTE, which encapsulates the exponential attenuation of a neutron emitted from some source at $\vec{r}\prime$ in some direction $\vec{\Omega}$ up to a point $\vec{r}$, where it contributes to the flux. Practically, only isotropic scatter is considered; although some anisotropy may be modelled by simply allowing for a transport correction of the cross-section data. The physical domain, V, is discretised into E-number of homogeneous regions, such that $V = \bigcup_{e=1}^{E} V_e$; and the probability that a neutron emitted uniformly and isotropically in region V_i will experience its first collision in region V_j may be calculated in a deterministic manner for all pairs (i, j); i.e. $P_{i \to j} \neq 0$ for $i, j = 1, \ldots, E$. Once all of these collision probabilities are known, a dense $(E \times E)$-system of algebraic equations may be solved for the region-averaged neutron fluxes; and the source may be updated for the next iteration.

The MCP may only be applied to those regions where the incoming neutron current is known; and therefore, the equations must often be formulated over the entire domain V and solved simultaneously. Reciprocity and complementarity, or conservation, relationships do allow for a significant reduction in the number of collision probabilities to be evaluated explicitly. Nonetheless, the method does become computationally prohibitive should more than few hundred regions need to be modelled in multiple dimensions. To remedy this, approximations are routinely made to simplify the geometry. For example in the investigation of the effect of self-shielding in an infinite lattice of infinitely long fuel pins, as presented as a repeated square-shaped unit pin-cell in 2D, one may employ the Wigner-Seitz cell approximation to model this instead as a one-dimensional (1D) system; whereby, instead, a white BC is prescribed over an equivalent cylindrical boundary [87].

The MCP has been implemented in the the THESEUS, FLURIG, PROCOL, PRIZE and PIJ modules of the WIMS lattice physics software [62, 88]. The first two solve the collision probabilities equations in 1D; and the third also solves in 1D for double heterogeneity problems, such as granular fuels within the prismatic fuel assembles of high-temperature reactors (HTR) or the fuel pebbles of pebble-bed modular reactors (PBMR) [89]. PRIZE solves the MCP equations over rz-geometries [90]. Finally, the PIJ module solves in 2D; e.g. for a cluster of pins for a Canadian Deuterium Uranium (CANDU) fuel assembly [91]. An alternative approach to improving the computational efficiency of the MCP is to combine the method with response surface approaches; an example being the multi-cell MCP,

which is employed within the PERSEUS module of WIMS [92, 93].

The Method of Characteristics: As the name suggests, the method of characteristics (MoC) solves the characteristic form of the NTE [94]. Specifically, the characteristic equations are solved along a set of *tracks* that run parallel to each of the characteristics, which are chosen by quadrature rules in a manner not dissimilar to that of the discrete ordinates method; q.v. Section 2.5.2 The physical domain, V, is, again, discretised into E-number of homogeneous regions and, now, the source over each of these regions is approximated as a polynomial expansion; yet often, this is simply taken to be constant, as per the step characteristic method [95]. The NTE is integrated along a *segment* of the track, which is bounded within a given region. For each region in turn, the angular neutron flux may be solved for in an iterative procedure of transport sweeps and source updates. The region-averaged scalar fluxes can then be evaluated by collecting all of the tracks that intersect a given region; and by integrating the average angular flux along each of these tracks over space and angle using quadrature rules. However, in general, the region volumes are not preserved; and often it is necessary to *renormalise* the length of the segments.

A major challenge lies in developing such scalable techniques in 3D. For example, CACTUS3D, a 3D MoC module in the WIMS10 code, employs a *cyclic tracking* algorithm, analogous to that of the 2D CACTUS solver, which has demonstrated poor convergence behaviour of track coverage under refinement and the track data that must be stored for each segment becomes very memory intensive for whole-core calculations [96]. Such problems were addressed in the CACTUSOT solver in WIMS11, which employs a *once-through* tracking algorithm as well as a 2D *slice-based geometry* scheme to optimise the track configuration and minimise the track storage requirements, respectively [97].

In another code, APOLLO3, *compound trajectories* and *chord classification* schemes have been developed to reduce the required amount of track information to be stored [98].

2.3 Second-Order Forms of the Neutron Transport Equation

As it has been introduced, the standard form of the NTE is a linear hyperbolic PIDE with first-order partial derivatives, which require the use of certain spatial discretisation schemes to ensure numerical stability for those problems with abrupt changes in their material properties. It may, however, be manipulated into alternative, yet equivalent, forms that can accommodate alternative discretisation schemes that offer certain attractive features over those standard techniques applied to Equation 2.8. Namely, these are the second-order forms of the NTE, which, as the name would suggest, exhibit second-order derivatives with respect to the spatial variables. The most widely exploited of the second-

order forms is the neutron diffusion approximation, whose introduction is deferred until later. Several other second-order forms will be introduced now; however, special interest is reserved for the SAAF formulation. Such elliptic PIDE allow for the use of efficient matrix solvers, such as the conjugate gradient (CG) method, with effective preconditioners, such as multi-grid algorithms.

The Even-Parity and the Odd-Parity Transport Equations: Together and separately, the even-parity and the odd-parity (EOP) transport equations are further examples of second-order formulations that have been studied extensively since the 1960s [99, 100]. As per the transformation provided by Vladimirov, the full angular flux, φ, may be decomposed into even (symmetric), φ^{+}, and odd (anti-symmetric), φ^{-}, components, such that $\varphi = \varphi^{+} + \varphi^{-}$ [101]. Where the NTE, Equation 2.8, is valid for all unit-vectors, $\vec{\Omega}$, it can be re-written instead for -$\vec{\Omega}$; adding and, separately, subtracting this from the original equation yields a pair of coupled first-order differential equations known as the *canonical* form of the NTE [102]. The second-order EOP equations are obtained by combining those first-order ones in the canonical form and eliminating either of the odd-parity or the even-parity components, respectively. From the definitions of φ^{+} and φ^{-}, the scalar neutron flux, ϕ, can be expressed solely in terms of the former, since $\int_{4\pi} \varphi^{-} \, d^2\vec{\Omega} = 0$; and the net current, $\vec{J}$, can be expressed solely in terms of the latter, since $\int_{4\pi} \vec{\Omega}\varphi^{+} \, d^2\vec{\Omega} = 0$. In fact, the net current may also be expressed solely in terms of the even-parity flux component. Due to symmetry in $\vec{\Omega}$, one must consider only half of the angular domain; i.e. $\varphi^{+}(\vec{\Omega}) = \varphi^{+}(-\vec{\Omega})$ [103]. In most practical stead-state neutron transport problems, it is sufficient to consider only the even-parity (EP) transport equation for φ^{+} [104]. However, the even and the odd-parity components of the angular neutron flux are coupled for the prescription of vacuum BC. Therefore, in the presence of voids, the full angular neutron flux must be calculated. Furthermore, both φ^{+} and φ^{-} are required for time-dependent problems [105]. Indeed, the odd-parity component can be evaluated as the gradient of the even-parity component; however, for polynomial representations of the latter, this can incur a loss in numerical accuracy; viz. φ^{-} appears to be constant for φ^{+} represented by linear basis functions [105].

Generally, Σ_t makes an appearance in the denominator of at least one of the terms of a second-order form of the NTE; e.g. the EP transport equation. A consequence of this is an inability to model, directly and accurately, the streaming of neutrons through a vacuum, where the governing equation is undefined; i.e. $1/\Sigma_t \to \infty$ as $\Sigma_t \to 0$. This becomes very problematic, for example, when modelling regions of steam in ABWR or any air-gaps in radiation shielding problems; and therefore poses a major challenge to the widely-adopted use of second-order forms.

As already mentioned, the even-parity and the odd-parity components of the angular neutron flux are coupled at vacuum boundaries. In principle, the EP transport equation could be solved up to the

vacuum boundary and the FONT equation could be solved within the void region. The EP and the FONT equation could be matched to produce consistent interface conditions by evaluating φ^- from φ^+. Accordingly, this would provide the correct interface conditions for the application of the FONT equation within the void; however, one notes that such an approach would only be feasible for the discrete ordinates equations; q.v. Section 2.5.2. An angular discretisation of the FONT equation by the spherical harmonics method faces a separate issue of singularity within a void region, where the discretised equations become over-determined [106].

One simple workaround would be to treat any void as a region of low absorption and thus, assign some sufficiently small value to the corresponding cross-section, as per some definition of numerical tolerance for a vanishing one $\Sigma_t = 10^{-4}$ cm^{-2}. However, it has been demonstrated that such an approach can lead to the poor conditioning of the resulting matrix equations for finite difference and finite element spatial discretisations, as well as unsatisfactory physical results [107].

The Least-Squares Formulation: Certain least-squares (LS) formulations do avoid the problems posed by voids, although they may become incompatible with standard angular discretisations and solution methods; viz. the discrete ordinates method and the SI procedure, respectively; q.v. Section 2.8 [108]. Hansen et al presented an LS formulation that was compatible with the SI procedure; and one that could benefit from partially-consistent diffusion acceleration schemes [109]. However, it was observed that neutrons were only conserved asymptotically in the fine-mesh limit. Ultimately, such a formulation would be unsuitable in the application to reactor physics; especially so for criticality calculations. It was later demonstrated by Peterson et al that conservation could be enforced by a partially-consistent non-linear diffusion acceleration (NDA) scheme; even in non-diffusive cases [110].

Hammer et al proposed a weighted least-squares (WLS) formulation of the form $L^\dagger w L = L^\dagger w q$, which, again, was compatible with the SI procedure, with $L^\dagger$ denoting the left-hand side of Equation 2.14 [111]. For a particular choice of the weighting function as $w = 1$, the LS transport equation presented by Hansen et al is recovered; and for $w = 1/\Sigma_t$, the WLS formulation reduces to that of the SAAF formulation, up to the BC; q.v. Section 2.3.2. Manteuffel and Ressel had previously reported that difficulties experienced by the LS formulation in the *thick* diffusion limit could be remedied by such a weighted least-squares principle [112]. In voided regions, an upper limit, w_{max}, must be imposed upon the weighting function to ensure well-defined values and thus, the invertibility of the resulting matrix equations; and it was determined that the *best* choice of w_{max} is problem-dependent [113]. Although still wholly non-conservative, it was observed that the introduction of a weighting function significantly reduces any problems due to *causality* [114].

2.3.1 Causality

Causality is a rigorous property of neutron transport; that is, information travels downstream, only, along characteristic lines. However, this can be violated numerically. Whilst such non-physical behaviour is generally considered to be a *coarse-mesh* problem that improves under refinement, it is one that plagues many of the second-order forms of the NTE [111].

A *characteristic* is a path of propagation; or rather, a path along which some entity, such as a physical disturbance, propagates [115]. One appreciates that the solutions to hyperbolic PIDE are typically non-smooth functions of the independent variables; such that pronounced variations in space, angle and energy may be observed in the presence of abrupt changes to the cross-section data; e.g. about material interfaces and physical boundaries. Considering only the spatial variables, the characteristics of a linear hyperbolic PIDE are the paths along which any discontinuities may propagate; and any numerically stable and convergent discretisation scheme must take this into account.

Each of the forms of the steady-state NTE is a linear PIDE with respect to space. In a manner analogous to that of classifying conic sections, one may classify a PIDE by the discriminant of its characteristic equation. However, for this qualitative discussion, it is sufficient to consider that a single first-order PIDE is *hyperbolic*, for which real distinct characteristics exist; whereas a single second-order PIDE is *elliptic*, for which no real characteristics exist because they are imaginary. In turn, this serves to characterise *two* types of problems, *propagation* and *equilibrium* ones. The distinction is made apparent by introducing the notion that the angular neutron flux at a given point in space depends upon anything and everything that may happen in the *domain of dependence*; and it will influence any point that lies within the *range of influence*.

Should real characteristics exist, then there exist preferred paths along which physical information will propagate; and thus, for every point in the solution domain there exist specific domains of dependence and ranges of influence, bounded by the characteristics. That is to say, that the angular neutron flux at any given point depends only upon that which lies upstream and will influence only that which lies downstream; and rather importantly, it does not depend upon that which lies downstream nor will it influence that which lies upstream. Physical problems governed by PIDE that have real characteristics are propagation problems; and thus, hyperbolic PIDE govern propagation problems [116].

Conversely, should no real characteristics exist, then there exist no preferred nor meaningful path along which physical information should propagate, which is characteristic of steady-state diffusion-like problems, where some perturbation is dissipated in all directions. Consequently, the domain of dependence and the range of influence of every point coincide with the entire solution domain; wherein lies the violation of causality. In other words, the solution at every point will depend upon

the solution at all other points and the solution at each point will influence the solution at all other points. An appropriate solution strategy should therefore ensure that all internal requirements and auxiliary conditions are satisfied simultaneously; although some numerical schemes may be devised to allow for some *disconnectedness*, such that a *weaker* solution is sought to approximate better the true solution; qq.v. Chapter 4 and Chapter 5. Physical problems governed by PIDE that have complex characteristics are equilibrium problems; and thus, elliptic PIDE govern equilibrium problems [116].

Now, it should be evident that the first-order and the second-order forms of the NTE constitute *two* very different classes of problems.

2.3.2 The Self-Adjoint Angular Flux Formulation

The SAAF formulation was first derived from a variational principle by Pomraning and Clark for 1D slab geometries; and then more generally from a maximum principle by Ackroyd in multiple dimensions [99, 117]. There are, however, many ways in which to arrive at the SAAF formulation; e.g. by simply adding together the EP and the OP transport equations. The most direct route is the algebraic derivation provided by Morel and McGhee; i.e. rearranging Equation 2.8 as [118]:

$$\varphi(\vec{r},E,\vec{\Omega}) = \frac{1}{\Sigma_t(\vec{r},E)}\big[q(\vec{r},E,\vec{\Omega}) - \vec{\Omega}\cdot\vec{\nabla}\varphi(\vec{r},E,\vec{\Omega})\big]; \tag{2.17}$$

and then substituting the above back into the streaming term in Equation 2.8 to yield:

$$-\vec{\Omega}\cdot\vec{\nabla}\frac{1}{\Sigma_t(\vec{r},E)}\vec{\Omega}\cdot\vec{\nabla}\varphi(\vec{r},E,\vec{\Omega}) + \Sigma_t(\vec{r},E)\varphi(\vec{r},E,\vec{\Omega}) = q(\vec{r},E,\vec{\Omega}) - \vec{\Omega}\cdot\vec{\nabla}\frac{1}{\Sigma_t(\vec{r},E)}q(\vec{r},E,\vec{\Omega}). \tag{2.18}$$

One must solve for the angular neutron flux over the entire angular domain; with the scalar flux evaluated as per Equation 1.9. Although this does imply more computational work when compared to the EP transport equation, it does allow for a more direct implementation of the incoming BC; in a manner similar to that for the FONT equation [118]. In addition to the prescription of incoming BC, as Equation 2.18 is a second-order PIDE, it will be necessary to prescribe outgoing BC, namely, the outgoing angular neutron flux should satisfy Equation 2.8, which can be derived from a maximum variational principle [119].

One remarks that the cross-section of total collision, Σ_t, appears in the denominator of a couple of terms. Therefore, Equation 2.18 becomes increasingly undefined for increasingly thin materials; viz. it becomes increasingly singular as $\Sigma_t \to 0$. By introducing a non-symmetric first-order stabilisation term, Wang et al manipulated the SAAF formulation such that for a specific choice of the stabilisation parameter the overall formulation reduced to the steady-state NTE in a void [120]. Whilst void-compatible methods are useful, such exotic treatments are not considered within the scope of this

work. Pragmatically, one opts to attribute a sufficiently small value of Σ_t to any void regions, where necessary; q.v. Section 5.5.2. One acknowledges that in such situations it may be better to avoid such second-order forms of the NTE since some neutrons are unwontedly absorbed where they would otherwise not be if the FONT equation were employed instead.

2.4 A Discretisation in Energy

In some instances, where the effects of energy may not be of primary importance, it is useful to consider a simplified *one-speed*, or *mono-energetic*, model of neutron transport. That is to integrate Equation 2.8, or Equation 2.18, over the entire energy spectrum and to solve for an energy-integrated angular neutron flux, with energy-averaged parameters, similar to those described in Chapter 1. Likewise, one may apply such a crude treatment of the energy variable over just some part of the energy spectrum; e.g. in the study of either very fast or very slow moving neutrons [121, 122]. In which case, any contributions due to the transfer of neutrons from energies outside of the range of interest are simply regarded as sources, which are assumed to be known, *a priori*.

Indeed, in the absence of energy-dependence, it is much easier to establish the mathematical properties of existence, uniqueness and completeness for elementary solutions to the *one-speed* NTE [123]. Furthermore, analytical *mono-energetic* solutions may be derived for certain geometric configurations, which may then serve as benchmarks against which to verify numerical methods [77]. Of greater importance, however, is the use of such a simplified transport model in the development of a more sophisticated energy-dependent one.

It is well understood that the neutron population in a typical thermal reactor spans a large energy range, over which the microscopic cross-section data varies greatly. For example, fission-born neutrons must be *slowed down* before inducing thermal fission events that then emit fast neutrons; and so on and so forth, *ad nauseam*. Therefore, for any accurate description of neutron transport, it is therefore necessary to consider any energy-dependence in detail. A universally adopted approach is to pose the problem as a set of coupled *one-speed* NTE.

2.4.1 The Multi-Group Approximation

Credit for the current implementation of the multi-group energy approximation is due to Hurwitz and Ehrlich; who reduced the age-diffusion equations to a finite set of coupled linear differential equations called the *multi-group* equations [124]. First, one assumes an upper and a lower limit for the practical range of neutron energies. In the consideration of a typical thermal reactor, one may ignore any neutron

that possesses more than ~ 20 MeV; viz. $E_{max} = 20$ MeV and $E_{min} = 0$ MeV. The continuous energy range, $(E_{min}, E_{max}]$, is then partitioned into distinct values, $E_{max} = E_0 > E_1 > \ldots > E_G = E_{min}$, forming G-number of energy groups; such that the gth-group spans the half-closed interval $(E_g, E_{g\text{-}1}]$. One has chosen a *backward* indexing of the groups to reflect the fact that most neutrons will scatter progressively towards lower energies [77].

Any integrals over the entire energy spectrum can now be expressed as a sum of individual integrals over each energy group; i.e.:

$$\int_0^{\infty} dE \approx \int_{E_G}^{E_0} dE \Rightarrow \sum_{g=1}^{G} \int_{E_g}^{E_{g\text{-}1}} dE. \tag{2.19}$$

The steady-state NTE, Equation 2.8, can then integrated over each energy group to give a system of multi-group equations; i.e. for $g = 1, \ldots, G$:

$$\vec{\Omega} \cdot \vec{\nabla} \varphi^g(\vec{r}, \vec{\Omega}) + \Sigma_t^g(\vec{r}) \varphi^g(\vec{r}, \vec{\Omega}) = q^g(\vec{r}, \vec{\Omega}); \tag{2.20}$$

and similarly for the SAAF formulation, Equation 2.18; i.e. $g = 1, \ldots, G$ [100]:

$$\text{-}\vec{\Omega} \cdot \vec{\nabla} \frac{1}{\Sigma_t^g(\vec{r})} \vec{\Omega} \cdot \vec{\nabla} \varphi^g(\vec{r}, \vec{\Omega}) + \Sigma_t^g(\vec{r}) \varphi^g(\vec{r}, \vec{\Omega}) = q^g(\vec{r}, \vec{\Omega}) - \vec{\Omega} \cdot \vec{\nabla} \frac{1}{\Sigma_t^g(\vec{r})} q^g(\vec{r}, \vec{\Omega}). \tag{2.21}$$

The above *within-group* equations are solved for the multi-group components of the angular neutron flux, φ^g, which is no longer a distribution in energy, but rather an energy-integrated quantity; i.e. for $g = 1, \ldots, G$:

$$\varphi^g(\vec{r}, \vec{\Omega}) = \int_{E_g}^{E_{g-1}} \varphi(\vec{r}, E, \vec{\Omega}) \, dE; \tag{2.22}$$

and the multi-group components of the neutron scalar flux, ϕ^g, are given for $g = 1, \ldots, G$:

$$\phi^g(\vec{r}) = \int_{E_g}^{E_{g\text{-}1}} \int_{4\pi} \varphi(\vec{r}, E, \vec{\Omega}) \, d^2\vec{\Omega} dE = \int_{E_g}^{E_{g\text{-}1}} \phi(\vec{r}, E) \, dE; \tag{2.23}$$

and similarly, the multi-group components of the net neutron current, $\vec{J}^g$, are given for $g = 1, \ldots, G$:

$$\vec{J}^g(\vec{r}) = \int_{E_g}^{E_{g\text{-}1}} \int_{4\pi} \vec{\Omega} \varphi(\vec{r}, E, \vec{\Omega}) \, d^2\vec{\Omega} dE = \int_{E_g}^{E_{g\text{-}1}} \vec{J}(\vec{r}, E) \, dE. \tag{2.24}$$

Accounting for all contributions from all groups, the multi-group emission source densities, q^g, are given for $g = 1, \ldots, G$:

$$\begin{aligned} q^g(\vec{r}, \vec{\Omega}) = & \sum_{g\prime=1}^{G} \int_{4\pi} \Sigma_s^{g\prime \to g}(\vec{r}, \vec{\Omega}\prime \cdot \vec{\Omega}) \varphi^{g\prime}(\vec{r}, \vec{\Omega}\prime) \, d^2\vec{\Omega}\prime \\ & + \frac{1}{4\pi} \chi^g(\vec{r}) \sum_{g\prime=1}^{G} \nu \Sigma_f^{g\prime}(\vec{r}) \phi^{g\prime}(\vec{r}) \\ & + q_x^g(\vec{r}, \vec{\Omega}); \end{aligned} \tag{2.25}$$

where the group-wise macroscopic cross-sections of total collision, Σ_t^g, are defined as per Equation 1.17 and the group-wise differential cross-sections of scatter, $\Sigma_s^{g' \to g}$, are defined as per Equation 1.18; such that they are now constant parameters within a given energy group and the reaction rates have been preserved [87]. In the consideration of large reactor shielding problems, it may not be necessary to model any up-scattering of neutrons; accordingly, the lower limit of the summation of the multi-group approximation for scatter sources may be modified: $\sum_{g' \leq g}$. For a suitable choice of group spacing, a further simplification can be made, such that neutrons scatter only into a neighbouring group; in which case, the multi-group NDE are said to be *directly coupled*.

One has also defined an energy-averaged product of $\nu\Sigma_f$ for $g = 1, \ldots, G$:

$$\nu\Sigma_f^g(\vec{r}) = \frac{\int_{E^g}^{E^{g-1}} \nu(\vec{r}, E)\Sigma_f(\vec{r}, E)\hat{\phi}(\vec{r}, E)\, dE}{\int_{E^g}^{E^{g-1}} \hat{\phi}(\vec{r}, E)\, dE}; \tag{2.26}$$

accompanied by an energy-integrated fission spectrum, χ^g, which is the fraction of all those neutrons emitted from fission events that are born into the gth-group; i.e. for $g = 1, \ldots, G$:

$$\chi^g(\vec{r}) = \int_{E^g}^{E^{g-1}} \chi(\vec{r}, E)\, dE; \tag{2.27}$$

and an energy-integrated extraneous or fixed source of neutrons, q_x^g; defined for $g = 1, \ldots, G$:

$$q_x^g(\vec{r}, \vec{\Omega}) = \int_{E^g}^{E^{g-1}} q_x(\vec{r}, E, \vec{\Omega})\, dE. \tag{2.28}$$

The structure of the multi-group problem closely resemble that of the original one; even though one is now solving for an energy-integrated quantity that depends upon energy-averaged cross-section data.

The semi-discrete problem is posed as a system of G-number of equations coupled weakly through the scatter and fission integral terms; where the group-wise, or within-group, equations are given in the form of effective *one-speed* NTE. In the vast majority of neutron transport problems, one may assume that a neutron will lose energy as it collides with a target nucleus; this suggests a natural manner in which to use the best available information to solve the multi-group equations. Starting with the group with the highest energy, one solves the within-group equations for φ^g in a successive sequence of $g = 1, \ldots, G$; such that the down-scatter source for higher groups are provided for upon converging the *self-scatter* source for the current group. In what is known as the SI procedure, or the power-iteration (PI) procedure in the case of eigenvalue problems, one continues to cycle through each of the groups and to update each of the *group-to-group* sources until they have converged; and whilst the down-scatter sources are known from the current iteration of the SI procedure, any up-scatter or fission sources may only be known from the previous iteration; q.v. Section 2.8 [79].

Likewise, the multi-group approximation may be applied to the adjoint problem [124]. Rather cautiously, however, one notes that the different possible permutations of the order of discretisation and

adjoint operations do not generally commute [125]. This would not be a problem if the direct and adjoint energy-averaged parameters were weighted by the exact direct and adjoint solutions, respectively; in which case, the multi-group approximation would not actually incur any errors. However, such adjoint-weighted cross-sections are not commonly available; and even so, any commutativity would be lost since approximate spectra are used in the evaluation of group-wise cross-section data.

Therefore, with practical consideration, one will always discretise the primal problem with respect to energy before formulating the dual problem.

2.5 A Discretisation in Angle

Various methods have been developed for the discretisation of the angular variable; the most popular being the methods of discrete ordinates (S_N) and spherical harmonics (P_N). The former approximates the angular dependence by a discrete set of angular directions; whereas the latter does this by a truncated expansion of Legendre polynomials in 1D, or spherical harmonic functions in 2D and 3D. The S_N method is chosen for the angular discretisation of Equation 2.20 and Equation 2.21; nevertheless, the P_L method will be presented to express anisotropy in the scattering kernel; such a discretisation is more rigorously dubbed the $S_N P_L$ method, where by convention, N and L are even and odd integers, respectively. Moreover, the P_N method is useful when deriving the NDE.

2.5.1 The Spherical Harmonics Method

Firstly, the angular dependence of the differential macroscopic multi-group cross-sections of scatter, $\Sigma_s^{g\prime\to g}$, is expanded in a basis of spherical harmonics, as per the P_L method; and secondly, this same expansion is applied to the angular neutron flux as per the P_N method.

The P_L method: On the whole, the slowing down process in a thermal nuclear reactor occurs during the elastic scattering of neutrons with light target nuclides in the moderator. For isotropic media, it is enough to assume that the differential macroscopic cross-section of scatter does not explicitly depend upon the angle of incidence, $\vec{\Omega}\prime$, but simply upon the angle of deviation, or rather the cosine of, $\mu = \vec{\Omega}\prime \cdot \vec{\Omega}$; q.v. Equation 1.19. Although the scattering of neutrons may be presumed to be isotropic in the CoM reference system, such an assumption cannot be made in the LAB frame. In fact, forward-peaked scatter becomes increasingly prevalent for increasingly fast neutrons incident upon increasingly small nuclides; where the mean cosine of the scattering angle, $\bar{\mu}_L$, is approximated by Equation 1.20.

The angular dependence in Equation 1.19, may be expressed in a series expansion of Legendre polynomials, P_l, with respect to $\vec{\Omega}\prime \cdot \vec{\Omega}$, as measured in the LAB frame, truncated to the Lth term, which can then substituted into the scattering integral term of the emission source density, Equation 2.25; i.e. for $g = 1, \ldots, G$:

$$\sum_{g\prime=1}^{G} \int_{4\pi} \Sigma_s^{g\prime\to g}(\vec{r}, \vec{\Omega}\prime \cdot \vec{\Omega}) \varphi^{g\prime}(\vec{r}, \vec{\Omega}\prime)\, d^2\vec{\Omega}\prime \approx \sum_{g\prime=1}^{G} \sum_{l=0}^{L} \frac{2l+1}{4\pi} \Sigma_{sl}^{g\prime\to g}(\vec{r}) \int_{4\pi} P_l(\vec{\Omega}\prime \cdot \vec{\Omega}) \varphi^{g\prime}(\vec{r}, \vec{\Omega}\prime)\, d^2\vec{\Omega}\prime; \quad (2.29)$$

where the terms of $l = 0$ and $l = 1$ correspond to those of isotropic and linearly anisotropic scattering, respectively. The Legendre polynomials satisfy their normal recursive formula and the expansion coefficients in the above, $\Sigma_{sl}^{g\prime\to g}$, are defined as:

$$\Sigma_{sl}^{g\prime\to g}(\vec{r}) = \int_{4\pi} P_l(\vec{\Omega}\prime \cdot \vec{\Omega}) \Sigma_s^{g\prime\to g}(\vec{r}, \vec{\Omega}\prime \cdot \vec{\Omega})\, d(\vec{\Omega}\prime \cdot \vec{\Omega}). \quad (2.30)$$

Collecting the above, expansion coefficients of total scatter out of a given group, Σ_{sl}^{g}, may be expressed in a manner similar to that of applying an energy condensation to Equation 1.14; i.e. for $g, = 1, \ldots, G$:

$$\Sigma_{sl}^{g}(\vec{r}) = \sum_{g\prime=1}^{G} \Sigma_{sl}^{g\to g\prime}(\vec{r}); \quad (2.31)$$

and more rigorously, the multi-group approximation of the mean cosine of the scattering angle, Equation 1.20, may be deduced for $g, = 1, \ldots, G$ [47]:

$$\bar{\mu}_L^g(\vec{r}) = \frac{\sum_{g\prime=1}^{G} \int_{4\pi} (\vec{\Omega}\prime \cdot \vec{\Omega}) \Sigma_s^{g\to g\prime}(\vec{r}, \vec{\Omega}\prime \cdot \vec{\Omega})\, d(\vec{\Omega}\prime \cdot \vec{\Omega})}{\sum_{g\prime=1}^{G} \int_{4\pi} \Sigma_s^{g\to g\prime}(\vec{r}, \vec{\Omega}\prime \cdot \vec{\Omega})\, d(\vec{\Omega}\prime \cdot \vec{\Omega})} = \frac{\Sigma_{s1}^{g}(\vec{r})}{\Sigma_{s0}^{g}(\vec{r})} \approx \frac{2}{3A}; \quad (2.32)$$

where A is the atomic mass number of the target nuclide relative to the incident neutron.

The Legendre polynomials in Equation 2.29 projected onto a basis of Laplace's spherical harmonic functions, Y_l^m, as per the following addition property [80]:

$$P_l(\vec{\Omega}\prime \cdot \vec{\Omega}) = \frac{4\pi}{2l+1} \sum_{m=-l}^{l} Y_l^m(\vec{\Omega}\prime) Y_l^m(\vec{\Omega}). \quad (2.33)$$

The real (tesseral) spherical harmonics form a basis over the unit sphere and are defined $\forall l \geq 0$ and $\forall 0 \leq |m| \leq l$ as:

$$Y_l^m(\vec{\Omega}) \equiv Y_l^m(\mu, \psi), \quad (2.34a)$$

$$= \sqrt{\frac{(2l+1)}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^{|m|}(\mu) \begin{matrix} \cos(m\psi) \\ \sin(m\psi) \end{matrix}; \quad (2.34b)$$

where P_l^m are the associated Legendre polynomials of major and minor indices, l and m, respectively; and by convention, $m \geq 0$ and $m < 0$ designate the cosine series and the sine series, respectively.

Equation 2.34b is normalised such that Y_l^m belong to an orthonormal set according to:

$$\int_{4\pi} Y_l^m(\vec{\Omega}) Y_{l\prime}^{m\prime}(\vec{\Omega})\, d^2\vec{\Omega} = \delta_{ll\prime}\delta_{mm\prime}; \tag{2.35}$$

where the Kronecker delta, δ_{ij}, retains its usual definition; i.e.:

$$\delta_{ij} = \begin{cases} 1, & \text{if } i = j; \\ 0, & \text{otherwise.} \end{cases} \tag{2.36}$$

The neutron scatter sources in the Equation 2.25 may now be re-expressed for $g = 1, \ldots, G$:

$$\sum_{g\prime=1}^{G} \sum_{l=0}^{L} \Sigma_{sl}^{g\prime\to g}(\vec{r}) \sum_{m=-l}^{l} Y_l^m(\vec{\Omega}) \Phi_l^{g\prime m}(\vec{r}); \tag{2.37}$$

where the multi-group components of the flux moments, Φ_l^{gm}, are defined as a projection onto the spherical harmonics for $g = 1, \ldots, G$ and for $0 \leq l \leq L$ and for $-l \leq m \leq l$:

$$\Phi_l^{gm}(\vec{r}) = \int_{4\pi} Y_l^m(\vec{\Omega}) \varphi^g(\vec{r}, \vec{\Omega})\, d^2\vec{\Omega}. \tag{2.38}$$

A similar expansion may be performed in the treatment of the fission sources, which is always truncated at zeroth term because it may be assumed that those fission-born neutrons are always emitted isotropically, in either of the reference systems. In fact, in the sole consideration of all other sources being isotropic, it will be sufficient to specify only the strength of any extraneous fixed sources of neutrons, q_x; i.e. for $g = 1, \ldots, G$:

$$\begin{aligned} q^g(\vec{r}, \vec{\Omega}) = & \sum_{g\prime=1}^{G} \sum_{l=0}^{L} \Sigma_{sl}^{g\prime\to g}(\vec{r}) \sum_{m=-l}^{l} Y_l^m(\vec{\Omega}) \Phi_l^{g\prime m}(\vec{r}) \\ & + \chi^g(\vec{r}) \sum_{g\prime=1}^{G} \nu\Sigma_f^{g\prime}(\vec{r}) Y_0^0(\vec{\Omega}) \Phi_0^{g\prime 0}(\vec{r}) \\ & + \frac{1}{4\pi} q_x^g(\vec{r}). \end{aligned} \tag{2.39}$$

The Transport Correction: Some degree of anisotropy can be reclaimed even when formally modelling for isotropic scatter in the LAB frame. For brevity, the derivation is omitted here; however, it is sufficient to know that a linearly-anisotropic transport correction can be made to the macroscopic cross-sections of isotropic scatter, Σ_{s0}, which are then substituted for effective transport-corrected ones, $\tilde{\Sigma}_{s0}$, that include a forward-peaked component; i.e. for $g, g\prime = 1, \ldots, G$ [126]:

$$\tilde{\Sigma}_{s0}^{g\to g\prime}(\vec{r}) = \Sigma_{s0}^{g\to g\prime}(\vec{r}) - \delta_{gg\prime} \sum_{g\prime=1}^{G} \Sigma_{s1}^{g\to g\prime}(\vec{r}). \tag{2.40}$$

Suppressing any spatial dependence to simply the notation, one may use the above in an expression for the effective transport-corrected macroscopic cross-sections of total collision; i.e. for $g = 1, \ldots, G$:

$$\tilde{\Sigma}_t^g = \Sigma_a^g + \sum_{g\prime=1}^{G} \tilde{\Sigma}_{s0}^{g\to g\prime}, \tag{2.41a}$$

$$= \Sigma_a^g + \sum_{g\prime=1}^{G} \Sigma_{s0}^{g\to g\prime} - \sum_{g\prime=1}^{G} \Sigma_{s1}^{g\to g\prime}, \tag{2.41b}$$

$$= \Sigma_a^g + \Sigma_{s0}^g - \Sigma_{s1}^g, \tag{2.41c}$$

$$= \Sigma_t^g - \bar{\mu}_L \Sigma_{s0}^g, \tag{2.41d}$$

$$= \Sigma_{tr}^g; \tag{2.41e}$$

where the first line recalls an energy condensation of Equation 1.15 with the macroscopic cross-section of isotropic scatter substituted for effective transport-corrected one, which is then expanded in the second line as per Equation 2.40. The third line is expressed in terms of total isotropic and total linearly anisotropic scatter; as per Equation 2.31. Equation 1.15 is employed again in the fourth line to reveal Σ_t and Equation 2.32 is used to replace Σ_{s1}; all of which is encapsulated in the fifth line in what is more commonly referred to as the macroscopic transport cross-section, denoted by Σ_{tr}.

The linear anisotropic effects are thus, incorporated into effective transport-corrected cross-sections; and it must be emphasised that the NTE is now rendered no more challenging to solve than it is to solve for isotropic scattering only. This is particularly useful and important for the integral forms of the transport equation, which would otherwise be limited to isotropic scatter [74].

The $\mathbf{P_N}$ method: One can extend the same P_L method beyond the scattering kernel by expanding in a truncated series of spherical harmonics the angular dependence in the angular neutron flux; such a projection is referred to as the P_N method; i.e. for $g = 1, \ldots, G$ [75]:

$$\varphi^g(\vec{r}, \vec{\Omega}) = \sum_{l=0}^{N} \sum_{m=-l}^{l} Y_l^m(\vec{\Omega}) \Phi_l^{gm}(\vec{r}); \tag{2.42}$$

where the multi-group components of the flux moments, Φ_l^{gm}, are those defined in Equation 2.38; and the spherical harmonics, Y_l^m, are those defined in Equation 2.34a.

Only the first two terms in the expansion of the major index, $l = 0, 1$, in Equation 2.42 bear any

physical significance; i.e. for $g = 1, \ldots, G$:

$$\varphi^g(\vec{r}, \vec{\Omega}) = \sum_{l=0}^{N} \sum_{m=-l}^{+l} Y_l^m(\vec{\Omega}) \Phi_l^{gm}(\vec{r}), \tag{2.43a}$$

$$= Y_0^0(\vec{\Omega}) \Phi_0^{g0}(\vec{r}) + Y_1^{-1}(\vec{\Omega}) \Phi_1^{g(-1)}(\vec{r}) + Y_1^0(\vec{\Omega}) \Phi_1^{g0}(\vec{r}) + Y_1^1(\vec{\Omega}) \Phi_1^{g1}(\vec{r}) + \ldots, \tag{2.43b}$$

$$= \frac{1}{4\pi} \left[\phi^g(\vec{r}) + 3 \left[\Omega_x J_x^g(\vec{r}) + \Omega_y J_y^g(\vec{r}) + \Omega_z J_z^g(\vec{r}) \right] \right] + \ldots, \tag{2.43c}$$

$$= \frac{1}{4\pi} \left[\phi^g(\vec{r}) + 3 \vec{\Omega} \cdot \vec{J}^g(\vec{r}) \right] + \ldots \tag{2.43d}$$

From the definition of the multi-group components of the flux moments, Equation 2.38, and the definition of the spherical harmonics, Equation 2.34a, one notes that the first two coefficients in the expansion above for the spherical harmonics of the zeroth and the first degree, correspond to the isotropic part and the linearly anisotropic part of the associated angular distribution of the angular neutron flux, respectively; as made apparent by the scalar neutron flux, Equation 2.23, and the net neutron current, Equation 2.24, respectively.

Substituting the expansion in Equation 2.43 into Equation 2.20 or Equation 2.21, and then multiplying through by Y_l^m, integrating over the unit sphere and exploiting the orthogonality property of the spherical harmonics, Equation 2.35, yields a system of $N+1$ coupled equations. For each major index l, there are $2l+1$ linearly independent spherical harmonics of degree l, which are coupled to those spherical harmonics of degrees above and below. To ensure solvability, the components of the spherical harmonics of degree $l \leq N$ must be decoupled from those of degree $l > N$ by a closure relationship that simply neglects them [127].

For an increasingly large N, the P_N approximation, Equation 2.42, converges quickly for small flux gradients and less so for increasingly large ones; nonetheless, the series does tend towards the exact solution as $N \to \infty$ [128]. However, for general geometries in multiple dimensions, the P_N equations become too complicated to solve for $N > 1$. The number of unknowns to solve for can be reduced by assuming azimuthal invariance in the angular flux distribution; in which case, for $m = 0$, the spherical harmonics reduce to Legendre polynomials. Moreover, approximations can be made to simplify the geometry of the problem; e.g. should it be possible to reduce the system to a 1D slab geometry problem, there exists a formal equivalence between the $P_{N\text{-}1}$ method and the discrete ordinates (S_N) method, to be presented [129]. However, no such equivalence exists in multiple spatial dimensions [130].

Alternatively, the heuristically derived simplified spherical harmonics (SP_N) method serves as an efficient angular discretisation technique that may be readily employed in whole-core calculations [131]. However, as it is based upon an expansion in an incomplete basis of orthogonal functions, it is not guaranteed that the numerical solution will converge as $N \to \infty$ [132]. That being said, the SP_3 and the SP_5 equations have been shown to be more accurate than the P_1, or diffusion, approximation [133].

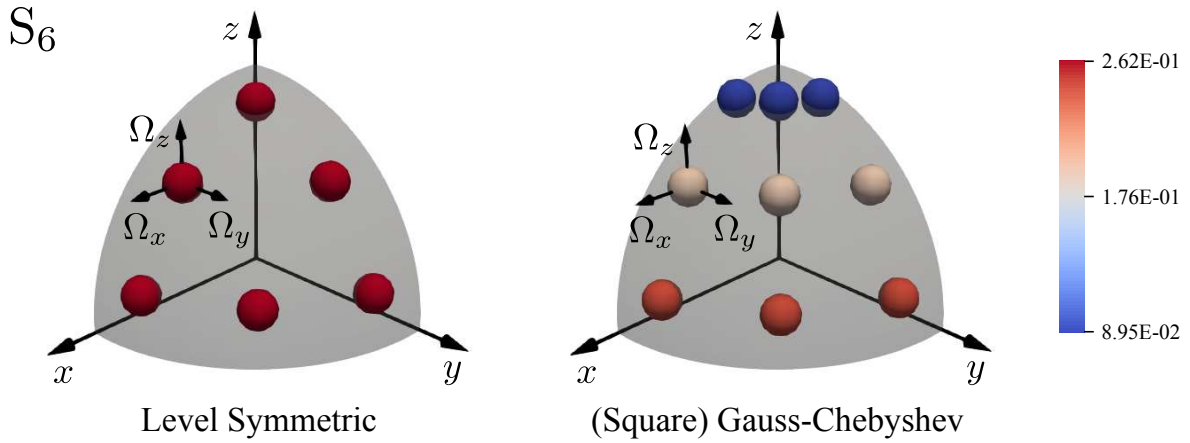


Figure 2.1: An example of *two* different S_6 angular quadrature sets. The LS quadrature set (left) contains $N(N+2)/8$ discrete angular directions per octant of the unit sphere; whereas the product (square) GC quadrature set (right) contains $N^2/4$ [136, 137]. Any discrete angular direction may be represented as a vector whose components are the direction cosines: $\vec{\Omega} = (\Omega_x, \Omega_y, \Omega_z)$ and the associated weight is indicated by the colour map; the weights of a given quadrature set are normalised to $\sum_k w_k = 4\pi$ steradians over the unit sphere. (*Vide* the web-based version for reference to colour.)

Generally, the truncation of the series in Equation 2.37 Equation 2.42 are chosen such that $L \leq N$ and N is odd. The latter allows for a certain symmetry in the implementation of $N+1$ BC that are required by $N+1$ coupled equations; especially for those incoming BC which may apply only to half of the angular domain [75]. For those boundaries where there is no incoming neutron current, the vacuum BC is rather difficult to implement because no finite sum of spherical harmonics could represent such a step function exactly; practically, it is treated approximately as a Mark condition or a Marshak condition [74, 81, 87]. Alternatively, for the half-range expansion in Legendre polynomials in 1D, Yvon's double-spherical harmonics (DP_N) method exactly satisfies the vacuum BC of no-incoming angular neutron flux; this approach has been extended to 2D Cartesian grids [134, 135].

2.5.2 The Discrete Ordinates Method

The method of discrete ordinates (S_N) and much of the preliminary work thereof may be credited to Carlson, Lathrop and Lee [55, 138–140]. This angular discretisation technique has been implemented within industrial 3D radiation transport codes such as PENTRAN and PARTISN; as well as more recent ones, such as Rattlesnake within the MOOSE framework [25, 26, 141–143]. The premise is that any integral over the angular domain can be performed numerically and sufficiently using quadrature rules [144, 145]. Consequently, the integrand of the multi-group components of the flux moments needs only to be evaluated at discrete points; i.e. for $g = 1, \ldots, G$ and for $0 \leq l \leq L$ and -$l \leq m \leq l$:

$$\Phi_l^{gm}(\vec{r}) = \int_{4\pi} Y_l^m(\vec{\Omega}) \varphi^g(\vec{r}, \vec{\Omega}) \, d^2\vec{\Omega}, \tag{2.44a}$$

$$\approx \sum_{k=1}^{K} w_k Y_l^m(\vec{\Omega}^k) \varphi^{gk}(\vec{r}); \tag{2.44b}$$

where for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$\varphi^{gk}(\vec{r}) \coloneqq \varphi^{g}(\vec{r}, \vec{\Omega}^{k}); \tag{2.45}$$

and:

$$q^{gk}(\vec{r}) \coloneqq q^{g}(\vec{r}, \vec{\Omega}^{k}). \tag{2.46}$$

Accordingly, the angular quadrature set consists of a set of pairs of quadrature points (QP)over the unit sphere, or rather the discrete ordinates, and their associated weights, $\{\vec{\Omega}^{k}, w_{k}\}_{k=1}^{K}$. The dimension of the set, $K = K(N)$, depends upon the order of the S_N approximation, N, which refers to the symmetry-preserving latitudinal point arrangement of the discrete ordinates; e.g. an example of *two* different types of S_6 quadrature sets are presented in Figure 2.1.

Naturally, the weights should be normalised such that their summation gives the total area of the unit sphere; whereby one may consider them as being proportional to the areas subtended by the solid angles associated with each of the discrete ordinates; i.e.:

$$\int_{4\pi} d^{2}\vec{\Omega} = \sum_{k=1}^{K} \int_{\vec{\Omega}^{k}} d^{2}\vec{\Omega} = \sum_{k=1}^{K} w_{k} = 4\pi. \tag{2.47}$$

In the evaluation of the multi-group components of the flux moments in Equation 2.44, one chooses simply to solve Equation 2.20 for φ^{gk} at the each of the same QP, $\vec{\Omega}^{k}$, in a manner not dissimilar to that of the collocation method; i.e. for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$ [146]:

$$\int_{4\pi} \delta(1 - \vec{\Omega} \cdot \vec{\Omega}^{k}) \left[L^{g}\varphi^{g}(\vec{r}, \vec{\Omega}) - q(\vec{r}, \vec{\Omega}) \right] d^{2}\vec{\Omega} = L^{gk}\varphi^{gk}(\vec{r}) - q^{gk}(\vec{r}) = 0. \tag{2.48}$$

Weighted by the Dirac delta function, $\delta(\cdot)$, the S_N method minimises the residual at the QP and thus, the differential equation is enforced in a point-wise manner only in those discrete angular directions as chosen by the quadrature set; where the operators L and q represent the left-hand side and the emission source density, respectively, of Equation 2.20. Conversely, however, no attempt is made to approximate any distribution across the angular domain and the angular neutron flux assumes discrete values only at the QP; i.e. for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$ [147]:

$$\varphi^{g}(\vec{r}, \vec{\Omega}) = \begin{cases} \varphi^{g}(\vec{r}, \vec{\Omega}^{k}) =: \varphi^{gk}(\vec{r}), & \text{if } \vec{\Omega} = \vec{\Omega}^{k}; \\ 0, & \text{otherwise.} \end{cases} \tag{2.49}$$

This allows for the implementation of a relatively simple fixed-point inner-iteration scheme for the convergence of the within-group, or self-scatter sources, which itself is nested within the SI procedure for the convergence of the group-to-group sources in an outer-iteration scheme, to be presented in Section 2.8. However, it is precisely this discontinuous definition that leads to a loss of rotational invariance in the solution, which manifests itself in the so-called *ray effects*, to be discussed in Section

2.5.2.1. Besides, with regard to the numerical integration in Equation 2.44, it may be more intuitive to consider the discretised angular neutron flux, φ^{gk}, as well as the discrete angular direction, $\vec{\Omega}^k$, as being piece-wise constants within the associated solid angle of weight w_k; rather than by the point-wise definition in Equation 2.49.

The semi-discrete problem is now posed as a system of $G \times K$ equations coupled weakly through the scattering and fission integral terms; where the multi-group FONT-S_N equations are given explicitly for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk}(\vec{r}) + \Sigma_t^g(\vec{r})\varphi^{gk}(\vec{r}) = q^{gk}(\vec{r}). \tag{2.50}$$

Similarly, for the multi-group SAAF-S_N equations for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$-\vec{\Omega}^k \cdot \vec{\nabla}\frac{1}{\Sigma_t^g(\vec{r})}\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk}(\vec{r}) + \Sigma_t^g(\vec{r})\varphi^{gk}(\vec{r}) = q^{gk}(\vec{r}) - \vec{\Omega}^k \cdot \vec{\nabla}\frac{1}{\Sigma_t^g(\vec{r})}q^{gk}(\vec{r}). \tag{2.51}$$

Just as any collocation method is strongly dependent upon the choice of the quadrature, or collocations, points, the accuracy of the S_N method depends strongly upon the choice of the angular quadrature set. Whilst one may suggest various different quadrature sets, they are all guided by the desire to integrate as accurately and as *cheaply* as possible the multi-group components of the flux moments in Equation 2.44, whilst also incorporating any constraints on symmetry. Moreover, should one be interested in modelling for anisotropic scattering, the order of anisotropic scatter, L, in Equation 2.37, would have to be truncated to the maximum degree of polynomial for which the quadrature would be exact [81]. Unsurprisingly, there are some restrictions in choosing the sets of discrete angular directions and corresponding weights; the first being that posed in Equation 1.4.

As it will be appreciated in the ensuing chapters, there is a great deal of flexibility available in the choice of spatial discretisation and thus, the subsequent spatial refinement of the NTE; yet, it is often customary to apply the same angular resolution, as embodied by the angular quadrature set, to each energy group and to each spatial mesh. Therefore, a second requirement would be to ensure that the S_N method maintains the principle of *optical reciprocity* [55]. In general, it may not be possible to preserve computational invariance of the solution under *all* geometric transformations. However, it will be sufficient to choose an even-numbered, K, angular quadrature set that is invariant under rotations of $\pi/2$ steradians about about each of the axes in an orthogonal Cartesian coordinate system. In practice, the distribution of the QP over the entire unit sphere and their associated weights can therefore be described in reference to that over a single principal octant; and allowing for all possible permutations of sign changes to the components of the discrete angular directions [140]. For those geometries whose physical boundaries align with the coordinate axes, this greatly simplifies the implementation of (specular) reflective BC; such that, for every pair of $\vec{\Omega}^k$ and w_k, the angular quadrature set also contains a pair of $\vec{\Omega}^{k\prime} = \vec{\Omega}^k - 2\vec{n}(\vec{\Omega}^k \cdot \vec{n})$, and $w_{k\prime} = w_k$, where $\vec{n}$ is the outward

facing unit-normal vector on the boundary.

Since the spherical harmonic flux moments are themselves polynomials of the direction cosines, additional constraints are applied in the following *moment conditions*; viz. the quadrature formula should integrate as accurately as possible the following, for $i = x, y, z$ [55]:

$$\sum_{k=1}^{K} w_k (\mu_i^k)^n = \begin{cases} 0, & \text{for } n\text{-odd;} \\ \frac{4\pi}{n+1}, & \text{for } n\text{-even.} \end{cases} \tag{2.52}$$

In particular, Equation 2.47 is recovered for $n = 0$; and the *flow condition* and the *diffusion theorem condition* conditions are revealed for $n = 1$ and for $n = 2$, respectively; where, in fact, any odd-moment condition is guaranteed by the property of optical reciprocity [136, 148].

The level-symmetric (LS) angular quadrature set is one such triangular set that exhibits full rotational-reflection invariance; and hence, there is no bias in any particular direction of neutron flight [140]. In any given octant, for the same distribution of direction cosines on each of the axes, the QP lie on $N/2$ lines of latitude, or polar levels, such that the ith polar level bears $N/2 - i + 1$ QP; thus, accounting for $N(N+2)/8$ discrete ordinates over the surface of an octant. An example of the S_6 LS quadrature set is given in Figure 2.1. Such a latitudinal point arrangement allows for a normally 2D quadrature to be performed as a *direct sum evaluation*; q.v. Equation 2.44 [136]. However, LS-based numerical integration is often inaccurate, even for low-order polynomial functions; and $N \geq 18$ yields negative weights and non-physical solutions [149].

In general, it is sufficient to consider only 2^d octants for a d-dimensional problem; and therefore, conditions of complete symmetry are only required by those angular quadrature sets in the application to full 3D problems. As such, the angular quadrature set for a 2D problem is only required to be invariant under rotations of $\pi/2$ steradians about the principal directional axis; compared to a fully symmetric set that exhibits such invariance about every axis. A relaxation of symmetry requirements in lower-dimensional geometries entails a relaxation in the latitudinal point arrangement, by which one gains in the flexibility of choosing alternative angular quadrature sets, whilst retaining the polar level arrangement and the property of direct sum evaluation.

Such an example of half-symmetric angular quadrature sets are the *product* Gauss-Chebyshev (GC) sets, which combine a 1D Gauss-Legendre polar quadrature rule with an equally-weighted Gauss-Chebyshev azimuthal one [137, 150]. An example of a S_6 square GC quadrature set is given in Figure 2.1; where each of the $N/2$ polar levels has $N/2$ QP. Therefore, accounting for $N^2/4$ discrete angular directions over the surface of any given octant in the hemi-sphere. The set of associated weights is formed simply as the product of those from the well-established Gauss-Legendre and Gauss-Chebyshev quadrature sets; and thus, are guaranteed to be positive for an arbitrarily large N [137].

In the case of 1D slab and spherical geometries, that require no conditions of symmetry, the 2D integration over the unit sphere reduces to a 1D quadrature with respect to the polar cosine, μ, over the interval [-1, 1], for which polynomials of order $2N$-1 or less can be integrated exactly by an N-point Gauss-Legendre rule. For QP chosen as the zeros of the Legendre polynomials, the S_N equations are entirely equivalent to the P_{N-1} approximation, with Mark BC [81, 129]. Howbeit, such an equivalence does not extend to multiple dimensions [151].

It would appear that no single optimal quadrature set exists for all problems; and in fact, the choice of quadrature set is often problem-dependent. For example, deep-penetration problems may favour biased quadratures, which cluster the QP about the preferred direction of forward-peaked scatter [136, 139]. In which case, it may be that such non-symmetric quadrature sets are tailored specifically to the geometry of the problem [80]. An eventual spatial discretisation introduces additional complexity when considering that the final linear system of algebraic equations will depend equally upon angular as well as spatial quadrature rules; where experience has demonstrated that the errors in both are interdependent [55]. Nonetheless, within the scope of the current project, one is content in developing the spatial discretisation techniques in isolation by simply assuming either LS or GC angular quadrature sets [136, 137]. To facilitate a more mathematically rigorous and more fair comparison between the two, triangular GC quadrature sets, with second-order 1D Gauss-Legendre polar quadrature, will be generated such that they contain the same number of QP as the LS set; however, it should be emphasised, that they are entirely different sets.

Finally, it is worth noting that, following the energy discretisation by the multi-group approximation, it is arbitrary the order in which one performs the adjoint operation or the angular discretisation by the S_N method because these two *do* commute [125]. In other words, discretisation of the dual NTE yields the adjoint of the discretised primal NTE; this is the property of discrete adjoint consistency.

2.5.2.1 The Ray Effects

The S_N method yields a system of equations that bears resemblance to the (semi-)analytical multi-group NTE; albeit evaluated in discrete angular directions of neutron flight, as chosen by some quadrature rules. Subject to a spatial discretisation, the resulting system of weakly coupled PIDE may be solved in a relatively straightforward manner; q.v. Section 2.8. For an identical treatment of BC and for any of the discrete directions in which the source does not depend upon the neutron flux, one can obtain an identical solution to the analytical one; viz. the *uncollided* flux distribution [75, 153]. On the contrary, owing to an inability to approximate reasonably any scalar quantity by angular quadrature, even though the integrand may be exact at the QP, the numerical solution may exhibit aberrations and anomalies in the presence of scatter and fission sources; these are the *ray effects*, as manifested by

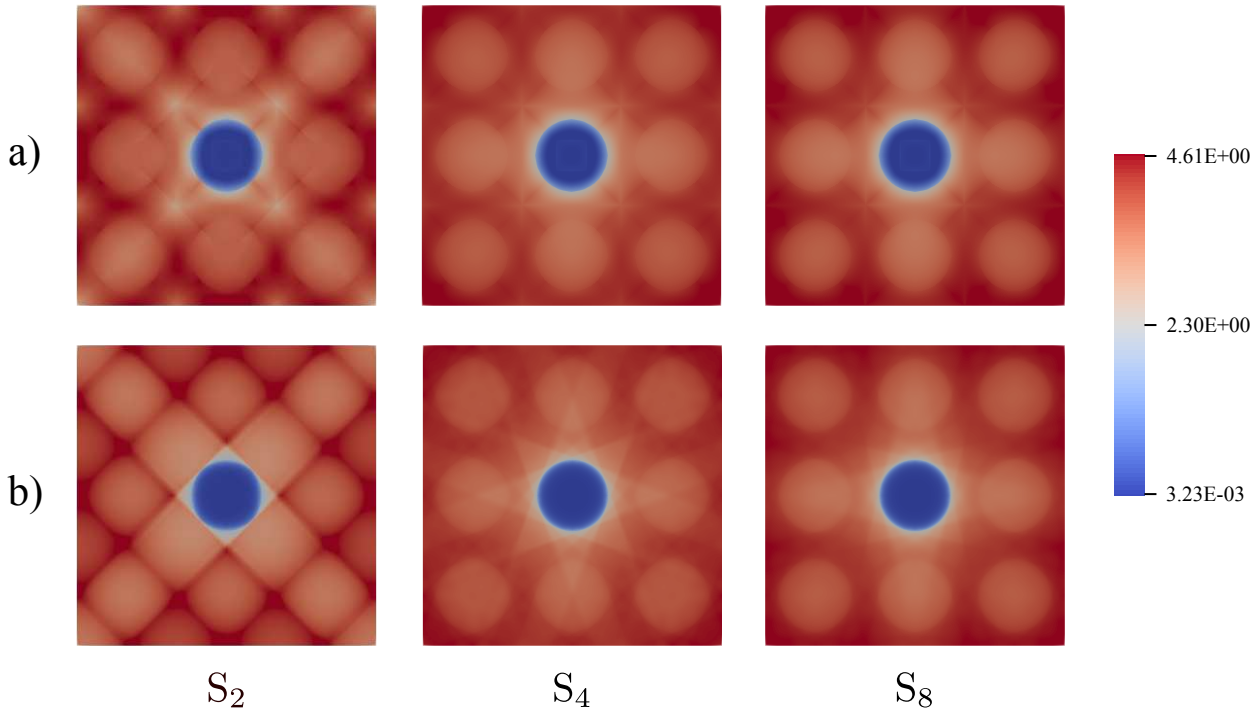


Figure 2.2: A demonstration of the ray effects for the example of the 2D IAEA CEA-1 benchmark and the solution to the 1G sweep-based DG-IGA FONT-S_N equations over a coarse mesh (row a) of 81 bi-quadratic NURBS patches, or 729 DoF per ordinate; and a more refined mesh (row b) of $8,100$ bi-quadratic NURBS patches, or $72,900$ DoF per ordinate; q.v. Table 7.25 [152]. For an increasingly large LS angular quadrature set, $N = \{2, 4, 8\}$ (left to right), the scalar neutron flux ($\text{cm}^{-2}\text{s}^{-1}$) distribution is indicated by the colour map. (*Vide* the web-based version for reference to colour.)

regions of non-physical high and low flux densities [151]. As the continuum of characteristic directions in which neutrons are allowed to stream is restricted to a finite set, so too are the contributions to any scalar quantity, at a given point, restricted to those sources, or sinks, that are intercepted by the discrete ordinates [154].

These ray effects may appear less pronounced in large diffusive systems where distributed sources of scattered neutrons are often isotropic and important, $c = \Sigma_s/\Sigma_t \to 1$; and where the angular distribution of neutrons tends to be smooth [75]. As the spectral radius of the iteration matrix is approximately equal to the scattering ratio, c, the SI procedure can become slow to converge in such scattering-dominated scenarios, which provides the motivation for the acceleration techniques applied in Chapter 7 [40].

It must be emphasised that the ray effects are an inherent *defect* of the the S_N method itself; and exist independently of the discretisation of other independent variables [155]. That being said, once subject to a spatial discretisation, a coarse mesh may *interrupt* the streaming of neutrons along the characteristics; and thus, to some extent, mitigate the ray effects [151]. This is demonstrated in Figure 2.2, where the *one-speed* (1G) sweep-based discontinuous Galerkin FONT-S_N equations are solved for the 2D IAEA CEA-1 benchmark over both a coarse spatial mesh (row a), as well as a more refined one (row b), for increasing orders of the LS angular quadrature set, $N = \{2, 4, 8\}$ [152]. All simulations

were terminated upon convergence of the scalar neutron flux to a tolerance of machine precision in the relative L_2-error in successive iterations. In each instance, the true behaviour of the S_N approximation is only masked by a poor resolution of the spatial mesh.

Somewhat representative of the heterogeneities found in a typical LWR, the 2D IAEA CEA-1 benchmark models the thermal neutron flux over an infinite three-by-three lattice arrangement of *eight* fuel pins that surround a central pure absorber, with a moderator region in between; this benchmark shares the same material compositions and nuclear data as the 2D IAEA CEA-2 benchmark presented in Section 7.4.4 [152]. Although the moderator is representative of a highly scattering material, $c \approx 0.99$, the ray effects are still observed; where it may happen that there exist uninterrupted streams of neutrons in certain discrete directions that do not intercept the sink, thus, resulting in an uneven depression in the scalar neutron flux about the central poison. At the same time, such a defect is equally apparent in the application of the S_N method to the second-order forms of the NTE; q.v. Section 5.5.3. However, due to their elliptic, or numerically diffusive or dissipative nature, the effect may appear less pronounced, qualitatively [156].

Clearly, the severity of the ray effects depends upon the choice of the angular quadrature set and whilst persistent, they are diminishingly so for the added computational cost of an increasingly large N; v. Figure 2.2 [154]. Adaptive angular refinement techniques are challenging to implement because the standard angular quadrature sets are not hierarchically defined. For this reason, specially designed ones have been devised [149]. Various other attempts have been made to mitigate the ray effects, whilst leveraging the simple structure of underlying the S_N method; including the phase-space finite element methods (FEM) and wavelets [157–164]. Such angularly *smeared* methods may mitigate the ray effect to some extent. However, only the rotationally invariant spherical harmonics methods have the potential to eliminate them altogether [156]. Indeed, one should expect the ray effects to persist in any angular discretisation that retains the hyperbolic approximation of the streaming operator [165].

To reiterate, the ray effects are due to a lack of rotational invariance in the discretised streaming operator in the S_N equations, which is not exhibited by the corresponding $P_{N\text{-}1}$ equations [166]. Lathrop went further to demonstrate that this general non-equivalence between the *two* angular discretisation techniques may be attributed to a retention of additional moments in the S_N equations [151]. It was then devised that this could be compensated for by establishing $S_N \rightarrow P_{N\text{-}1}$ equivalence through the inclusion of some *fictitious source* term [146, 154, 167, 168]. Such an approach was implemented in the TWOTRAN code [169]. However, in the transformation to the spherical harmonics equations, any desirable convergence properties of *matrix-free* sweep-based solvers are lost; as is any potential to benefit from standard acceleration schemes that are compatible with the original S_N equations.

Such exotic ray-effect mitigation techniques are beyond the scope of this work. One recalls the aim

of developing refinement strategies to minimise the error due to the spatial discretisation; as well as to apply convergence acceleration techniques that are compatible with the standard S_N equations. Whilst the ray effects may feature, it is usually the spatially-integrated scalar quantities such as the fluence, detector response or general reaction rates that are specific QoI; in which case, the S_N method may still allow for reasonable accuracy [75]. Should it be necessary, however, one is content in simply increasing the size of the angular quadrature set; whereby the associated computational cost is expected to increase linearly with the number of discrete angular directions; i.e. $\mathcal{O}(K)$. Therefore, it becomes increasingly important to implement some form of HPC parallelisation within the numerical code; such that each discrete ordinate is not simply solved for in rude succession. The investigation of such parallel sweep-based algorithms on regular Cartesian grids lead to the (modified) Koch-Baker-Alcouffe (KBA) algorithm [170, 171]. That being said, such scalable algorithms over general unstructured meshes remain an active field of research [172].

2.6 Diffusion Theory

Evidently, thus far, the discretisation of the NTE with respect to each of its independent variables has become an increasingly arduous task. By neglecting any detailed directional aspect of neutron transport, a simplified model can be found in neutron diffusion theory, which is readily employed in whole-core reactor physics codes such as the nodal code, PANTHER [63].

As it happens, diffusion theory is an asymptotic limit of transport theory; whereby, despite presenting itself formally as a hyperbolic PIDE, Equation 2.50 behaves like an elliptic one in the asymptotic diffusion limit; i.e. where $\Sigma_t \to \infty$ and $\Sigma_a \to 0$ [173, 174]. Equally, the (semi-)analytical second-order forms should retain the same asymptotic properties in the diffusion limit because one expects the same analytical solution to the NTE, irrespective of its form [109]. Ideally, this should hold for the spatially discretised equations; however it is not always guaranteed in the intermediate and thick diffusion limits; and especially so in the vicinity of unresolved boundary layers [148].

Although the full derivation is omitted here for brevity, the multi-group NDE are somewhat equivalent to the multi-group P_1 equations; wherein one assumes a rather weak and linearly anisotropic expansion of the angular neutron flux by truncating Equation 2.43 to $N = 1$ and all higher-order moments, $l \geq N+1$, are ignored; i.e. for $g = 1, \ldots, G$:

$$\begin{cases} \vec{\nabla} \cdot \vec{J}^g(\vec{r}) + \Sigma_t^g(\vec{r})\phi^g(\vec{r}) = \sum_{g\prime=1}^{G} \Sigma_{s0}^{g\prime\to g}(\vec{r})\phi^{g\prime}(\vec{r}) + \chi^g(\vec{r}) \sum_{g\prime=1}^{G} \nu\Sigma_f^{g\prime}(\vec{r})\phi^{g\prime}(\vec{r}) + q_x^g(\vec{r}); & \text{(2.53a)} \\ \frac{1}{3}\vec{\nabla}\phi^g(\vec{r}) + \Sigma_t^g(\vec{r})\vec{J}^g(\vec{r}) = \sum_{g\prime=1}^{G} \Sigma_{s1}^{g\prime\to g}(\vec{r})\vec{J}^{g\prime}(\vec{r}). & \text{(2.53b)} \end{cases}$$

The first equation ($l = 0$), Equation 2.53a, is simply the neutron conservation equation, or rather a balance that relates the divergence of the net neutron current to the scalar neutron flux and any isotropic sources. The second equation ($l = 1$), Equation 2.53b, is the corresponding conservation equation for the net neutron current, or rather a balance that relates the gradient of the scalar neutron flux to the net currents of any anisotropic sources due to scatter.

The micro-reversibility approximation (MRA), or the out-flow approximation, can be used to simplify Equation 2.53b, whereby one assumes that the total anisotropic scatter of neutrons into a given group is equal to that out of the same group; i.e. for $g = 1, \ldots, G$ [60]:

$$\sum_{g\prime=1}^{G} \Sigma_{s1}^{g\prime\to g}(\vec{r})\vec{J}^{g\prime}(\vec{r}) = \sum_{g\prime=1}^{G} \Sigma_{s1}^{g\to g\prime}(\vec{r})\vec{J}^{g}(\vec{r}). \tag{2.54}$$

Strictly speaking, such a detailed balance would only be valid at thermal energies, where the projectile neutrons are assumed to be in (quasi-)equilibrium with the target nuclei. With slight modification, the above *row sum* approach can extend to epi-thermal energies, where down-scatter processes dominate [175]. At higher energies, however, such an approximation does not hold; and from this point on, the multi-group equations depart from the multi-group P_1 equations [54, 176].

Although it was not detailed explicitly, the outflow approximation was implicit in deriving the transport-corrected cross-sections in Equation 2.40 and Equation 2.41a; and they reappear now in defining the transport-corrected diffusion coefficient, D (cm); viz. by substitution of Equation 2.54 into Equation 2.53b and then rearranging to yield Fick's (first) law as it applies to neutron diffusion for $g = 1, \ldots, G$ [175]:

$$\vec{J}^{g}(\vec{r}) = -\frac{1}{3\Sigma_{tr}^{g}(\vec{r})}\vec{\nabla}\phi^{g}(\vec{r}) = -D^{g}(\vec{r})\vec{\nabla}\phi^{g}(\vec{r}). \tag{2.55}$$

Under steady-state conditions, the above postulates that the diffusive flux, or rather the net neutron current, tends in the direction of low neutron density across the concentration gradient, or the spatial derivative of the scalar neutron flux. This may be a reasonable assumption for optically thick media where the mean free path (MFP), $\sim 1/\Sigma_{tr}$, is of the same size as the collision length; and where the scattering events dominate any streaming behaviour. Accordingly, any densely populated regions of neutrons, as characterised by frequent scatter and random trajectories, will exhibit higher-rates of collision and the neutrons will tend to scatter more frequently towards less populated regions [47].

The rate at which the neutrons diffuse depends upon the diffusion coefficient, D, which is generally defined as a (3×3)-tensor whose entries are directional diffusion coefficients. The full tensor-product form of D is useful for representing the anisotropic nature of neutron diffusion in problems that have been spatially homogenised by lattice physics calculations; e.g. neutron streaming in AGR, where the radial and axial diffusion coefficients can be different [177]. However, it has been simplified somewhat

by the outflow approximation, in Equation 2.54; whereby any directional dependence was lost in removing the need for specific neutron current spectra. For most PWR lattice calculations, the effects due to isotropic streaming can be sufficiently represented by non-directional and transport-corrected diffusion coefficients [74].

By substitution of Equation 2.55, into the first of the P_1 equations, Equation 2.53a, the multi-group NDE are obtained for $g = 1, \ldots, G$:

$$-\nabla \cdot D^g(\vec{r})\nabla\phi^g(\vec{r}) + \Sigma_r^g(\vec{r})\phi^g(\vec{r}) = \sum_{g\prime\neq g}^{G} \Sigma_{s0}^{g\prime\to g}(\vec{r})\phi^{g\prime}(\vec{r}) + \chi^g(\vec{r}) \sum_{g\prime=1}^{G} \nu\Sigma_f^{g\prime}(\vec{r})\phi^{g\prime}(\vec{r}) + q_x^g(\vec{r}); \tag{2.56}$$

where the within-group scatter, or self-scatter, terms have been subtracted from either side; and the effective macroscopic cross-section of removal, Σ_r, is defined for $g = 1, \ldots, G$:

$$\Sigma_r^g = \Sigma_t^g - \Sigma_s^{g\to g} = \Sigma_a^g + \sum_{g\prime\neq g}^{G} \Sigma_s^{g\prime\to g}. \tag{2.57}$$

It is markedly *cheaper* to solve the multi-group NDE than it is to solve any higher-order approximation of the NTE. However, neutron diffusion theory only provides a valid mathematical description of the behaviour of neutrons when the assumptions implicit in its derivation are satisfied [80, 178]. The most stringent being that the angular neutron flux distribution varies *gently* in space [87, 179]. This is conceptually equivalent to the linearly anisotropic behaviour of the angular flux assumed in P_1 approximation, which, it must be reminded, is itself a *low-order* approximation of the NTE [180]. Importantly, this assumption of weak angular dependence is violated wherever sharp gradients appear in the neutron flux distribution, such that strong spatial variations infer strong angular dependence and thus, anisotropic behaviour; viz. in the vicinity of localised sources and strong absorbers; and within a couple MFP of physical boundaries and material interfaces [39, 181].

With regard to continuity in the angular neutron flux, the interface condition, Equation 2.7, must be enforced inexactly since the P_1 equations consider only the first two moments in Equation 2.42. In order to enforce the conservation of neutron transport across an interface it will be sufficient to ensure continuity in the scalar neutron flux as well as continuity of the normal component of the net neutron current, respectively, across that same interface; e.g. for some interface ∂V_{12} that separates two regions, V_1 and V_2, and for $g = 1, \ldots, G$:

$$\left.\begin{aligned} \phi^g(\vec{r})|_{V_1} &= \phi^g(\vec{r})|_{V_2}; && \text{(2.58a)} \\ -D^g(\vec{r})\vec{\nabla}\phi^g(\vec{r}) \cdot \vec{n}(\vec{r})|_{V_1} &= -D^g(\vec{r})\vec{\nabla}\phi^g(\vec{r}) \cdot \vec{n}(\vec{r})|_{V_2} && \text{(2.58b)} \end{aligned}\right\} \quad \forall \vec{r} \in \partial V_{12};$$

where $\vec{n}$ is the unit-normal vector located on the interface pointing from $V_1 \to V_2$; and the net neutron current has been expressed as per Equation 2.55. From the second of the interface conditions above, one remarks that for *two* different materials the diffusion coefficient will be discontinuous in $\vec{r}$ across ∂V_{12} and so too will the gradient of the scalar neutron flux discontinuous across the interface.

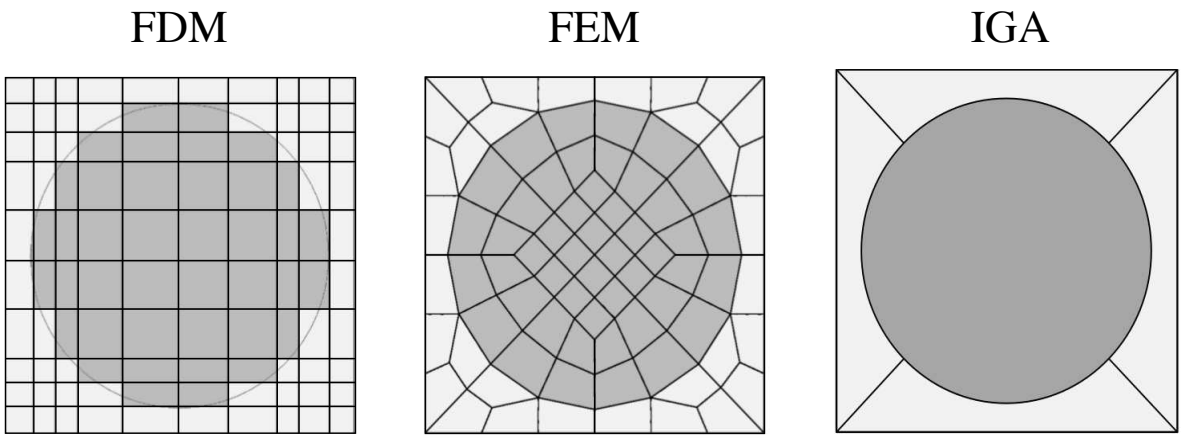


Figure 2.3: A comparison of the typical meshes employed by different spatial discretisation techniques for the example of a 2D pin-cell: FDM, FEM and NURBS-based IGA, respectively. With respect to FDM and linear FEM, the underlying geometry of the central fuel pin, given in dark grey, may be better approximated under refinement of the mesh; however, this is exactly represented in IGA by the coarsest mesh of bi-quadratic NURBS patches. (*Vide* the web-based version for reference to colour.)

Similarly from Equation 2.42 and Equation 2.55, the outgoing and incoming partial currents in Equation 1.11 may be approximated by diffusion theory as [74]:

$$\vec{J}^{g\pm}(\vec{r}) \approx \frac{1}{4}\phi^g(\vec{r}) \pm \frac{1}{2}\vec{J}^g(\vec{r}) \cdot \vec{n}(\vec{r}) \approx \frac{1}{4}\phi^g(\vec{r}) \mp \frac{1}{2}D^g(\vec{r})\vec{\nabla}\phi^g(\vec{r}) \cdot \vec{n}(\vec{r}). \tag{2.59}$$

In spite of its limitations and despite the complex core configurations of most reactors, neutron diffusion theory serves as a vital computational tool for reactor analysis in providing *equivalent* homogenised core calculations, which rely heavily upon detailed, heterogenous, lattice calculations provided by transport theory [87, 182]. Typically, however, deep-penetration reactor shielding problems are less amenable to a direct application of diffusion theory because the often highly energetic neutrons tend to stream over large distances between interactions in regions with small cross-sections; and in the limit of vanishingly small cross-sections in voided regions, the diffusion coefficient is undefined as per Equation 2.55; i.e. $D \to \infty$ as $\Sigma_t \to 0$.

A better approach would be to employ diffusion theory in tandem with transport theory; where a compatible and consistent, or at least partially consistent, DSA can be devised to *precondition*, or accelerate the fixed-point convergence of the latter [183]. In essence, intermediate-iteration diffusion corrections are made to at least the zeroth-degree flux moments; where upon convergence, the accelerated solution is the same as that obtained without acceleration, just in a fewer number of iterations [80]. Such acceleration strategies form the topic of Chapter 7.

2.7 A Discretisation in Space

Analytical solutions to the (semi)-discrete FONT-S_N equations do exist for certain simplified problems that typically assume *one* single energy group (1G) in either 1D slab geometries or homogeneous 2D Cartesian geometries [184–186]. With some manipulation, the 1G NDE can be expressed as a Helmholtz equation, to which the fundamental solution and geometric buckling for bare and uniform reactors of simple shape are separable and well known, analytically, in eleven different coordinate systems [187–189]. Whilst such analytical solutions are useful for code verification and benchmarking purposes, they are often intractable for general engineering problems.

Following the discretisation of the NTE with respect to energy and angle, all that remains is to perform that with respect to space. Accordingly, this current section allows for an informal and non-exhaustive discussion of spatial discretisation techniques, including the finite differences methods (FDM) and the finite elements methods (FEM). A comprehensive discussion of NURBS-based isogeometric analysis (IGA) is reserved for Chapter 3.

The goal is to approximate a continuous field by a set of discrete values at discrete points, or nodes; as defined by some appropriately constructed grid or physical mesh; whereby, the defining continuous linear operator equation is formulated locally at each of the discrete points, yielding a system of algebraic equations that can be solved, as per Equation 1.2; where the standard modern approach is to employ Galerkin-based FEM [190]. The spatial discretisation is said to be *consistent* under mesh refinement if the sequence of numerical solutions converges in some sense to the true solution [191].

2.7.1 The Finite Differences Method

The partial-differential operator in Equation 2.50 can be approximated by a discrete (central-)difference relation; and for a grid defined by a set of cells of characteristic size, h, one can then formulate a system of equations in volume-averaged and surface-averaged quantities [74]. Provided some relationship exists between these *two* different quantities over a given cell, their calculation can proceed as a transport sweep across the grid in each direction of propagation, $\vec{\Omega}^k$, that belongs to the angular quadrature set [55]. The solution to any outgoing face-averaged values of a given cell are completely determined by that of all incoming face-averaged values, as well as the volume-averaged values; and once these outgoing values are known, they can then serve as incoming ones for those cells downwind. One of the earliest and most common closure relationships applied to the FONT-S_N equations was that of the diamond-difference (DD) scheme; whereby the volume-averaged value is taken simply as the the arithmetic average of the surface-averaged values [138].

The FDM continues to garner attention because of its relative simplicity and ease of implementation [192, 193]. In the application to second-order elliptic systems, the resulting matrices are banded by dint of the *nearest-neighbour* coupling; yielding diagonal, tridiagonal and pentadiagonal matrices in 1D, 2D and 3D, respectively [87]. One can take advantage of this by employing compressed diagonal storage (CDS) data structures for sparse matrices.

A poor choice of the grid spacing, h, however, especially in optically thick regions, can result in negative surface-averaged values and non-physical flux oscillations, which would require some form of *fix up* or stabilisation; e.g. the step method or a weighted DD scheme, respectively [54, 55]. More importantly, FDM is unsuitable in its application to unstructured meshes and complex geometries [194]. For example, the *stepped* approach employed to model the central fuel pin in Figure 2.3 is inaccurate; instead, it may prove more effective to simply homogenise the entire pin-cell region [186]. Reed did attempt to capture greater geometric detail by using triangular meshes; however, this does not appear to have gained much traction [195].

2.7.2 The Method of Weighted-Residuals

In a manner similar to that of the P_N method in Section 2.5.1, the spatial dependence in the angular neutron flux, φ, can be expanded in a basis of linearly independent analytical functions, u, dubbed the *trial functions*, whose functional dependence on position $\vec{r}$ is known *a priori*; i.e. $g = 1, \ldots, G$ and for $k = 1, \ldots, K$ [196]:

$$\varphi^{gk}(\vec{r}) \approx \hat{\varphi}^{gk}(\vec{r}) = \sum_{j=1}^{N} \hat{\varphi}_j^{gk} u_j(\vec{r}), \quad \forall \vec{r} \in V. \tag{2.60}$$

For practical reasons, only the first N members are retained and the trial solution, $\hat{\varphi}$, only approximates the true angular neutron flux. Accordingly, the trial functions, u_j, are chosen such as to ensure an increasingly better approximation of φ for an increasingly large N [197].

Inserting Equation 2.60 into whichever form of the NTE is of interest, Equation 2.50 or Equation 2.51, the unknown expansion coefficients, or control-variables, can be determined by first recasting the governing PIDE in some natural integral statement and then requiring it to be stationary with respect to independent variations in $\hat{\varphi}_j$ for $j = 1, \ldots, N$. To elucidate, the domain of this abstract functional, F, is a linear vector space of functions of position, $\vec{r}$, amongst which is the solution to the PIDE of interest, $\hat{\varphi}$, as characterised as a stationary point, $\delta F = 0$ [198]. The original problem has now become one of minimising a system of linear algebraic equations with adjustable parameters; generally, this is known as Ritz's synthesis method, as it pertains to stationary principles; or rather the Rayleigh-Ritz method if it pertains more specifically to extrema principles [199–201].

For $\hat{\varphi}$ characterised variationally, it is possible to deduce the *strong form* of a set governing PIDE, including BC. In fact, the Euler-Lagrange equations can be established for any variational principle; viz. those equations that characterise a stationary point of the variational statement, or the *weak form*, as the solution to the strong form of the governing PIDE [196]. However, this does not typically hold for the contrary as only certain forms of PIDE can be identified as the Euler-Lagrange equations of a variational problem; namely, those that are second-order and self-adjoint [101]. In other words, Ritz's method can be applied to any variational problem, yet variational principles cannot be derived for every PIDE [196].

Due to the group-to-group sources, the different forms of the multi-group NTE are generally not self-adjoint. Therefore, it may be more conducive to reformulate a given problem in the more general context of the method of weighted-residuals (MWR), which can be applied to a larger class of problems without the need for explicit variational principles [202, 203].

By substitution of φ for $\hat{\varphi}$, the governing PIDE defined over a general space of functions, are mapped to approximate ones defined over a trial space spanned by the trial functions, to which the trial solution also belongs. Having assumed its form in Equation 2.60, one must ensure, at least in a weak sense, that the trial solution satisfies the NTE along with any BC, the extent to which can be quantified by the residual. Where, obviously, the residual of the projected equations is vanishing for an exact solution that also belongs to the same trial space; i.e. $\hat{\varphi} = \varphi$. However, such a feat is rarely achievable.

Alternatively, yet equivocally, one can demand that the residual be weighted by some test function, w, integrated over the entire real physical domain, V, and then be made to be vanishing for each member of the set $\{w\}_{i=1}^{M}$; viz. dropping the superscripts g and k that denote energy-group and discrete angular direction, respectively, and for $i = 1, \ldots, M$:

$$\int_V w_i(\vec{r}) \big[L\hat{\varphi}(\vec{r}) - q(\vec{r}) \big] \, dV = 0; \tag{2.61}$$

where the group and angle-dependent operators L and q may represent the left-hand side and the emission source density, respectively, of Equation 2.50, for example.

In practice, there exist many different possible sets of test functions, each one corresponding to its own numerical method; and although definitively there appears to be no best choice, the convergence of successive approximations of increasingly large N and M may give some indication [201]. One remarks that the collocation method, for example, is one such derivation of the MWR; whereby the test functions are chosen as Dirac delta functions and the residual is made to be vanishing in a point-wise manner at the collocation points [204].

2.7.2.1 The Bubnov-Galerkin Method

More formally, one can denote the finite-dimensional vector spaces to which belong the trial and test functions as $\hat{\mathcal{U}}$ and $\hat{\mathcal{V}}$, respectively; such that $N = \dim \hat{\mathcal{V}}$ and $M = \dim \hat{\mathcal{V}}$. By containing all of the unknowns, φ, on the left-hand side of the weighted-residual statement and the known given data on the right-hand side of Equation 2.61, the discrete weak problem may be posed concisely as finding $\hat{\varphi}^{gk} \in \hat{\mathcal{U}}$, such that for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$a^{gk}(\hat{\varphi}^{gk}, \hat{w}) = b(\hat{w}), \quad \forall \hat{w} \in \hat{\mathcal{V}}; \tag{2.62}$$

where for simplicity, each of the multi-group components of the discrete angular neutron flux is sought in the same solution space; with the discrete bilinear form $a^{gk}(\cdot,\cdot)\colon \hat{\mathcal{V}} \times \hat{\mathcal{U}} \to \mathbb{R}$, and the discrete linear functional $b(\cdot)\colon \hat{\mathcal{V}} \to \mathbb{R}$, whose properties depend upon whichever form of the NTE is under investigation. As would be the case of any chosen discretisation technique, the bilinear form must be analysed for its mathematical convergence properties; viz. in terms of consistency, continuity, or boundedness, and coercivity, or ellipticity.

The total number of unknowns, N, or DoF, corresponds to the truncated expansion in Equation 2.60; and since the weighted-residual statement must be satisfied for each member of the set of test functions, there are as many conditions to solve for as there are test functions. To determine the control-variables, φ_j, uniquely, there must be as many algebraic equations as there are DoF; i.e. $M = N$. As such, one may suggest to remove any need for a separate calculation of the weighting functions by simply taking them from the same space as the trial functions, or *vice versa*; i.e. $\hat{\mathcal{U}} = \hat{\mathcal{V}}$. In particular, the residual is made to be vanishing for the choice of a set of orthogonal functions; since it will itself be made orthogonal with respect to each of the test functions of this chosen set. So therein lies the basic principle of the Bubnov-Galerkin method; where the distinction is made from the Petrov-Galerkin method for the general case $\hat{\mathcal{U}} \neq \hat{\mathcal{V}}$ [205, 206].

In most applications of Galerkin-based methods the weighted-residual statement is reformulated by Green's identities, or integration by parts, such as to express any suitably differentiable linear operators as those of reduced orders of derivatives; whence justifies its appellation the *weak form*, since any required smoothness in the weak solution is relaxed; whilst the strong form calls explicitly for a solution that is differentiable to the order in which it appears in the governing PIDE [194]. In relaxing the requirements for continuity in the trial functions, tighter ones are imposed on the test functions; viz. an appropriate choice would ensure $C^{s\text{-}1}$ continuity in the basis of the functions for the strong form of a set of governing PIDE that exhibit at most s-order derivatives [197]. Moreover, in applying the divergence theorem, Equation 2.5, surface integrals over the physical boundary appear in the weak form, which lends itself to the introduction of those BC that may be satisfied *naturally* by the

weighted-residual statement [207]. Those that are not are of the essential type; for it is essential that they be built directly into the solution space or weakly enforced in some sense [197].

One recalls that self-adjoint second-order forms may be formulated as variational principles; and should the variational statement be of no great significance and account for the natural BC, then Ritz's method will yield a variational form that is equivalent to the weak form of the Bubnov-Galerkin method [200,201]. As such, it may be seen that the terms *variational* and *weak* are used interchangeably; even when variational principles can be derived.

2.7.3 The Finite Elements Method

The finite elements methods (FEM) are a class of numerical technique used to model a variety of scalar, vector and tensor field problems; including, heat transfer, fluid flow, electric and magnetic fields and structural analysis [194, 208–210]. Within the field of nuclear reactor physics, FEM were first applied to the static multi-group NDE, for which considerable gains in accuracy and efficiency were demonstrated over standard FDM, especially in 2D [211–213]. Most first-order and second-order forms of the NTE have since benefitted from a discretisation by the FEM; including for the application to complete phase-space [157,214–217]. A good example of the commercial use of FEM is that implemented in the general-purpose ANSYS software [218].

Rather than implicitly formulating the integral statement over the entire physical domain, V, with some single expression for the basis functions throughout V, a more systematic approach would be to construct the approximation spaces locally over simple contiguous subdomains, V_e, called the *finite elements*, which, it can be assumed, retain their respective homogeneity and well-defined properties and can be arranged in some convenient manner [219]. In doing so, it becomes easier to treat any problem that presents irregular boundaries and material heterogeneity [191]. Moreover, refinement of the numerical technique can proceed simply by enlarging the number of elements; such that the total DoF of the discrete field increases whilst the finite element function spaces themselves remain relatively small [196]. Essentially, the FEM encapsulate Galerkin's method in a piece-wise manner; and clearly, continuity of the numerical approximation must be enforced over V in some manner.

An automated ancillary mesh generator is needed to decompose the problem geometry into a tessellation of E-number of non-overlapping regular geometric primitives such as triangles, quadrilaterals, tetrahedra and hexahedra; whose vertices connect together at common nodal points to form the physical mesh, $\mathcal{T}(V) = \{V_e\}_{e=1}^{E}$, over which the numerical analysis is performed [220]. It is possible for FEM to model a wide range of complex geometries. However, in the case of general non-Cartesian geometries, for example in Figure 2.3, $\mathcal{T}$ may represent only an approximation of V, albeit an increas-

ingly better one under refinement; i.e. $V \leftarrow \bigcup_{e=1}^{E} V_e$ as $E \rightarrow \infty$. One notes that the preservation of geometry and thus, conservation of mass, is of critical importance in the application to nuclear reactor physics; and although volume and surface conservation algorithms do exist, it is often not possible to perform both in tandem [221].

Nevertheless, the strength of FEM lies is in the union of the elements; whereby the algebraic equations can be formulated locally and the global weighted-residual statement obtained in a systematic *assembly* of the element-wise contributions; i.e. [197]:

$$\int_V dV = \sum_{e=1}^{E} \int_{V_e} dV. \tag{2.63}$$

By their piece-wise definition and compact local support, the basis functions are assigned a *narrow base* and the element-wise assembly of dense *blocks*, as per Equation 2.63, produces banded matrices, at least in the case of linear elliptic BVP; whose sparsity pattern is directly related to the global enumeration of the basis functions [222]. As the computational effort varies proportionally with the bandwidth of the matrix, it is desirable to minimise this as much as possible; wherefore permutation algorithms, such as the reverse Cuthill-McKee algorithm, can be employed to optimise the bandwidth of symmetric sparse matrices [223].

Within the scope of this thesis, new spatial discretisation techniques and strategies have been developed within a new numerical code that is based upon a standard FEM framework; it is therefore worthwhile to recount the *four* basic stages required for analysis:

1. Pre-processing:
    - to read in the geometry, material data and all other problem specifications
    - to identify the bases of the trial and test functions
    - to generate the physical mesh
    - to assign the BC
2. Assembly:
    - to establish the global as well as the local mesh connectivity information
    - to build the dense element contributions
    - to assemble the global integral statement
3. Processing:
    - to solve the discretised system of algebraic equations
4. Post-processing:
    - to compute any secondary results
    - to write out and display all relevant results

Any higher-level decisions are made at the pre-processing stage; viz. *which* problem to solve for is determined by supplying an appropriate description of the geometry, the material data and any other problem specifications; the user must also determine *how* the finite element mesh is generated and assign the BC. The global and the local ordering of the basis functions, as per the discrete connectivity and the discrete decomposition, respectively, must be established and the equations over the elements formulated before the *connected model* is assembled as a simple matrix sum. The processing stage is generally the most computationally intensive part of the analysis; where the manner in which the matrix equations are *inverted*, or *smoothed*, depends upon the spatial discretisation, as governed by the nature of the problem. Assuming that no further refinement of the model is required, any secondary results are computed and all relevant analysis results can be visualised; otherwise, one returns to the pre-processing stage, where a new physical mesh is generated and the analysis repeated.

2.7.3.1 Isoparametric Elements

Typically, the finite elements are defined in terms of interpolatory polynomials; e.g. those belonging to the classical Lagrange and Hermite families of polynomials [222]. Accordingly, the basis functions are more commonly known as *interpolation* or *shape* functions [224]. The *element library* may contain many different types of regular element shapes such as triangles and quadrilaterals; however, it is desirable that the elements be able to assume irregular shapes in order to allow for the modelling of structures with curved surfaces or edges that meet at obtuse angles, for example. This can be achieved by means of some coordinate transformation, whereby a model element, as described by its local coordinates in a regular reference space, (ξ, η, ζ), is mapped into curvilinear form, as described by its global coordinates in the real physical space, $\vec{r} = (x, y, z)$; i.e. $\vec{r} = \vec{r}(\xi, \eta, \zeta)$.

The transformation mapping itself is arbitrary and general in its application to any element topology; where a natural choice is to approximate the geometry, $\mathcal{T}(V)$, in the same basis as that chosen to approximate the solution; and such is the *isoparametric* concept [225, 226]. Thus, the shape functions describe not only the displacements of the field within a given element, V_e, but also the geometric shape of that element; i.e. for N_e-number of shape functions, u_i, with local support over V_e:

$$\vec{r} = \sum_{i=1}^{N_e} \vec{r}_i u_i(\xi, \eta, \zeta), \quad \forall V_e \in \mathcal{T}(V); \tag{2.64}$$

where $\vec{r}_i$ are the global coordinates of the node that is associated to the ith shape function.

Although versatile, the isoparametric elements are not arbitrarily irregular since they must adhere to certain conditions; viz. those of compatibility and non-degeneracy [227]. More obviously, an element can only assume those geometric forms that are permitted by the shape functions as they interpolate

the nodes; and often the interpolant does not exactly coincide with the underlying geometry [228]. For example, conic sections, such as the central pin in Figure 2.3, cannot be represented exactly by any number of piece-wise polynomial functions, even higher-order ones [229]. In practice, it is often much simpler to apply mass preservation techniques to a mesh of straight-sided, or planar-faced, (linear) elements than it is to one of higher-order curvilinear elements because the placement of any mid-side or mid-face nodes would need to be determined in some pre-processing step; thereby preserving neither volume nor surface area exactly at an additional cost [230, 231]. This is unfortunate because higher-order methods are capable of yielding enhanced rates of convergence; and more specifically, they may benefit those sweep-based solvers, for which the computational effort is proportional to the number of elements that belong to the mesh; q.v. Section 2.8.1. That being said, the use of linear finite elements is further motivated by the unrelated issue of Runge's, or Gibbs', phenomena; whereby the approximation of certain functions by increasingly higher-order polynomial interpolation can become divergent over a set of equally spaced nodes [232].

Fortunately, the isoparametric concept is general enough to incorporate easily bases more exotic than those introduced here. Indeed, this is envisaged in developing NURBS-based IGA as a convergent spatial discretisation technique for neutron transport; q.v. Chapter 3. Sufficient conditions for convergence are often satisfied by a reasonably smooth and complete basis; and this last property is guaranteed for any isoparametric basis that also forms a partition of unity [224].

2.7.3.2 The Discontinuous Galerkin Method

As has been alluded to, the piece-wise definition of the finite element approximation spaces does naturally introduce some degree of discontinuity into the numerical solution. Assuming that the finite elements satisfy conditions of compatibility, the standard approach, as per the Bubnov-Galerkin method, is to globally enforce C^0 continuity across the domain by essentially *gluing* together the shape functions across the element interfaces. Practically, this is achieved by assigning a unique global number to each of the nodes that appear in the connected model. Such a *strong* enforcement of continuity, however, may not be suitable for those problems in which the solution is known to exhibit sharp gradients or discontinuities because a continuous Bubnov-Galerkin (CBG) spatial discretisation will try to smoothen these out [196]. In extreme cases, a CBG spatial discretisation may even introduce spurious numerical artefacts for those problems that present strong material heterogeneity; where abrupt changes to the cross-section data may induce non-physical oscillations in the numerical solution. Indeed, the typical finite element interpolation functions are known to oscillate when attempting to fit points of discontinuous data, this is the Gibbs phenomenon, for which increasingly violent oscillations can be observed for increasingly higher-order polynomials [233].

A CBG spatial discretisation of the first-order form of the neutron transport equation is particularly unstable in the presence of strong material heterogeneity; Martin and Duderstadt demonstrated this for the example of Reed's 1D problem, where the code FTRAN gave extremely poor results for a discretisation of the phase-space by continuous linear elements [207]. The geometry configuration and the nuclear data for Reed's 1D problem are presented in Section 5.5.2 [234]. For this same problem, Martin and Duderstadt then employed discontinuous angular elements; whereby a decoupling of the forward-travelling and the backward-travelling angular components saw a significant improvement to the overall solution [207]. A qualitatively similar result is presented in Section 6.3 for the CBG-IGA spatial discretisation of the 1G FONT-S_8 equations; where one recalls that an angular discretisation by the discrete ordinates method, Section 2.5.2, is inherently discontinuous by nature [235]. The same authors for the same problem then took the natural step of allowing for discontinuities in the solution across the spatial mesh; whereby jump conditions across a discontinuity were incorporated into the integral law, or weak form, as natural interface conditions in the direction of neutron travel [207]. Weakly enforcing continuity in the angular neutron flux in such a manner was found to be equivalent to that of the standard numerical upwinding scheme employed by most discrete ordinates codes; in either case, the coupling of variables between elements is achieved solely through surface boundary integral terms [194,236]. The ensuing results justify the current practice of treating the FONT equation by discontinuous methods; that is, an application of discontinuous Galerkin (DG-)FEM to phase-space effectively mitigated the spurious oscillations in almost every region of Reed's 1D problem [207].

In reducing continuity across the spatial mesh, the support of the basis functions is localised; as are any possible numerical pollution effects [229]. Moreover, by dint of the local formulation, the elementary *stencil* is reduced to a given element and every neighbour that shares an interface, only. This simplifies the implementation of the AMR strategies in the chapters that follow [237]. However, this comes at the cost of effectively doubling the DoF of the connected system since the control-variables across an element interface are no longer unique [194]. As a further matter, problems of unresolved boundary layers may persist, although perhaps to a lesser extent when compared to a similar CBG spatial discretisation; such that non-physical *bumps* may still present themselves about strong material discontinuities where the numerical solution is unable to resolve sharp gradients in the field. This is generally a coarse mesh problem that usually resolves itself under refinement of the spatial mesh; viz. for elements that are small with respect to the MFP within a boundary layer [109]. Accordingly, optically thick elements may only be appropriate in the asymptotic diffusion limit.

Reed and Hill first introduced DG methods for solving the FONT equation over triangular meshes of varying polynomial degree; as implemented in the code TRIPLET [235, 238]. Their work was then extended to a multi-group formalism, as implemented in the TRIDENT code; however, this was limited to linear basis functions only [239]. For similar reasons to those discussed in Section 2.7.3.1, much of

the work on DG-FEM discretisations over unstructured meshes has focused on straight-sided (linear) finite elements [236, 240, 241]. Although continuous and discontinuous methods alike employ similar types of basis, or shape, functions, their respective approximation spaces are fundamentally different. In essence, CBG methods make use of a single piece-wise continuous function space that is local to the whole domain; the DG methods, however, make use of several piece-wise discontinuous function spaces that are local to the finite elements. Consequently, DG methods are locally conservative over individual elements and thus, by extension, globally conservative over the whole domain [242]. Conversely, CBG methods are globally conservative over the whole domain, as well as locally conservative over those elements that share a common node; viz. they are locally conservative over a tent function of that shared node [243].

Again, due to their local formulation, the discontinuous elements are not linked together in any parametric sense; and they must be weakly coupled across their interfaces. In the case of the multi-group FONT-S_N equations, this coupling is achieved via a numerical upwinding scheme, q.v. Chapter 6; and whilst there exist myriad coupling schemes for elliptic problems, one opts for an interior penalty (IP) scheme in the application to the multi-group NDE and the multi-group SAAF-S_N equations; qq.v. Chapter 4 and Chapter 5, respectively.

2.8 The Source-Iteration Procedure

Reserving any formal spatial discretisation for the coming chapters, the (semi-)analytical steady-state multi-group FONT-S_N and SAAF-S_N equations have been derived as a system of equations weakly coupled in energy and angle; q.v. Equation 2.50 and Equation 2.51, respectively. A solution strategy must be devised to determine the order in which they are to be solved; a task that is somewhat more complicated than simply solving $(G \times K)$-number of discrete equations, since unknown variables appear on both sides of the equations due to the coupling of φ in the scatter and fission sources; q.v. 2.39. The most basic of strategies is that of the (scattering) SI procedure, which under steady-state conditions, *ceteris paribus*, can be implemented as *two* embedded stages of iteration; v. Figure 2.4 [55]. Effectively, an inner-iteration procedure resolves the spatial and angular dependence of the angular neutron flux within a given energy group; whereas an outer-iteration procedure resolves the overall energy dependence.

Inner Iterations: It is natural to start with the inner-most part of the problem and work outwards. That is, for a given group, g, the SI procedure resolves the angular neutron flux in all discrete directions of neutron flight; assuming that the group-to-group sources are known *a priori*, this can be treated as

an effective mono-energetic problem [80]. Using the best available information at each outer iteration of index (j), the solution to a given within-group equation proceeds in a succession of inner iterations of index (i); e.g. letting the operators L^g and q^g denote the left-hand side and the emission source density of Equation 2.50, respectively, in the restriction to the gth group, for $i \geq 0$:

$$L^g \varphi^{g(i+1,j+1)} = q^{g(i,j+1)}, \tag{2.65a}$$

$$= S^g_{\downarrow} \varphi^{(j+1)} + S^g_o \varphi^{g(i,j+1)} + S^g_{\uparrow} \varphi^{(j)} + q^g_x; \tag{2.65b}$$

where the scattering contributions to the right-hand side have been split into the down-scattering, self-scattering and up-scattering components, respectively; and sources due to fission have been ignored. The spherical harmonic flux moments, Φ, are successively updated at each inner iteration, as per Equation 2.44, which are used to evaluate $q^{(i+1)}$, in the next iteration.

The inner-iteration procedure is terminated upon either attaining the maximum number of inner iterations, i_{max}, or upon convergence of the within-group (self-scatter) source; where the error due to any spurious sources introduced into the self-scattering source must be vanishing with respect to the actual sources. A test for convergence would be for the difference in the flux moments between successive inner iterations to be less than some criterion, τ_i; i.e. $\epsilon^{(i)} = \|\Phi^{(i)} - \Phi^{(i\text{-}1)}\| \leq \tau_i$.

As is the case for any iterative method, one must supply an initial guess to generate a sequence of increasingly better approximations; for the particular choice of $\varphi^{(0)} = 0$, one may interpret $\varphi^{(i)}$ as the angular neutron flux due to those neutrons that have been scattered at most $(i$-1)-number of times [40]. Consequently, the SI procedure can be slow to converge in optically-thick regions where the neutrons will undergo many scattering events during their life time; viz. an important source of within-group scattering neutrons will require a large number of inner iterations to suppress sufficiently any spurious errors introduced therewithin [55, 183]. These are precisely the situations in which diffusion theory serves as a good approximation to transport theory; q.v. Section 2.6. In fact, the NDE, Equation 2.56, has no need for any inner-iteration procedure since the self-scatter source is removed from either side of the equation and the *unknowns* are contained only on the left-hand side.

Outer Iterations: The inner-iteration procedure assumes prior knowledge of the group-to-group sources; however, this is actually just the best available information for the scatter and fission sources. That is, they can only be known *a posteriori*, since they themselves depend upon the multi-group components of the flux moments, which may have yet to be (re-)evaluated [74]. This suggests a succession of outer iterations of index (j), not dissimilar to a Gauss-Seidel iterative method, in which one cycles through each of the energy groups in turn, resolving the respective multi-group component of the flux moments via the inner-iteration procedure; and (re)computing the group-to-group sources [54].

Starting with the highest-energy group, for $g = 1, \ldots, G$, the down-scatter sources, $g \to g\prime$ for $g < g\prime$, can be evaluated upon convergence of the respective within-group inner-iteration procedures; and thus, it can be used by those lower-energy groups in the same outer iteration. Yet, any up-scatter sources, $g \to g\prime$ for $g > g\prime$, could only be determined for the following $(j{+}1)^{th}$ iteration; i.e. for $j \geq 0$ [244]:

$$(L - S_{\downarrow} - S_o)\varphi^{(j+1)} = S_{\downarrow}\varphi^{(j)} + q_x. \tag{2.66}$$

The error due to any spurious sources introduced into the group-to-group sources must be vanishing with respect to the true sources; and a test for convergence would be for the difference in the multi-group components of the flux moments between successive outer iterations to be less than some criterion, τ_o; i.e. $\epsilon^{(j)} = \|\Phi^{(j)} - \Phi^{(j\text{-}1)}\| \leq \tau_o$. In the case of criticality calculations, the eigenvalue, $1/\mathrm{k_{eff}}$, must be updated between outer iterations using the Rayleigh quotient, for example; and this is more commonly known as *power iteration* [245]. Appropriately, an additional stopping condition would be $|\mathrm{k_{eff}}^{(j)} - \mathrm{k_{eff}}^{(j\text{-}1)}| < \tau_o$; although this convergence criterion would generally be stricter than that imposed to test the convergence of the flux moments [76]. In the absence of up-scatter and fission sources, the outer iteration will simply terminate after a single iteration.

2.8.1 The Transport Sweep

Although their discretisation with respect to the independent spatial variables has only been discussed rather informally, one can expect that applying FEM-like methods to Equation 2.50 or Equation 2.51 will yield $(G \times K)$-number of large, sparse, linear systems of algebraic equations over the spatial domain. Otherwise, a corresponding matrix equation, Equation 1.2, will need to be solved for each within-group equation in a given direction of neutron travel; corresponding to process P-1 on Figure 2.4. Due to its general size and sparsity, (inexact) iterative methods must be employed to *invert* a given coefficient matrix, $\underline{\underline{A}}$; often requiring only simple matrix-vector multiplication operations. Nevertheless, the processing stage is often the most computationally demanding phase of the analysis; excluding, perhaps, that of mesh (re)generation.

The SI procedure will not change under different spatial discretisations; however its convergence properties are strongly affected by it, as are the properties of any potential acceleration techniques; q.v. Chapter 7. This is in contrast to the rather benign effect of the $\mathrm{S_N}$ angular discretisation for the choice of standard symmetric angular quadrature sets [40]. More importantly, the actual manner in which $\underline{\underline{A}}$ is inverted is strongly dependent upon the spatial discretisation, which itself, is determined by the nature of the governing PIDE.

Marching Schemes: Hyperbolic PIDE govern propagation problems; viz. initial-value problems posed in open domains in which an initial *disturbance* propagates at finite speeds along characteristic lines; q.v. Section 2.3. Since the steady-state multi-group FONT-S_N equations are first-order PIDE, only *one* (incoming) BC is required to describe the *initial state* of the angular neutron flux, from which it *marches* forward in the time-like direction.

For reasons given in Section 2.7.3, only DG(-FEM-like) methods will be considered for the spatial discretisation of Equation 2.50, subject to which an arbitrary assembly of the globally connected model will yield *reducible* matrices, $\underline{\underline{A}}$. That is, permutations of a given $\underline{\underline{A}}$ can be expressed in block lower-triangular form, which would also be obtained in assembling the element *blocks*, $\{\underline{\underline{A}}_e\}_{e=1}^{E}$, in the same order in which the neutrons march across the physical mesh in a given direction of travel.

Owing to the discontinuous approximation space, the discrete problem admits local formulations; and since information propagates from the upwind direction the local streaming and total collision operators can be inverted in sequential fashion by *following the neutron*, as per block forward substitution [55]. Consequently, there is no need for an explicit *assembly* stage, since for a given energy group, one now solves for E-number of local systems in each $\vec{\Omega}^k$ for $k = 1, \ldots, K$; thus, *one* within-group *transport sweep* corresponds to *one* inner iteration in Figure 2.4, after which the respective multi-group component of the flux moments can be (re-)evaluated. Generally, the asymmetric local blocks themselves are relatively small and can be efficiently inverted using direct methods; e.g. via LU decomposition. Such is the matrix-free approach of a numerical upwinding scheme [236].

Clearly, it must be known *a priori*, the order in which to sweep the elements to be solved for in $\vec{\Omega}^k$; this may be determined by some topological sorting algorithm, to be developed more specifically in Chapter 6. Ideally, it should be possible to resolve the spatial distribution of the neutron flux in a given group in one transport sweep; viz. in the case of 2D straight-sided triangular meshes, the sweeping order is *acyclic* and there are no dependencies in the *digraph* associated with $\vec{\Omega}^k$ [246]. However, in the general case of skewed or curvilinear elements, the sweeping order will be *cyclic* and permutations of $\underline{\underline{A}}$ will not be strictly block lower-triangular; in which case, any cycles must be resolved within the strongly coupled components (SCC); q.v. Chapter 6 [247].

Relaxation Methods: Elliptic PIDE govern equilibrium problems; viz. BVP posed in closed domains in which a *disturbance* will dissipate in all directions; q.v. Section 2.3. Since the multi-group SAAF-S_N equations are second-order PIDE, *two* BC will be required to define the steady-state system; these are the incoming and outgoing BC.

In contrast to the above, the domain of dependence and the region of influence do coincide for equi-

librium problems; and the globally connected model requires simultaneous satisfaction of all internal requirements and BC. This translates as *irreducible* matrices, $\underline{\underline{A}}$, that must be inverted by iterative relaxation methods, or *smoothing*, of which the *two* main classes are the stationary (Richardson) iterative methods and the more general Krylov subspace methods [116]. Interestingly enough, as it is implemented in Figure 2.4, the SI procedure is itself a fixed-point iteration scheme, which can be *preconditioned*; reformulated as the GMRes method; and both in tandem; q.v. Chapter 7.

Since the enumeration of the basis functions and the connectivity of the mesh are determined in the pre-processing stage, any data structures should be pre-allocated to avoid any dynamic (re-)allocations of new blocks of memory; likewise, any data structures that share the same non-zero sparsity pattern should be reused, wherever possible [248,249]. Fortunately, Galerkin FEM-type methods for the spatial discretisation of each of the considered second-order forms of the NTE, Equation 2.51 and Equation 2.56, yield symmetric positive definite (SPD) matrix equations, which can be solved efficiently by the CG algorithm, with either one-level or multi-level preconditioners [250, 251].

Until recently, sweep-based algorithms for the FONT-S_N equations were significantly more computationally efficient to solve for than preconditioned CG algorithms for second-order forms of the NTE [252]. However, with the advent of modern HPC hardware architecture and software technologies, this performance gap between the *two* approaches is now closing. Indeed, the latest exascale HPC hardware architectures are primarily limited by memory bandwidth. In which case, there may be an eventual shift towards favouring the development of massively parallel multi-level matrix solution algorithms for higher-order spatial discretisations of second-order forms of the NTE.

2.9 A Discussion

This chapter introduced the NTE in several of its most common PIDE forms. Several legacy techniques were presented for their discretisation with respect to each of the independent variables; yet, details of the IGA spatial discretisation are reserved for Chapter 3. Where standard FEM rely upon ancillary mesh generators in order to perform the analysis over what is typically just an approximation of the true geometry, there is a clear desire to model exactly those geometric forms that are ubiquitous in reactor physics problems and radiation shielding problems; v. Figure 2.3. It would be more desirable still to perform the numerical analysis over the same description of the geometry, as provided generally by some CAD software. In the aspiration of coupling with other codes, it would be convenient for each of the different numerical analyses to be performed over the same representation of the underlying geometry. This provides the initial motivation for NURBS-based IGA spatial discretisations, which will be developed in a new numerical code written in modern Fortran [73]. Since IGA and classical

FEM are both isoparametric implementations of Galerkin's method, they may share very similar code architectures; q.v. Section 2.7.3 [229]. Whilst an existing code may offer facilities to automate or ignore much of the code development, it may not allow for the identification of those important decisions that could be made to improve its overall quality. In writing a new code for research, one gains full control of *what* is implemented, along with a sense of personal achievement; perhaps, admittedly, to the detriment of performance.

There is also the motivation to develop AMR strategies to treat unresolved boundary layers; whilst avoiding any superfluous (over-)refinement of other regions. It is envisaged that the error due to the spatial discretisation can be controlled in an iteration of adaptive mesh refinements that can be wrapped around the outer-iteration loop of the SI procedure; v. Figure 2.4. Moreover, NURBS provides a good framework within which to implement such a simple cell-based AMR strategy as that based upon the hierarchical nesting of meshes from previous iterations.

On a final note, the discretisation of *two* different types of PIDE has unsurprisingly resulted in *two* different methods for resolving their respective linear systems of algebraic equations. On the one hand, the discretisation of the steady-state multi-group FONT-S_N equations yields a system that is weakly coupled in energy, angle and space, for which the matrix equations are best solved in a numerical upwinding scheme. On the other hand, the discretisation of the steady-state multi-group SAAF-S_N equations yields a system that is weakly coupled in energy and angle but tightly coupled in space, for which the matrix equations are best solved via iterative relaxation methods. Since there is no argument to exclude either of these *two* forms of the NTE, the current research proceeds in the development of NURBS-based IGA spatial discretisation techniques and refinement strategies for both.

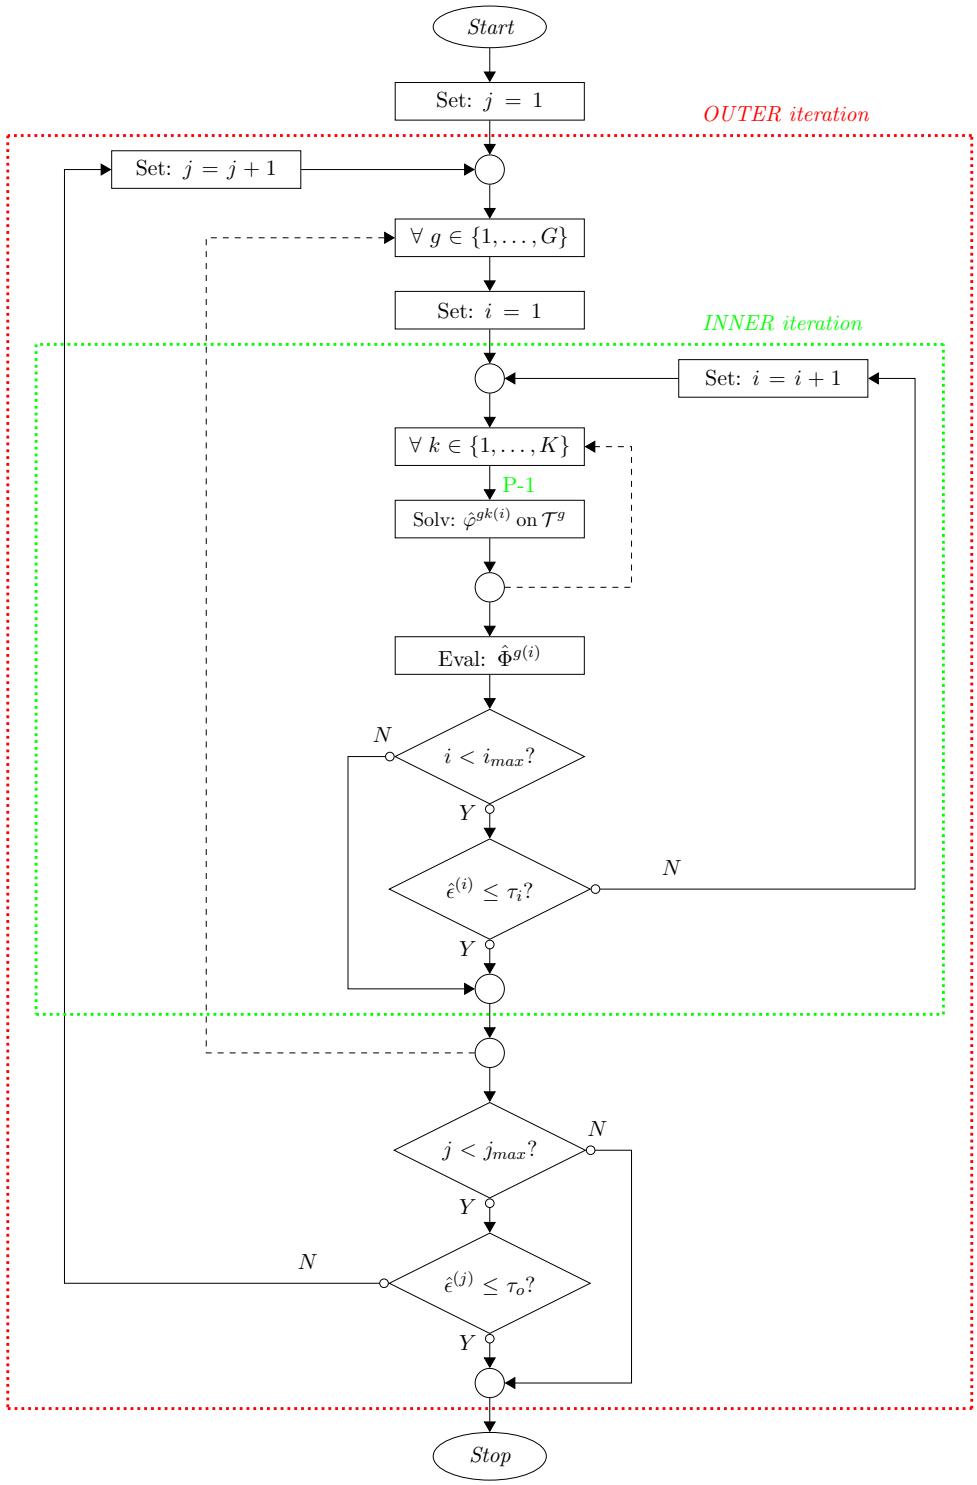


Figure 2.4: The SI procedure depicted as a flowchart. The black dashed process-flow arrows distinguish *looping* operations, which take priority. The within-group spatial and angular dependence of the angular neutron flux is resolved in an inner-iteration procedure, enclosed in the green dashed line; whilst the overall energy dependence is resolved in an outer-iteration procedure, enclosed in the red dashed line. The *inversion* of the within-group matrix equations is performed by Process P-1, the manner in which depends upon the spatial discretisation, as per the nature of the problem. (*Vide* the web-based version for reference to colour.)

Chapter 3

NURBS-Based Isogeometric Analysis

This chapter introduces NURBS-based isogeometric analysis (IGA) as the spatial discretisation of choice for the development of subsequent chapters; the basic concepts are discussed, the process of describing geometries explained and a global uniform refinement strategy presented. Finally, *two* different interpolation-based error measures are devised to drive *two* different AMR-iteration procedures; each with the capacity for energy-dependent meshes.

3.1 An Introduction

Galerkin-based FEM serve as powerful tools for numerical analysis because they can model complex behaviours over multi-dimensional curvilinear geometries. This is achieved by decomposing the geometry into a collection of basic geometric primitives; e.g. triangles and quadrilaterals in 2D; and tetrahedra, prisms, pyramids and hexahedra in 3D. Both straight-sided geometries, as well as curvilinear ones, can be modelled using the isoparametric concept where the solution field within an element and the geometry of the element employ the same higher-order basis functions; q.v. Section 2.7.3.1. There are, however, *two* major shortcomings of existing FEM technologies.

Firstly, the vast majority of time and effort spent in the overall engineering design-and-analysis process is devoted to creating analysis-suitable meshes; viz. mesh (re)generation can account for up to 80% of the total time spent on large-scale engineering problems [233]. Secondly, most general designs cannot be represented exactly by low-order piecewise polynomials. In which case, the finite element mesh will not coincide with the underlying geometry; and the analysis is performed only on an approximation thereof. Furthermore, any difficulties encountered in the meshing of small and detailed features often justifies ignoring them; viz. *feature removal*, or geometrical de-featuring, is common practice when generating low-order finite element meshes [229]. Such an approach does not readily lend itself to

the study of nuclear reactor physics or criticality calculations where the preservation of fissile mass could not be assured without an appropriate smearing of the cross-section data or some other mass preservation technique [230]. Even still, it may not be possible to model with great precision the neutron leakage across a mere approximation of a curved boundary; v. Figure 2.3 [253].

Where its purpose is to provide (spatial) support for numerical simulation, the quality of the mesh should be measured only by the quality, or numerical accuracy, of the solution; accordingly, the most obvious condition for *good* mesh quality is that the (spatial) domain, where the problem is formulated, should be represented as precisely as possible [220]. Otherwise, it is irrelevant how *nice* the mesh looks. The numerical accuracy of the solution field, for a given number of elements or nodes, is a rigorous metric against which the quality of the mesh can be assessed. However, in selecting the mesh resolution, one must also be cognisant of the requirement for the solution field to be free from any spurious numerical artefacts. Indeed, the resolution of the mesh is intimately related to the numerical stability of the solution; e.g. the Courant-Friedrichs-Lewy criteria determines a constraint on the mesh resolution for a given time-step size for certain, usually hyperbolic, PIDE. A similar criteria exists upon the mesh resolution and the angular resolution for certain discrete ordinates schemes, such as the DD-S_N method [75].

Typically, the geometries of most engineering problems are built with some CAD software. Most often, this CAD description must first be made *watertight* before some ancillary program can generate the finite element mesh; it will then need to be revisited by the mesh generator for every subsequent refinement of the mesh that incurs any change to its representation of the underlying geometry [254]. Routinely, however, this may not be possible at analysis run-time; and may suggest why AMR strategies tend to be of academic interest only [255–258]. It would therefore be pragmatic to focus on a single geometric description that is common to both the design and the analysis. This is the approach envisaged by IGA. In summary of the work by Hughes et al, IGA reverses the classical *flow* of the isoparametric concept; viz. the same basis used to describe the *known* geometry is then used to approximate the *unknown* distribution of some desired field; q.v. Equation 2.60 [233].

There exist several different technologies that can be employed by CAD software for the mathematical description of geometric forms; e.g. those based upon A-patches, subdivision surfaces and other related spline-based descriptions [259, 260]. The industry standard, however, is to use Non-Uniform Rational B-Splines (NURBS) surfaces in the boundary representation of 3D objects; cf. popular commercial software packages such as Rhinoceros and AutoCAD [261–263]. The premise of IGA is such that the NURBS-based CAD description should be passed directly to the analysis program to serve as the initial coarsest mesh; without the need for some intermediate mesh generator; viz. to be read in standard file formats such as the standard for the exchange of product data (STEP) or the initial graphics

exchange specification (IGES) [264, 265]. Since the coarsest mesh encapsulates the exact geometry, any subsequent refinement thereof can be performed simply by enriching the parametric space; for the reason that the parametrisation of the geometrical mapping is invariant under refinement, thus, again, obviating the need to (re-)communicate with any ancillary mesh generator [233].

The NURBS themselves are constructed from B-Splines, piecewise linear combinations of Bernstein Polynomials, and can be expressed by convenient recursion formulae [266]. The approximation space has been more than somewhat improved since NURBS can exactly represent a wide range of complex geometric forms, in particular, those that are common in nuclear reactor physics and shielding problems; e.g. the annular cladding region of fuel pins, fuel pellets and critical spheres; q.v. Section 3.3. In this sense, IGA is an inherently higher-order discretisation technique by the fact that conic sections in 2D and quadric surfaces in 3D can only be exactly represented by at least quadratic and quartic rational polynomials, respectively. This is in contrast to the standard applications of FEM to neutron transport, which typically rely upon unstructured triangular meshes and linear shape functions; q.v. Section 2.7.3.1.

An initial investigation into NURBS-based IGA saw it successfully applied to continuous Bubnov-Galerkin (CBG) discretisations for solid mechanics problems and to streamline-upwind Petrov-Galerkin (SUPG) discretisations for fluid dynamics problems [233]. It has since been employed in many other fields of research; e.g. structural vibrations, acoustics, electromagnetics, fluid-structure interactions and thermal-hydraulics [267–275]. Of particular interest, however, are the applications of IGA to the field of neutronics, where it was first applied to the CBG-IGA spatial discretisation of the *one-speed* NDE by Hall [221]. This was extended to the multi-group formalism by Welch et al who demonstrated the superiority of C^0-continuous NURBS-based IGA over mass-preserved polygonal standard FEM and non-mass-preserved higher-order FEM in 2D; and better yet, allowing for higher-order continuity in the basis saw improved accuracy per DoF [253]. CBG-IGA spatial discretisations have also been applied to second-order forms of the NTE in 2D; cf. the multi-group EP-P_N equations, the multi-group SAAF-S_N equations and the multi-group WLS-S_N equations [113, 119, 276]. Wilson et al applied a symmetric interior penalty (SIP) scheme to a discontinuous Galerkin (DG) discretisation of the multi-group NDE in 2D [277]. Owens et al developed sweep-based solvers for DG-IGA discretisations of the multi-group FONT-S_N equations in 2D; and separately, *eidem* developed heuristic and goal-based AMR algorithms [231, 247, 278].

The rest of the chapter is organised as follows: The mathematical and theoretical concepts of IGA are presented, briefly; as is the particular process for generating NURBS-based descriptions of the geometries to be used by the numerical code developed for this work. Within the framework of IGA, analogues of the standard h-refinement and p-refinement mechanisms from classical FEM are

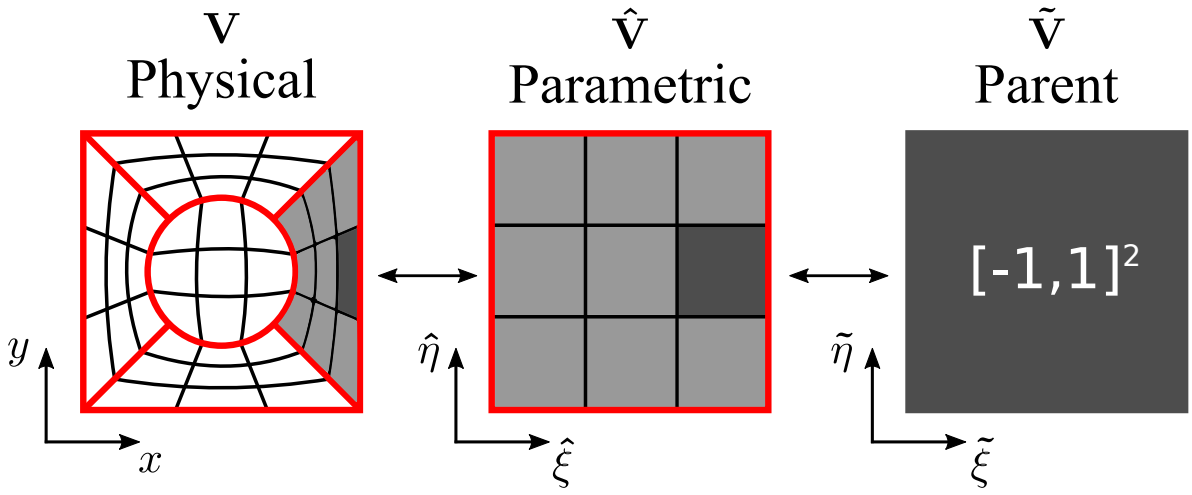


Figure 3.1: The *domain*, as represented in the parametric space, $\hat{V}$, formed from the tensor-product of *two* 1D spaces, as denoted by the parameters ξ and η; and its codomain, as represented, for the example of a 2D pin-cell, by a bi-quadratic NURBS patch (macro-element) in the real physical space, V, as denoted by the parameters x and y. Through KI-refinement, the parametric space has been further decomposed into more non-vanishing knot-spans (micro-elements), over which numerical integration is performed by an affine mapping from the parent domain, $\tilde{V}$, as denoted by the parameters $\tilde{\xi}$ and $\tilde{\eta}$. The patch boundaries are highlighted in red; whereas the non-vanishing knot-spans within a patch are separated by black lines. The image of given non-vanishing knot-span, shaded in dark grey, can be represented in all *three* spaces. (*Vide* the web-based version for reference to colour.)

introduced as part of a global uniform refinement strategy. The ability to interpolate conservatively a scalar field between *volume* meshes is then developed in the guise of energy-dependent meshes, which will find further use when devising interpolation-based error measures to drive *two* different AMR strategies in the following section. This same general interpolation procedure will then be extended to interpolate conservatively a scalar field between *surface* meshes when enforcing continuity in a CBG-IGA spatial discretisation over irregular meshes via constraint equations; qq.v. Chapter 4 and Chapter 5; and yet again, finally, when enforcing continuity in a sweep-based DG-IGA spatial discretisation as a Galerkin projection of upwinded fluxes; q.v. Chapter 6.

The AMR strategies are general in the sense that they can be employed for each of the different discretisations of each of the different forms of the NTE considered in this work. One strategy aims to minimise the global discretisation error in the solution, as measured in an appropriate norm, dubbed the energy-driven (NRG-)AMR strategy; whilst the other aims to minimise the error as measured in some specific goal, dubbed the goal-based (DWR-)AMR strategy.

3.2 The Rudiments

To reiterate, the geometries of most standard engineering problems are typically built with CAD software in a basis of NURBS, or some other related spline-based technology. The simplest description of the geometry is often based upon the decomposition of the spatial domain into its constituent

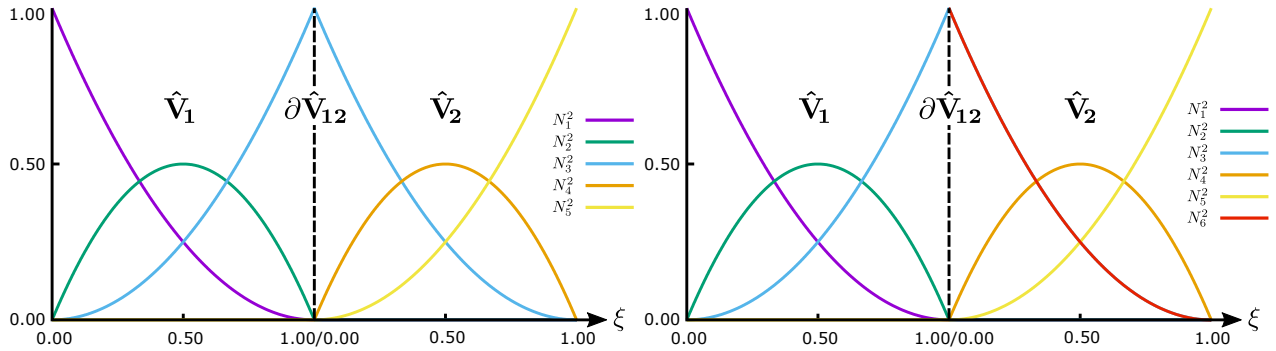


Figure 3.2: The left and right-hand sides are the result of inserting knots at ξ=0.5 with multiplicity $\dot{m} = p, p{+}1$, respectively, into a basis of quadratic (p=2) B-splines with no initial internal refinement structure; i.e. Ξ_ξ=$\{0,0,0,1,1,1\}$. In either case, continuity in the basis has been reduced enough to form a new patch boundary, $\partial\hat{V}_{12}$, and the parametric space has been renormalised to $[0,1]$ in each of the newly formed patches, $\hat{V}_1$ and $\hat{V}_2$. C^0 continuity across $\partial\hat{V}_{12}$ is naturally enforced in a strong sense for $\dot{m}$=p; however, the basis becomes discontinuous, or C^{-1} continuous across $\partial\hat{V}_{12}$ for $\dot{m}$=$p{+}1$, where it must be explicitly enforced in either a strong sense or a weak one. (*Vide* the web-based version for reference to colour.)

materials. Regardless, one assumes that the geometry is represented as a tessellation of NURBS patches, the union of which forms the coarsest physical mesh, $\mathcal{T}(V)$; where V denotes the domain consisting of all possible sets of values that the independent spatial variables can assume; i.e. $\vec{r} = (x, y, z) \in V$. Since the computational domain exactly represents the physical domain, the one may be interchangeably referred to as the other; i.e. $V = \bigcup_{e=1}^{E} V_e$ for $E \geq 1$.

The geometrical mapping is the parametrisation of $\vec{r}$: $\hat{V} \rightarrow V$ that takes a point in the parametric domain, $\vec{r}^{\,-1} \in \hat{V} \subset \mathbb{R}^{\hat{d}}$, where the NURBS are defined, and maps it to a point in Euclidean space, $\vec{r} \in V \subset \mathbb{R}^d$, for $\hat{d}, d = 1, 2$ or 3. A NURBS function defined over the real physical space, $R(\vec{r})$, is the composition of that function, as defined in the parametric domain, $R(\vec{r}^{\,-1})$, with the inverse of the geometrical mapping; i.e. $R(\vec{r}) = \{R(\vec{r}^{\,-1}) \circ \vec{r}^{\,-1} : \forall\, \vec{r}^{\,-1} \in \hat{V}\}$. The parametric domain is local to a given patch; and consequently, as per the bijective mapping, $\vec{r}$, a NURBS function has finite local support on a single given patch; i.e.:

$$\text{supp}\,(R) = \left\{\vec{r} \in V_e : R(\vec{r}) \neq 0, \forall\, \vec{r}^{\,-1} \in \hat{V}_e\right\}, \quad \forall V_e \in \mathcal{T}(V). \tag{3.1}$$

3.2.1 B-Splines

Apropos the 1D parametric domain, denoted by some parameter $0 \leq \xi \leq 1$, the description begins with a *knot-vector*, Ξ_ξ, a set of non-decreasing coordinates, $\Xi_\xi = \{\xi_i\}_{i=1}^{N+p+1}$; where $\xi_i \in \mathbb{R}$ denote the *knots* of index $i = \{1, 2, \ldots, N + p + 1\}$; p denotes the polynomial degree of the basis; and N the total number of functions that span the parametric domain. A knot-vector is *open* if its first and last knot-values are repeated $p{+}1$ times; and is *uniform* if the knots-values are equally spaces [266].

Multi-dimensional, $d > 1$, NURBS are constructed in tensor-product fashion from univariate B-splines,

which themselves are a piecewise linear combinations of Bernstein Polynomials that are ofttimes simply expressed via the Cox-de Boor recursion formula [279, 280]. In defining a basis of B-splines, N_i^p, one begins with piecewise constants, i.e. for $p = 0$ and for $i = 1, \ldots, N$:

$$N_i^0(\xi) = \begin{cases} 1, & \text{if } \xi_i \leq \xi < \xi_{i+1}; \\ 0, & \text{otherwise;} \end{cases} \tag{3.2}$$

from which the higher-order B-splines are evaluated; i.e. for $p > 0$ and for $i = 1, \ldots, N$:

$$N_i^p(\xi) = \frac{\xi - \xi_i}{\xi_{i+p} - \xi_i} N_i^{p\text{-}1}(\xi) + \frac{\xi_{i+p+1} - \xi}{\xi_{i+p+1} - \xi_{i+1}} N_{i+1}^{p\text{-}1}(\xi). \tag{3.3}$$

Should the denominator become vanishing, then the whole fraction is made to be vanishing; i.e. for $\xi_{i+p} = \xi_i$ and for $\xi_{i+p+1} = \xi_{i+1}$. The B-splines have a number important properties; e.g. they are point-wise non-negative over the domain; i.e. $N_i^p(\xi) \geq 0$, $\forall \xi \in [0, 1]$; and the basis constitutes a partition of unity and is therefore capable of representing any constant value; i.e.

$$\sum_{i=1}^{N} N_i^p(\xi) = 1, \quad \forall \xi \in [0, 1]. \tag{3.4}$$

Knots are points of reduced continuity that partition the domain into *knot-spans*; where the ith knot-span is the half-closed interval $[\xi_i, \xi_{i+1})$. The insertion of new knots splits existing knot-spans, whilst the repetition of old ones further decreases continuity in the basis. For $\dot{m}$-number of knots that take on the same value, that knot-value has *multiplicity* of $\dot{m}$ knots; and the continuity across this knot-value is $C^{p\text{-}\dot{m}}$. The basis remains C^∞ continuous, or infinitely differentiable, within the non-vanishing knot-spans; i.e. for $\xi_i \neq \xi_{i+1}$. Increasing the multiplicity of a knot to $\dot{m} = p$ renders the basis interpolatory at that knot-value, effectively forming a new patch boundary, across which C^0 continuity is naturally enforced in a strong sense; v. Figure 3.2. For $\dot{m} = p{+}1$, again, a new patch boundary is formed; however, the basis is now discontinuous, or $C^{\text{-}1}$ continuous, and continuity can be enforced in either a strong sense or a weak one, according to the discretisation scheme; q.v. Chapter 4.

The support of B-splines, $\{N_i^p\}_{i=1}^N$, is compact and extends over $p{+}1$ knot-spans; i.e. for $i = 1, \ldots, N$:

$$\text{supp}\left(N_i^p\right) = \left\{\xi \in \hat{V} \subset \mathbb{R}^1 \colon N_i^p(\xi) \neq 0, \forall\, \xi_i \leq \xi < \xi_{i+p+1}\right\}. \tag{3.5}$$

This support also extends over those vanishing knot-spans; and $p{+}1$ functions always have support within any given knot-span. This property of *locality* will be important for constructing geometries; q.v. Section 3.3. The NURBS inherit all of these properties and many others from the B-splines; e.g. the convex hull property and the variation diminishing property [266]. Owing to the latter, the B-splines and the NURBS by extension are less susceptible to Gibbs' phenomenon and may therefore allow for an increasingly better representation of sharp layers than their Lagrange interpolatory polynomial counterparts for a basis of increasingly higher-order polynomials; q.v. Section 2.7.3 [233].

By analogy to CBG-FEM discretisations, the patches can be thought of as macro-elements, since the basis is C^0 continuous across the patch interfaces. The non-vanishing knot-spans can be thought of as the (micro-)elements, since they are the smallest physical entity dealt with in IGA, over which numerical integration is performed by an affine mapping from the parent domain, $\tilde{V} = [\text{-}1, 1]^{\hat{d}}$; q.v. Section 3.2.3. Unlike FEM, however, the basis in IGA can exhibit higher-order continuity between elements; viz. up to $C^{p\text{-}1}$ continuity across knots. For elliptic PIDE characterised by smooth solutions this has been shown to yield greater accuracy per DoF [253, 270].

3.2.2 NURBS

Although useful, B-splines are, nevertheless, polynomials and thus, could only exactly represent those geometries described as such, which exclude conic sections, such as circles and ellipses, and quadric surfaces, such as cylinders, spheres and ellipsoids. Equation 3.6a and Equation 3.6b, respectively, express the tensor-product fashion in which bi-variate NURBS, $R^{p,q}_{i,j}(\xi,\eta)$, and tri-variate NURBS, $R^{p,q,r}_{i,j,k}(\xi,\eta,\zeta)$, are formed from sets of uni-variate B-splines, $\{N^p_i(\xi)\}^N_{i=1}$, $\{M^q_j(\eta)\}^M_{j=1}$ and $\{L^r_k(\zeta)\}^L_{k=1}$; each defined by their respective polynomial degrees, p, q and r; and their respective knot-vectors, $\Xi_\xi = \{\xi_i\}^{N+p+1}_{i=1}$, $\Xi_\eta = \{\eta_j\}^{M+q+1}_{j=1}$ and $\Xi_\zeta = \{\zeta_k\}^{L+r+1}_{k=1}$.

Each of the NURBS corresponds to a vector-valued coefficient of the same index called a control-point, the set of which serves as the *scaffolding* for the problem geometry; i.e. the control-net, $\mathcal{B} = \{\vec{B}_{i,j}\}^{N,M}_{i,j=1} \subset \mathbb{R}^2$, and the control-lattice, $\mathcal{B} = \{\vec{B}_{i,j,k}\}^{N,M,L}_{i,j,k=1} \subset \mathbb{R}^3$, in 2D and 3D, respectively. A NURBS object, $V \in \mathbb{R}^d$, is the piecewise projective transformation of a B-spline-object in $\mathbb{R}^{d+1}$; the latter is defined by a set of projective control-points, to each of which is attributed a weight that belongs to a set that forms the $(d{+}1)^{th}$-component; i.e. $\mathcal{W} = \{W_{i,j}\}^{N,M}_{i,j=1} \subset \mathbb{R}$, and $\mathcal{W} = \{W_{i,j,k}\}^{N,M,L}_{i,j,k=1} \subset \mathbb{R}$. The control-points are obtained by dividing each of the projective control-points by its respective weight; such that the weights convey a sense of *affinity* for the control-point in the real physical space. The same transformation is applied to every point in the projective space and thus, the NURBS are rendered rational by the division of the polynomial weighting function, the denominator of Equation 3.6a and Equation 3.6b; i.e. for $i, j = 1, \ldots, N, M$ and for $i, j, k = 1, \ldots, N, M, L$, respectively:

$$R^{p,q}_{i,j}(\xi,\eta) = \frac{N^p_i(\xi) M^q_j(\eta) W_{i,j}}{\sum\limits_{i\prime=1}^{N}\sum\limits_{j\prime=1}^{M} N^p_{i\prime}(\xi) M^q_{j\prime}(\eta) W_{i\prime,j\prime}}; \tag{3.6a}$$

$$R^{p,q,r}_{i,j,k}(\xi,\eta,\zeta) = \frac{N^p_i(\xi) M^q_j(\eta) L^r_k(\zeta) W_{i,j,k}}{\sum\limits_{i\prime=1}^{N}\sum\limits_{j\prime=1}^{M}\sum\limits_{k\prime=1}^{L} N^p_{i\prime}(\xi) M^q_{j\prime}(\eta) L^r_{k\prime}(\zeta) W_{i\prime,j\prime,k\prime}}. \tag{3.6b}$$

The NURBS are point-wise non-negative for non-negative weights, which, in most instances, will be normalised such that $0 < W \leq 1$. One notes that it is not the absolute value of the weight that counts

but rather its value relative to that of the other control-points. The parametrisation, $\vec{r}\colon \hat{V} \to V$, is now given as a linear combination of NURBS and control-points, such as to define a surface, S, or a volume, V, over the domain, $\hat{V}$, in the physical space, respectively:

$$S(\xi,\eta) = \sum_{i=1}^{N}\sum_{j=1}^{M} \vec{B}_{i,j} R_{i,j}^{p,q}(\xi,\eta), \qquad \forall\, \xi,\eta \in \hat{V}; \tag{3.7a}$$

$$V(\xi,\eta,\zeta) = \sum_{i=1}^{N}\sum_{j=1}^{M}\sum_{k=1}^{L} \vec{B}_{i,j,k} R_{i,j,k}^{p,q,r}(\xi,\eta,\zeta), \qquad \forall\, \xi,\eta,\zeta \in \hat{V}. \tag{3.7b}$$

By application of the quotient rule to the above, the first derivative with respect to the parametric space can be expressed simply in terms of the partial derivatives of the B-splines; whereas the gradient over the domain in the real physical space is expressed as:

$$\underline{\underline{\mathbf{J}}}_V \cdot \vec{\nabla} R_i^{p,q}(\xi,\eta) = \hat{\nabla} R_i^{p,q}(\xi,\eta), \qquad \text{for } i = 1,\ldots,N\times M; \tag{3.8a}$$

$$\underline{\underline{\mathbf{J}}}_V \cdot \vec{\nabla} R_i^{p,q,r}(\xi,\eta,\zeta) = \hat{\nabla} R_i^{p,q,r}(\xi,\eta,\zeta), \qquad \text{for } i = 1,\ldots,N\times M\times L; \tag{3.8b}$$

where the convention for the nabla operator is retained, with the unit-vectors defined parallel to each of the axes in a rectangular Cartesian coordinate system; i.e. for $\vec{e}_\xi = \vec{e}_x = [1,0,0]^T$, $\vec{e}_\eta = \vec{e}_y = [0,1,0]^T$ and $\vec{e}_\zeta = \vec{e}_z = [0,0,1]^T$:

$$\hat{\nabla} = \vec{e}_\xi \frac{\partial}{\partial \xi} + \vec{e}_\eta \frac{\partial}{\partial \eta} + \vec{e}_\zeta \frac{\partial}{\partial \zeta}; \tag{3.9a}$$

$$\vec{\nabla} = \vec{e}_x \frac{\partial}{\partial x} + \vec{e}_y \frac{\partial}{\partial y} + \vec{e}_z \frac{\partial}{\partial z}; \tag{3.9b}$$

and the volumetric Jacobian $(d\times d)$-matrix, is defined below:

$$\underline{\underline{\mathbf{J}}}_V \coloneqq \left(\begin{array}{ccc:c} \partial x/\partial\xi & \partial y/\partial\xi & \partial z/\partial\xi & \\ \partial x/\partial\eta & \partial y/\partial\eta & \partial z/\partial\eta & \\ \hdashline \partial x/\partial\zeta & \partial y/\partial\zeta & \partial z/\partial\zeta & \end{array}\right). \tag{3.10}$$

Stability of the discretisation necessitates the mapping to be invertible and non-singular almost everywhere, such that the NURBS and their derivatives are square-integrable; viz. the gradient of the mapping should be well defined and $\det \underline{\underline{\mathbf{J}}}_V > 0$ almost everywhere; q.v. Section 3.3.1.

For the example of a 2D pin-cell constructed from *five* bi-quadratic NURBS patches, the schematic in Figure 3.1 presents the domain of a single NURBS patch as that represented in the *three* classical spaces of IGA; viz. the real physical, the parametric and the parent spaces. Since the coarsest physical mesh is ideally read in as some CAD description of the geometry of the problem, the only contribution to the discretisation error is an inability in the finite basis of NURBS to represent the true solution and not an inability to represent the underlying geometry. This is in contrast to standard FEM, where the latter can becoming dominating, especially when modelling for rational geometric forms such as conic sections; q.v. Section 2.7.3 [221, 253]. Moreover, as per Figure 3.1, this exact representation of the underlying geometry is maintained as the mesh is refined; viz. the geometrical mapping and

the Jacobian are invariant under refinement; however, this should not be confused for being piecewise constant [266]. All subsequent refinement of the coarsest mesh can therefore be performed by simply enriching the parametric space at run-time, without the need to revisit some ancillary mesh generator.

3.2.3 Gaussian Quadrature

In order to formulate the weighted-residual statement, as per Equation 2.61, one needs to evaluate integral quantities over the spatial domain, for which it is assumed that analytical expressions are not tractable, especially over general non-Cartesian d-dimensional geometries. As per Equation 2.63, the integral quantities can be constructed as the sum of contributions over individual elements, or rather the knot-spans, since they delimit points of reduced continuity in the basis.

Owing to the tensor-product structure of the parametric space the integration with respect to each direction, ξ, η and ζ, can proceed by means of standard (Legendre-)Gaussian quadrature rules, which are well tabulated for lower-order quadrature sets, otherwise, there exist robust and efficient algorithms; i.e. for some function, f, expanded in a basis of NURBS, as per Equation 2.60 [281]:

$$\int_{V_e} f(\vec{r})\, dV = \int_{\hat{V}_e} f(\hat{r}) \times |\frac{dV}{d\hat{V}}|\, d\hat{V}, \tag{3.11a}$$

$$= \int_{\tilde{V}} f(\tilde{r}) \times |\frac{dV}{d\hat{V}}| \times |\frac{d\hat{V}}{d\tilde{V}}|\, d\tilde{V}, \tag{3.11b}$$

$$\approx \sum_{q=1}^{Q} w_q f(\tilde{r}_q) \times \det \underline{\underline{\mathbf{J}}}_V|_q \times \det \underline{\underline{\mathbf{J}}}_{\hat{V}}|_q, \tag{3.11c}$$

$$\approx \sum_{q_i=1}^{Q_\xi} \sum_{q_j=1}^{Q_\eta} \sum_{q_k=1}^{Q_\zeta} w_{q_i} w_{q_j} w_{q_k} f(\tilde{\xi}_{q_i}, \tilde{\eta}_{q_j}, \tilde{\zeta}_{q_k}) \times \det \underline{\underline{\mathbf{J}}}_V|_{q_i,q_j,q_k} \times \det \underline{\underline{\mathbf{J}}}_{\hat{V}}|_{q_i,q_j,q_k}. \tag{3.11d}$$

A change of variable facilitates the numerical integration above, such that the NURBS basis is sampled locally at Q-number of QP, $\hat{r}_q = \hat{r}_q(\tilde{\xi}_{q_i}, \tilde{\eta}_{q_j}, \tilde{\zeta}_{q_k})$, over a given knot-span, $\hat{V}_e = [\xi_i, \xi_{i+1}) \times [\eta_j, \eta_{j+1}) \times [\zeta_k, \zeta_{k+1})$, mapped from a *canonical* parent, or reference, space, $\tilde{V} = [\text{-}1, 1]^d$; v. Figure 3.1. The sets of weights are normalised such that $\sum_{q_i} w_{q_i} = 2$; and the Jacobian determinant of the mapping from the parent to the parametric spaces, $\det \underline{\underline{\mathbf{J}}}_{\hat{V}}$, must be calculated at each QP, as well as that from the parametric to the real physical spaces, $\det \underline{\underline{\mathbf{J}}}_V$ [233]. Since both the parametric and the parent spaces are topologically Cartesian in structure, the gradient of the mapping yields a tridiagonal matrix; i.e. $\underline{\underline{\mathbf{J}}}_{\hat{V}} = \text{diag}(d\xi/d\tilde{\xi}, d\eta/d\tilde{\eta}, d\zeta/d\tilde{\zeta})$. In which case, the determinant is simply the product of the entries on the diagonal. A simple relation can be derived to express the parametric coordinates of the QP for the mapping $\hat{r}\colon \tilde{V} \to \hat{V}$, from which the gradient of the mapping can be determined; e.g. for $\xi \in [\xi_i, \xi_{i+1})$ mapped from $\tilde{\xi} \in [\text{-}1, 1]$:

$$\xi = \xi_i + (\tilde{\xi} + 1)\frac{\xi_{i+1} - \xi_i}{2} \implies \frac{d\xi}{d\tilde{\xi}} = \frac{\xi_{i+1} - \xi_i}{2}; \tag{3.12}$$

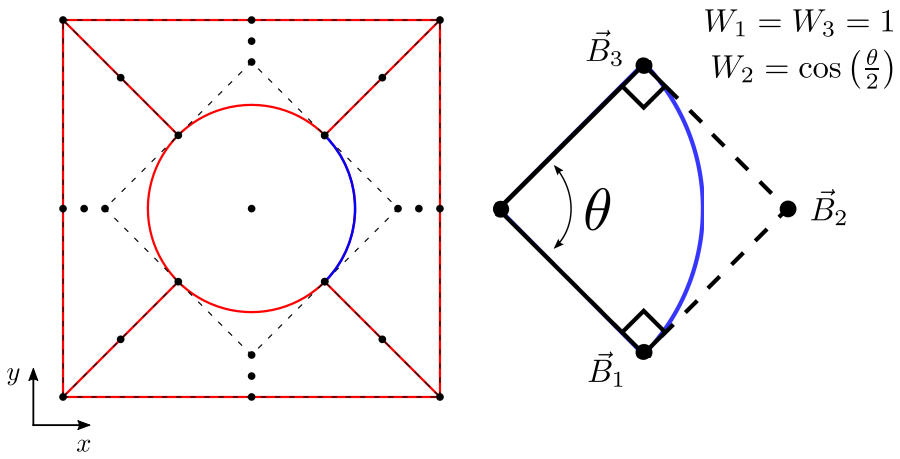


Figure 3.3: An illustration of a multi-patch configuration for the example of a 2D unit pin-cell with a fuel pin of radius 0.25 cm; the patch boundaries are highlighted in red; and the control-points indicated by black circles and the control-net by black dashed lines. Owing to the use of open knot-vectors, the basis is interpolatory at the patch boundaries, therefore any *surfaces* can themselves be described as NURBS objects of dimension one-less than the *volume*; e.g. the eastern boundary of the central pin, or rather its image as mapped from the parametric space for ξ=1, highlighted in blue, is a conic section represented by a 1D quadratic NURBS edge. (*Vide* the web-based version for reference to colour.)

with analogous expressions in the other parametric directions, η and ζ.

Gaussian quadrature rules will not exactly integrate rational polynomials, yet by *rule of thumb*, the NURBS will be integrated using the lowest-order rule that would exactly integrate B-splines of the same polynomial degree; viz. Q-number of Legendre-Gauss QP exactly integrate $2Q$-1-polynomials; although this may not always be sufficient for integration over partial boundaries. Despite an inexact numerical integration, the non-polynomial contributions to the integrand, the polynomial weighting function in the denominator of Equations 3.7 and the Jacobian, Equation 3.10, are invariant under refinement of the parametric space; therefore, their approximation by non-rational polynomials will improve as the error due to numerical integration becomes negligible [282].

3.3 Geometric Design and Mesh Generation

Before constructing any coarsest mesh upon which the analysis is to be performed, one must first identify the computational domain and any approximations thereof that may simplify the model; e.g. by identifying simple planes of symmetry to reduce the number of dimensions in which the problem is posed; viz. in consideration of the computation time as well as the model creation itself, a 3D solid analysis will generally require an order of magnitude more effort than a 2D one. Moreover, it may be possible to identify periodic, or cyclic, symmetry in some problems; e.g. an entire reactor core, or an entire assembly, can be modelled as an infinite lattice by a small array of pin-cells and the prescription of reflective BC. In this sense, the general advice is to model the problem in only as much detail as is

```
!
2D_patch_description_start
!
2D_patch_ID:
n
!
knot_vector_xi:
0.0000000000000E+00
0.0000000000000E+00
0.0000000000000E+00
1.0000000000000E+00
1.0000000000000E+00
1.0000000000000E+00
!
knot_vector_eta:
0.0000000000000E+00
0.0000000000000E+00
0.0000000000000E+00
1.0000000000000E+00
1.0000000000000E+00
1.0000000000000E+00
!
control_points:
1.7677669529663E-01  -1.7677669529663E-01   0.0000000000000E+00   1.0000000000000E+00
3.3838834764831E-01  -3.3838834764831E-01   0.0000000000000E+00   1.0000000000000E+00
5.0000000000000E-01  -5.0000000000000E-01   0.0000000000000E+00   1.0000000000000E+00

3.5355339059327E-01   0.0000000000000E+00   0.0000000000000E+00   7.0710678118654E-01
4.2677669529663E-01   0.0000000000000E+00   0.0000000000000E+00   8.5355339059327E-01
5.0000000000000E-01   0.0000000000000E+00   0.0000000000000E+00   1.0000000000000E+00

1.7677669529663E-01   1.7677669529663E-01   0.0000000000000E+00   1.0000000000000E+00
3.3838834764831E-01   3.3838834764831E-01   0.0000000000000E+00   1.0000000000000E+00
5.0000000000000E-01   5.0000000000000E-01   0.0000000000000E+00   1.0000000000000E+00
!
material_name:
XX
!
2D_patch_description_stop
!
```

Figure 3.4: A snippet of the geometry (.geo) input file to be read in by the numerical code for the declaration of the easternmost *volume* NURBS patch of the 2D unit pin-cell illustrated in Figure 3.3, centred about the origin; where n and XX are an appropriate integer value and string, respectively. The polynomial degree of the basis can be deduced from the open knot-vectors; and the coordinates, (x, z, y, w), for each control-point are declared on separate lines, varying fastest in the $x(\xi)$-direction.

necessary, without introducing redundant complexities and any associated errors.

As it is the ultimate goal of IGA, the geometry of a given problem would ideally be created with some CAD software and its mathematical description passed directly to the analysis program; e.g. in neutral IGES file format [233]. Whilst one could envisage several CAD programs with which it would be possible to do so, this would incur additional complexity in developing an analysis program that would be capable of reading in such arbitrary geometries, which is beyond the scope of this work. The focus of CAD lies primarily with the *design* of engineering components, almost independently of any analysis thereof, so, understandably, most CAD descriptions are not readily suitable for analysis, often being described as *dirty*, since they most likely contain *gaps*, both geometrically and parametrically, and thus, would first need to be *cleaned up* [254, 261, 284]. Furthermore, commercial programs such as Rhinoceros rely heavily upon an ability to *trim* NURBS objects, by which the tensor-product structure of the parametric space is lost and it becomes unclear how to best define a basis of NURBS functions for analysis [262]. It is apparent that an ability to *fix up* effectively CAD descriptions is vital to the

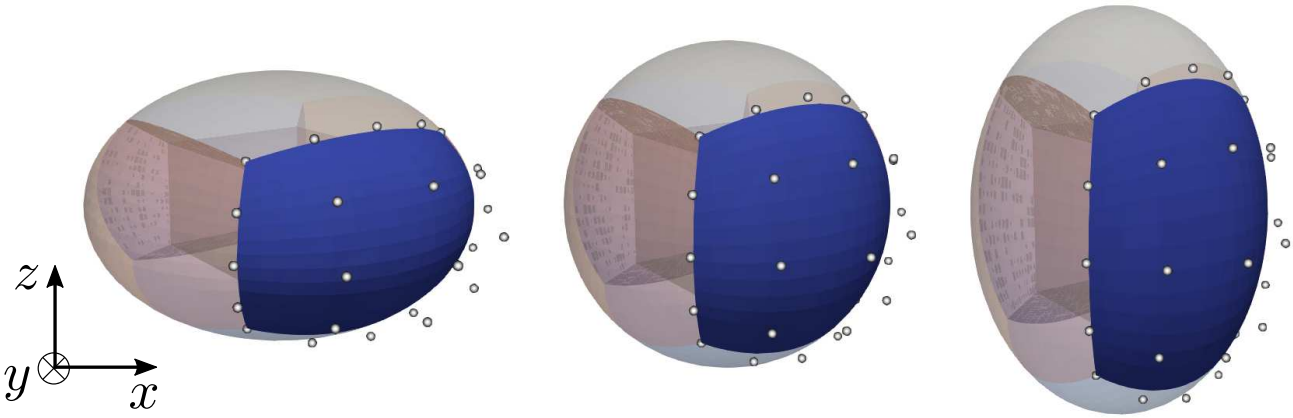


Figure 3.5: Based upon the topology of a cube, as per the tensor-product structure of the parametric space, the surface of a general spheroid can be represented exactly by quartic NURBS patches by the stereographic mapping of *tiles* from the plane [283]. The control-points, denoted by grey orbs, are visible for *one* of *six* faces, highlighted in blue, for each of an oblate spheroid (left) with $e = 0.749$ for $r_x > r_z$; a sphere (centre) with $e = 0.000$ for $r_x = r_z$; and a prolate spheroid (right) with $e = 0.698$ for $r_x < r_z$; where, assuming rotational symmetry about the z-axis, the parameter e denotes the eccentricity of the spheroid; and where $r_x = r_y$ and r_z denote the major and minor radii, respectively, of an oblate spheroid, or the minor and major radii, respectively, of a prolate spheroid. (*Vide* the web-based version for reference to colour.)

overall advancement of IGA technology, such as to ensure that *watertight* analysis-suitable NURBS models are readily tractable; cf. the work of Sederberg et al on converting trimmed NURBS surfaces into untrimmed and unstructured T-splines [285]. Again, this is beyond the scope of this work.

Instead, bespoke bits of software have been developed specifically for the purpose of generating the NURBS-based descriptions of the geometries of those problems considered in this work, which are simple enough always to be defined by patches that are at least star-convex, or star-shaped, and thus, simply connected; and the mesh itself will always forms a convex set; v. Figure 3.3 [222]. Indeed, one remarks that only a relatively few number of different geometric forms will need to be modelled for; viz. conic sections and quadric surfaces. This includes, for example, an ability to create, duplicate, translate and scale 2D pin-cell arrangements; as well as include for an annular cladding region about the fuel pins and vary the radii arbitrarily. A simple format was developed to encapsulate the NURBS description of the problem geometry in a *.geo* file, which forms part of the input deck, along with the *.mat* and *.spx* files that declare all nuclear data and general problem specifications, respectively.

To simplify the connectivity information for the mesh and other properties of the *elements*, all NURBS objects considered are based upon topologically equivalent hypercubes, or rather, a d-dimensional NURBS object has *four* or *six* (d-1)-dimensional boundaries for $d = 2, 3$, respectively; viz. a consistent local enumeration of the basis functions is adopted by varying fastest in ξ, then η and ζ; the ordering of the vertices, edges and faces is determined in a similar manner, running in a positive sense. To establish the initial mesh connectivity information, the coarsest mesh is made to be as uniform and as *regular* as possible in the sense that conditions of compatibility are satisfied; although *irregular* meshes that contain *hanging-nodes* will be permitted in subsequent AMR iterations [222]. In either

case, an interface is always shared by a maximum of *two* elements.

One recalls that unlike the typical interpolation functions of classical FEM, there is no notion of the Kronecker delta property in IGA since the NURBS basis is generally not interpolatory, except for at the boundaries of a patch described by open knot-vectors [286]. This simplifies the manipulation of any boundaries or the extraction of any values at boundaries because they are only influenced by the control-points or the control-variables at the boundary. In this sense, the surfaces of a d-dimensional NURBS volume are themselves (d-1)-dimensional NURBS objects; v. Figure 3.3.

A snippet of the *.geo* file used to generate the geometry for the 2D unit pin-cell with fuel pin of radius 0.25 cm in Figure 3.3 is presented in Figure 3.4. In the declaration of the NURBS patches, a subtle choice of vocabulary is made to distinguish them by their functionality and not the space in which they are defined or used; viz. for a d-dimensional problem, one must declare a number of d-dimensional *volume* NURBS patches and associated material; along with a number of (d-1)-dimensional *surface* NURBS patches and associated BC. One remarks that the weight assigned to any central control-points is somewhat arbitrary since they do not affect the overall geometric form; however, the best accuracy in the analysis is achieved for the best representation of the tensor-product structure of the parametric space and thus, allowing for a smooth mapping, or Jacobian, throughout the domain; e.g. the weight of middle control-point in Figure 3.4 is taken simply as the average of that either side.

In the field of neutronics there is certainly the motivation to provide full 3D reactor core calculations in order to circumvent the problems and inaccuracies associated with the current standard practices of homogenisation and reconstruction. Unfortunately, most CAD programs do not generally support solid 3D NURBS objects since boundary representations are sufficient for design purposes; and thus, the generation of analysis-suitable trivariate parameterisations does face many similar problems as those encountered by typical FEM mesh generators [287]. Fortunately, however, the rendering solid 3D volume NURBS objects is trivial enough for simple geometries and the bespoke bits of software allow for the extrapolation of 2D NURBS objects in the axial direction as well as the full 3D rendering of general spheroids from their surface representations; v. Figure 3.5. Although omitted here for brevity, the general procedure for tiling the surface of a spheroid with quartic NURBS patches is based upon that for the topology of a cube [283].

3.3.1 Degeneracy

Multi-patch configurations are clearly convenient for assigning different material properties to the different regions of a problem. Moreover, the tensor-product structure of the parametric space can complicate the analysis over free-form geometries constructed from too few NURBS patches because

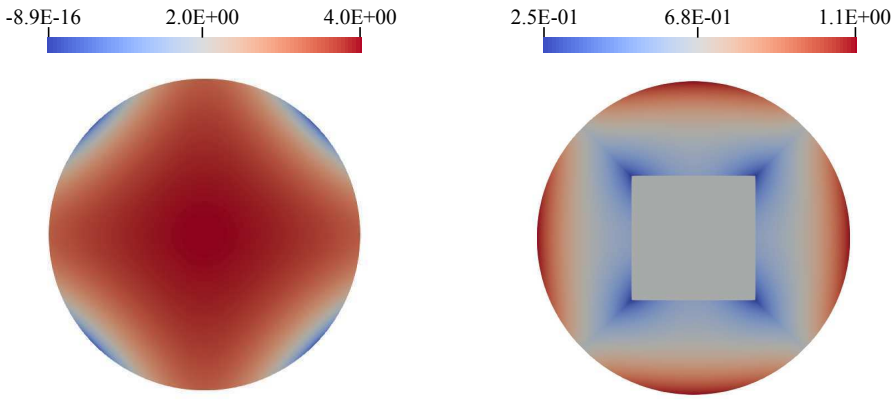


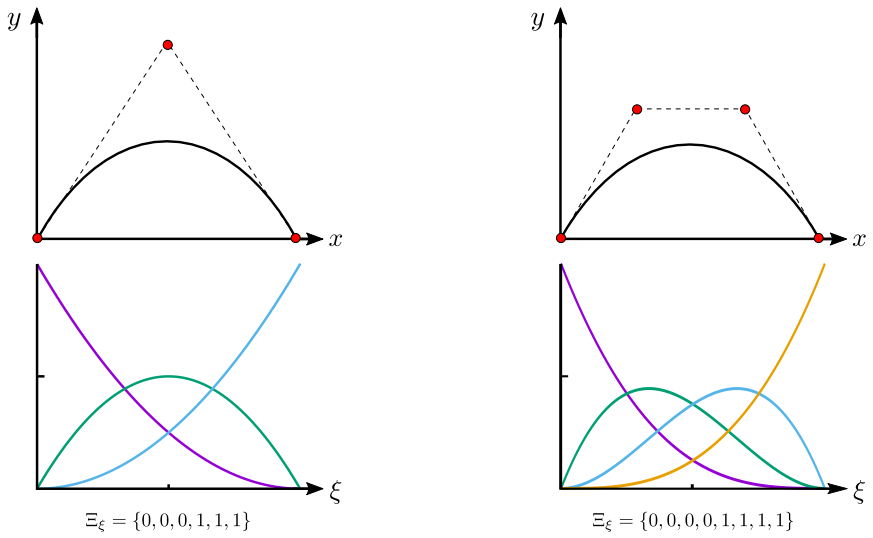
Figure 3.6: The Jacobian determinant for the example of a 2D unit circle as represented by a single degenerate bi-quadratic NURBS patch (left) and *five* non-degenerate bi-quadratic NURBS patches (right). The geometrical mapping is non-invertible at the *corners* of the circle on the left-hand side where the Jacobian determinant becomes vanishing. (*Vide* the web-based version for reference to colour.)

the elements thereof can become very distorted; and in extreme cases, the geometrical mapping can become singular, or at least in its inverse. Any derivatives of the NURBS with respect to the real physical space can be determined from the chain rule, or rather by rearranging Equation 3.8; in either case, one must call upon entries from the inverse of the volumetric Jacobian matrix, $[\underline{\underline{\mathbf{J}}}_V]^{-1}$. Fortunately, there exist explicit expressions for the inversion of 2×2 and 3×3 matrices; e.g. in 2D:

$$[\underline{\underline{\mathbf{J}}}_V]^{-1} = \begin{bmatrix} \frac{d\xi}{dx} & \frac{d\eta}{dx} \\ \frac{d\xi}{dy} & \frac{d\eta}{dy} \end{bmatrix} = \frac{1}{\det \underline{\underline{\mathbf{J}}}_V} \begin{bmatrix} \frac{dy}{d\eta} & -\frac{dy}{d\xi} \\ -\frac{dx}{d\eta} & \frac{dx}{d\xi} \end{bmatrix} . \tag{3.13}$$

Clearly, $[\underline{\underline{\mathbf{J}}}_V]^{-1}$ is not defined when the gradient of the mapping, or rather its determinant, is vanishing; v. Figure 3.6. Yet, this is not always to the detriment of the numerical method; viz. when formulating the integral quantities by Gaussian quadrature rules, the QP may never actually coincide with any points of singularity [227]. Accordingly, however, the coefficient matrix is quadrature dependent and for an increasingly large quadrature set, the evaluation points will lie ever close to the singularity and some entries may become divergent; yet, CBG-IGA discretisations are typically well conditioned and the resulting volumetric mass matrices, or Gram matrices of category 0, and the stiffness matrices remain diagonally dominant for all practical numerical quadrature sets [229].

On the contrary, such degenerate patches would not be suitable should the discretisation require $\vec{\nabla} R$ to be (square-)integrable over the boundary, ∂V_e, which is exactly where the singularities lie in the case of a 2D circle represented by a single bi-quadratic NURBS patch; v. Figure 3.6 and q.v. Section 4.4.2. Since the parametrisation is invariant under refinement, any distortion and such problems associated with degeneracy can only be rectified by a better initial decomposition of the mesh; viz. a circle should be represented by *five* non-degenerate NURBS patches instead of a single degenerate one. Indeed, an isoparametric element performs best when it resembles its unmapped shape and accuracy tends to deteriorate the more it is distorted; viz. upper and lower bounds are generally imposed upon the



(a) Unrefined Quadratic NURBS Curve

(b) Degree Elevation (DE-1)

(c) ($\dot{m} = 1$) Knot Insertion (KI-1)

(d) ($\dot{m} = p$) Knot Insertion (KI-h)

Figure 3.7: The different mechanisms for refinement for the example of a quadratic NURBS curve; whose coarsest description is given in Figure 3.7a; refinement by DE-refinement in Figure 3.7b; and the insertion of a single knot and p-number of knots at $\xi = 0.5$ in Figure 3.7c and Figure 3.7d, respectively. The blue or black segments represent the image of a knot-span in (x, y); and the red circles represent the control-points. (*Vide* the web-based version for reference to colour.)

interior angles or the aspect ratio of triangular and quadrilateral elements, respectively [288]. In a similar fashion, the coarsest mesh of any problem considered in this work must take into consideration the regularity of the NURBS patches, in the sense that they must possess distinct edges and surfaces as per an equivalent hypercube topology in the parametric space; e.g. this translates to representing all circles and spheroids by *five* and *seven* non-degenerate NURBS patches, respectively.

3.4 The Global Uniform Refinement Strategy

As it has already been made apparent, an important feature of IGA is that the parametrisation of the domain in the real physical space is invariant under refinement; viz. the representation of the underlying geometry is always exact for any subsequent refinement of the coarsest mesh description and there is no need to (re-)communicate with any ancillary mesh generator [233].

There are many combinations of mechanisms by which to enrich the solution space, one of which is suggested by Figure 3.1, that is to split existing elements by knot insertion (KI-)refinement. New knots can be inserted with multiplicity of *one* up to multiplicity of p+1, which are referred to as KI-1-refinement and KI-h-refinement, respectively. The latter is so called because it is analogous to the h-refinement of standard FEM since the NURBS basis is reduced to C^0 continuity across newly-formed element boundaries, once continuity has been enforced; this is in contrast to KI-1-refinement where the basis is reduced to $C^{p-\dot{m}}$ continuity across the newly-formed element boundaries. Both of these types of KI-refinement are demonstrated in Figure 3.7 for the example of a quadratic NURBS curve. In both cases, the geometry of the curve is unchanged and the mesh is partitioned in the same place; however, they differ by the number of basis functions and thus, control-points; viz. KI-h-refinement becomes increasingly more expensive to perform for bases of increasingly large polynomial degree.

Another way in which to enrich the basis is by degree elevation (DE-)refinement, which is somewhat analogous to p-refinement of standard FEM, whereby the polynomial degree of the basis is increased. Importantly, however, any DE-refinement must preserve the inherent smoothness of the original basis; therefore the multiplicity of each existing knot-value must be increased by *one* for each degree the basis is raised by. This is demonstrated in Figure 3.7 where the curve is now represented by cubic NURBS and the basis remains interpolatory at the boundaries. Generally, the operations of KI-1-refinement and DE-refinement are not communicative and this can be exploited to maximise smoothness, C^p continuity, in the basis by performing DE-1-refinement followed by KI-1-refinement; this mechanism of refinement is dubbed k-refinement for which there is no known analogue in classical FEM [229].

Owing to the tensor-product structure of the NURBS and by the compatibility conditions imposed on the coarsest description of the geometry, any refinement to a CBG-IGA discretisation scheme must be allowed to propagate throughout the mesh. That is, to define single piecewise continuous function space local to the entire computational domain, it was assumed that the mappings and parameterisations on adjoining patch faces were identical and by extension it was assumed that the control-variables and control-points on adjoining patch faces were in one-to-one correspondence with each other. Such a strong enforcement of continuity in the basis is indistinguishable from the operation of KI-h-refinement since they both result in C^0 continuity across element boundaries. This relationship must be preserved

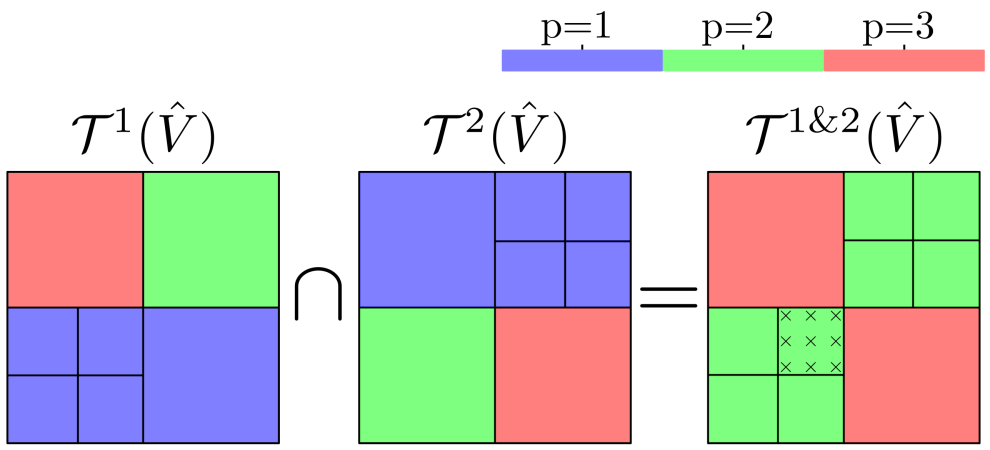


Figure 3.8: An example of the *implied* parametric super-mesh, $\mathcal{T}^{1\&2}$, formed as the intersection of *two* volume meshes, $\mathcal{T}^1$ and $\mathcal{T}^2$, for the conservative interpolation of a scalar field from one to the other. The crosses represent the appropriate number of QP required to integrate that element of the super-mesh, such its boundaries are defined as per $\mathcal{T}^1$, whilst the polynomial degree of the basis is defined as per $\mathcal{T}^2$. (*Vide* the web-based version for reference to colour.)

under refinement, unless continuity can be enforced by other means. One of the attractive features of discontinuous discretisations is that they naturally allow for arbitrary refinement of the mesh because the approximation space is discontinuous by definition and continuity must always be enforced in a weak sense, regardless of the *irregularity* of the mesh.

In consideration of the above, one may devise a global uniform refinement strategy whereby the characteristic mesh spacing is decreased, $h \to 0$; or the polynomial degree of the basis is increased, $p \to \infty$; or both in tandem, in the same consistent and isotropic manner over each NURBS patch that belongs to the physical mesh. Such global refinement strategies are useful for verifying consistency in the numerical code since theoretical *a priori* rates of convergence are well established [67,69]. However, they do not generally allow for an *optimal* allocation of resources.

3.5 Energy-Dependent Meshes

Within nuclear-related industries that employ deterministic neutron transport codes, the current practice is to solve each of the multi-group components of the neutron flux distribution over a single spatial mesh; e.g. PANTHER and ATILLA [236, 289]. However, the neutron cross-section data is known to vary significantly across the range of possible neutron energies and thus, the flux distribution will equally vary significantly across the energy groups. Faster neutrons tend to have longer MFP and typically exhibit smoother flux profiles than slower ones that tend to have shorter MFP and thus, typically exhibit more pronounced variations in their flux profiles, especially about material interfaces and boundaries. *Optimality* could not be achieved over a single energy-independent mesh, $\mathcal{T}(V)$, and the associated finite dimensional function space, $\hat{\mathcal{V}}_{h,p}$, as characterised by the char-

acteristic mesh spacing, h, and the polynomial degree of the basis, p. Instead, the multi-group NTE should be solved over separate energy-dependent meshes, $\mathcal{T}^g(V)$, with respective function spaces, $\hat{\mathcal{V}}^g_{h,p}$, that should be refined separately to maximise the accuracy per total DoF; viz. it is expected that $G \times \dim \hat{\mathcal{V}}_{h,p} > \sum_g \dim \hat{\mathcal{V}}^g_{h,p}$.

The problem has been rendered somewhat more complex by requiring the storage of at least *one* spatial mesh per energy group; yet, the multi-group method does lend itself to the use of energy-dependent meshes because the global linear system of algebraic equations that result from Galerkin-type discretisations is only loosely coupled in energy. In practice the within-group equations are solved separately, iterating on the within-group (self-scattering) sources, whilst the group-to-group sources are fixed using the best available information. Sweeping through the energy groups, the discrete within-group systems are not required to be of the same size; however, it must be possible to project the components of the multi-group fluxes between the different meshes in order to evaluate the group-to-group scattering and fission sources.

When Wang and Ragusa developed adaptive *hp*-refinement strategies for a FEM discretisation of the multi-group NDE over energy-dependent meshes in multiple dimensions, they employed an adaptive integration technique to reduce the error in the coupled group-to-group source terms [290]. For FEM discretisations of the FONT-P_N equations, Goffin et al achieved the transfer of information between energy-dependent meshes as a conservative Galerkin projection between volume meshes [291]. As proposed by Farrell and Madison, the information passed between *two* different volume meshes must be interpolated onto a *super-mesh*, formed as the intersection of the *two* meshes [292]. Constructing the energy-dependent meshes independently of one another saw a significant increase in the performance times involved when compared to similar analyses performed over a single energy-independent one. In the application to the NDE in 3D, Ragusa and separately Wang and Ragusa sought to remedy this by limiting their adaptive *h*-refinement strategies to a sequence of adapted meshes belonging to a hierarchy of nested Cartesian meshes, forming a tree-like structure stemming from a common coarsest mesh shared by all energy groups [293, 294]. Although the numerical integration procedure was exact, the problem geometries were restricted to those that could be exactly modelled, at the coarsest level, by straight-sided finite elements. Applying such a strategy to complex rational geometries would see the construction of an initial polygonal, or polyhedral, mesh that could not benefit from any subsequent improvement to its representation of the underlying geometry; thus, capping the attainable accuracy of the scheme.

One designates the interpolation $\prod^t_{h,p} v^d \colon \mathcal{V}^d(V) \to \mathcal{V}^t(V)$ as the projection of some function, $v^d \in \mathcal{V}^d_{h,p}$, defined on a donor mesh, $\mathcal{T}^d(V)$, onto a target mesh, $\mathcal{T}^t(V)$. The optimum target interpolant,

$v^t \in \mathcal{V}^t_{h,p}$, is approximated such that:

$$\|v^t - v^d\| = \min_{v \in \mathcal{V}^t_{h,p}} \|v - v^d\|.$$

In order to preserve locality and global continuity, the residual, as measured in some norm, is made to be vanishing by solving the appropriate Sobolev minimisation problem, from which conservation follows naturally [295]. The functions are those taken from the same spaces spanned by the NURBS functions used to define each of the donor mesh and the target mesh, denoted by the superscripts d and t, respectively. By letting $N^i = \dim \hat{\mathcal{V}}^i_{h,p}$, the NURBS functions defined over a given mesh are also denoted by a similar superscript; i.e. for $i = d, t$ and for $j = 1, \ldots, N^i$:

$$R^i_j(\vec{r}) = \left(R(\vec{r})|_{\mathcal{T}^i}\right)_j, \qquad \vec{r} \in V.$$

The weighted-residual formulation is chosen to be optimal in the H^1-norm and the residual is minimised if the first variation in the control-variables of the target function, v^t_i, is vanishing for $i = 1, \ldots, N^t$:

$$\begin{aligned}
&\frac{\partial}{\partial v^t_i} \sum_{V_e \in \mathcal{T}^t} \left[\int_{V_e} \left[\sum_{j=1}^{N^t} v^t_j \vec{\nabla} R^t_j(\vec{r}) - \vec{\nabla} \hat{v}^d(\vec{r}) \right]^2 + \left[\sum_{j\prime=1}^{N^t} v^t_{j\prime} R^t_{j\prime}(\vec{r}) - \hat{v}^d(\vec{r}) \right]^2 dV \right], \\
&= \sum_{V_e \in \mathcal{T}^t} \left[\int_{V_e} 2\vec{\nabla} R^t_i(\vec{r}) \sum_{j=1}^{N^t} v^t_j \vec{\nabla} R^t_j(\vec{r}) + 2R^t_i(\vec{r}) \sum_{j\prime=1}^{N^t} v^t_{j\prime} R^t_{j\prime}(\vec{r}) \, dV \right] \\
&- \sum_{V_e \in \mathcal{T}^t} \left[\int_{V} 2\vec{\nabla} R^t_i(\vec{r}) \sum_{\hat{j}=1}^{N^d} v^d_{\hat{j}} \vec{\nabla} R^d_{\hat{j}}(\vec{r}) - 2R^t_i(\vec{r}) \sum_{\hat{j}\prime=1}^{N^d} v^d_{\hat{j}\prime} R^d_{\hat{j}\prime}(\vec{r}) \, dV \right] = 0.
\end{aligned} \tag{3.14}$$

This is analogous to an orthogonal Galerkin projection of a discrete field onto the function space spanned by the target basis, where the projection is the function that minimises the interpolation error in the desired norm [292]. Equation 3.14 can be re-written in matrix form:

$$\left[\underline{\underline{S}}^{tt}_V + \underline{\underline{M}}^{tt}_V\right] \cdot \underline{v}^t = \left[\underline{\underline{S}}^{td}_V + \underline{\underline{M}}^{td}_V\right] \cdot \underline{v}^d; \tag{3.15a}$$

$$\implies \underline{\underline{K}}^{tt}_V \cdot \underline{v}^t = \underline{\underline{K}}^{td}_V \cdot \underline{v}^d; \tag{3.15b}$$

where the consistent mass and stiffness matrices, $\underline{\underline{M}}^{tt}_V$ and $\underline{\underline{S}}^{tt}_V$, respectively, are formed by integrating the product of those functions, defined over $\mathcal{T}^t$, that belong to the target function space, $\hat{\mathcal{V}}^t_{h,p}$. Whereas the projective mass and stiffness matrices, $\underline{\underline{M}}^{td}_V$ and $\underline{\underline{S}}^{td}_V$, respectively, require that in the same integrand are evaluated the product of functions defined over different spatial meshes, $\mathcal{T}^t$ and $\mathcal{T}^d(V)$, that belong to different function spaces, $\hat{\mathcal{V}}^t_{h,p}$ and $\hat{\mathcal{V}}^d_{h,p}$, respectively. This is achieved by integrating over an implied *super-mesh* formed by the intersection of the target mesh and the donor mesh, $\mathcal{T}^g \cap \mathcal{T}^{g\prime}$, which decomposes the parametric domain into subdomains appropriate for numerical integration; v. Figure 3.8. Although it has not yet been detailed, an ability to perform local refinement has been assumed. Looping over the elements formed by the intersection of *two* meshes, the bases are sampled

at a number of QP in each parametric direction, Q, determined by:

$$Q = \max_{i=t,d} Q|_{\mathcal{T}^i}, \quad \forall V_{\hat{e}} \in \bigcap_i \mathcal{T}^i(V). \tag{3.16}$$

It was found that employing p+1 QP in each parametric direction over the canonical representation of a given element in the parent domain did provide sufficient accuracy in the integration over volumes for practical purposes; q.v. Section 3.2.3. Of course, this will be verified; qq.v. Section 4.5.1, Section 5.5.1 and Section 6.5.1.

Complexity can be largely circumvented if the energy-dependent meshes are derived from a common coarsest-mesh description such that any subsequent refinement results in a hierarchy of meshes nested therewithin. For the sake of convenience, this mesh is assumed to be regular and conforming. The intersection of two elements from two different meshes will always form either an empty (disjoint) set, a shared interface, $V_{i'}^g \subseteq V_{j'}^{g'}$ or $V_{i'}^g \supseteq V_{j'}^g$. The tensor-product structure of the parametric space of NURBS-based IGA lends itself to such an approach and since the coarsest-mesh description is exactly the original CAD description, each subsequent mesh nested within a hierarchy of meshes would also exactly represent the underlying geometry; thus, avoiding any complications faced by standard FEM in matching different geometric approximations. Moreover, the intersection of *two* meshes is simplified in the parametric space, which is initially defined by a single parent patch common to both meshes. Welch et al computed the group-to-group source terms for a CBG-IGA spatial discretisation of the multi-group EP-S_N equations on energy-dependent meshes and their mesh-to-mesh interpolation was based upon an L_2-minimisation that preserved the expected theoretical rates of convergence of relatively simple test cases with smooth analytical solutions [276]. Although conservative, an interpolation optimal in the L_2-norm may not be sufficient in a wider application; especially for those problems that exhibit strong material heterogeneity over which the source function changes abruptly [295].

Some function may undergo a change of basis transformation by simply pre-multiplying the control-variables, $\underline{v}^d$, by the combined transfer matrix, $\underline{\underline{K}}_V^{t\leftarrow d}$, for which the interpolation of the function from a donor mesh onto a target mesh is optimal with respect to the H^1-norm. Although $\underline{\underline{S}}_V^{tt}$ is singular, $\underline{\underline{K}}_V^{tt}$ is symmetric-positive definite and therefore readily invertible using standard iterative smoothers:

$$\underline{v}^t = [\underline{\underline{K}}_V^{tt}]^{-1} \cdot \underline{\underline{K}}_V^{td} \cdot \underline{v}^d, \tag{3.17a}$$

$$= \underset{(N^t \times N^d)\ (N^d \times 1)}{\underline{\underline{K}}_V^{t\leftarrow d} \cdot \underline{v}^d}. \tag{3.17b}$$

Owing to the local formulation of discontinuous Galerkin discretisations, it will be possible to avoid the assembly of large matrices since they can be formulated locally and $\underline{\underline{K}}_V^{tt}$ can be inverted locally; e.g. using efficient LAPACK routines such as `DGETRF` and `DGETRI` [296]. Otherwise, for continuous discretisations, it will be necessary to (re)use sparse matrix storage structures and efficient smoothers

to perform the interpolation [249].

Group-to-group scattering and fission contributions may not exist for some ordered pairs of energy groups, $(g\prime, g)$, such that it would be unnecessary to perform the interpolation $\prod {}^{g}_{h,p}\hat{\phi}^{g\prime}_{h,p} : \hat{\mathcal{V}}^{g\prime}_{h,p} \to \hat{\mathcal{V}}^{g}_{h,p}$; this is particularly pertinent to radiation shielding problems that exhibit mostly down-scatter for $G \gg 1$. For those ordered pairs in which information is transferred in both directions, additional savings can be made by $\underline{\underline{K}}_V^{g\leftarrow g\prime} = [\underline{\underline{K}}_V^{g\prime\leftarrow g}]^T$; in which case, only $\underline{\underline{K}}_V^{g\prime\rightarrow g}$ must be stored for $g\prime \geq g$ [278].

In the wider context of the current work, the interpolation of a discrete scalar field between *volume* meshes is not unique in its application to energy-dependent meshes; it will also be employed in the next section when devising interpolation-based error measures to drive the adaptive refinement strategies. This same general interpolation procedure will then be extended to interpolate conservatively a scalar field between *surface* meshes when enforcing continuity in a CBG-IGA spatial discretisation over irregular meshes via constraint equations; qq.v. Chapter 4 and Chapter 5; and yet again, finally, when enforcing continuity in a sweep-based DG-IGA spatial discretisation as a Galerkin projection of upwinded fluxes; q.v. Chapter 6.

3.6 The Local Adaptive Refinement Strategies

In the pursuit of solving PIDE to a high degree of accuracy at minimal costs, an AMR strategy is favoured over a global uniform one; viz. the latter becomes impractical for large complex multi-dimensional problems in terms of memory and performance [297]. Whilst the former behoves some way in which to enforce continuity in the solution across any inevitably irregular meshes, the details are reserved for the ensuing chapters since the approach depends upon the nature of the discretisation as well as the governing equations; viz. constraints based upon a Galerkin projection for CBG-IGA discretisations of the multi-group NDE, Section 4.3, and the multi-group SAAF-S_N equations, Section 5.3; IP-DG-IGA discretisations of the multi-group NDE, Section 4.4, and the multi-group SAAF-S_N equations, Section 5.4; and a numerical upwinding scheme for the DG-IGA discretisation of the multi-group FONT-S_N equations, Section 6.4.

Should the behaviour of the solution for a particular problem be known, possibly from previous experience, an appropriately designed initial mesh could be constructed and uniform p-refinement strategies, or adaptive ones, could be employed to yield exponential rates of convergence, with some dependence upon the regularity of the solution [298]. Notwithstanding that such information is not generally known *a priori*; and such a p-refinement strategy that hinges upon an initial (quasi-)optimal mesh would suffer from pollution effects were any non-smooth components of the field not already addressed by an appropriate geometric grading of the mesh; e.g. those due to boundary layers about

material discontinuities and physical boundaries. Accordingly, adaptive h-refinement techniques were further developed to rectify automatically any inefficiencies in the initial mesh. Both uniform and AMR-h strategies are capable only of achieving algebraic rates of convergence that are restricted by the polynomial degree of the basis, p; the convergence of the former is also limited by the regularity of the solution, whereas that of the latter is not [299, 300]. Whilst separately, both AMR-h and AMR-p strategies may perform better than their respective and corresponding global refinement strategies, a combined AMR-hp strategy addresses both the non-smooth and the smooth parts of the solution; and thus, is capable of attaining (quasi-)optimal exponential rates of convergence, regardless of the regularity of the solution [298, 301, 302].

For global error estimates that can be established *a posteriori*, it is possible to minimise the overall discretisation error in an automated succession of computations performed on locally adapted meshes; and the procedure can be terminated once the solution has converged to some user-defined tolerance [303]. Otherwise, the next *optimised* mesh is determined by maximising the rate of decrease in the local error contributions per increase in the number of unknowns by selecting some optimal local refinement strategy for a given element that has been *flagged*. Although simple enough to implement in 1D, complexities arise when constructing non-conforming, or irregular, meshes in multiple dimensions; viz. the optimal refinement of a given element may directly conflict with that of its neighbour. Thus, to facilitate their usage and reduce complexity in the strategy itself, heuristics such as the *minimum rule* and the *1-irregularity rule* were conceived, thereby placing certain restrictions on *where* and *how* the meshes can be refined; whilst this does often allow for a better conditioning of the connected model, it may be to the detriment of mesh *optimality* by the accruement of additional DoF [290, 304].

To ascertain the error in the solution over the current mesh, such AMR strategies will also require some knowledge of the true solution itself; or rather, some reference solution that is more readily available, obtained over some more-refined reference mesh that provides a better approximation to the true solution than the coarser current mesh. The difference between the *two* solutions drives the AMR procedure such that relatively cheap cell-based error measures can be devised in an appropriate norm to discern those elements that should be flagged for further refinement.

Wang and Ragusa developed such an AMR-$h(p)$ technique for the multi-group NDE in 1D, taking their reference mesh to be that of one-level of global uniform hp-refinement of the coarser current mesh [305]. However, the explicit computation over both was deemed to be too expensive and impractical in multiple dimensions; therefore, *eidem* redeveloped their AMR-hp strategy, this time employing interpolation-based error estimates for problems in 1D; and then applied it more generally to multiple dimensions [290, 306]. This approach had previously been adopted by Rachowicz and Demkowicz et al and applied successfully to other elliptic systems; whereby invoking Céa's Lemma, or Galerkin's

best approximation for SPD systems, the *coarser* solution can be approximated by projection of the *more-refined* reference solution onto the coarser current mesh [257, 299, 307]. The discretisation error can be controlled by minimising the interpolation error, thereby circumventing the need to compute *two* different solutions on *two* different meshes.

Alternatively, elements may be flagged by less-strict error measures used to gauge *smoothness* of the solution, typically by computing estimates of higher-order derivatives and measuring their jumps across element boundaries [71, 231, 294, 308]. Indeed, the use of such computationally inexpensive heuristics can provide more accuracy per DoF over global uniform refinement strategies, especially when the important regions correspond to those in which the solution exhibits pronounced spatial variations; although this is not guaranteed.

Upon reflection, it is perhaps worth revisiting the definition of *optimality* and what this means to the user, since it may not, in fact, be at all necessary to resolve finely the problem throughout the entire computational domain. Often in practice, the ultimate goal of the numerical method is to provide, up to some given tolerance in accuracy, some QoI. The latter may be the solution itself, or some functional of the solution in some specific region of the problem; e.g. volumetric detector responses, or the absorption rate in a given material. Rannacher et al presented the original framework for goal-based AMR strategies for the linear NTE for a fixed source, whereby local contributions to the error in a specific goal can be expressed in terms of the local error measures weighted by a measure of the importance of their respective elements to the functional of interest [309]. Lathouwers extended this work to optimise the effective neutron multiplication factor, k_{eff} for eigenvalue problems [255]. Appropriately, elements with large errors may be ignored if they are not deemed to be *important*; conversely, elements with small errors may be flagged if they have a strong influence on the desired goal. The importance is computed by solving the appropriately formulated physical adjoint (dual) of the direct (primal) governing equations; whence stems the preferred appellation dual-weighted residual (DWR-)AMR strategies.

Solin and Demkowicz modified their DWR error measures by invoking Galerkin's orthogonality property such as to express them instead as the product of a local error measure for the primal solution, μ_e, and a similar error measure for the dual solution, $\mu_e^\dagger$; rather than the importance function itself, $\phi^\dagger$ [310]. Wang and Ragusa argued that weighting energy-norm error estimators for the primal problem in such a way will yield greater accuracy in the results because trends are generally more reliable than absolute values [290]. In any case, when compared to a standard NRG-AMR strategy, a DWR-AMR one is expected to deliver better accuracy per DoF by favouring the refinement of those regions that are important to a specific goal.

IGA provides a good framework within which to implement AMR-$h(p)$ strategies because the parametri-

sation, Equation 3.7 and Equation 3.10, does not vary under refinement and thus, the underlying geometry is always exactly represented. Adhering to the IGA paradigm, the current work employs the standard tensor-product NURBS patches, which both simplify the analysis and are readily found in most CAD environments; although it is appreciated that more exotic techniques do exist; cf. the development of other spline-based descriptions such as the truncated (T-)splines and the locally refined (LR)B-splines [311]. Welch briefly investigated a heuristic-based AMR-h strategy for the NDE; and Welch et al investigated a standard AMR-h strategy for the EP-$\mathrm{P_N}$ equations, chasing the global error as measured in the L_2-norm on energy-dependent meshes [71, 276]. Within a sweep-based solver for DG-IGA discretisations of the FONT-$\mathrm{S_N}$ equations, Owens et al investigated a heuristic-based AMR-h strategy; *eidem* then extended their work to a goal-based DWR-AMR-h strategy over energy-dependent meshes [231, 278]. In each instance, the conservative mesh interpolation between energy-dependent meshes was based upon hierarchy of meshes nested within a common coarsest mesh; q.v. Section 3.5.

Although increased smoothness is something to be desired when solving elliptic problems, especially as higher-order continuity, $C^{p\text{-}1}$, between elements is achievable within a NURBS-based IGA framework, it will not be exploited here. At the most, C^0 continuity is considered in order to facilitate a wholly-local hp-refinement strategy; whereby any new elements formed by h-refinement will become, in their own right, new NURBS patches lacking any internal refinement structure. Allowing for reduced continuity in the solution has been shown to favour better numerical dispersion and diffusion properties in the solution, at least for the FONT-$\mathrm{S_N}$ equations [231, 247]. Furthermore, in limiting the basis to C^0 continuity, one can exploit the efficient and well-established preconditioners that have been developed for low-order continuity FEM, which are not generally well suited for bases that exhibit higher-order continuity [312, 313]. The tensor-product structure of the NURBS patches lends itself to the simple cell-based approach adopted by Wang and Ragusa; whereby any local isotropic h-refinement consists simply of bisecting an element in each parametric direction forming 2^d new elements, for $d = 2, 3$; and the AMR procedure can proceed as a sequence of adapted meshes that belong to a hierarchy of nested meshes, with the lowest level coinciding with the initial coarsest mesh [294].

In a newly-developed multi-dimensional multi-group neutron diffusion and transport code, the successive convergence of the within-group sources, using the best available information for the group-to-group sources, constitutes the outer-(source)-iteration loop over all groups; where, for a given group, any spurious sources introduced into the within-group scatter source are required to become vanishing upon exiting the respective (within-group) inner-iteration loop; q.v. Section 2.8. In the presence of neutron up-scatter or fission events, the total source must be recomputed after each outer iteration until it has converged; where, in the case of an eigenvalue problem, this outer-iteration sequence is dubbed the PI procedure and the $\mathrm{k_{eff}}$ must also be updated. The proposed AMR-$h(p)$ strategy is

encapsulated within an iterative mesh refinement procedure that wraps around the outer-most loop in Figure 2.4, such that all of the group-to-group sources must be fully converged, to some prescribed tolerance, before refining the spatial mesh(es) for the next AMR iteration; v. Figure 3.10. Furthermore, all of the multi-group components of the numerical solution would need to be known before evaluating, *a posteriori*, the global error in the QoI, $\epsilon^{(i)}$, as measured in an appropriate norm:

$$\epsilon^{(i)} = \|QoI(\phi) \text{ - } QoI(\hat{\phi}^{(i)}_{h,p})\|. \tag{3.18}$$

The QoI represents a physical quantity of the solution; e.g. this could be, but is not restricted to, the solution itself, averaged and point values, or some functional of the solution. Where $\hat{\phi}^{(i)}_{h,p}$ is the discrete solution obtained from the numerical analysis performed on the mesh(es), $\mathcal{T}^{(i)}(V)$, characterised by some combination of h and p, at the ith iteration of the AMR procedure, as denoted by the parenthesised index (i). The relative global error, $\bar{\epsilon}^{(i)}$, is defined as:

$$\bar{\epsilon}^{(i)} = \frac{\|QoI(\phi) \text{ - } QoI(\hat{\phi}^{(i)}_{h,p})\|}{\|QoI(\phi)\|} = \frac{\epsilon^{(i)}}{\|QoI(\phi)\|}; \tag{3.19}$$

where the AMR iteration is terminated once $\bar{\epsilon}^{(i)} < \tau$, for some prescribed tolerance, τ; or once the maximum number of iterations has been attained.

Mesh optimality is striven for by maximising the decrease in the global error, $\Delta\epsilon^{(i)}$, between successive iterations, per increase in the number of unknowns, $\Delta DoF^{(i)}$; or rather, borrowing the terminology from Ceze et al, a *merit* function, divided by any associated *costs* of refinement, is maximised through the selection of some combination of h^*, p^* [314]:

$$\max \frac{\Delta\epsilon^{(i)}}{\Delta DoF^{(i)}} = \max \frac{\epsilon^{(i)} \text{ - } \epsilon^{(i+1)}}{|DoF^{(i)} \text{ - } DoF^{(i+1)}|}, \tag{3.20a}$$

$$= \max_{h^*,p^*} \frac{\|QoI\big(\phi\big) \text{ - } QoI\big(\hat{\phi}^{(i)}_{h,p}\big)\| - \|QoI\big(\phi\big) \text{ - } QoI\big(\hat{\phi}^{(i+1)}_{h^*,p^*}\big)\|}{|DoF^{(i)}_{h,p} \text{ - } DoF^{(i+1)}_{h^*,p^*}|}, \tag{3.20b}$$

$$= \max_{h^*,p^*} \bar{\Delta}\epsilon^{(i)}_{h^*,p^*}. \tag{3.20c}$$

Employing a similar approach to Jeffers et al, who investigated DWR-AMR-*hp* for cell-vertex and cell-centred DD-S_N methods in 1D, the *merit* is calculated as the difference in the error in solving over the current mesh(es), $\epsilon^{(i)}$, and that in solving over the same mesh(es) after refinement, $\epsilon^{(i+1)}$, as per some choice of h^*, p^* [193]. In the application to higher-dimensions, however, the number of different ways in which a cell can be refined becomes large, it is therefore practical to limit the mechanisms for refinement. The current work restricts this choice to isotropic DE-1-refinement and isotropic KI-*h*-refinement; i.e. $h^*, p^* = \{h/2, p \text{ or } h, p{+}1\}$, which is comparable to a standard AMR-$h(p)$ strategy for classical FEM. As an aside, one notes that with respect to Equation 3.20, the cost of performing KI-*h*-refinement becomes increasingly more expensive as the polynomial degree of the basis, p, increases;

this is due to the fact that knots must be inserted with a multiplicity of $\dot{m} = p$ in order to render the basis C^0 continuous and the number of unknowns is increased for every knot inserted.

Two different AMR strategies are introduced in the following subsections: one that aims to minimise, almost everywhere, the global discretisation error in the solution, as measured in an appropriate norm, dubbed the NRG-AMR strategy, although the error does not necessarily have to be measured in the energy-norm; viz. the L_2-norm or the H^1-(semi-)norm. The other strategy aims to minimise the error in some specific goal, dubbed the DWR-AMR strategy. Neither strategy considers mesh coarsening; viz. those elements that contribute the least to the error in the quantity of interest are not *un-refined*.

Within a multi-group framework, the notion of optimality naturally extends to the use of energy-dependent meshes; q.v. Section 3.5. Owing to the often significant variation of nuclear data with respect to space and more importantly with respect to energy, the smoothness of the multi-group flux components can vary drastically between the energy groups; viz. faster groups tend to exhibit smoother flux profiles, whereas the thermal group typically exhibits more pronounced variations in its spatial distribution. An effective refinement strategy should thus be able to meet the demands of each group whilst refraining from a superfluous expenditure of resources. A mesh adapted to the higher-energy groups is likely to be too coarse for the lower-energy groups; whereas a mesh tailored to the latter is likely to be excessive in regards to the former. A single energy-independent mesh could not therefore yield the greatest accuracy per total DoF.

Before continuing, one must add the following caveats. For the discretisation error defined as the difference between the true solution and the numerical approximation, establishing any sharp error bounds thereof is generally an arduous task and problem dependent; and especially so for hyperbolic systems, for which a lack of smoothness is reflected by weaker error estimates and often it is simpler to derive any error measures heuristically [69, 237, 315–317]. For this reason, the error measures adopted in this work are based upon mathematical proofs and error bounds that are well studied and well established for elliptic systems; namely, the multi-group NDE. To this end, one can establish relatively cheap and flexible error measures that are based upon angular-integrated quantities; viz. this simplifies the analysis since the number of flux moments will be far less than the number of discrete ordinates and the angular neutron flux can be discarded after a full transport sweep. The same projection-based error measures will be employed for each of the discretisations of each form of the NTE considered in this work; and their efficacy will be judged by numerical experimentation.

Whilst it is appreciated that the multi-group NDE have yet to be formally discretised, it is sufficient for now to assume that the weak problem can be posed as Equation 2.62; i.e. to find $\hat{\phi} \in \hat{\mathcal{V}}$, such that:

$$a(\hat{\phi}, \hat{w}) = b(\hat{w}), \quad \forall \hat{w} \in \hat{\mathcal{V}}; \tag{3.21}$$

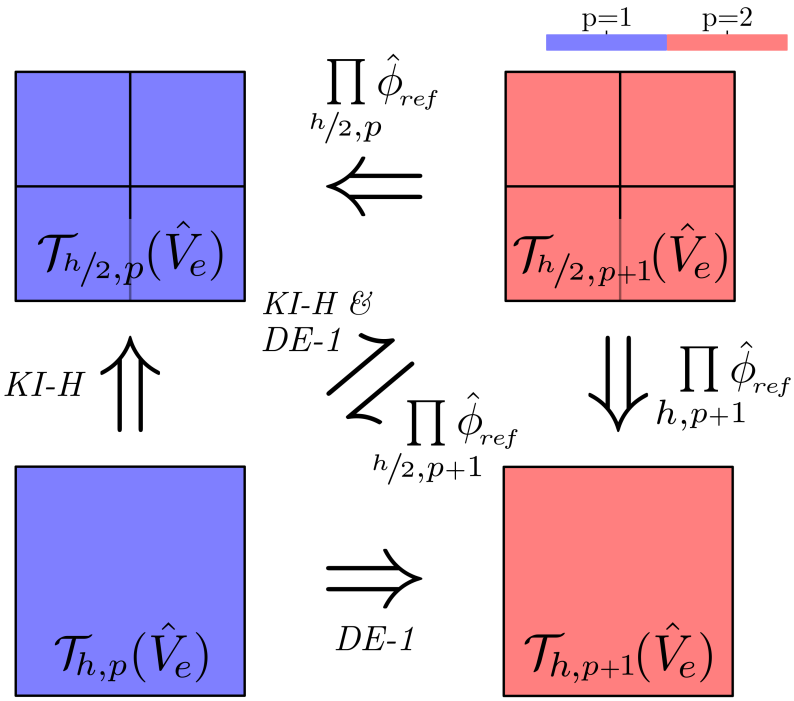


Figure 3.9: A schematic of the interpolation between a more-refined reference mesh and a coarser current mesh; and the projections to prospective subspaces for the next iteration of the AMR procedure. (*Vide* the web-based version for reference to colour.)

where it is assumed that the bilinear form, $a(\cdot,\cdot) = \sum_g a^g(\cdot,\cdot)$, bears the properties of *continuity* and *coercivity*; qq.v. Section 4.3 and Section 4.4 for further details for continuous and discontinuous Galerkin discretisations of the multi-group NDE, respectively.

3.6.1 Energy-Driven Error Measures

The NRG-AMR strategy aims to minimise the residual, or rather the discretisation error, $e = \phi - \hat{\phi}_{h,p}$, over the entire computational domain by locally refining those regions that contribute significantly to e, as measured in the energy-norm; viz. from Equation 3.23, $\epsilon = \|e\|$. By the equivalence of norms on the chosen function space and the assumption that the first-order derivative term is dominant, the discretisation error can be simply measured in the H^1-seminorm, $|\cdot|_{H^1}$; i.e.:

$$|\phi - \hat{\phi}_{h,p}|^2_{H^1} = |e|^2_{H^1} = \|\vec{\nabla} e\|^2_{L_2} = \int_V \sum_i^d \left[\frac{\partial}{\partial x_i} e(\vec{r})\right]^2 dV. \tag{3.22}$$

An approximate reference solution, $\hat{\phi}_{ref}$, is proposed in lieu of the true solution, ϕ. Although still an approximation, the reference solution, calculated over some more-refined reference mesh(es), $\mathcal{T}_{ref}$, is a more accurate approximation of the true solution. With respect to Equation 3.20, the reference solution must sufficiently resolve the local scales of either of the *two* proposed local refinement options, KI-h-refinement and DE-1-refinement, in order to justify the use of one over the other in the next iteration of mesh refinement; i.e. $\hat{\mathcal{V}}_{h^*,p^*}(V_e) \subset \hat{\mathcal{V}}_{ref}(V_e)$, [299]. Accordingly, the reference mesh should

be at least one-level of global uniform KI-h-refinement and DE-1-refinement above that of the coarser current mesh; v. Figure 3.9, where $\mathcal{T}_{ref} = \mathcal{T}_{\frac{1}{2},p+1}$ and therefore $\hat{\phi}_{ref} = \hat{\phi}_{\frac{1}{2},p+1} \in \hat{\mathcal{V}}_{\frac{1}{2},p+1}$. All measurements of error are now estimates made with respect to the reference solution; i.e. $\hat{\epsilon} = |\hat{e}|_{H^1}$, where $\hat{e} = \hat{\phi}_{ref} - \hat{\phi}_{h,p}$.

For $\hat{\phi}_{ref} \in \hat{\mathcal{V}}_{ref}$ and for $\hat{\phi}_{h,p} \in \hat{\mathcal{V}}_{h,p} \subset \hat{\mathcal{V}}_{ref}$ it can be shown that the global error estimate, $\hat{\epsilon}$, provides an upper bound to the true global error, under the assumption that the reference solution is a much better approximation of the true solution than is $\hat{\phi}_{h,p}$ calculated on $\mathcal{T}_{h,p}$; i.e. $\|\phi - \hat{\phi}_{ref}\| \ll \|\hat{\phi}_{ref} - \hat{\phi}_{h,p}\|$:

$$C_\beta \|\phi\|^2 \leq a(\phi, \phi); \tag{3.23a}$$

$$\implies C_\beta \|\phi - \hat{\phi}_{h,p}\|^2 \leq a(\phi - \hat{\phi}_{h,p}, \phi - \hat{\phi}_{h,p}), \tag{3.23b}$$

$$\leq a(\phi - \hat{\phi}_{h,p}, \phi - \hat{\phi}_{ref}) + a(\phi - \hat{\phi}_{h,p}, \hat{\phi}_{ref} - \hat{\phi}_{h,p}), \tag{3.23c}$$

$$\leq C_\alpha \|\phi - \hat{\phi}_{ref}\| \|\phi - \hat{\phi}_{h,p}\| + C_\alpha \|\hat{\phi}_{ref} - \hat{\phi}_{h,p}\| \|\phi - \hat{\phi}_{h,p}\|; \tag{3.23d}$$

$$\implies \|\phi - \hat{\phi}_{h,p}\| \leq \frac{C_\alpha}{C_\beta} \Big[\|\phi - \hat{\phi}_{ref}\| + \|\hat{\phi}_{ref} - \hat{\phi}_{h,p}\| \Big], \tag{3.23e}$$

$$\|e\| \lesssim \frac{C_\alpha}{C_\beta} \|\hat{e}\|, \tag{3.23f}$$

$$\epsilon \lesssim \frac{C_\alpha}{C_\beta} \hat{\epsilon}. \tag{3.23g}$$

The first line expresses the property of coercivity in the bilinear form, with some constant, C_β; and the second line expresses the first in terms of the discretisation error; where the linearity of each argument of a is exhibited in the third line by introducing $\hat{\phi}_{ref}$; and the fourth line expresses the property of continuity, with some constant, C_α. The constants of coercivity and continuity depend upon the nature of the discretisation; viz. whether it be *conforming* in the case of continuous discretisations or *non-conforming* in the case of discontinuous ones; qq.v Section 4.3 and Section 4.4, respectively.

The performance of *a posteriori* error estimation is measured by the (reciprocal) effectivity index [309]:

$$\theta = \frac{\hat{\epsilon}}{\epsilon}. \tag{3.24}$$

Again, as the true solution is not available on the left-hand side of Equation 3.23, this too is approximated by some reference solution in its stead; one that is much more accurate than $\hat{\phi}^g_{\frac{1}{2},p+1}$. It is clearly desirable that $\theta_1 \to 1$ as $h \to 0$ and $p \to \infty$; however, it is more important that $\hat{\epsilon}$ provides a sharp upper bound to the true remaining error such that it may be used as a criterion to terminate the AMR procedure; viz. for $\theta > 1$.

Owens et al investigated the effectiveness of using error estimates based upon the explicit calculation of a current coarse solution, as well as a refined reference solution, to drive their AMR-h strategy

for the DG-IGA discretisation of the FONT-S_N equations in 2D; where the reference mesh was taken to be that resulting from a global uniform DE-1-refinement of the current coarser mesh [278]. For Equation 3.23 to hold, the reference solution must always provide a much better approximation of the true solution than that of the coarser solution; however, it was observed that the accuracy in the estimation of the error degraded as the polynomial degree of the NURBS basis increased. As per the tensor-product structure of the NURBS basis, it was understood that the ratio of DoF per element in the reference mesh to that of the current mesh, $(p+2)^d/(p+1)^d$, is a monotonically decreasing function for positive integer values of p; viz. quadratics solutions are significantly more accurate than linear ones, yet, cubics are not much more accurate than quadratics. Should the reference mesh be taken as $\mathcal{T}_{h/2,p+1}$, for arbitrarily constructed meshes, $\mathcal{T}_{h,p}$, it would be expected that a more robust exit-criterion, $\hat{\epsilon}$, could be obtained. Nevertheless, the relative sizes in the element contributions to the overall error would still be useful in driving the AMR procedure; and significant accuracy can be gained over global uniform refinement strategies for comparable computational efforts [318].

It would be computationally expensive to have to perform *two* explicit calculations over *two* different levels of mesh refinement for each energy group. Fortunately, one can invoke Céa's Lemma, Equation 4.36, for conforming and consistent discretisations or Strang's second Lemma, Equation 4.81, for non-conforming and consistent discretisations, to demonstrate that the chosen approximation space will yield a solution, $\hat{\phi}_{h,p} \in \hat{\mathcal{V}}_{h,p}$, whose error is bounded by the quasi-best approximation error, which itself is bounded by the difference in the exact solution and any function that belongs to the finite-dimensional *hp* function space. A particular choice would consist of choosing the projection of the exact solution onto this approximation space, $\prod_{h,p}\phi$. The entire AMR procedure can therefore be driven by estimating and then controlling the interpolation error that bounds the quasi-best approximation error, which in turn bounds the true discretisation error; i.e. [319]:

$$\|\phi - \hat{\phi}_{h,p}\|_{H^1} \leq \frac{C_\alpha}{C_\beta} \inf_{\hat{\phi}_{h,p} \in \hat{\mathcal{V}}_{h,p}} \|\phi - \hat{\phi}_{h,p}\|_{H^1} \leq \frac{C_\alpha}{C_\beta} \|\phi - \prod_{h,p} \phi\|_{H^1};$$

where again, the constants of coercivity and continuity will depend upon the nature of the discretisation; qq.v Section 4.3 and Section 4.4.

Of course, the exact solution must be substituted for the reference solution, which yields an iterative procedure that requires an explicit solution over the more-refined reference mesh and its projection onto the coarser current mesh. This projection, $\prod_{h,p}\colon \hat{\mathcal{V}}_{h/2,p+1} \to \hat{\mathcal{V}}_{h,p}$, is performed as a conservative interpolation between volume meshes, similar to that introduced for energy-dependent meshes; q.v. Section 3.5. As an approximation, $\hat{\phi}$, is sought to minimise the discretisation error as that measured in the energy-norm, or the H^1-(semi)norm by the equivalence of norms, unsurprisingly, the interpolation over the reference solution onto the coarser mesh is chosen to be optimal in the H^1-norm. It was noted

by Wang and Ragusa that the explicit calculation of both a reference solution and a solution over the coarser current mesh did not yield any advantage in terms of accuracy over an interpolation-based approach; especially in the asymptotic range of convergence [290].

The estimation of the global discretisation error, based upon the interpolation of a reference solution, is summarised below:

$$\epsilon^2 = \sum_{g=1}^{G} (\epsilon^g)^2, \tag{3.25a}$$

$$= \sum_{g=1}^{G} \| \phi^g - \hat{\phi}^g_{h,p} \|^2_{H^1}, \tag{3.25b}$$

$$\leq \sum_{g=1}^{G} \| \phi^g - \prod_{h,p} \phi^g \|^2_{H^1}, \tag{3.25c}$$

$$\lesssim \sum_{g=1}^{G} \| \hat{\phi}^g_{h/2,p+1} - \prod_{h,p} \hat{\phi}^g_{h/2,p+1} \|^2_{H^1}, \tag{3.25d}$$

$$\lesssim \sum_{g=1}^{G} \sum_{V_e \in \mathcal{T}^g_{h,p}} \| \hat{\phi}^g_{h/2,p+1} - \prod_{h,p} \hat{\phi}^g_{h/2,p+1} \|^2_{H^1(V_e)}, \tag{3.25e}$$

$$\lesssim \sum_{g=1}^{G} \sum_{V_e \in \mathcal{T}^g_{h,p}} \mu^g_e. \tag{3.25f}$$

The global discretisation error estimate, $\hat{\epsilon}$, is now a projection-based error estimate of the global discretisation error, bounded by the sum of element-local interpolation-based error measures, μ^g_e; i.e.:

$$\hat{\epsilon} = \left[\sum_{g=1}^{G} (\hat{\epsilon}^g)^2 \right] \frac{1}{2} = \left[\sum_{g=1}^{G} \sum_{V_e \in \mathcal{T}^g_{h,p}} \mu^g_e \right] \frac{1}{2}; \tag{3.26}$$

where, by the Poincaré inequality, μ^g_e are defined in terms of the first-order Sobolev space seminorm, for $g = 1, \ldots, G$:

$$\mu^g_e = | \hat{\phi}^g_{h/2,p+1} - \prod_{h,p} \hat{\phi}^g_{h/2,p+1} |^2_{H^1(V_e)}, \quad \forall \, V_e \in \mathcal{T}^g_{h,p}(V). \tag{3.27}$$

In the event that only a single energy-independent mesh is employed, the error measure will also be made to be energy-independent; i.e. $\mu_e = \sum_g \mu^g_e$.

Driven by the residual in the interpolation-based error estimate, successive meshes at each iteration of the AMR procedure are obtained by minimising local contributions to the error. At a global level, those elements that contribute the most to the overall error are deemed eligible for further refinement; where at a local level, an optimal refinement strategy is determined for each, as per Equation 3.20. This approach differs from that of Demkowicz et al who first determined an optimal refinement strategy for all elements and then refined those whose potential to reduce the overall error was greatest [257].

One could choose to refine a certain proportion of those elements that contribute the most to the overall error at each iteration on the mesh; alternatively, an element could be flagged should its local contribution to the overall error lie above a certain threshold. The former approach may not always allow for sufficient refinement of the mesh at each iteration of the AMR procedure, therefore, the current work adopts the latter; where the threshold is taken to be some fraction, $\beta_1 \in [0,1]$, of the maximum error over a single element belonging to any of the current coarser meshes. For energy-dependent meshes, $g = 1, \ldots, G$, at a some given iteration count in the AMR procedure, (i), the first decision in the AMR procedure, D-1 in Figure 3.10, is thus to determine those elements that satisfy:

$$w^{g(i)}\mu_e^{g(i)} \geq \beta_1 \max_{g\prime=1,\ldots,G} \max_{V_{e\prime}\in\mathcal{T}_{h,p}^{g\prime(i)}} w^{g\prime(i)}\mu_{e\prime}^{g\prime(i)} \quad , \quad \forall\, V_e \in \mathcal{T}_{h,p}^{g(i)}(V). \tag{3.28}$$

Although somewhat arbitrary, experience suggests that $\beta_1 = \frac{1}{3}$ provides a good compromise of not having to perform too many AMR iterations whilst not refining too many elements at any one time; where the contrary to the latter would see the AMR procedure tend towards a global uniform refinement strategy, which is certainly not optimal in terms of the greatest accuracy per DoF [307].

Non-smooth multi-group components of the solution are targeted by choosing to refine all of the energy-dependent meshes at once. However, one must prevent the AMR procedure from simply refining those elements over which the magnitude of the error is small in comparison to the solution over the respective mesh, yet, may be large in comparison to the magnitude of the errors over other meshes; whence the apparition of the group-importance factor, $w^{g(i)}$, in Equation 3.28, which simply normalises the group error measures with respect to their respective multi-group component of the scalar neutron flux, as measured in some appropriate norm; i.e. for $g = 1, \ldots, G$:

$$w^g = \frac{1}{\|\hat{\phi}^g_{h/2,p+1}\|}. \tag{3.29}$$

For those elements that satisfy Equation 3.28, an optimal local refinement strategy is determined by maximising the rate at which the local contribution to the interpolation-based error estimate can be decreased, Equation 3.20. As stated previously, the options for local refinement are restricted to 1×KI-h-refinement or 1×DE-1-refinement; and as the the reference mesh is chosen to be $\mathcal{T}_{h/2,p+1}(V)$, it is possible to interpolate the reference solution onto the prospective subspaces, $\hat{\mathcal{V}}_{h^*,p^*}(V_e)$, and evaluate, $\bar{\Delta}\epsilon^{(i)}_{h^*,p^*}|_{V_e}$; v. Figure 3.9. For $h^*, p^* = \{h/2, p \text{ or } h, p{+}1\}$, the optimal combination of h^*, p^* on $V_e \in \mathcal{T}^g(V)$ is chosen to be that which maximises the following:

$$\max_{h^*,p^*} \bar{\Delta}\epsilon^{g(i)}_{h^*,p^*}|_{V_e} = \max_{h^*,p^*} \frac{\mu_e^{g(i)} - |\hat{\phi}^{g(i)}_{h/2,p+1} - \prod_{h^*,p^*} \hat{\phi}^{g(i)}_{h/2,p+1}|^2_{H^1(V_e)}}{|DoF^{g(i)}_{h,p} - DoF^{g(i)}_{h^*,p^*}|_{V_e}}. \tag{3.30}$$

Two conditions are attached to Equation 3.30 in settling the dispute between competing refinement strategies over a given element, such that certain constraints are imposed upon the use of either DE-1-

refinement or KI-h-refinement. The first corresponds to that employed at D-2 in Figure 3.10, where an upper limit is placed upon the maximum polynomial degree of the NURBS basis; i.e. DE-1-refinement will only be considered should:

$$p|_{V_e} \leq \beta_2. \tag{3.31}$$

The second corresponds to that employed at D-3 in Figure 3.10, where in the case of regular smooth solutions, it is demanded that $\Delta\epsilon^{g(i)}_{\mathrm{h}^1_2,p}|_{V_e}$ must be relatively large, with respect to the error estimated over the element, $\mu^{g(i)}_e$, for KI-h-refinement to be considered, since it reduces continuity in the basis. For some parameter, $\beta_3 \in [0,1]$, KI-h-refinement will only be considered should:

$$\Delta\epsilon^{g(i)}_{\mathrm{h}^1_2,p}|_{V_e} \geq \beta_3 \mu^{g(i)}_e; \tag{3.32}$$

where $\beta_3 = 0$ implies that no restriction is placed upon the use of KI-h-refinement; whereas $\beta_3 = 1$ implies that KI-h-refinement is strictly prohibited. It has been shown, for an AMR-$h(p)$ strategy applied to 1D FEM-discretisations of the multi-group NDE, that values of $\beta_3 \to 1$ are preferred [318]. Without being too restrictive, the current work employs $\beta_3 = ½$.

To recapitulate the parameters employed in the constraints imposed upon local refinement are:

$$\begin{cases} \beta_1 = ⅓ & \text{(Equation 3.28);} \\ \beta_2 = p_{max} & \text{(Equation 3.31);} \\ \beta_3 = ½ & \text{(Equation 3.32).} \end{cases} \tag{3.33}$$

3.6.2 Goal-Based Error Measures

The DWR-AMR strategy aims to minimise the residual in some specific QoI by locally refining those regions who significantly contribute to the error in the specified goal. Although in practice there exists an assortment of different types of QoI, the current work considers only *two* of the most common in their application to reactor physics problems: volumetric detector responses and the effective neutron multiplication factor, $\mathrm{k_{eff}}$.

Under the assumption that each of the primal and the dual problems have been discretised in the same manner, it is assumed in the following that the discrete variational statements for each are adjoint-consistent; i.e. $a(\hat{u},\hat{v}) - a^{\dagger}(\hat{v},\hat{u}) = 0 \ \forall \hat{u},\hat{v} \in \hat{\mathcal{V}}_{h,p}$. Where adjoint-consistency of the discrete bilinear forms signifies a consistent discretisation of the (semi-)analytical adjoint problem, from which the solution to the discrete adjoint problem inherits its properties of smoothness. Consequently, (quasi-)optimal rates of convergence can be attained for the error measured in some specified goal [320].

3.6.2.1 A Volumetric Detector Response

The response of a detector uniformly distributed within some given detector material, denoted by the subscript d, is measured as some reaction rate over the subdomain, V_d; i.e.:

$$QoI(\phi) = \sum_{g=1}^{G} \int_{V_d} \Sigma_d^g(\vec{r}) \phi^g(\vec{r}) \, dV. \tag{3.34}$$

The reaction rate could be that measured by Uranium fission chambers, for example; however, more generally, the reactions that are considered are those characterised by the effective macroscopic cross-section of absorption; and Σ_d^g is defined such that, for $g = 1, \ldots, G$:

$$\Sigma_d^g(\vec{r}) = \begin{cases} \Sigma_a^g(\vec{r}), & \vec{r} \in V_d; \\ 0, & \vec{r} \notin V_d. \end{cases} \tag{3.35}$$

Fixed-Source Problems By the arbitrariness of the fixed adjoint source, Equation 2.15, and letting $q_x^\dagger = \Sigma_d$, the QoI is made equal to the adjoint linear functional for some test function ϕ; i.e. for $\phi^\dagger \in \mathcal{V}$:

$$a^\dagger(\phi^\dagger, \phi) = b^\dagger(\phi), \quad \forall \phi \in \mathcal{V}, \tag{3.36a}$$

$$= QoI(\phi), \quad \forall \phi \in \mathcal{V}. \tag{3.36b}$$

The DWR-AMR strategy aims to minimise the absolute error in the goal, $|QoI(\phi) - QoI(\hat{\phi}_{h,p})|$, Equation 3.23. Again, one can establish a global error estimate, $\hat{\epsilon}$, measured with respect to a reference solution, that provides an upper bound to the true global error, ϵ; this is summarised in the following lines for the QoI as defined in Equation 3.34:

$$|QoI(\phi) - QoI(\hat{\phi}_{h,p})| = |QoI(\phi - \hat{\phi}_{h,p})|, \tag{3.37a}$$

$$= |a^\dagger(\phi^\dagger, \phi - \hat{\phi}_{h,p})|, \tag{3.37b}$$

$$= |a(\phi - \hat{\phi}_{h,p}, \phi^\dagger)|, \tag{3.37c}$$

$$= |a(\phi - \hat{\phi}_{h,p}, \phi^\dagger - \hat{\phi}_{h,p}^\dagger)|, \tag{3.37d}$$

$$\begin{aligned} &= |a(\phi - \hat{\phi}_{ref}, \phi^\dagger - \hat{\phi}_{ref}^\dagger)| \quad + |a(\phi - \hat{\phi}_{ref}, \hat{\phi}_{ref}^\dagger - \hat{\phi}_{h,p}^\dagger)| \\ &+ |a(\hat{\phi}_{ref} - \hat{\phi}_{h,p}, \phi^\dagger - \hat{\phi}_{ref}^\dagger)| + |a(\hat{\phi}_{ref} - \hat{\phi}_{h,p}, \hat{\phi}_{ref}^\dagger - \hat{\phi}_{h,p}^\dagger)|, \end{aligned} \tag{3.37e}$$

$$\begin{aligned} &\leq C_\alpha \Big[\|\phi^\dagger - \hat{\phi}_{ref}^\dagger\| \|\phi - \hat{\phi}_{ref}\| \quad + \|\hat{\phi}_{ref}^\dagger - \hat{\phi}_{h,p}^\dagger\| \|\phi - \hat{\phi}_{ref}\| \\ &+ \quad \|\phi^\dagger - \hat{\phi}_{ref}^\dagger\| \|\hat{\phi}_{ref} - \hat{\phi}_{h,p}\| + \|\hat{\phi}_{ref}^\dagger - \hat{\phi}_{h,p}^\dagger\| \|\hat{\phi}_{ref} - \hat{\phi}_{h,p}\| \Big], \end{aligned} \tag{3.37f}$$

$$\lesssim C_\alpha \|\hat{\phi}_{ref}^\dagger - \hat{\phi}_{h,p}^\dagger\| \|\hat{\phi}_{ref} - \hat{\phi}_{h,p}\|, \tag{3.37g}$$

$$\lesssim C_\alpha \|\hat{e}^\dagger\| \|\hat{e}\|; \tag{3.37h}$$

$$\implies \epsilon \lesssim C_\alpha \hat{\epsilon}; \tag{3.37i}$$

where, by the linearity of each argument in the linear functional, the error in the goal is re-expressed in the first line, Equation 3.37a, which is now taken to be the right-hand side of the variational form of the adjoint problem, Equation 3.36. The adjoint bilinear form, $a^\dagger$, satisfies the property of adjoint-consistency, from which the error in the QoI is given in the third line, Equation 3.37c, as the residual of the direct problem weighted by the importance function, as measured in the bilinear inner-product, a; and as per Galerkin's property of orthogonality, the fourth line is equal to that above it. A numerical reference solution to each of the direct and the adjoint problems, $\hat{\phi}_{ref} = \hat{\phi}_{h/2,p+1}$ and $\hat{\phi}^\dagger_{ref} = \hat{\phi}^\dagger_{h/2,p+1}$, respectively, is introduced in the fifth line, which is expanded out, by the linearity of each argument in the bilinear form. The property of continuity is expressed in the sixth line, Equation 3.37f; and in the seventh line, Equation 3.37g, it is assumed that the error in the more-refined reference solution with respect to the true solution is negligible in comparison to the residual formed by the coarser current solution with respect to the former. Again, the effectivity index, Equation 3.24, can be employed to measure the performance of estimating the error in the goal, $\hat{\epsilon}$.

For a given energy group, the reference solutions to the primal and the dual problems are solved on the same reference mesh, $\mathcal{T}_{h/2,p+1}$; and rather than effectuate the explicit calculation of *four* different solutions over *two* different levels of mesh refinement, only the reference solutions need to be computed. The DWR-AMR strategy proceeds in a similar manner to that of the NRG-AMR strategy since it is driven by estimating the error in the QoI by interpolation of the reference solutions onto the coarser current meshes; i.e:

$$\epsilon^2 = \sum_{g=1}^{G} (\epsilon^g)^2, \tag{3.38a}$$

$$\leq \sum_{g=1}^{G} \| \phi^{g\dagger} - \hat{\phi}^{g\dagger}_{h,p} \|^2_{H^1} \| \phi^g - \hat{\phi}^g_{h,p} \|^2_{H^1}, \tag{3.38b}$$

$$\leq \sum_{g=1}^{G} \| \phi^{g\dagger} - \prod_{h,p} \phi^{g\dagger} \|^2_{H^1} \| \phi^g - \prod_{h,p} \phi^g \|^2_{H^1}, \tag{3.38c}$$

$$\lesssim \sum_{g=1}^{G} \| \hat{\phi}^{g\dagger}_{h/2,p+1} - \prod_{h,p} \hat{\phi}^{g\dagger}_{h/2,p+1} \|^2_{H^1} \| \hat{\phi}^g_{h/2,p+1} - \prod_{h,p} \hat{\phi}^g_{h/2,p+1} \|^2_{H^1}, \tag{3.38d}$$

$$\lesssim \sum_{g=1}^{G} \sum_{V_e \in \mathcal{T}^g_{h,p}} \| \hat{\phi}^{g\dagger}_{h/2,p+1} - \prod_{h,p} \hat{\phi}^{g\dagger}_{h/2,p+1} \|^2_{H^1(V_e)} \| \hat{\phi}^g_{h/2,p+1} - \prod_{h,p} \hat{\phi}^g_{h/2,p+1} \|^2_{H^1(V_e)}, \tag{3.38e}$$

$$\lesssim \sum_{g=1}^{G} \sum_{V_e \in \mathcal{T}^g_{h,p}} \mu^g_e. \tag{3.38f}$$

The global error estimate in the QoI, $\hat{\epsilon}$, is now a projection-based error estimate of ϵ, bounded by the sum of element-local interpolation-based error measures, μ^g_e, Equation 3.26, which can be measured

in the H^1-seminorm; i.e. for $g = 1, \dots, G$:

$$\mu_e^g = |\hat{\phi}^g_{\text{h}\!/\!2,p+1} - \prod_{h,p} \hat{\phi}^g_{\text{h}\!/\!2,p+1}|^2_{H^1(V_e)} |\hat{\phi}^{g\dagger}_{\text{h}\!/\!2,p+1} - \prod_{h,p} \hat{\phi}^{g\dagger}_{\text{h}\!/\!2,p+1}|^2_{H^1(V_e)}, \quad \forall\, V_e \in \mathcal{T}^g_{h,p}(V). \tag{3.39}$$

Although one must now solve the adjoint problem in tandem with the direct one at P-1 in Figure 3.10, one can expect faster rates of convergence and greater mesh optimality with respect to a specified goal, when compared to the NRG-AMR strategy, since one now considers both the smoothness of the solution as well as the influence of specific regions on the QoI.

Those elements eligible for further refinement at D-1 in Figure 3.10 are determined by Equation 3.28; yet, one recognises that the *importance* of each energy group is already accounted for by the error measures, Equation 3.39; therefore $w^g = 1$ for $g = 1, \dots, G$. The same options for local refinement are available, as in the case of the NRG-AMR strategy, and those same constraints in Equation 3.31 and Equation 3.32, made at D-2 and D-3 in Figure 3.10, respectively, still apply; q.v. Equations 3.33.

One subtle difference, however, exists in choosing the refinement option that maximises the rate at which the local contribution to the interpolation-based error estimate decreases, Equation 3.20; i.e.:

$$\max_{h^*,p^*} \bar{\Delta}\epsilon^{g(i)}_{h^*,p^*}|_{V_e} = \max_{h^*,p^*} \frac{\mu_e^{g(i)} - |\hat{\phi}^{g(i)}_{\text{h}\!/\!2,p+1} - \prod\limits_{h^*,p^*} \hat{\phi}^{g(i)}_{\text{h}\!/\!2,p+1}|_{H^1(V_e)} |\hat{\phi}^{g\dagger(i)}_{\text{h}\!/\!2,p+1} - \prod\limits_{h^*,p^*} \hat{\phi}^{g\dagger(i)}_{\text{h}\!/\!2,p+1}|_{H^1(V_e)}}{|DoF^{g(i)}_{h,p} - DoF^{g(i)}_{h^*,p^*}|_{V_e}}. \tag{3.40}$$

Eigenvalue Problems In the interest of detector responses for fixed-source problems the QoI is a linear quantity. However, in the consideration of eigenvalue problems, or rather fission source problems in the absence of extraneous fixed sources of neutrons, the QoI becomes a non-linear quantity and the fission sources must be known before the detector response can be determined. In which case, one must first converge the fission sources, as well as any other group-to-group sources, and the eigenvalue in the dual problem as per the PI procedure, after which any discretisation errors therein can be assumed to be negligible; the dual problem can then be resolved with the QoI as per Equation 3.34 and Equation 3.36 [321]. Overall, this translates as *one* additional iteration when resolving the dual problem; indeed, the k_{eff} determined for the primal problem can be used to expedite the power-iteration convergence of the eigenvalue for the dual problem [291, 294].

3.6.2.2 The Effective Neutron Multiplication Factor

In the consideration of multiplying systems and in the absence of any extraneous fixed sources of neutrons, $q_x = 0$ and hence $b = 0$, the weak primal problem is posed as finding $\hat{\phi} \in \hat{\mathcal{V}}_{h,p}$ such that:

$$\tilde{a}(\hat{\phi}, \hat{w}) = \frac{1}{\text{k}_{\text{eff}}} a_F(\hat{\phi}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}; \tag{3.41}$$

and similarly, the weak dual problem is posed as finding $\hat{\phi}^\dagger \in \hat{\mathcal{V}}_{h,p}$ such that:

$$\tilde{a}^\dagger(\hat{\phi}^\dagger, \hat{w}) = \frac{1}{\mathrm{k_{eff}}} a_F^\dagger(\hat{\phi}^\dagger, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}; \tag{3.42}$$

where $\tilde{a}$ and $\tilde{a}^\dagger$ are the original bilinear forms in Equation 3.21 and Equation 3.36, respectively, less the fission-source terms; i.e.: $\tilde{a} = a + a_F$ and $\tilde{a}^\dagger = a^\dagger + a_F^\dagger$, respectively.

When solving eigenvalue problems, experience suggests that the solution should be normalised after each outer iteration, such as to avoid problems of underflow or overflow; and good practice suggests that the primal and the dual problems should not be normalised independently. Accordingly, the direct solution is normalised such that $a_F(\phi, \phi) = 1$; and the adjoint solution is normalised such that $a_F^\dagger(\hat{\phi}^\dagger, \hat{\phi}) = a_F(\hat{\phi}, \hat{\phi}^\dagger) = 1$; in which case, the QoI can be expressed as [255, 256]:

$$QoI(\phi) = \frac{1}{\mathrm{k_{eff}}} = \frac{1}{\mathrm{k_{eff}}} a_F^\dagger(\hat{\phi}^\dagger, \hat{\phi}). \tag{3.43}$$

With respect to a reference solution, one can establish a global error estimate in the eigenvalue, $\hat{\epsilon}$, one that provides an upper bound to the true global error, ϵ, which is summarised as:

$$\begin{aligned}
|QoI(\phi) \text{ - } QoI(\hat{\phi}_{h,p})| &= |QoI(\phi \text{ - } \hat{\phi}_{h,p})|, && \text{(3.44a)}\\
&= |\frac{1}{\mathrm{k_{eff}}} a_F^\dagger(\phi^\dagger, \phi \text{ - } \hat{\phi}_{h,p})|, && \text{(3.44b)}\\
&= |\frac{1}{\mathrm{k_{eff}}} a_F(\phi \text{ - } \hat{\phi}_{h,p}, \phi^\dagger)|, && \text{(3.44c)}\\
&= |\tilde{a}(\phi \text{ - } \hat{\phi}_{h,p}, \phi^\dagger)|; && \text{(3.44d)}\\
\implies \epsilon &\lesssim \tilde{C}_\alpha \hat{\epsilon}; && \text{(3.44e)}
\end{aligned}$$

where by the linearity of each argument in the linear functional the error in the goal is re-expressed in the first line, which is taken as the right-hand side of the variational form of the dual problem, Equation 3.42. The adjoint bilinear form of the fission sources, $a_F^\dagger$, is assumed to satisfy the property of adjoint-consistency, from which the error in the QoI is given in the third line, Equation 3.44c, as the product of the eigenvalue and the residual of the direct problem weighted by the importance function, as measured in the bilinear inner-product, a_F; which can then be expressed as the left-hand side of Equation 3.41. The derivation proceeds as that from Equation 3.37c onwards, whereby, again, the same error measures as those defined in Equation 3.39 can be employed.

3.7 A Discussion

This chapter introduced NURBS-based IGA as the spatial discretisation of choice for the development of subsequent chapters, which will focus on the actual spatial discretisation of the different forms of

the NTE. In contrast to the nodal values of classical FEM, one appreciates that the non-interpolatory nature of the NURBS basis does not lead to a concrete interpretation of the unknowns being solved for in the resulting matrix equations, in essence, they are merely parameters that serve as the *scaffolding* to the numerical solution just as the control-points provides the scaffolding for the geometry.

The process for generating the NURBS-based descriptions of the geometries that will serve as the initial coarsest mesh for analysis was detailed. Ideally, this pre-processing step would be performed by some CAD software and the information passed in some standard file format. Within the scope of this work, however, this is achieved by bespoke bits of software. Future work would naturally extend to allow for the analysis to be performed over more general geometries that could be described in bases other than the standard NURBS and *read in* as STEP or IGES files, for example.

General global uniform refinement strategies were discussed by the mechanisms knot insertion and degree elevation, which are somewhat analogous to those of standard h-refinement and p-refinement of classical FEM. The process for conservative interpolation of a scalar field between volume meshes was also developed in the guise of energy-dependent meshes; and this same projection procedure was used to devise *two* different types of relatively cheap and flexible interpolation-based error measures. Although based upon well-established mathematical proofs and error bounds for elliptic systems and derived specifically for the multi-group NDE, these same error measures will be employed for each of the discretisations of each form of the NTE considered in this work; and their efficacy will be judged by numerical experimentation.

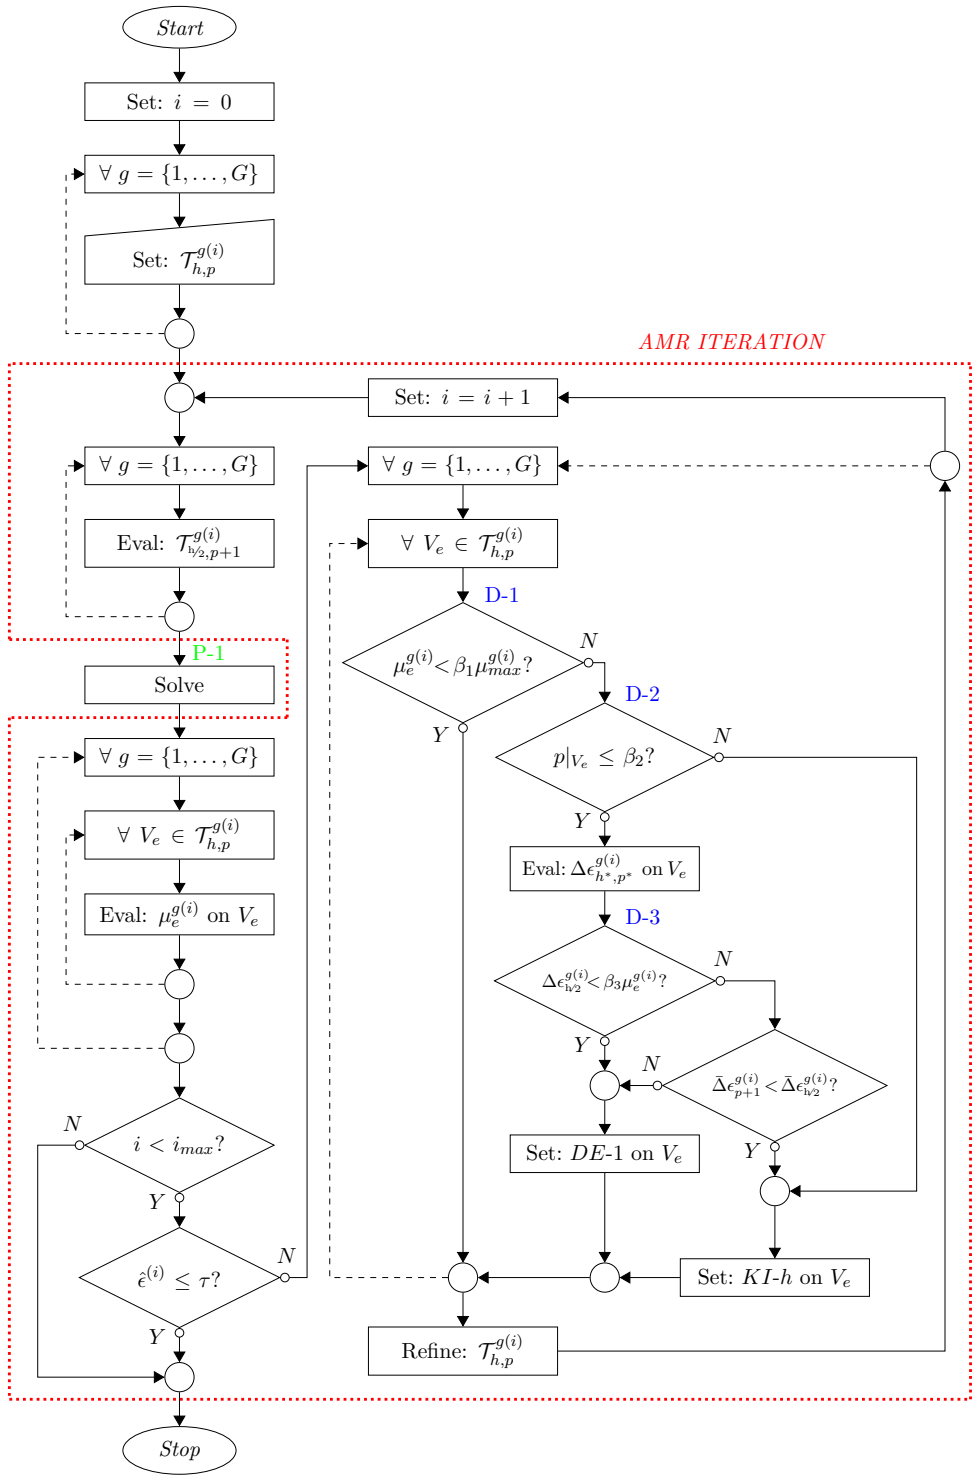


Figure 3.10: The general AMR-iteration procedure depicted as a flowchart, enclosed in the red dotted line. The procedure terminates upon either attaining the maximum number of iterations or upon satisfying some convergence criterion. The black dashed process-flow arrows distinguish *looping* operations, which take priority. For standard NRG-AMR, the interpolation-based error measures are weighted by some appropriate energy-dependent function. Otherwise, in the case of DWR-AMR, they are weighted by the error in the importance function; Process P-1 must therefore resolve the direct (primal) problem as well as the adjoint (dual) one, for some user-specified QoI. *Which* elements are eligible for further refinement is determined at Decision D-1; whereas constraints are placed on *how* they are refined at Decisions D-2 and D-3. (*Vide* the web-based version for reference to colour.)

Chapter 4

Elliptic Systems: Part 1

This chapter presents the application of automised NRG-AMR-$h(p)$ and DWR-AMR-$h(p)$ strategies to separate CBG-IGA and SIP-DG-IGA spatial discretisations of the multi-group NDE.

4.1 An Introduction

The NTE behaves like an elliptic second-order PIDE in optically thick and scattering-dominated regions, away from material discontinuities and physical boundaries; q.v. Section 2.6. Serving as a low-order approximation to the full transport equation at a reduced number of independent variables, the steady-state NDE finds its use in a variety of practical situations; e.g. whole-core systems calculations [54]. Another keen interest lies in its application to diffusion acceleration schemes; q.v. Chapter 7.

Following a standard multi-group discretisation in energy, the PIDE is reduced to a system of BVP weakly coupled in energy that have yet to be spatially discretised or prescribed BC. Working within *two* different frameworks, the discrete problem is posed as that resulting from either a *conforming* CBG-IGA spatial discretisation or a *non-conforming* DG-IGA spatial discretisation. Statements of quasi-best approximation may be established for either asymptotically consistent discretisation.

Within the continuous setting, the basis is defined globally across the computational domain and continuity is generally enforced in a *strong* sense; C^0 continuity is trivially enforced across *regular* compatible meshes because of the one-to-one correspondence that exists between the control-variables at the patch interfaces. Local refinement strategies, however, will inevitably give rise to *irregular* meshes where the patches may no longer be connected in a parametric sense and therefore must be coupled explicitly for analysis. Cottrell et al devised constraint equations based upon the fundamental relationship between the NURBS that define a given object and those of a more refined description of

the same object, expressed via operators that encapsulate the process of refinement [229, 270]. This necessitates at any given irregular interface an ability to determine the more-refined (slave) patch and the less-refined (master) patch, such as to constrain the former to the latter, or some other coarsest refinement shared by the two. This approach was adopted by both Welch and Latimer et al in their separate studies of AMR-h strategies for second-order forms of the neutron transport equation [119, 276]. Although restricted to C^0 continuity between the patches, such a strategy does exploit the potential for higher-order continuity within the basis; in which case, the refinement is not strictly *local* as it propagates within a patch, incurring DoF unnecessarily. Moreover, it was not demonstrated whether it was possible to incorporate p-refinement capabilities. In the application to the Helmholtz equation for harmonic acoustics in 2D and 3D, Coox et al generalised the constraint-based method and were able to enforce higher-order continuity, C^0 and C^1, across irregular patch interfaces by considering the constraints in a least-square sense via the method of virtual uncommon knot-inserted master-slave (VUKIMS) [322]. Lei et al derived their continuity constraints in a similar manner; however, their interface constraints were imposed directly by means of a penalty formulation [323].

In general, CBG-IGA discretisations perform well on a per DoF basis and the C^0 continuity between patches facilitates the modelling of heterogeneous problems to a high degree of accuracy. Yet, in the development of local refinement strategies, the implementation of constraint equations can become rather involved; especially in multiple dimensions where a given control-variable can be common to several indirectly neighbouring elements. More complex still are the problem of multiply-constrained variables; and often some degree of regularity must be imposed over the mesh in order to avoid such cycles of compounded constraints; this does come at the cost of accruing additional DoF [290, 299]. Furthermore, the constraint equations themselves introduce errors and can render the connected model ill conditioned. Fortunately, although not wholly suited to discretisations by NURBS, the conditioning of the matrix may be improved by the application of any traditional preconditioner developed for those general bases of classical FEM, with varying degrees of effectiveness [324].

With regard to DG methods, the locally defined basis is left to be discontinuous and an interior penalty (IP) scheme can be devised to enforce weakly continuity across the domain; however, this does come at the cost of at least twice as many DoF. Nitsche introduced penalty schemes to enforce weakly essential Dirichlet BC by appending penalty terms to the bilinear form [325]. This penalisation was then applied to the interior of the finite element mesh by Wheeler and by Percell and Wheeler as a means of weakly enforcing continuity across element interfaces [326, 327]. IP-DG-FEM was also developed independently by Arnold for general non-linear second-order elliptic and parabolic problems [288]. More recently, IP schemes have been applied to IGA discretisations of elliptic systems [277, 328].

In the case of DG methods, the BVP is formulated locally over a given element and a projection of

residuals is made to be vanishing across its boundaries [242]. IP-DG methods are often compared to stabilised mixed FEM for second-order differential equations, since residual-based penalty terms are introduced to stabilise the discrete problem [194]. There exist myriad IP schemes, each employing its own combination of terms; typically, however, the penalty terms take the form of weighted L_2-inner products of jumps in the trial and test functions at element interfaces; and the coupling of control-variables exists only between direct neighbours that share an interface [329, 330]. From a practical point of view, the inevitable costs incurred in the implementation of an IP-DG scheme may be compensated by its compact discretisation stencil; as well as the flexibility that it offers in terms of mesh refinement [237].

The choice of penalty parameters is crucial to the effectiveness of the numerical method. For arbitrarily large values, iterative smoothers become slow to converge as the conditioning of the iteration matrix worsens [331]. For values not large enough, the residual-based jumps across element boundaries may not be sufficiently penalised and coercivity of the bilinear could not be guaranteed. In the literature, it was found that penalty parameters were assigned values based upon experience and numerical experiments [332]. More often though, the penalty parameters are taken as some user-dependent sufficiently large constant that scales under refinement of the mesh [333].

The rest of the chapter is organised as follows: First, the governing equations are presented, the multi-group NDE, along with their adjoint equations and appropriate BC. This is followed by a CBG-IGA discretisation, where higher-order continuity within a patch is sacrificed in order to facilitate the wholly-local nature of an AMR-$h(p)$ strategy. Any eventual *hanging-nodes* are addressed by establishing constraint equations across irregular patch interfaces, which are based upon master-slave relationships and the conservative interpolation between surface meshes. A similar Galerkin projection was presented in Chapter 3 for the conservative interpolation between volume meshes when evaluating the group-to-group source terms on energy-dependent meshes and the interpolation-based error measures. A separate SIP-DG-IGA discretisation follows, where *optimised* penalty parameters are calculated for general element types, based upon the stiffness-matrix method by Owens et al [334], derived from a coercivity analysis of the bilinear form. Finally, the discretisations are verified and the convergence properties analysed by the MMS; and numerical results are presented for a couple of nuclear benchmarks.

4.2 The Multi-Group Neutron Diffusion Equations

4.2.1 The Direct (Primal) Equations

The direct (primal) multi-group NDE are posed on the closure of a connected and bounded domain, $V \subset \mathbb{R}^d$, in d-dimensional Euclidean space for $d = 2, 3$. The strong form of the problem is cast as solving $\phi^g \colon V \to \mathbb{R}$ for $g = 1, \ldots, G$:

$$\begin{cases} -\vec{\nabla} \cdot D^g(\vec{r})\vec{\nabla}\phi^g(\vec{r}) + \Sigma_r^g(\vec{r})\phi^g(\vec{r}) = q^g(\vec{r}), & \vec{r} \in V; \quad (4.1a) \\ \phi^g(\vec{r}) = g_D^g, & \vec{r} \in \partial V_D; \quad (4.1b) \\ \vec{\nabla}\phi^g(\vec{r}) \cdot \vec{n}(\vec{r}) = g_N^g, & \vec{r} \in \partial V_N; \quad (4.1c) \end{cases}$$

where Dirichlet and Neumann-type BC, Equation 4.1b and Equation 4.1c, respectively, are prescribed on the external physical boundaries; i.e. $\partial V_{ext} = \partial V_D \cup \partial V_N$, respectively; and the normal vector, $\vec{n}$, is outward-pointing. The defining equations can be recast in block-like group operator notation:

$$L\phi = q, \quad (4.2a)$$

$$= S\phi + F\phi + q_x; \quad (4.2b)$$

with the scalar neutron flux, $\phi = [\phi^1, \ldots, \phi^G]^T$; the extraneous fixed sources, $q_x = [q_x^1, \ldots, q_x^G]^T$; the diagonal loss operator, L, comprised of within-group losses due to diffusive-leakage and total collision, L^g, Equation 4.3a; the scatter operator, S, comprised of group-to-group scatter terms, S^g, Equation 4.3b; and the fission operator, F, comprised of multi-group fission terms, F^g, Equation 4.3c.

$$L^g\phi := -\vec{\nabla} \cdot D^g\vec{\nabla}\phi^g + \Sigma_r^g\phi^g; \quad (4.3a)$$

$$S^g\phi := \sum_{g' \neq g} \Sigma_s^{g' \to g}\phi^{g'}; \quad (4.3b)$$

$$F^g\phi := \chi^g \sum_{g'} \nu^{g'}\Sigma_f^{g'}\phi^{g'}. \quad (4.3c)$$

Defining the operator, A, such that:

$$A\phi = (L - S)\phi, \quad (4.4)$$

one can distinguish fixed-source problems:

$$(A - F)\phi = q_x, \quad (4.5)$$

from eigenvalue ones:

$$A\phi = \frac{1}{\mathrm{k_{eff}}}F\phi. \quad (4.6)$$

4.2.2 The Adjoint (Dual) Equations

Goal-based (DWR-)AMR relies upon duality arguments, therefore, one introduces the adjoint (dual) problem; q.v. Section 2.2.3. Mathematically, the dual problem should be adjoint-consistent with the primal one; q.v. Equation 2.12. By linearity of the inner-product, adjoint-consistency can be satisfied in each of the operators defined in Equations 4.3; i.e. [75]:

$$\sum_{g=1}^{G}\left[\int_{V}\phi^{g\dagger}(\vec{r})B^{g}\phi^{g}(\vec{r})\,dV\right]\equiv\sum_{g=1}^{G}\left[\int_{V}\phi^{g}(\vec{r})B^{g\dagger}\phi^{g\dagger}(\vec{r})\,dV\right]. \tag{4.7}$$

The dual problem is posed in the same domain, $V\in\mathbb{R}^{d}$, as the primal one; its strong form cast in multi-group operator notation as finding $\phi^{g\dagger}\colon V\to\mathbb{R}$ for $g=G,\ldots,1$:

$$L^{\dagger}\phi^{\dagger}=q^{\dagger}, \tag{4.8a}$$

$$=S^{\dagger}\phi^{\dagger}+F^{\dagger}\phi^{\dagger}+q_{x}^{\dagger}; \tag{4.8b}$$

with the scalar neutron importance, $\phi^{\dagger}=[\phi^{1\dagger},\ldots,\phi^{G\dagger}]^{T}$; the adjoint source, $q_{x}^{\dagger}=[q_{x}^{1\dagger},\ldots,q_{x}^{G\dagger}]^{T}$; and the loss, scatter and fission adjoint operators, $L^{\dagger}$, $S^{\dagger}$ and $F^{\dagger}$, respectively:

$$L^{g\dagger}\phi^{\dagger}:=-\vec{\nabla}\cdot D^{g}\vec{\nabla}\phi^{g\dagger}+\Sigma_{r}^{g}\phi^{g\dagger}; \tag{4.9a}$$

$$S^{g\dagger}\phi^{\dagger}:=\sum_{g'\neq g}\Sigma_{s}^{g\to g'}\phi^{g'}; \tag{4.9b}$$

$$F^{g\dagger}\phi^{\dagger}:=\nu^{g}\Sigma_{f}^{g}\sum_{g'}\chi^{g'}\phi^{g'}. \tag{4.9c}$$

The scatter and the fission operators are not self-adjoint; essentially their adjoint operators are obtained by transposition and the interchanging of arguments of the energy-integral kernels [125]. The loss operator is formally self-adjoint, $(L\phi,\phi^{\dagger})=(\phi,L\phi^{\dagger})$; however, the diffusive-leakage term necessitates an appropriate prescription of adjoint BC. Equation 4.8 is multiplied by the scalar neutron flux and subtracted from Equation 4.2, which is first multiplied by the scalar neutron importance; and by linearity of the inner-product:

$$(L\phi,\phi^{\dagger})-(\phi,L^{\dagger}\phi^{\dagger})=(q,\phi^{\dagger})-(\phi,q^{\dagger}). \tag{4.10}$$

The inner-product formed with the loss operator is written explicitly in the first line of Equation 4.11; where Green's first identity of the first term is revealed and the spatial dependence suppressed to give the second line, to which, again, the Divergence Theorem is applied to the second term to give the

third line, on which the integrals over the surface and those over the volume are collected:

$$(L\phi, \phi^\dagger) = \sum_{g=1}^{G} \left[- \int_V \phi^{g\dagger}(\vec{r}) \vec{\nabla} \cdot D^g(\vec{r}) \vec{\nabla} \phi^g(\vec{r}) \, dV + \int_V \phi^{g\dagger}(\vec{r}) \Sigma_r^g(\vec{r}) \phi^g(\vec{r}) \, dV \right], \tag{4.11a}$$

$$= \sum_{g=1}^{G} \left[- \int_{\partial V} \phi^{g\dagger} D^g \vec{\nabla} \phi^g \cdot \vec{n} \, dS + \int_V \vec{\nabla} \phi^{g\dagger} D^g \vec{\nabla} \phi^g \, dV + \int_V \phi^{g\dagger} \Sigma_r^g \phi^g \, dV \right], \tag{4.11b}$$

$$= \sum_{g=1}^{G} \left[\int_{\partial V} \phi^g D^g \vec{\nabla} \phi^{g\dagger} \cdot \vec{n} - \phi^{g\dagger} D^g \vec{\nabla} \phi^g \cdot \vec{n} \, dS \right] + (\phi, L^\dagger \phi^\dagger). \tag{4.11c}$$

For a set of BC for the primal problem, there is a set for the dual one:

$$\sum_{g=1}^{G} \left[\int_{\partial V} \phi^g D^g \vec{\nabla} \phi^{g\dagger} \cdot \vec{n} - \phi^{g\dagger} D^g \vec{\nabla} \phi^g \cdot \vec{n} \, dS \right] = 0. \tag{4.12}$$

In which case, it simplifies the problem to restrict the general Dirichlet and Neumann BC, Equation 4.1b and Equation 4.1c, respectively, to a certain few that are commonly used in reactor physics and radiation shielding problems. For the set of zero-flux, reflective and vacuum BC, prescribed on the physical boundaries, ∂V_0, ∂V_R and ∂V_V respectively:

$$\begin{cases} \phi^g(\vec{r}) = 0, & \vec{r} \in \partial V_0; \quad (4.13a) \\ -D^g(\vec{r}) \vec{\nabla} \phi^g(\vec{r}) \cdot \vec{n} = 0, & \vec{r} \in \partial V_R; \quad (4.13b) \\ -D^g(\vec{r}) \vec{\nabla} \phi^g(\vec{r}) \cdot \vec{n} = \frac{1}{2} \phi^g(\vec{r}), & \vec{r} \in \partial V_V. \quad (4.13c) \end{cases}$$

The corresponding set of adjoint BC are prescribed over the same boundaries:

$$\begin{cases} \phi^{g\dagger}(\vec{r}) = 0, & \vec{r} \in \partial V_0; \quad (4.14a) \\ D^g(\vec{r}) \vec{\nabla} \phi^{g\dagger}(\vec{r}) \cdot \vec{n} = 0, & \vec{r} \in \partial V_R; \quad (4.14b) \\ D^g(\vec{r}) \vec{\nabla} \phi^{g\dagger}(\vec{r}) \cdot \vec{n} = \frac{-1}{2} \phi^{g\dagger}(\vec{r}), & \vec{r} \in \partial V_V; \quad (4.14c) \end{cases}$$

where the vacuum BC has been considered as a general Robin BC of the form: $\alpha\phi + \vec{\nabla}\phi \cdot \vec{n} = g$, which assumes that the incoming partial current of neutrons as well as the outgoing partial current of neutron importance are both vanishing. In turn, the BC from the sets of Equations 4.13 and Equations 4.14 satisfy Equation 4.12; they are compatible for use in an eigenvalue problem; and indeed, a physical interpretation of Equations 4.14 can be deduced when deriving Equation 4.8 from the FONT equation [335].

In a similar manner to Equation 4.4 for the primal problem, the operator $A^\dagger$ is defined such that:

$$A^\dagger \phi^\dagger = (L^\dagger - S^\dagger) \phi^\dagger. \tag{4.15}$$

Again, one distinguishes fixed-source problems:

$$(A^\dagger - F^\dagger)\phi^\dagger = q_x^\dagger, \tag{4.16}$$

from eigenvalue problems:

$$A^\dagger\phi^\dagger = \frac{1}{\mathrm{k}_{\mathrm{eff}}^\dagger}F^\dagger\phi^\dagger. \tag{4.17}$$

Under the assumption that the property of (semi-discrete) adjoint-consistency holds, it should be clear that $\mathrm{k}_{\mathrm{eff}}^\dagger = \mathrm{k}_{\mathrm{eff}}$; i.e. the Rayleigh quotient:

$$\mathrm{k}_{\mathrm{eff}}^\dagger = \frac{(\phi, F^\dagger\phi^\dagger)}{(\phi, A^\dagger\phi^\dagger)} = \frac{(F\phi, \phi^\dagger)}{(A\phi, \phi^\dagger)} = \mathrm{k}_{\mathrm{eff}}; \tag{4.18}$$

where $|1/\mathrm{k}_{\mathrm{eff}}|$ is determined after each $(i)^{th}$ *power* iteration; q.v. Section 2.8:

$$\mathrm{k}_{\mathrm{eff}}^{(i)} = \frac{(F\phi^{(i)}, \phi^\dagger)}{(A\phi^{(i)}, \phi^\dagger)} = \frac{(F\phi^{(i)}, \phi^\dagger)}{\frac{1}{k_{\mathrm{eff}}^{(i+1)}}(F\phi^{(i+1)}, \phi^\dagger)}; \tag{4.19a}$$

$$\implies \frac{1}{\mathrm{k}_{\mathrm{eff}}^{(i+1)}} = \frac{1}{\mathrm{k}_{\mathrm{eff}}^{(i)}}\frac{(F\phi^{(i)}, \phi^\dagger)}{(F\phi^{(i+1)}, \phi^\dagger)}; \tag{4.19b}$$

where $\phi^\dagger$ is not obliged to be the adjoint solution in the above iteration; and $\lim_{(i)\to\infty} k_{eff}^{(i)} = \mathrm{k}_{\mathrm{eff}}$.

4.3 A Continuous Bubnov-Galerkin Isogeometric Discretisation

First, one introduces formally some abstract function spaces. The Lebesgue space, $L_2(V)$, of some measurable functions, u, that are square-integrable on the bounded domain, V, is defined as:

$$L_2(V) \coloneqq \{u \text{ measureable} : \|u\|_{LE(V)} < \infty\}. \tag{4.20}$$

This is a real linear Hilbert space furnished with an inner-product and norm, respectively:

$$(u, v)_{L_2(V)} \coloneqq \int_V vu\, dV; \tag{4.21a}$$

$$\|u\|_{L_2(V)} \coloneqq (u, u)^{1/2}_{L_2(V)}. \tag{4.21b}$$

The s-order Sobolev space, $H^s(V) \subset L_2(V)$, is defined for those functions belonging to $L_2(V)$, that also admit all weak partial derivatives, D^α, up to order s, belonging to $L_2(V)$; for $\vec{r} = (x_1, x_2, x_3) = (x, y, z)$ and $D^\alpha = \prod_{i=1}^d D_i^{\alpha_i}$ for $D_i^{\alpha_i} = \partial^{\alpha_i}/\partial_{x_i}^{\alpha_i}$; and $|\alpha| = \sum_i^d \alpha_i$ for some multi-index $\alpha \in \mathbb{R}^d$:

$$H^s(V) \coloneqq \{u \in L_2(V) \colon D^\alpha u \in L_2(V), \forall\, 0 \leq |\alpha| \leq s\}. \tag{4.22}$$

The Sobolev space, $H^s(V)$, is a Hilbert space equipped with an inner-product and norm, respectively:

$$(u, v)_{H^s(V)} \coloneqq \sum_{|\alpha|\leq s} (D^\alpha v, D^\alpha u)_{L_2(V)}; \tag{4.23a}$$

$$\|u\|_{H^s(V)} \coloneqq (u, u)^{1/2}_{H^s(V)}. \tag{4.23b}$$

The restriction of such functions and their normal derivatives with support over the domain, V, are defined over the boundary, ∂V, by trace operators, $\gamma_i \colon H^s(V) \to H^{s-i-1/2}(\partial V)\ \forall\ 0 \leq i < s$; viz. the Dirichlet trace, $\gamma_0 u = u|_{\partial V}$, is well defined in $L_2(\partial V)$ for $s > 3/2$ [222].

Now, the point of departure is the semi-discrete multi-group NDE, Equation 4.2, and their adjoint equations, Equation 4.8. Both the primal and the dual problem constitute systems of equations that are solved for the multi-group components $\phi^g \colon V \to \mathbb{R}$ and $\phi^{g\dagger} \colon V \to \mathbb{R}$, respectively; with BC, such as those in Equations 4.13 and Equations 4.14, respectively, prescribed over separate sections of the external boundary.

Without loss of generality, an emphasis is placed upon manipulating the direct equations; and the same discretisation is applied to the adjoint equations. The strong form of the defining equations is reformulated into a suitable abstract variational boundary value problem (VBVP), more generally dubbed the *weak* form, which itself is obtained by multiplying Equation 4.2 by a test function, $w \in \mathcal{V}(V)$, and then by integrating over the entire domain in the real physical space, V. For $q_x \in L_2(V)$, rather than searching for a smooth solution in the space of continuous functions, a solution is sought in the s-order Sobolev space, Equation 4.22, where derivatives are understood in a weak distributional sense. The (semi-)analytical solution is approximated by some infinite linear expansion of trial functions, $u(\vec{r}) \in \mathcal{U}(V) \subseteq H^s(V)$, and control-variables, ϕ_i; q.v. Equation 2.60:

$$\phi(\vec{r}) \approx \sum_{i=1}^{\infty} \phi_i\ u_i(\vec{r}), \quad \forall \vec{r} \in V. \tag{4.24}$$

In the integral statement defined over V, substitution of the approximation made above into the integrand, Equation 4.2, yields a weighted-residual, which is forced to vanish everywhere and thus, the residual is made orthogonal to every member of the set of test functions. As it stands, the weighted-residual expression necessitates that the trial functions are strictly second-order differentiable, yet this is relaxed by application of the divergence theorem to the double-differential diffusion term. Now, the test and trial functions are only required to be first-order differentiable and as per a Bubnov-Galerkin approach, they are both taken from the same space and the residual is made orthogonal to the approximate solution; i.e. $\mathcal{U}(V), \mathcal{V}(V) \subseteq H^1(V)$ [336].

Appearing in the weak form is a surface integral term over the entire physical boundary, which lends itself to the introduction of *natural* and *essential* BC. Naturally satisfied, in the present case, is the homogeneous Neumann-type BC, or the reflective BC; q.v. Equation 4.13b. The vacuum BC is implemented in a weak sense by the simple substitution of Equation 4.13c into the VBVP. Nitsche's penalty method is employed to enforce weakly Equation 4.13a at the Dirichlet boundaries by the

following penalisation, for $\mu_f > 0$ [325]:

$$\sum_{\partial V_f \in \partial V_D} \mu_f \int_{\partial V_f} w(\vec{r})\phi(\vec{r})\, dS = \sum_{\partial V_f \in \partial V_D} \mu_f \int_{\partial V_f} w(\vec{r})\, g_D\, dS. \tag{4.25}$$

The domain in the real physical space, V, is decomposed into E-number of non-degenerate, non-overlapping NURBS patches, V_e; i.e. $V = \bigcup_{e=1}^{E} V_e$. Within any given patch, the material is assumed to be homogeneous and its properties smeared and uniform; viz. the nuclear data is assigned as piecewise constants; e.g. $\Sigma_{r,e} = \Sigma_r(\vec{r})|_{V_e}$, $\forall \vec{r} \in V_e \in \mathcal{T}$. The finite union of macro-elements forms the general physical mesh, $\mathcal{T}(V) = \{V_e\}_{e=1}^{E}$. For the sake of simplicity in developing the spatial discretisation, one assumes a single energy-independent mesh. The intersection of *two* patches belonging to the same physical mesh forms either an empty set (disjoint); an edge shared by indirect neighbours, this being a NURBS vertex in 2D or a NURBS edge in 3D; or a face shared by direct neighbours, this being a NURBS edge in 2D or a NURBS surface in 3D. A patch interface formed at the boundaries of *two* adjoining patches that share a common face thus belongs to the set of interior faces, ∂V_{int}; and therefore, an arbitrary patch face, ∂V_f, belongs to the complete set $\partial V_f \in \partial V = \partial V_{int} \cup \partial V_{ext}$.

The integral terms over the entire domain are obtained as a summation of dense local contributions from over each patch, $V_e \in \mathcal{T}$, in a systematic *assembly* process. A virgin patch with no internal refinement structure can be decomposed into more non-vanishing knot-spans, $V_{\hat{e}}$, via knot insertion; i.e. $V_e = \bigcup_{\hat{e}} V_{\hat{e}}$. Each new micro-element retains its own homogeneity and well-defined material properties, as inherited from the macro-element; viz. $\Sigma|_{V_{\hat{e}}} = \Sigma|_{V_e}\ \forall\ V_{\hat{e}} \in V_e$. Integration is always performed over the smallest entity and more rigorously the assembly process loops over all non-vanishing knot-spans belonging to the mesh. Without loss of generality, however, the subscript e is retained, as pertaining to the element over which the integral quantity is evaluated; and the distinction between macro-element and the micro-element will be made when necessary.

The weak form of Equation 4.5 is to find $\phi \in \mathcal{V}(V) \subseteq H^1(V)$, such that:

$$a(\phi, w) = b(w), \quad \forall\ w \in \mathcal{V}(V). \tag{4.26}$$

The bilinear form, $a(\cdot,\cdot)\colon \mathcal{V}(V) \times \mathcal{V}(V) \to \mathbb{R}$, retains its block-like structure and comprises of a weak diffusive-leakage and group-removal term, a_L, defined over volumes; a group-to-group scattering term, a_S; a fission term, a_F; as well as a term that incorporates the BC prescribed on the external physical boundary, a_{ext}:

$$a(\phi, w) = a_L(\phi, w) - a_S(\phi, w) - a_F(\phi, w) + a_{ext}(\phi, w). \tag{4.27}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$\begin{aligned} a_L(\phi, w) &= \sum_{g=1}^{G} \left[\sum_{V_e \in \mathcal{T}} \int_{V_e} \vec{\nabla} w \cdot D_e^g \vec{\nabla} \phi^g + w \Sigma_{r,e}^g \phi^g \, dV \right], \\ &= \sum_{g=1}^{G} a_L^g(\phi^g, w); \end{aligned} \tag{4.28a}$$

$$\begin{aligned} a_S(\phi, w) &= \sum_{g=1}^{G} \left[\sum_{g' \neq g}^{G} \sum_{V_e \in \mathcal{T}} \int_{V_e} w \Sigma_{s,e}^{g' \to g} \phi^{g'} \, dV \right], \\ &= \sum_{g=1}^{G} a_S^g(\phi, w); \end{aligned} \tag{4.28b}$$

$$\begin{aligned} a_F(\phi, w) &= \sum_{g=1}^{G} \left[\sum_{g'=1}^{G} \sum_{V_e \in \mathcal{T}} \chi_e^g \int_{V_e} w \nu_e^{g'} \Sigma_{f,e}^{g'} \phi^{g'} \, dV \right], \\ &= \sum_{g=1}^{G} a_F^g(\phi, w); \end{aligned} \tag{4.28c}$$

$$\begin{aligned} a_{ext}(\phi, w) &= \sum_{g=1}^{G} \left[\sum_{V_e \in \mathcal{T}} \sum_{\partial V_f \in V_e \cap \partial V_0} \mu_f \int_{\partial V_f} w \phi^g \, dS + \sum_{\partial V_f \in V_e \cap \partial V_V} \frac{1}{2} \int_{\partial V_f} w \phi^g \, dS \right], \\ &= \sum_{g=1}^{G} a_{ext}^g(\phi^g, w). \end{aligned} \tag{4.28d}$$

The linear functional, $b(\cdot)\colon \mathcal{V}(V) \to \mathbb{R}$, comprises only of fixed sources, since $g_D = 0$:

$$\begin{aligned} b(w) &= \sum_{g=1}^{G} \left[\sum_{V_e \in \mathcal{T}} \int_{V_e} w q_{x,e}^g \, dV \right], \\ &= \sum_{g=1}^{G} b^g(w). \end{aligned} \tag{4.29}$$

For multiplying systems, $q_x = 0$, the weak form of Equation 4.6 is to find $\phi \in \mathcal{V}(V)$, such that:

$$\tilde{a}(\phi, w) = \frac{1}{\mathrm{k_{eff}}} a_f(\phi, w), \quad \forall \, w \in \mathcal{V}(V); \tag{4.30}$$

where $\tilde{a}$ is the original bilinear form in Equation 4.27, a, less the fission term; i.e.:

$$\tilde{a} = a + a_F = a_L - a_S + a_{ext}. \tag{4.31}$$

The *discrete* variational problem is formulated by restricting the function space to a finite-dimensional subset; i.e. $\hat{\mathcal{V}}_{h,p}(V) \subset \mathcal{V}(V)$. As per the isoparametric concept, this space is chosen as that spanned by the same N-number of NURBS used to describe the mesh, $\mathcal{T}$; where the NURBS, R, Equations 3.7, are characterised by the mesh spacing, h, as well as the polynomial degree of the basis, p; i.e.:

$$\begin{cases} \left[\hat{\mathcal{V}}_{h,p}(V)\right]^T = \mathrm{span}\left\{R_i^p\right\}_{i=1}^N, \\ \text{with } \mathrm{supp}\left(R^p\right) = \left\{\vec{r} \in V\colon R_i^p(\vec{r}) \neq 0 \text{ for } i = 1, \ldots, N\right\}. \end{cases} \tag{4.32}$$

Indeed, the tensor-product structure of the parametric domain permits the NURBS to be constructed from different polynomials in different parametric directions; for simplicity, however, they will be characterised by a single p and unless necessary, it will be suppressed as a superscript. Accordingly, Equation 4.24 is further approximated by truncating the expansion to the Nth term:

$$\phi(\vec{r}) \approx \hat{\phi}(\vec{r}) = \sum_{i=1}^{N} \phi_i R_i(\vec{r}). \tag{4.33}$$

Clearly $\hat{\mathcal{V}}_{h,p}(V)$ is chosen such that the weighted-residual becomes vanishing with respect to variations in the control-variables, ϕ_i, as the dimension of the finite-dimensional space increases; i.e. $\hat{\phi} \to \phi$ as $N \to \infty$. The weighted-residual statement must be satisfied for each member in the weighting set; and since the space of trial and test functions coincide, this forms the same number of algebraic equations as there are unknowns per energy group; i.e. $DoF^g = \dim \hat{\mathcal{V}}_{h,p}(V) = N$ for $g = 1, \ldots, G$; and hence, $DoF = \sum_g DoF^g = G \times N$. Consequently, for the analysis over a single energy-independent mesh, the discrete problem is posed as a system of $(GN \times GN)$-matrix equations.

The space $\hat{\mathcal{V}}_{h,p}(V)$ is the collection of global functions whose restriction to a patch lives in the tensor-product parametric domain $\hat{V}_e$. Consequently, NURBS in the real physical space also have local support over a given patch; i.e. $\hat{\mathcal{V}}^e_{h,p}(V_e) \subset \hat{\mathcal{V}}_{h,p}(V)\ \forall V_e \in \mathcal{T}$. This piecewise definition assigns the functions a narrow base and results in the banded structure of the assembled matrices. It also introduces discontinuity into the basis; yet the $H^1(V)$ space demands strict conformity of the numerical solution; viz. global continuity across V. Dirichlet traces are well defined in $L_2(\partial V_e)$ for $s > ½$; i.e. $\gamma_0 \colon H^1(V_e) \to \mathrm{H}^{½}(\partial V_e)$; and the extension of functions defined up to the boundary of a patch are made continuous across an interface by virtually *gluing* them together. This is trivial in the case of regular meshes where the control-variables at patch interfaces exist in a one-to-one correspondance; i.e. $R(\vec{r})|_{V_e} = R\prime(\vec{r})|_{V_{e\prime}}\ \forall \vec{r} \in \partial V_f \in V_e \cap V_{e\prime}$ and $\forall V_e, V_{e\prime} \in \mathcal{T}$. The basis is reduced to C^0 continuity across any patch interface; v. Figure 3.2.

Alternatively, continuity across the domain could be enforced weakly in a discontinuous setting; q.v. Section 4.4. Akin to the weak imposition of the Dirichlet BC, Nitsche's penalty method is applied to the interior of the mesh as per an IP-DG scheme; where the errors across patch interfaces are penalised in a similar manner to that in Equation 4.25. Bazilevs et al investigated the weak imposition of Dirichlet BC, in the application to fluid dynamics and found that such an approach did aid in eliminating some of the spurious oscillations associated with unresolved boundary layer phenomena encountered where BC had previously been imposed in a strong sense [337]. In the same way, for both elliptic and hyperbolic systems, allowing for reduced continuity in the solution can aid in reducing spurious oscillations, especially in the presence of strong material heterogeneity [234]. This provides the rationale for the investigation of wholly-local AMR-$h(p)$ strategies that consider at most C^0 continuity.

Casting back to the (semi-)analytical problem, it is evident that the scattering and the fission terms, Equation 4.28b and Equation 4.28c, respectively, are asymmetric; e.g. $a_S(u,v) \neq a_S(v,u)$. Nevertheless, the bilinear form, a, does possess the properties of continuity, or boundedness, Equation 4.34, and coercivity, or stability and ellipticity, Equation 4.35. The latter is often simply guaranteed by nature of the problem; viz. physical non-negative nuclear cross-section data.

For $u, v \in \mathcal{V}(V)$ and some constant, $C_\alpha > 0$, the bilinear form is bounded if:

$$a(u,v) \leq C_\alpha \|v\| \|u\|; \tag{4.34}$$

and for $u \in \mathcal{V}(V)$ and some constant, $C_\beta > 0$, the bilinear form is coercive if:

$$a(u,u) \geq C_\beta \|u\|^2. \tag{4.35}$$

As per the Lax-Milgram Lemma, these properties of continuity and more specifically coercivity are sufficient conditions to assert well-posedness; viz. the admission of a unique and stable solution [338]. Since $\hat{\mathcal{V}}_{h,p}(V) \subset \mathcal{V}$, the discrete problem inherits well-posedness and other properties from the (semi-)analytical weak-problem [222]. Céa's Lemma of quasi-best approximation for conforming and consistent discretisations, Equation 4.36, states that up to a constant, as those defined by continuity and coercivity, C_α and C_β, respectively, the set $\hat{\mathcal{V}}_{h,p}(V)$ is that chosen from all possible finite sets of function spaces that minimises the error in the numerical solution, $\hat{\phi} \in \hat{\mathcal{V}}_{h,p}(V)$, with respect to the true solution, $\phi \in \mathcal{V}(V)$, as measured in the energy-norm; i.e. for some $\hat{v} \in \hat{\mathcal{V}}_{h,p}(V)$ [339]:

$$\begin{aligned}
C_\beta \|\phi - \hat{\phi}\|^2 &\leq a(\phi - \hat{\phi}, \phi - \hat{\phi}), && \text{(4.36a)}\\
&\leq a(\phi - \hat{\phi}, \phi - \hat{v}) + a(\phi - \hat{\phi}, \hat{v} - \hat{\phi}), && \text{(4.36b)}\\
&\leq a(\phi - \hat{\phi}, \phi - \hat{v}), && \text{(4.36c)}\\
&\leq C_\alpha \|\phi - \hat{v}\| \|\phi - \hat{\phi}\|; && \text{(4.36d)}\\
\implies \|\phi - \hat{\phi}\| &\leq \frac{C_\alpha}{C_\beta} \inf_{\hat{v} \in \hat{\mathcal{V}}_{h,p}(V)} \|\phi - \hat{v}\|; && \text{(4.36e)}
\end{aligned}$$

where the first line, Equation 4.36a, expresses in the bilinear form, a, the property of coercivity, Equation 4.35; and the linearity of each argument of a is exhibited in the second line by introducing $\hat{v}$; where the second term is made to be vanishing because the error in the numerical solution is orthogonal to the chosen subspace; i.e. since $\hat{\mathcal{V}}_{h,p}(V) \subset \mathcal{V}(V)$:

$$\begin{aligned}
a(\phi - \hat{\phi}, w) &= a(\phi, w) - a(\hat{\phi}, w), \quad \forall w \in \mathcal{V}(V), && \text{(4.37a)}\\
&= b(w) - b(w), \quad \forall w \in \mathcal{V}(V), && \text{(4.37b)}\\
&= 0. && \text{(4.37c)}
\end{aligned}$$

The fourth line, Equation 4.36d, is obtained by expression of continuity, Equation 4.34, in the third;

and Céa's Lemma is obtained, Equation 4.36e, by taking the infimum over all possible $\hat{v} \in \hat{\mathcal{V}}_{h,p}(V)$. A strong definition of consistency has been enforced, yet Legendre-Gauss quadrature rules will be employed to evaluate all integral quantities; q.v. Section 3.2.3. Consequently, one expects the discretisation to be asymptotically consistent.

The system of equations is solved via the SI procedure; q.v. Section 2.8. Weakly coupled in energy, the within-group equations are solved sequentially with the most up-to-date available information for the group-to-group sources; whereby, the source terms are now linear functionals that can be moved to the right-hand side of the equation. In the presence of extraneous fixed sources, $\hat{\phi}^g \in \hat{\mathcal{V}}_{h,p}(V)$ is sought for $g = 1, \ldots, G$:

$$a_L^g(\hat{\phi}^g, \hat{w}) + a_{ext}^g(\hat{\phi}^g, \hat{w}) = a_S^g(\hat{\phi}, \hat{w}) + a_F^g(\hat{\phi}, \hat{w}) + b^g(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V). \tag{4.38}$$

The left-hand side, $a_L^g + a_{ext}^g$, is SPD and Equation 4.38 can be solved using efficient iterative smoothers, such as the PCG method; and additional savings can be made when storing symmetric matrices.

Employing the PI procedure to converge $|1/\mathrm{k_{eff}}|$, the eigenvalue problem is cast as finding $\hat{\phi}^g \in \hat{\mathcal{V}}_{h,p}(V)$ for $g = 1, \ldots, G$:

$$a_L^g(\hat{\phi}^g, \hat{w}) + a_{ext}^g(\hat{\phi}^g, \hat{w}) = a_S^g(\hat{\phi}, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a_F^g(\hat{\phi}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V). \tag{4.39}$$

In kind, the discrete variational statement of the adjoint problem in the presence of fixed adjoint sources, $q_x^\dagger \neq 0$, is cast as finding $\hat{\phi}^{g\dagger} \in \hat{\mathcal{V}}_{h,p}(V)$ for $g = G, \ldots, 1$:

$$a_L^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) + a_{ext}^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) = a_S^{g\dagger}(\hat{\phi}^\dagger, \hat{w}) + a_F^{g\dagger}(\hat{\phi}^\dagger, \hat{w}) + b^{g\dagger}(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V); \tag{4.40}$$

and the adjoint eigenvalue problem is cast as finding $\hat{\phi}^{g\dagger} \in \hat{\mathcal{V}}_{h,p}(V)$ for $g = G, \ldots, 1$:

$$a_L^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) + a_{ext}^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) = a_S^{g\dagger}(\hat{\phi}^\dagger, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a_F^{g\dagger}(\hat{\phi}^\dagger, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V); \tag{4.41}$$

where the adjoint bilinear form, $a^\dagger(\cdot,\cdot)\colon \mathcal{V}(V) \times \mathcal{V}(V) \to \mathbb{R}$, is defined by:

$$a^\dagger(\phi^\dagger, w) = a_L^\dagger(\phi^\dagger, w) - a_S^\dagger(\phi^\dagger, w) - a_F^\dagger(\phi^\dagger, w) + a_{ext}^\dagger(\phi^\dagger, w). \tag{4.42}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$\begin{aligned} a_L^\dagger(\phi^\dagger, w) &= \sum_{g=1}^{G}\left[\sum_{V_e\in\mathcal{T}}\int_{V_e}\vec{\nabla}w\cdot D_e^g\vec{\nabla}\phi^{g\dagger}+w\Sigma_{r,e}^g\phi^{g\dagger}\,dV\right], \\ &= \sum_{g=1}^{G} a_L^{g\dagger}(\phi^{g\dagger}, w); \end{aligned} \tag{4.43a}$$

$$\begin{aligned} a_S^\dagger(\phi^\dagger, w) &= \sum_{g=1}^{G}\left[\sum_{g\prime\neq g}^{G}\sum_{V_e\in\mathcal{T}}\int_{V_e} w\Sigma_{s,e}^{g\rightarrow g\prime}\phi^{g\prime\dagger}\,dV\right], \\ &= \sum_{g=1}^{G} a_S^{g\dagger}(\phi^\dagger, w); \end{aligned} \tag{4.43b}$$

$$\begin{aligned} a_F^\dagger(\phi^\dagger, w) &= \sum_{g=1}^{G}\left[\sum_{g\prime=1}^{G}\sum_{V_e\in\mathcal{T}}\nu_e^g\Sigma_{f,e}^g\int_{V_e} w\chi_e^{g\prime}\phi^{g\prime\dagger}\,dV\right], \\ &= \sum_{g=1}^{G} a_F^{g\dagger}(\phi^\dagger, w); \end{aligned} \tag{4.43c}$$

$$\begin{aligned} a_{ext}^\dagger(\hat{\phi}^\dagger, w) &= \sum_{g=1}^{G}\left[\sum_{V_e\in\mathcal{T}}\sum_{\partial V_f\in V_e\cap\partial V_0}\mu_f\int_{\partial V_f} w\hat{\phi}^{g\dagger}\,dS+\sum_{\partial V_f\in V_e\cap\partial V_V}\frac{1}{2}\int_{\partial V_f} w\hat{\phi}^{g\dagger}\,dS\right], \\ &= \sum_{g=1}^{G} a_{ext}^{g\dagger}(\hat{\phi}^{g\dagger}, w); \end{aligned} \tag{4.43d}$$

and the adjoint linear functional, $b^\dagger(\cdot)\colon \mathcal{V}(V)\rightarrow\mathbb{R}$, is defined as:

$$\begin{aligned} b^\dagger(w) &= \sum_{g=1}^{G}\left[\sum_{V_e\in\mathcal{T}}\int_{V_e} wq_{x,e}^{g\dagger}\,dV\right], \\ &= \sum_{g=1}^{G} b^{g\dagger}(w). \end{aligned} \tag{4.44}$$

For eigenvalue problems, one recalls that the solution is normalised after each outer iteration such that $a_F(\phi,\phi)=1$; and the adjoint solution is normalised such that $a_F(\phi,\phi^\dagger)=1$; q.v. Section 3.6.2.2.

4.3.1 Enforcing C^0 Continuity on Irregular Meshes

In the case of a standard CBG-IGA discretisation, any refinement made to the mesh propagates throughout the domain in order to satisfy conditions of compatibility. Such a global refinement strategy over large geometries induces a superfluous expenditure of computational resources since some regions will be unnecessarily resolved; in which case, IGA loses its competitive edge against the capability of FEM to simply adjust the local density of elements over unstructured meshes. The extension to include T-spline descriptions does overcome the limitation of the tensor-product refinement of NURBS; however, they also present their own limitations; e.g. their usage becomes difficult to generalise in 3D, especially where additional DoF may need to be introduced at T-junctions in order to adequately

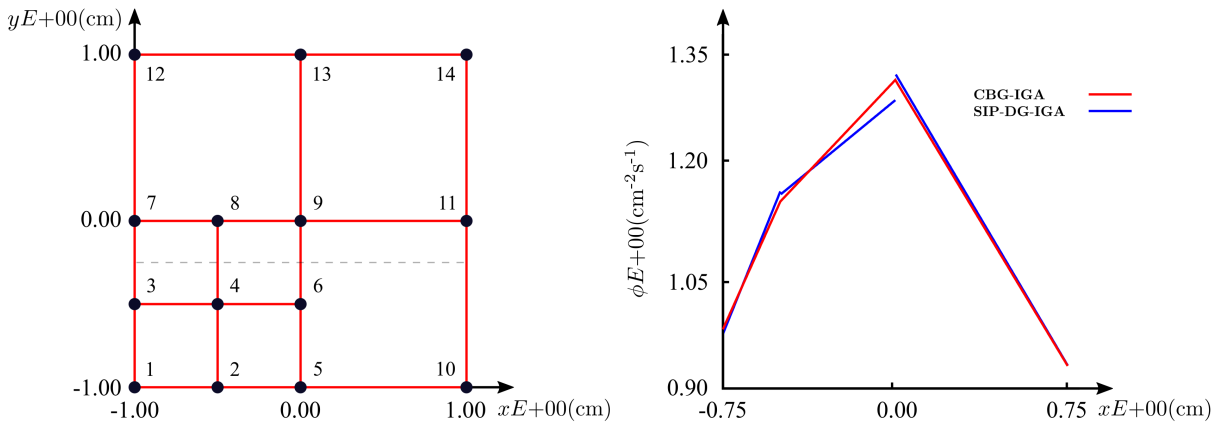


Figure 4.1: A comparison between the solutions (right) of a CBG-IGA discretisation and its DG-IGA counterpart along the dashed line at y = -0.25 cm over the irregular mesh (left) of a toy problem; a *one-speed*, (2 cm × 2 cm) homogeneous bare-square with isotropic unit-source, $\Sigma_t = 1.00$ cm^{-1}, $c = 0.7$. Continuity in the former is enforced by constraining the *hanging-nodes* 6 & 8 via a Galerkin projection; whereas continuity in the latter is enforced by means of a penalty scheme. The global enumeration of the control-variables for the CBG discretisation is given over the mesh on the left-hand side; where the patch boundaries are highlighted in red. (*Vide* the web-based version for reference to colour.)

define the basis; and it becomes unclear how T-splines should be refined [272,340]. Seeking a solution in the broken Sobolev space grants total flexibility in terms of local refinement; however, this comes at the price of significantly more DoF. Moreover, in the case of an IP scheme, the stability of the analysis depends upon the assignment of some penalty parameter; q.v. Section 4.4.1 [277]. Sticking to the theme of NURBS-based CBG-IGA spatial discretisations, one must develop a local refinement technique that can accommodate an AMR-$h(p)$ strategy and any subsequent mesh irregularities.

Figure 4.1 depicts a toy problem which has been solved for $\hat{\phi} \in \hat{\mathcal{V}}_{h,p}(V)$, where C^0 continuity has been enforced in a strong sense; and separately for $\hat{\phi}\prime \in \hat{\mathcal{V}}\prime_{h,p}(\mathcal{T})$, where continuity has be enforced in a weak sense. In the case of the latter, the basis is discontinuous and continuity has been enforced across patch interfaces by means of a penalty scheme; q.v. Section 4.4. In the case of the former, constraint equations have been applied over all irregular patch interfaces, ∂V_{sm}, those formed between a more-refined (slave) patch, V_s, and a less-refined (master) patch, V_m; i.e.:

$$\partial V_{sm} = \{\partial V_f \in V_s \cap V_m \backslash \emptyset \subseteq \partial V_{int} : \dim \hat{\mathcal{V}}_{h,p}(V_s) > \dim \hat{\mathcal{V}}_{h,p}(V_m), \forall\, V_s, V_m \in \mathcal{T}(V)\}.$$

For the ease of notation and the analysis, no distinction will be made between the constraints formulated over edges or faces, unless necessary; and to reiterate, any single intersection of patches that does not belong to an empty set comprises of a maximum of *two* patches; i.e. for $\bigcap_e V_e \backslash \emptyset \ni \partial V_f \in \partial V_{int}$ then card $(\bigcap_e V_e \backslash \emptyset) = 2$. The set of interfaces thus includes shared edges. As mentioned previously, a natural consequence of open knot-vectors is that the NURBS basis is interpolatory at patch boundaries. As such, the boundary of a NURBS surface or a NURBS volume is itself a NURBS edge or a NURBS surface, respectively, influenced solely by the NURBS functions at that given boundary and their associated control-points; the latter forms its own subset of control-points, $\mathcal{B}_f \subset \mathcal{B}$, separate

from those that do not influence the boundary, $\mathcal{B}_n \subset \mathcal{B}$, such that $\mathcal{B} = \mathcal{B}_n \cup \mathcal{B}_f$. Accordingly, those functions that have support on the boundary of a slave patch, $\partial V_f \in V_s \cap \partial V_{sm}$, span the proper subset $\hat{\mathcal{V}}^s_{h,p}(\partial V_f) \subset \hat{\mathcal{V}}_{h,p}(V_s)$; and similarly, those functions that have finite support on the boundary of a master patch, $\partial V_f \in V_m \cap \partial V_{sm}$, span $\hat{\mathcal{V}}^m_{h,p}(\partial V_f) \subset \hat{\mathcal{V}}_{h,p}(V_m)$; such that, for $e = s, m$ and $\forall V_e \in \mathcal{T}(V)$ and $\forall \partial V_f \in V_e \cap \partial V_{sm}$:

$$\begin{cases} \big[\hat{\mathcal{V}}^e_{h,p}(\partial V_f)\big]^T = \text{span}\left\{(R|_{\partial V_f})_i\right\}_{i=1}^{N^e_f}, \\ \text{with } \text{supp}\left(R|_{\partial V_f}\right) = \left\{\vec{r} \in \partial V_f \colon (R(\vec{r})|_{\partial V_f})_i \neq 0 \text{ for } i = 1, \dots, N^e_f\right\}; \end{cases} \tag{4.45}$$

where the NURBS defined locally over a given element, $V_e \in \mathcal{T}(V)$, have also been denoted by a similar superscript; i.e. by letting $N^e = \dim \hat{\mathcal{V}}_{h,p}(V_e)$, where $\hat{\mathcal{V}}_{h,p}(V_e) \subset \hat{\mathcal{V}}_{h,p}(V)$, then for $i = 1, \dots, N^e$:

$$R^e_i(\vec{r}) = \big(R(\vec{r})|_{V_e}\big)_i, \qquad \vec{r} \in V_e.$$

Continuity in the geometry across an irregular patch interface is enforced simply by applying a change of basis to the linear combination of NURBS and control-points, such that the control-points on one side of the interface, $\{\vec{B}^s_f\}$, can be represented by $\{\vec{B}^m_f\}$ on an adjoining side that exhibits a different refinement structure to the former. As per the isoparametric concept, continuity in the solution field at the irregular interface, $\{\vec{\phi}_f\}$, can be constrained in a similar fashion; i.e.:

$$\underline{\vec{B}}^s_f = \underline{\underline{K}}^{s\leftarrow m}_{\partial V} \cdot \underline{\vec{B}}^m_f; \tag{4.46a}$$

$$\underline{\phi}^s_f = \underline{\underline{K}}^{s\leftarrow m}_{\partial V} \cdot \underline{\phi}^m_f. \tag{4.46b}$$

The combined transfer matrix above, $\underline{\underline{K}}^{s\leftarrow m}_{\partial V}$, is exactly that which was introduced for the conservative interpolation between volume meshes, Equation 3.17b; however, in this instance, the H^1-Sobolev minimisation is performed over an irregular patch interface, $\forall \partial V_f \in \partial V_{sm}$, or a surface mesh. The interpolation $\prod^s_{h,p} \hat{v}^m \colon \hat{\mathcal{V}}^m(\partial V_f) \to \hat{\mathcal{V}}^s(\partial V_f)$ designates the projection of some function, $\hat{v}^m \in \hat{\mathcal{V}}^m_{h,p}(\partial V_f)$, defined on the surface mesh, $\mathcal{T}(\partial V_f)$, of a master patch onto that of a slave patch, each characterised by some different combination of h and p. For some $\partial V_f \in \partial V_{sm}$, the optimum target interpolant, $\hat{v}^s \in \hat{\mathcal{V}}^s_{h,p}(\partial V_f)$, is approximated such that:

$$\|\hat{v}^s - \hat{v}^m\|_{H^1(\partial V_f)} = \min_{\hat{v} \in \hat{\mathcal{V}}^s_{h,p}(\partial V_f)} \|\hat{v} - \hat{v}^m\|_{H^1(\partial V_f)}.$$

The basis of the slave patch should sufficiently represent any function that belongs to that of the master patch, $\dim \hat{\mathcal{V}}^s_{h,p}(\partial V_f) > \dim \hat{\mathcal{V}}^m_{h,p}(\partial V_f)$; however, it is not required that $\hat{\mathcal{V}}^s_{h,p}(\partial V_f) \supset \hat{\mathcal{V}}^m_{h,p}(\partial V_f)$; in which case, the continuity would be enforced exactly. For clarity, the combined surface transfer matrix, for which the interpolation between surface meshes is optimal in the H^1-norm, is defined by:

$$\underset{(N^s_f \times N^m_f)}{\underline{\underline{K}}^{s\leftarrow m}_{\partial V}} = \underset{(N^s_f \times N^s_f)}{[\underline{\underline{K}}^{ss}_{\partial V}]^{-1}} \cdot \underset{(N^s_f \times N^m_f)}{\underline{\underline{K}}^{sm}_{\partial V}};$$

with the consistent matrix formed over the slave-patch boundary, ∂V_s, $\underline{\underline{K}}^{ss}_{\partial V} = \underline{\underline{S}}^{ss}_{\partial V} + \underline{\underline{M}}^{ss}_{\partial V}$; and the projective matrix formed over the entire irregular interface, ∂V_{sm}, $\underline{\underline{K}}^{sm}_{\partial V} = \underline{\underline{S}}^{sm}_{\partial V} + \underline{\underline{M}}^{sm}_{\partial V}$.

Although a surface mesh may exist in d-dimensions in the real physical space, it is mapped from a (d-1)-dimensional parametric domain. The gradient of NURBS over a boundary of the domain in the real physical space is thus expressed, for 2D and 3D, respectively, via:

$$\underline{\underline{\mathbf{J}}}_{\partial V} \cdot \vec{\nabla} R^{p}_{i}(\xi) = \hat{\nabla} R^{p}_{i}(\xi), \qquad \text{for } i = 1, \ldots, n; \tag{4.47a}$$

$$\underline{\underline{\mathbf{J}}}_{\partial V} \cdot \vec{\nabla} R^{p,q}_{i}(\xi, \eta) = \hat{\nabla} R^{p,q}_{i}(\xi, \eta), \qquad \text{for } i = 1, \ldots, n \times m; \tag{4.47b}$$

and the surface Jacobian (d-1 $\times$ d)-matrix, is defined below:

$$\underline{\underline{\mathbf{J}}}_{\partial V} \coloneqq \begin{pmatrix} \partial x/\partial \xi & \partial y/\partial \xi & \partial z/\partial \xi \\ \hdashline \partial x/\partial \eta & \partial y/\partial \eta & \partial z/\partial \eta \end{pmatrix}; \tag{4.48}$$

where the non-square matrix is *inverted* by evaluating its pseudo-inverse [295]:

$$\left[\underline{\underline{\mathbf{J}}}_{\partial V}\right]^{-1} = \left[\underline{\underline{\mathbf{J}}}_{\partial V}\right]^{T} \cdot \left[\underline{\underline{\mathbf{J}}}_{\partial V}\right]^{-T} \cdot \left[\underline{\underline{\mathbf{J}}}_{\partial V}\right]^{-1}, \tag{4.49a}$$

$$= \underset{(d\times d\text{-}1)}{\left[\underline{\underline{\mathbf{J}}}_{\partial V}\right]^{T}} \cdot \underset{(d\text{-}1\times d\text{-}1)}{\left[\underline{\underline{\mathbf{J}}}_{\partial V} \cdot \left[\underline{\underline{\mathbf{J}}}_{\partial V}\right]^{T}\right]^{-1}}. \tag{4.49b}$$

The derivatives of the NURBS with respect to the boundary must be evaluable in the real physical space, Equations 4.47; this requires the surface mesh to be non-degenerate and the surface Jacobian non-singular, irregardless of the volumetric Jacobian; q.v. Section 3.3.1. Again, the notion of an implied *super-mesh* is employed to ensure that the integral quantities evaluated over an irregular interface are constructed using a quadrature set that sufficiently samples the bases either side of the interface. This is especially important for maintaining consistency in the space of control-variables that approximate the field and thus maintaining the conservation of neutrons. For those irregularities induced by h-refinement and the subsequent integration over element boundaries that contain such hanging-nodes, the use of p+1 QP would be inadequate because the integrand will not generally be polynomial; even for the underlying geometry represented by B-splines. Nonetheless, the integrand over partial boundaries remains a continuous function over smooth NURBS objects and it is possible to gain in accuracy by supplying additional QP [70].

The manner in which the constraint equations are incorporated into the globally assembled system is based upon the work of Cottrell et al, which was built upon that of Kagan et al for restoring continuity of a basis in the presence of hanging-nodes on volumetric NURBS objects and B-spline surfaces [270, 341]. The coupling relationships depend only upon the mesh itself and can therefore be generated in a pre-processing step; along with the other dense element matrices. Unlike the proposed method, however, their constraints were based upon some common coarsest refinement structure to

which both the slave patch and the master patch must be constrained in the event that the former cannot be obtained from some combination of knot-insertion in the former. Cottrell et al were able to exploit higher-order continuity of the basis within a patch by allowing for internal refinement and applying the constraints at a patch level. However, it was not obvious how local refinement could in fact be performed, such that any non-zero knot-span flagged for p-refinement would then result in the degree elevation of the entire patch; in which case, the coarsest description of the mesh would thus require a sufficient number of initial patches to allow for any form of local p-refinement.

A wholly-local AMR-$h(p)$ strategy is proposed by considering, at most, C^0 continuity across element interfaces. The roles of *master* and *slave* are assigned to those patches that intersect at an irregular patch interface and the constraint equations are enforced via a Galerkin projection, otherwise agnostic of the underlying refinement structure of either patch. In the assembly process, the local linear systems are constructed as integral quantities over the elements, which now correspond to entire patches, V_e. The discrete system of linear equations local to a given slave patch, $\underline{\underline{A}}^s \cdot \underline{\phi}^s = \underline{b}^s \ \forall V_s \in \mathcal{T}(V)$, for any number of constrained sides, $\forall \partial V_f \in V_s \cap \partial V_{sm}$, is partitioned into those integral quantities that account for the behaviour within the patch and separately for any interaction with its neighbours:

$$\underline{\underline{A}}^s_{nn} \cdot \underline{\phi}^s_n + \sum_j \left[\underline{\underline{A}}^s_{nf_j} \cdot \underline{\phi}^s_{f_j} \right] = \underline{b}^s_n; \tag{4.50a}$$

$$\sum_{\hat{\imath}} \left[\underline{\underline{A}}^s_{f_{\hat{\imath}}n} \cdot \underline{\phi}^s_n + \sum_{\hat{\jmath}} \left[\underline{\underline{A}}^s_{f_{\hat{\imath}}f_{\hat{\jmath}}} \cdot \underline{\phi}^s_{f_{\hat{\imath}}} \right] = \underline{b}^s_{f_{\hat{\imath}}} \right]; \tag{4.50b}$$

where $\underline{\underline{A}}^s_{nn}$ is the left-hand-side partial matrix with coefficients that support only those internal functions that have no support at an irregular interface; $\underline{\underline{A}}^s_{ff}$ is the partial matrix with coefficients that support only those functions with support at an irregular interface; and $\underline{\underline{A}}^e_{nf}$ and $\underline{\underline{A}}^e_{fn}$ are the partial matrices with coefficients that support those functions both internal to the patch and at an irregular interface. The linear functional, $\underline{b}^s$, and the solution vector, $\underline{\phi}^s$, are decomposed in a similar manner; however, the latter must also be constrained to the master patch that adjoins the slave patch at a given irregular interace, $\partial V_f \in V_m \cap \partial V_{sm}$, by Equation 4.46b; i.e.:

$$\underline{\underline{A}}^s_{nn} \cdot \underline{\phi}^s_n + \sum_j \left[\underline{\underline{A}}^s_{nf_j} \cdot \underline{\underline{K}}^{s \leftarrow m_j}_{\partial V_j} \cdot \underline{\phi}^{m_j}_{f_j} \right] = \underline{b}^s_n; \tag{4.51a}$$

$$\sum_{\hat{\imath}} \left[\underline{\underline{A}}^s_{f_{\hat{\imath}}n} \cdot \underline{\phi}^s_n + \sum_{\hat{\jmath}} \left[\underline{\underline{A}}^s_{f_{\hat{\imath}}f_{\hat{\jmath}}} \cdot \underline{\underline{K}}^{s \leftarrow m_{\hat{\jmath}}}_{\partial V_{\hat{\jmath}}} \cdot \underline{\phi}^{m_{\hat{\jmath}}}_{f_{\hat{\jmath}}} \right] = \underline{b}^s_{f_{\hat{\imath}}} \right]. \tag{4.51b}$$

The above equations are tested against the larger slave-patch space, yet the trial functions that belong to the same space are restricted to act only in a linear combination, as defined by the Galerkin projection, that results in those functions existing in the master-patch function space. Without wishing to over or under-constrain the DoF at an irregular interface, so too must those test functions defined at the slave-patch boundary act only in a linear combination such as to replicate those functions

in the weighting space as defined over the master patch. Therefore, those equations over irregular interfaces must be pre-multiplied by $[\underline{\underline{K}}^{s\leftarrow m}_{\partial V_f}]^T$; where clearly, constraining a more-refined subspace to a less-refined one will inevitably induce some error at the irregular interface:

$$\underline{\underline{A}}^s_{nn}\cdot\underline{\phi}^s_n+\sum_j\underset{(N^s_n\times N^s_{f_j})\,(N^s_{f_j}\times N^{m_j}_{f_j})\,(N^{m_j}_{f_j}\times 1)}{\Big[\underline{\underline{A}}^s_{nf_j}\cdot\underline{\underline{K}}^{s\leftarrow m_j}_{\partial V_j}\cdot\underline{\phi}^{m_j}_{f_j}\Big]}=(\underline{b}^s_n); \tag{4.52a}$$

$$\sum_{\hat{\imath}}\underset{(N^{m_{\hat{\imath}}}_{f_{\hat{\imath}}}\times N^s_{f_{\hat{\imath}}})\,(N^s_{f_{\hat{\imath}}}\times N^s_n)\,(N^s_n\times 1)}{\Big[[\underline{\underline{K}}^{s\leftarrow m_{\hat{\imath}}}_{\partial V_{\hat{\imath}}}]^T\cdot\underline{\underline{A}}^s_{f_{\hat{\imath}}n}\cdot\underline{\phi}^s_n}+\sum_{\hat{\jmath}}\underset{(N^{m_{\hat{\imath}}}_{f_{\hat{\imath}}}\times N^s_{f_{\hat{\imath}}})\,(N^s_{f_{\hat{\imath}}}\times N^s_{f_{\hat{\jmath}}})\,(N^s_{f_{\hat{\jmath}}}\times N^{m_{\hat{\jmath}}}_{f_{\hat{\jmath}}})\,(N^{m_{\hat{\jmath}}}_{f_{\hat{\jmath}}}\times 1)}{\Big[[\underline{\underline{K}}^{s\leftarrow m_{\hat{\imath}}}_{\partial V_{\hat{\imath}}}]^T\cdot\underline{\underline{A}}^s_{f_{\hat{\imath}}f_{\hat{\jmath}}}\cdot\underline{\underline{K}}^{s\leftarrow m_{\hat{\jmath}}}_{\partial V_{\hat{\jmath}}}\cdot\underline{\phi}^{m_{\hat{\jmath}}}_{f_{\hat{\jmath}}}\Big]}=\underset{(N^{m_{\hat{\imath}}}_{f_{\hat{\imath}}}\times N^s_{f_{\hat{\imath}}})\,(N^s_{f_{\hat{\imath}}}\times 1)}{[\underline{\underline{K}}^{s\leftarrow m_{\hat{\imath}}}_{\partial V_{\hat{\imath}}}]^T\cdot\underline{b}^s_{f_{\hat{\imath}}}\Big]}. \tag{4.52b}$$

These equations can be assembled with those established on the master patch adjoining the slave patch at a given irregular interface, which thus far have remained unchanged; i.e. $\forall\forall\partial V_f\in V_m\cap\partial V_{sm}$:

$$\underline{\underline{A}}^m_{nn}\cdot\underline{\phi}^m_n+\underline{\underline{A}}^m_{nf}\cdot\underline{\phi}^m_f=\underline{b}^m_n; \tag{4.53a}$$

$$\underline{\underline{A}}^m_{fn}\cdot\underline{\phi}^m_n+\underline{\underline{A}}^m_{ff}\cdot\underline{\phi}^m_f=\underline{b}^m_f. \tag{4.53b}$$

Referring back to the toy problem in Figure 4.1, it should be evident that C^0 continuity has been enforced by constraining those control-variables at the slave-patch boundary, in the bottom left-hand corner, on its east-side, $\phi^s_{f_E}$, and its north side, $\phi^s_{f_N}$, to those master patches that adjoin the slave patch on its east-side, V_{m_E}, and its north-side, V_{m_N}, respectively; i.e. $\prod^{m_E}_{h,p}\phi^s_{f_E}\colon\{\phi_5,\phi_6,\phi_9\}\rightarrow\{\phi_5,\phi_9\}$ and $\prod^{m_N}_{h,p}\phi^s_{f_N}\colon\{\phi_7,\phi_8,\phi_9\}\rightarrow\{\phi_7,\phi_9\}$, respectively. Again, those control-variables at regular and compatible boundaries exist in a one-to-one correspondence between *all* adjoining patches. As such, *corner* control-variables can be shared by several patches; e.g. $\phi_9\in V_s\cap V_{m_E}\cap V_{m_N}$. Consequently, such shared *corner* control-variables are constrained to themselves across an irregular patch interface because they lie where the basis is interpolatory; e.g. $\phi_9=\underline{\underline{K}}^{s\leftarrow m_E}\cdot\phi_9$. Care must be taken not to over-constrain these *corner* control-variables; and the number of times that they are constrained is tallied so that the respective constraint equations can be corrected; e.g.:

$$\phi_9=\frac{1}{2}\big[\underline{\underline{K}}^{s\leftarrow m_E}_{\partial V}\cdot\phi_9\big]+\frac{1}{2}\big[\underline{\underline{K}}^{s\leftarrow m_N}_{\partial V}\cdot\phi_9\big]=\frac{1}{2}\big[\phi_9+\phi_9\big]=\phi_9.$$

A profound consequence of the proposed local refinement technique is the transpiration of *hanging-nodes* at the irregular interfaces. More specifically there is concern for those constrained control-variables that result from local h-refinement, whereby, it may occur that a set of (master) control-variables, to which another set of (slave) control-variables is constrained, is itself constrained to some other set of (master) control-variables; in which case, the constraints on the original set of (slave) control-variables are compounded, or the original set of (slave) control-variables is said to be multiply constrained [299]. The existence of multiply constrained control-variables greatly complicates the numerical analysis because graph dependencies would need to be resolved such as to apply the constraints and then recover the (slave) control-variables in a certain order. It is therefore desirable

to avoid such scenarios and instead enforce some degree of *regularity* over the mesh.

Where *corner* control-variables can be shared by several patches, the only way in which a compounded constraint could present itself would be if the master to some slave patch were both constrained to some other common master patch; e.g. the left-hand side of Figure 4.2 presents a mesh with such a dependency. In this case, *two* patches in the lower quadrants have been h-refined; and consequently, ϕ_8 is multiply constrained; i.e. $\{\phi_4, \phi_{10}, \phi_8\} \rightarrow \{\phi_4, \phi_8\}$, yet, $\{\phi_7, \phi_8, \phi_{13}, \phi_{14}\} \rightarrow \{\phi_7, \phi_{14}\}$. Despite its constraints also being compounded, ϕ_{14}, poses no problem because it is always constrained to itself; i.e. $\{\phi_{14}, \phi_{19}, \phi_{20}\} \rightarrow \{\phi_{14}, \phi_{20}\}$, yet $\{\phi_6, \phi_{12}, \phi_{14}\} \rightarrow \{\phi_6, \phi_{14}\}$. To alleviate the problem of multiply constraints, some regularity is enforced over the mesh; e.g. on the right-hand side of Figure 4.2, an additional patch has been h-refined, as highlighted in blue, and now $\{\phi_{11}, \phi_{12}, \phi_{13}, \phi_{17}, \phi_{18}\} \rightarrow \{\phi_{11}, \phi_{18}\}$. The regularity enforced in the current work is not as strict as the *1-irregularity* rule imposed by Wang and Ragusa; whereby, an element may only be h-refined if it does not already possess a hanging-node and in the case that it does, it must first be resolved [290].

With regard to the interpolation-based nature of the AMR strategy, an inherent bonus of defining the more-refined reference mesh, $\mathcal{T}_{ref} = \mathcal{T}_{½,p+1}$, as that which results from one uniform hp-refinement of the coarser current mesh is that some regularity is implicitly enforced over the mesh where the numerical analysis is performed; viz. h-refinement resolves some of the graph dependencies and the elements are assured to have at least *two* or *three* compatible neighbours in 2D and 3D, respectively; and p-refinement ensures a basis of at least quadratic NURBS and that at least *one* central control-variable will remain unconstrained since the support of interior functions vanishes at the boundary.

Whilst Figure 4.1 and Figure 4.2 are useful for demonstrative purposes, without loss of generality, the proposed method for enforcing C^0 continuity via projection-based constraint equations is applicable to arbitrarily irregular meshes. One admits, however, that the task of establishing constraints and resolving any dependencies does become one of good bookkeeping, which becomes evermore complicated in 3D when considering those constraints between *direct* neighbours that share a face, as well as those between *indirect* neighbours that can share an edge. In essence, the elementary stencil can become a lot larger than that for a DG discretisation over the same mesh; q.v. Section 4.4.

4.4 An Interior Penalty Scheme for a Discontinuous Galerkin Isogeometric Discretisation

Changing tack, one derives now an IP-DG-IGA discretisation for the multi-group NDE, Equation 4.2, and their adjoint equations, Equation 4.8, are solved for the multi-group components $\phi^g : V \rightarrow \mathbb{R}$ and

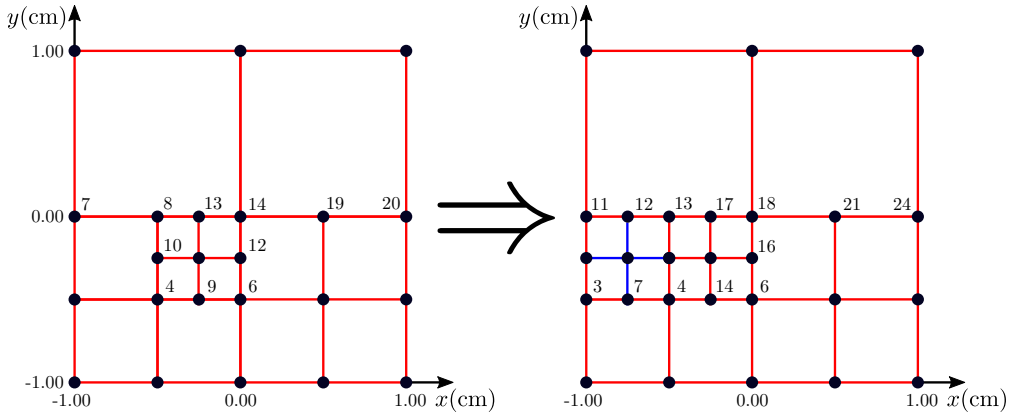


Figure 4.2: An example of regularity enforced over the irregular mesh of a toy problem; v. Figure 4.1. The initial mesh (left) presents a multiply constrained control-variable, ϕ_8. Dependencies are resolved (right) by enforcing some degree of regularity; viz. the h-refinement highlighted in blue. For ease of viewing, only the enumeration of those control-variables that lie at an irregular patch interface are indicated. (*Vide* the web-based version for reference to colour.)

$\phi^{g\dagger}\colon V \to \mathbb{R}$, respectively, over the same mesh as that adopted in Section 4.3; i.e. $\mathcal{T}(V) = \{V_e\}_{e=1}^{E}$ and $V = \bigcup_{e=1}^{E} V_e$; and $\partial V_f \in \partial V = \partial V_{int} \cup \partial V_{ext}$; where $\partial V_{ext} = \partial V_0 \cup \partial V_R \cup \partial V_V$. In the interest of discontinuous formulations, it is more appropriate to exploit the usual space of trial functions in a piecewise manner. The broken Sobolev space, $H^s(\mathcal{T})$, defines a set of spaces, $H^s(V_e)$, spanned by locally supported functions that belong to the regular Sobolev space, Equation 4.22, in the restriction to V_e; and one notes that continuity constraints are not imposed across element interfaces:

$$H^s(\mathcal{T}) \coloneqq \{u \in L_2(V)\colon u|_{V_e} \in H^s(V_e), \forall V_e \in \mathcal{T}(V)\}; \tag{4.54}$$

where the usual Sobolev embeddings rules apply; i.e. $H^{s+1}(\mathcal{T}) \subset H^s(\mathcal{T})$. The associated broken inner-product and broken norm, are given below in Equation 4.55a and Equation 4.55b, respectively:

$$(u,v)_{H^s(\mathcal{T})} \coloneqq \sum_{V_e \in \mathcal{T}(V)} (u,v)_{H^s(V_e)}; \tag{4.55a}$$

$$\|u\|_{H^s(\mathcal{T})} \coloneqq (u,u)^{1/2}_{H^s(\mathcal{T})}. \tag{4.55b}$$

The restriction of such functions and their normal derivatives with support over V_e are defined over ∂V_e by trace operators, $\gamma_i\colon H^s(V_e) \to H^{s-i-1/2}(\partial V_e)$ for $0 \leq i < s$ [222]. Of particular interest are the Dirichlet trace, γ_0, and the Neumann trace, γ_1, which are well defined in $L_2(\partial V_e)$ for $s > 1/2$:

$$\gamma_0 u \coloneqq u|_{\partial V_e}; \tag{4.56a}$$

$$\gamma_1 u \coloneqq \vec{\nabla} u \cdot \vec{n}|_{\partial V_e}. \tag{4.56b}$$

Any function $u \in H^s(\mathcal{T})$ is discontinuous, along with its normal derivatives, across element interfaces; one therefore defines the gradient operator in an element-wise manner; i.e. the broken gradient:

$$(\vec{\nabla} u)|_{V_e} \coloneqq \vec{\nabla}(u|_{V_e}), \quad \forall V_e \in \mathcal{T}(V). \tag{4.57}$$

For some function $u \in H^s(V)$, for $s \geq 1$, where any jump in u is always vanishing, the broken gradient coincides with the gradient defined in a distributional sense; i.e. $(\vec{\nabla} u)|_{V_e} = \vec{\nabla} u$. Such a function

also belongs to the broken Sobolev space, $u \in H^s(\mathcal{T})$, where it is not required for any jump in u to be vanishing. One can therefore assume that $H^s(V) \subset H^s(\mathcal{T})$; or rather, under some assumption of regularity, only sufficiently smooth solutions belong to $H^s(V)$ [237]. This may explain any observed increase in accuracy for IP-DG discretisations of elliptic problems over standard CBG ones over the same mesh; and because the former employs at least twice as many DoF as the latter [198, 277].

It is useful to provide a formal definition of the *jump* in a function, as well as the *average* across an interface, $\partial V_f \in \partial V_{int}$; where for $s > 3/2$, the domain is smooth enough to admit two-valued traces. Although arbitrary, the definitions are perspective-dependent; viz. the superscript (+) pertains to the *active*-element, V_+, whose outward-pointing normal vector over its boundary, $\vec{n}^+$, coincides with that, $\vec{n}$, at the interface, $V_+ \cap V_-$; whereas the superscript (-) pertains to the *neighbour*-element, V_-, whose outward-pointing normal vector over its boundary, $\vec{n}^-$, has the opposite orientation to that of the active-element; i.e. $\vec{n} = \vec{n}^+ = -\vec{n}^-$. The jump, $[\![\cdot]\!]_{\partial V_f}$, and the average, $\{\!\{\cdot\}\!\}_{\partial V_f}$, are defined for some scalar-valued function, u, in Equation 4.58a and Equation 4.58b, respectively:

$$[\![u]\!]_{\partial V_f} \coloneqq \begin{cases} u^+ - u^- \quad , & \forall \partial V_f \in \partial V_{int}; \\ u, & \forall \partial V_f \in \partial V_{ext}; \end{cases} \tag{4.58a}$$

$$\{\!\{u\}\!\}_{\partial V_f} \coloneqq \begin{cases} \frac{1}{2}(u^+ + u^-), & \forall \partial V_f \in \partial V_{int}; \\ u, & \forall \partial V_f \in \partial V_{ext}. \end{cases} \tag{4.58b}$$

To allow for a more compact notation in due course, the jumps and averages over those boundaries that lie adjacent to the physical boundary, ∂V_{ext}, have been included in the above.

Physically, the solution to the primal problem is assumed to be sufficiently smooth that any jumps in the scalar neutron flux, ϕ, and any jumps in the normal component of the net neutron current, $-D\vec{\nabla}\phi \cdot \vec{n}$, are vanishing across the interior of the domain; i.e.

$$\left.\begin{matrix} [\![\phi]\!]_{\partial V_f}; \\ [\![-D\vec{\nabla}\phi \cdot \vec{n}]\!]_{\partial V_f} \end{matrix}\right\} = 0, \quad \forall \partial V_f \in \partial V_{int}. \tag{4.59a}$$

Likewise, as inherited by the property of adjoint-consistency, the solution to the dual problem, $\phi^\dagger$, is equally smooth; i.e. $\phi, \phi^\dagger \in \mathcal{V}(V) \subset H^{s>3/2}(V)$. In the consideration of material heterogeneity and the piecewise definition of the diffusion coefficient, D, one assumes weaker regularity of the true solution; viz. its first derivative is not expected to be smooth across the domain [342].

Without loss of generality, an emphasis is placed upon manipulating the direct equations; and the same discretisation is applied to the adjoint equations. Again, the strong form of the defining equations is reformulated into the *weak* form by multiplying Equation 4.2 by a test function, $w \in \mathcal{V}(\mathcal{T})$, integrating over each V_e and then taking the sum of these contributions. A solution is also sought in the s-order

broken Sobolev space, Equation 4.54; and thus, the (semi-)analytical solution is approximated by some linear expansion of trial functions, $u \in \mathcal{U}(\mathcal{T}) \subseteq H^s(\mathcal{T})$, and control-variables, ϕ_i:

$$\phi(\vec{r}) \approx \hat{\phi}(\vec{r}) = \sum_{i=1}^{\infty} \phi_i u_i(\vec{r}), \quad \forall \vec{r} \in V_e \in \mathcal{T}(V). \tag{4.60}$$

Substitution of the above into the integrand, Equation 4.2, of the integral statements over V_e yields a weighted-residual, which is forced to vanish everywhere. Taking both the trial and the test functions from the same space, the residual is made to be orthogonal to the approximate solution; i.e. $\mathcal{U}(\mathcal{T}), \mathcal{V}(\mathcal{T}) \subseteq H^s(\mathcal{T})$. Smoothness requirements of the trial functions are relaxed by application of the divergence theorem to the double-differential diffusion term; yet $s > 3/2$ is still enforced.

The proposed spatial discretisation is non-conforming in the sense that the space in which an approximate solution is sought is not a subset of that to which the true solution belongs; i.e. $\hat{\phi} \in \mathcal{V}(\mathcal{T}) \not\subset \mathcal{V}(V) \ni \phi$. Unlike a standard CBG discretisation, it cannot be assumed that any properties of the (semi-)analytical solution will be simply inherited by the discrete one; therefore one must directly establish discrete stability and consistency. The discrete variational problem is posed by restricting the collection of infinite approximation spaces to finite-dimensional subsets; i.e. $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \subset \mathcal{V}(\mathcal{T})$. As per the isoparametric concept, these broken spaces are chosen as those spanned by the same N_e-number of NURBS used to describe the patches, V_e. Although the basis may be smooth within a given patch, it is discontinuous across the boundaries; v. Figure 3.2. The NURBS, R, Equations 3.7, are characterised by the mesh spacing, h, as well as the polynomial degree of the basis, p; i.e. $\forall V_e \in \mathcal{T}$:

$$\begin{cases} \left[\hat{\mathcal{V}}_{h,p}(V_e)\right]^T = \text{span}\left\{R_i^p\right\}_{i=1}^{N_e}, \\ \text{with } \text{supp}\left(R^p\right) = \left\{\vec{r} \in V_e \colon R_i^p(\vec{r}) \neq 0 \text{ for } i = 1, \ldots, N_e\right\}. \end{cases} \tag{4.61}$$

Accordingly, the expansion in Equation 4.60 is truncated to the N_e^{th} term $\forall V_e \in \mathcal{T}$. Again, the weighted-residual statement must be satisfied for each member in the weighting set and since the test and trial spaces coincide, there are as many algebraic equations as there are unknowns per group; i.e. $DoF_e^g = \dim \hat{\mathcal{V}}_{h,p}(V_e) = N_e$ and $DoF^g = \sum_{V_e} DoF_e^g = N$ for $g = 1, \ldots, G$; and hence, $DoF = \sum_g DoF^g = G \times N$. Consequently, for the analysis over a single energy-independent mesh, the discrete problem is posed as a $(GN \times GN)$-matrix system of equations.

Naturally satisfied is the reflective BC, Equation 4.13b; the vacuum BC is implemented in a weak sense by the simple substitution of Equation 4.13c into the VBVP; and Equation 4.13a is weakly enforced by Nitsche's penalty method; q.v. Equation 4.25. As it stands, the discrete variational statement of Equation 4.5 is to find $\hat{\phi} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, such that:

$$a(\hat{\phi}, \hat{w}) = b(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}). \tag{4.62}$$

The discrete bilinear form, $a(\cdot,\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \to \mathbb{R}$, is block-like in structure and comprises of a weak diffusive-leakage and group-removal term, a_L; a group-to-group scattering term, a_S; a fission term, a_F; and a diffusive-leakage term, a_{IP}, defined over the interfaces and those external boundaries where Dirichlet BC are prescribed, $\partial V_{IP} = \partial V_{int} \cup \partial V_0$; and a final term, a_{ext}, that incorporates those BC prescribed on the remainder of $\partial V_{ext} n \partial V_0$:

$$a(\hat{\phi},\hat{w}) = a_L(\hat{\phi},\hat{w}) - a_S(\hat{\phi},\hat{w}) - a_F(\hat{\phi},\hat{w}) + a_{IP}(\hat{\phi},\hat{w}) + a_{ext}(\hat{\phi},\hat{w}). \tag{4.63}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$\begin{aligned} a_L(\hat{\phi},\hat{w}) &= \sum_{g=1}^{G}\Bigg[\sum_{V_e\in\mathcal{T}}\int_{V_e}\vec{\nabla}\hat{w}\cdot D_e^g\vec{\nabla}\hat{\phi}^g + \hat{w}\Sigma_{r,e}^g\hat{\phi}^g\,dV\Bigg], \\ &= \sum_{g=1}^{G} a_L^g(\hat{\phi}^g,\hat{w}); \end{aligned} \tag{4.64a}$$

$$\begin{aligned} a_S(\hat{\phi},\hat{w}) &= \sum_{g=1}^{G}\Bigg[\sum_{g\prime\neq g}^{G}\sum_{V_e\in\mathcal{T}}\int_{V_e}\hat{w}\Sigma_{s,e}^{g\prime\to g}\hat{\phi}^{g\prime}\,dV\Bigg], \\ &= \sum_{g=1}^{G} a_S^g(\hat{\phi},\hat{w}); \end{aligned} \tag{4.64b}$$

$$\begin{aligned} a_F(\hat{\phi},\hat{w}) &= \sum_{g=1}^{G}\Bigg[\sum_{g\prime=1}^{G}\sum_{V_e\in\mathcal{T}}\chi_e^g\int_{V_e}\hat{w}\nu_e^{g\prime}\Sigma_{f,e}^{g\prime}\hat{\phi}^{g\prime}\,dV\Bigg], \\ &= \sum_{g=1}^{G} a_F^g(\hat{\phi},\hat{w}); \end{aligned} \tag{4.64c}$$

$$\begin{aligned} a_{IP}(\hat{\phi},\hat{w}) &= \sum_{g=1}^{G}\Bigg[\sum_{V_e\in\mathcal{T}}\sum_{\partial V_f\in V_e\cap\partial V_{IP}} -\int_{\partial V_f}\hat{w}D_e^g\vec{\nabla}\hat{\phi}^g\cdot\vec{n}\,dS + \sum_{\partial V_f\in V_e\cap\partial V_0}\mu_f^g\int_{\partial V_f}\hat{w}\hat{\phi}^g\,dS\Bigg], \\ &= \sum_{g=1}^{G}\Bigg[\sum_{\partial V_f\in\partial V_{IP}} -\int_{\partial V_f}[\![\hat{w}D^g\vec{\nabla}\hat{\phi}^g\cdot\vec{n}]\!]\,dS + \sum_{\partial V_f\in\partial V_0}\mu_f^g\int_{\partial V_f}[\![\hat{w}]\!][\![\hat{\phi}^g]\!]\,dS\Bigg], \\ &= \sum_{g=1}^{G} a_{IP}^g(\hat{\phi}^g,\hat{w}); \end{aligned} \tag{4.64d}$$

$$\begin{aligned} a_{ext}(\hat{\phi},\hat{w}) &= \sum_{g=1}^{G}\Bigg[\sum_{V_e\in\mathcal{T}}\sum_{\partial V_f\in V_e\cap\partial V_V}\frac{1}{2}\int_{\partial V_f}\hat{w}\hat{\phi}^g\,dS\Bigg], \\ &= \sum_{g=1}^{G} a_{ext}^g(\hat{\phi}^g,\hat{w}). \end{aligned} \tag{4.64e}$$

The discrete linear functional, $b(\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \to \mathbb{R}$, comprises only fixed sources, since $g_D = 0$:

$$\begin{aligned} b(\hat{w}) &= \sum_{g=1}^{G}\Bigg[\sum_{V_e\in\mathcal{T}}\int_{V_e}\hat{w}q_{x,e}^g\,dV\Bigg], \\ &= \sum_{g=1}^{G} b^g(\hat{w}). \end{aligned} \tag{4.65}$$

For multiplying systems, $q_x = 0$, the weak form of Equation 4.6 is to find $\hat{\phi} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, such that:

$$\tilde{a}(\hat{\phi}, \hat{w}) = \frac{1}{\mathrm{k_{eff}}} a_f(\hat{\phi}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}); \tag{4.66}$$

where $\tilde{a}$ is the original bilinear form in Equation 4.63, a, less the fission term; i.e.:

$$\tilde{a} = a + a_F = a_L - a_S + a_{IP} + a_{ext}. \tag{4.67}$$

Consistency The discretisation is said to be consistent if the true solution satisfies the discrete system of equations; however, one cannot simply substitute ϕ into Equation 4.62 because a has been defined on $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}_{h,p}(\mathcal{T})$; and similarly for $\tilde{a}$ in Equation 4.66. Therefore, one must first assume that $\phi \in \mathcal{V}^*(V) \subset \mathcal{V}(V)$; and second, that $\hat{\phi}^* \in \hat{\mathcal{V}}^*_{h,p}(\mathcal{T}) = \mathcal{V}^*(V) + \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, such that the bilinear form can be extended to be bounded on $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}^*_{hp}(\mathcal{T})$. The space $\hat{\mathcal{V}}^*_{h,p}(\mathcal{T})$ is equipped with the norm $\|\cdot\|_{\hat{\epsilon}^*}$, such that for some arbitrary function, u, and some constant C:

$$\|u\|_{\hat{\epsilon}^*} = \|u\|_{\hat{\epsilon}}, \qquad \forall u \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}); \tag{4.68a}$$

$$\|u\|_{\hat{\epsilon}^*} < C\|u\|_{\epsilon^*}, \qquad \forall u \in \mathcal{V}^*(V). \tag{4.68b}$$

The former implies an equivalence of the norms $\|\cdot\|_{\hat{\epsilon}^*}$ and $\|\cdot\|_{\hat{\epsilon}}$ on $\hat{\mathcal{V}}_{h,p}(\mathcal{T})$; and the latter implies that $\mathcal{V}^*(V) \subset \hat{\mathcal{V}}^*_{h,p}(\mathcal{T})$; where $\|\cdot\|_{\hat{\epsilon}}$ and $\|\cdot\|_{\epsilon^*}$ are the energy-norms associated with the discrete bilinear form, $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, and the analytical one, $\mathcal{V}^*(V) \times \mathcal{V}^*(V)$, respectively.

From the definition of jumps, Equation 4.58a, the first line of Equation 4.64d expressed as a sum of integrals over element boundaries can be re-expressed in the second line as a sum of integrals over interfaces and Dirichlet boundaries, which can then be developed further using the definition of averages, Equation 4.58b; i.e. for $g = 1, \ldots, G$:

$$\begin{aligned}
&\sum_{\partial V_f \in \partial V_{IP}} \int_{\partial V_f} [\![\hat{w} D^g \vec{\nabla} \hat{\phi}^{g*} \cdot \vec{n}]\!]\, dS \\
&= \sum_{\partial V_f \in \partial V_{int}} \int_{\partial V_f} [\![\hat{w} D^g \vec{\nabla} \hat{\phi}^{g*} \cdot \vec{n}]\!]\, dS + \sum_{V_e \in \mathcal{T}} \sum_{\partial V_f \in V_e \cap \partial V_0} \int_{\partial V_f} \hat{w} D^g_e \vec{\nabla} \hat{\phi}^{g*} \cdot \vec{n}\, dS,
\end{aligned} \tag{4.69a}$$

$$= \sum_{\partial V_f \in \partial V_{int}} \int_{\partial V_f} \{\!\{ \hat{w} \}\!\} \cancelto{0}{[\![D^g \vec{\nabla} \hat{\phi}^{g*} \cdot \vec{n}]\!]}\, dS + \sum_{\partial V_f \in \partial V_{IP}} \int_{\partial V_f} [\![\hat{w}]\!] \{\!\{ D^g \vec{\nabla} \hat{\phi}^{g*} \cdot \vec{n} \}\!\}\, dS. \tag{4.69b}$$

The average of the normal component of the net neutron current across an interface is taken simply as the arithmetic average; q.v. Equation 4.58b; alternatively, one could have employed a weighted average that depends upon the diffusion coefficients defined either side of the interface [343].

The strong form, Equation 4.1a, and the BC, Equations 4.13, are retrieved by substitution of ϕ into Equation 4.62; and by application of the divergence theorem to the first term in Equation 4.64a, the double-differential term is recouped. This is followed by the re-expression of a_{IP}, Equation 4.64d, in

terms of Equation 4.69b; and then accounting for the regularity in the true solution, Equations 4.59. Thus, the discrete problem is consistent provided that the numerical fluxes are consistent; and the following modification can be made for $g = 1, \ldots, G$:

$$\begin{aligned} \implies a^g_{IP}(\hat{\phi}^{g*}, \hat{w}) = & \sum_{\partial V_f \in \partial V_{IP}} - \int_{\partial V_f} [\![\hat{w}]\!] \{\{ D^g \vec{\nabla} \hat{\phi}^{g*} \cdot \vec{n} \}\} \, dS \\ & + \sum_{\partial V_f \in \partial V_0} \mu^g_f \int_{\partial V_f} [\![\hat{w}]\!] [\![\hat{\phi}^{g*}]\!] \, dS. \end{aligned} \tag{4.70}$$

Galerkin's property of orthogonality stems from consistency since $\mathcal{V}^*(V) \subset \hat{\mathcal{V}}^*_{h,p}(\mathcal{T})$:

$$a(\phi - \hat{\phi}, \hat{w}) = a(\phi, \hat{w}) - a(\hat{\phi}, \hat{w}), \quad \forall \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}), \tag{4.71a}$$

$$= b(\hat{w}) - b(\hat{w}), \quad \forall \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}), \tag{4.71b}$$

$$= 0. \tag{4.71c}$$

A strong definition of consistency has been enforced, yet the integral quantities are evaluated via numerical quadrature; q.v. Section 3.2.3. Consequently, one expects consistency only asymptotically. Equally important in upholding the properties of local conservation is an ability to evaluate sufficiently any boundary integrals. The use of p+1 QP may no longer be adequate because the integrands are polynomial constructs of combined maximal polynomial degree $p + 2 \times (p-1) = 3p - 2$; q.v. Equation 4.70. Where each component of the normal vector, $\vec{n}$, is itself a linear combination of derivatives of the same functions; viz. polynomial constructs of degree p-1. Using Legendre-Gauss quadrature rules, the boundary integral should be evaluated sufficiently with Q-number of QP in each of the (d-1)-directions over the canonical representation of the boundary in the parent space; i.e. for $p > 1$:

$$2Q - 1 \geq 3p - 2; \tag{4.72a}$$

$$\implies Q = \begin{cases} \dfrac{3p}{2}, & \text{for } p\text{-even}; \\ \dfrac{3p-1}{2}, & \text{for } p\text{-odd}. \end{cases} \tag{4.72b}$$

In the interest of AMR, asymptotic consistency and the conservation of neutrons must hold over irregular meshes. To satisfy Equations 4.59, one employs quadrature rules that adequately sample the bases either side of an interface; i.e. $\forall \partial V_f \in \partial V_{int}$:

$$Q = \max_{V_e \in \bigcap_{e\prime} V_{e\prime} \ni \partial V_f} Q|_{V_e}. \tag{4.73}$$

For those irregularities induced by h-refinement and the subsequent integration over element boundaries that contain such hanging-nodes, Equation 4.72b would still be inadequate because the integrand will not generally be a polynomial; even for the underlying geometry represented by B-splines. Nonetheless, the integrand over partial boundaries remains a continuous function over smooth NURBS objects and it is possible to gain in accuracy by supplying additional QP.

Symmetry Foreshadowing the call upon duality arguments and the desire to employ the CG method, the following is appended to a_{IP} for $g = 1, \dots, G$:

$$\sum_{\partial V_f \in \partial V_{IP}} \text{-}\theta \int_{\partial V_f} \{\{D^g \vec{\nabla} \hat{w} \cdot \vec{n}\}\} [\![\hat{\phi}^{g*}]\!] \, dS = \begin{cases} 0, & \forall \partial V_f \in \partial V_{int}; \\ \sum_{\partial V_f \in \partial V_D} \text{-}\theta \int_{\partial V_f} D_e^g \vec{\nabla} w \cdot \vec{n} \, g_D^g \, dS, & \forall \partial V_f \in \partial V_D. \end{cases} \tag{4.74}$$

Or rather, symmetry of the within-group discrete bilinear form is reclaimed by letting $\theta = 1$ in the following modification to a_{IP}^g, Equation 4.70, for $g = 1, \dots, G$:

$$\implies a_{IP}^g(\hat{\phi}^{g*}, \hat{w}) = a_{IPc}^g(\hat{\phi}^{g*}, \hat{w}) + a_{IPs}^g(\hat{\phi}^{g*}, \hat{w}) + a_{IPp}^g(\hat{\phi}^{g*}, \hat{w}), \tag{4.75a}$$

$$\begin{aligned} &= \sum_{\partial V_f \in \partial V_{IP}} \text{-} \int_{\partial V_f} [\![\hat{w}]\!] \{\{D^g \vec{\nabla} \hat{\phi}^{g*} \cdot \vec{n}\}\} \, dS \\ &- \sum_{\partial V_f \in \partial V_{IP}} \theta \int_{\partial V_f} \{\{D^g \vec{\nabla} \hat{w} \cdot \vec{n}\}\} [\![\hat{\phi}^{g*}]\!] \, dS \\ &+ \sum_{\partial V_f \in \partial V_0} \mu_f^g \int_{\partial V_f} [\![\hat{w}]\!] [\![\hat{\phi}^{g*}]\!] \, dS. \end{aligned} \tag{4.75b}$$

Consistency is maintained since $[\![\phi]\!] = 0$.

Neglected thus far, a consistent discretisation of the adjoint problem, Equation 4.8, is needed to establish (quasi-)optimal error estimates in the L_2-norm and more importantly in some specified goal. Such a discretisation is adjoint-consistent, such that $a(\hat{u}, \hat{v}) - a^\dagger(\hat{v}, \hat{u}) = 0$, $\forall \hat{u}, \hat{v} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, where it can be inferred from Equation 4.7 that a lack of symmetry is often the root of a lack of adjoint-consistency [338]. Disregarding any super-penalisation, the adjoint solution to any non-symmetric discretisation will generally exhibit some degree of non-smoothness, which, contrary to a similar symmetric discretisation, tends to linger under refinement. Such a numerical method could not benefit from the *order-doubling* of any error measured in some specified goal functional of the direct solution; whence the sub-optimal rates of convergence [320].

In regard to the assembly process, with the exception of a_{IP}, one notes that Equations 4.64 are given as a sum over elements or boundaries whose local contributions are the product of basis functions that both have local support over the same domain. Differently, Equation 4.75 comprises of a sum over interfaces, formed as the intersection of *two* elements, whose local contributions are the product of basis functions that may not necessarily have local support over the same element, such that a_{IP} can be block-partitioned [237]. For the example of $\mathcal{T}(V) = \{V_1, V_2\}$ and $\partial V_f = V_1 \cap V_2 = \partial V_{IP}$, then:

$$a_{IP}(\hat{u}, \hat{v}) = \underline{\underline{A}}_{IP} = \begin{bmatrix} \underline{\underline{A}}_{IP}^{11} & \underline{\underline{A}}_{IP}^{12} \\ \underline{\underline{A}}_{IP}^{21} & \underline{\underline{A}}_{IP}^{22} \end{bmatrix}. \tag{4.76}$$

As mentioned previously, a natural consequence of open knot-vectors is that the NURBS basis is

interpolatory at patch boundaries. Accordingly, those functions that have support over a portion of the boundary, ∂V_f, of a given element, V_e, span the proper subset $\hat{\mathcal{V}}^e_{h,p}(\partial V_f) \subset \hat{\mathcal{V}}_{h,p}(V_e)$; i.e.:

$$\begin{cases} \left[\hat{\mathcal{V}}^e_{h,p}(\partial V_f)\right]^T = \text{span}\left\{(R^e|_{\partial V_f})_i\right\}_{i=1}^{N^e_f}, \\ \text{with supp}\left(R^e|_{\partial V_f}\right) = \left\{\vec{r} \in \partial V_f : (R^e(\vec{r})|_{\partial V_f})_i \neq 0 \text{ for } i = 1, \dots, N^e_f\right\}; \end{cases} \tag{4.77}$$

where the NURBS defined locally over a given element, $V_e \in \mathcal{T}(V)$, Equation 4.61, have been denoted by a similar superscript; i.e. by letting $N^e = \dim \hat{\mathcal{V}}_{h,p}(V_e)$, then for $i = 1, \dots, N^e$:

$$R^e_i(\vec{r}) = \left(R(\vec{r})|_{V_e}\right)_i, \qquad \vec{r} \in V_e.$$

A block from Equation 4.76 can be defined $\forall \hat{u} \in \hat{\mathcal{V}}^\beta_{h,p}(\partial V_f)$ and $\forall \hat{v} \in \hat{\mathcal{V}}^\alpha_{h,p}(\partial V_f)$, for $\alpha, \beta = 1, 2$ by:

$$\left(\underline{\underline{A}}^{\alpha\beta}_{IP}\right)_{ij} = \left[a_{IP}(R^\beta_j, R^\alpha_i)\right] \in \mathbb{R}^{N^\alpha_f \times N^\beta_f}.$$

Furthermore, a block, $\underline{\underline{A}}^{\alpha\beta}_{IP}$, may be split into *three* parts that correspond to the consistency, symmetry and penalty terms in the second, third and fourth lines of Equation 4.75, respectively; i.e.: $\underline{\underline{A}}^{\alpha\beta}_{IP} = \underline{\underline{A}}^{\alpha\beta}_{IPc} + \underline{\underline{A}}^{\alpha\beta}_{IPs} + \underline{\underline{A}}^{\alpha\beta}_{IPp}$. From a computational perspective, savings can be made by considering that the symmetry-part of the block is simply the transpose of the consistency-part; or more rigorously:

$$\underline{\underline{A}}^{\alpha\beta}_{IPs} = \left[\underline{\underline{A}}^{\beta\alpha}_{IPc}\right]^T.$$

Penalty As per the Lax-Milgram Lemma, the properties of continuity, Equation 4.78, and more specifically coercivity, Equation 4.79, are sufficient conditions to assert well-posedness [338]. For $\hat{u} \in \hat{\mathcal{V}}^*_{h,p}(\mathcal{T})$ and $\hat{v} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ and some constant, $C_\alpha > 0$, the discrete bilinear form is bounded on $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}^*_{hp}(\mathcal{T})$ if:

$$a(\hat{u}, \hat{v}) \leq C_\alpha \|\hat{v}\|_{\hat{\epsilon}} \|\hat{u}\|_{\hat{\epsilon}^*}; \tag{4.78}$$

and for $\hat{u} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ and some constant, $C_\beta > 0$, the discrete bilinear form is coercive if:

$$a(\hat{u}, \hat{u}) \geq C_\beta \|u\|^2_{\hat{\epsilon}}. \tag{4.79}$$

Quite often the latter is simply guaranteed by the nature of the problem; however, at a second glance, the consistency and the symmetry terms in Equation 4.75 are negative; and in general, it is not possible to uphold a statement of coercivity without some (residual-based) stabilisation mechanism [330]. Akin to the weak imposition of the Dirichlet BC, Equation 4.25, Nitsche's penalty method can be applied to the interior of the mesh; although it is understood that the discrete *inf-sup* condition could be satisfied in other ways, without having to introduce a least-squares penalisation term [344].

Discontinuities in the scalar neutron flux, as well as the net current, result from the non-conformity in seeking a discrete solution in a broken function space. A final modification is made to a_{IP}, such that

appended to the bilinear form is a further residual-dependent term that *contributes* to the coercivity by penalising the error, or rather quashing any jumps $[\![\hat{\phi}]\!]$ across the discontinuous interfaces; q.v. Equation 4.80. Further terms may be added to enforce weakly higher-order continuity; however, are not be considered here [242]. Stability of the discretisation is enhanced and invertibility is ensured under refinement for some suitable selection of penalty parameters, $\mu_f^g > 0$, for $g = 1, \ldots, G$:

$$\begin{aligned} \implies a_{IP}^g(\hat{\phi}^{g*}, \hat{w}) = & \sum_{\partial V_f \in \partial V_{IP}} - \int_{\partial V_f} [\![\hat{w}]\!] \{\{ D^g \vec{\nabla} \hat{\phi}^{g*} \cdot \vec{n} \}\} \, dS \\ & - \sum_{\partial V_f \in \partial V_{IP}} \theta \int_{\partial V_f} \{\{ D^g \vec{\nabla} \hat{w} \cdot \vec{n} \}\} [\![\hat{\phi}^{g*}]\!] \, dS \\ & + \sum_{\partial V_f \in \partial V_{IP}} \mu_f^g \int_{\partial V_f} [\![\hat{w}]\!] [\![\hat{\phi}^{g*}]\!] \, dS. \end{aligned} \tag{4.80}$$

Choosing the numerical fluxes as those that connect the elements through their interfaces as per an interior penalty scheme appears to provide for a *good* compromise between the sparsity and the conditioning of the ensuing matrix equations [345]. Whilst the latter is dependent upon the fine tuning of the penalty parameters, the former is better understood from a perspective of implementation. Owing to the local formulation, the elementary stencil for the proposed IP-DG discretisation accounts for those volume contributions over elements as well as those face contributions over interfaces formed with *direct* neighbours that share a face; viz. as per the tensor-product structure of the parametric space, the elementary stencil contains at most $(1+d^2)$-elements.

In the application to DG discretisations, there exist myriad varieties of IP schemes, each proposing its own combination of numerical fluxes and stabilisation mechanisms [330, 345]. Limiting the scope somewhat, the *three* standard variations, the non-symmetric interior penalty (NIP), the incomplete interior penalty (IIP) and the symmetric interior penalty (SIP) schemes can be expressed by setting θ to -1, 0 and 1, respectively, in Equation 4.80. Whilst the NIP-bilinear form is always coercive for any $\mu_f > 0$, a lack of symmetry and a lack of super-penalisation generally results in a lack of adjoint-consistency and thus, sub-optimal error estimates based upon duality-arguments [346]. As such, the principal interest lies with the SIP-DG-IGA discretisation.

Owing to the choice of discretisation, C^0 continuity is enforced across the domain in a weak sense at the patch-level; that is, for some internal refinement structure, $C^{p \geq \dot{m}}$ continuity in the basis is strongly enforced across the knots within a patch. Alternatively, had the approximation space been restricted to a subset of globally continuous functions, $H^s(V)$, continuity would be enforced in a strong sense and the discretisation would revert to a conforming CBG-IGA discretisation; i.e. for $[\![\hat{w}]\!] = 0$ and $[\![\hat{\phi}]\!] = 0$. The proposed SIP-DG discretisation is locally conservative, since a valid balance equation is satisfied over each macro-element; the symmetry and penalty terms are understood as pure numerical *masses* that vanish under refinement [242]. Furthermore, consistency has been maintained.

Strang's Second Lemma of quasi-best approximation for non-conforming and consistent discretisations, Equation 4.81, states that up to a constant, as defined by those constants of continuity and coercivity, the set $\hat{\mathcal{V}}_{h,p}(\mathcal{T})$ is that chosen from all possible finite sets of function spaces that minimises the error in the numerical solution, $\hat{\phi} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, with respect to the true solution, $\phi \in \mathcal{V}^*_{h,p}(V)$, as measured in the $\|\cdot\|_{\hat{\epsilon}^*}$-norm; i.e. for some $\hat{v} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ [338]:

$$
\begin{aligned}
C^*_\beta \|\hat{\phi} - \hat{v}\|^2_{\hat{\epsilon}^*} &\leq a(\hat{\phi} - \hat{v}, \hat{\phi} - \hat{v}), && \text{(4.81a)}\\
&\leq a(\hat{\phi} - \phi, \hat{\phi} - \hat{v}) + a(\phi - \hat{v}, \hat{\phi} - \hat{v}), && \text{(4.81b)}\\
&\leq a(\phi - \hat{v}, \hat{\phi} - \hat{v}), && \text{(4.81c)}\\
&\leq C^*_\alpha \|\hat{\phi} - \hat{v}\|_{\hat{\epsilon}^*} \|\phi - \hat{v}\|_{\hat{\epsilon}^*}; && \text{(4.81d)}\\
\implies \|\hat{\phi} - \hat{v}\|_{\hat{\epsilon}^*} &\leq \frac{C^*_\alpha}{C^*_\beta} \|\phi - \hat{v}\|_{\hat{\epsilon}^*}; && \text{(4.81e)}\\
\implies \|\phi - \hat{\phi}\|_{\hat{\epsilon}^*} &\leq \frac{C^*_\alpha + C^*_\beta}{C^*_\beta} \inf_{\hat{v} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})} \|\phi - \hat{v}\|_{\hat{\epsilon}^*}; && \text{(4.81f)}
\end{aligned}
$$

where, by the equivalence of the norms on $\hat{\mathcal{V}}_{h,p}(\mathcal{T})$, Equation 4.68a, the first line expresses the property of coercivity, Equation 4.79, with some different constant, C^*_β; the linearity of each argument of a is exhibited in the second line by introducing ϕ; where the first term is made to be vanishing in the third, Equation 4.81c, since the error in the numerical solution is orthogonal to the chosen subspace, Equation 4.71. The fourth line, Equation 4.81d, is obtained by the property of continuity, Equation 4.78, with some different constant, C^*_α, as per the equivalence of norms on the chosen finite-dimensional function space, Equation 4.68a. Finally, Strang's Second Lemma is obtained in the sixth line, Equation 4.81f, by taking the infimum over all possible $\hat{v} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$; after having applied the triangle inequality:

$$
\|\phi - \hat{\phi}\| \leq \|\phi - \hat{v}\| + \|\hat{\phi} - \hat{v}\|.
$$

The system of equations is solved via the SI procedure; q.v. Section 2.8. Weakly coupled in energy, the within-group equations are solved sequentially with the most up-to-date available information for the group-to-group sources; whereby the source terms are now linear functionals that can be moved to the right-hand side of the equation. In the presence of extraneous fixed sources, $\hat{\phi}^g \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ is solved for $g = 1, \ldots, G$:

$$
a^g_L(\hat{\phi}^g, \hat{w}) + a^g_{IP}(\hat{\phi}^g, \hat{w}) + a^g_{ext}(\hat{\phi}^g, \hat{w}) = a^g_S(\hat{\phi}, \hat{w}) + a^g_F(\hat{\phi}, \hat{w}) + b^g(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}). \tag{4.82}
$$

The left-hand side, $a^g_L + a^g_{IP} + a^g_{ext}$, is SPD, for $\theta = 1$ and the above can be solved by the PCG method; and additional savings can be made for the storage of symmetric matrices.

Employing the PI procedure to converge $|1/\mathrm{k_{eff}}|$, the eigenvalue problem is cast as finding $\hat{\phi}^g \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$

for $g = 1, \ldots, G$:

$$a_L^g(\hat{\phi}^g, \hat{w}) + a_{IP}^g(\hat{\phi}^g, \hat{w}) + a_{ext}^g(\hat{\phi}^g, \hat{w}) = a_S^g(\hat{\phi}, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a_F^g(\hat{\phi}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}). \tag{4.83}$$

In kind, the discrete variational statement of the adjoint problem in the presence of fixed adjoint sources, $q_x^\dagger \neq 0$, is cast as finding $\hat{\phi}^{g\dagger} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ for $g = G, \ldots, 1$:

$$a_L^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) + a_{IP}^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) + a_{ext}^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) = a_S^{g\dagger}(\hat{\phi}^\dagger, \hat{w}) + a_F^{g\dagger}(\hat{\phi}^\dagger, \hat{w}) + b^{g\dagger}(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}); \tag{4.84}$$

And the adjoint eigenvalue problem is cast as finding $\hat{\phi}^{g\dagger} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ for $g = G, \ldots, 1$:

$$a_L^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) + a_{IP}^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) + a_{ext}^{g\dagger}(\hat{\phi}^{g\dagger}, \hat{w}) = a_S^{g\dagger}(\hat{\phi}^\dagger, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a_F^{g\dagger}(\hat{\phi}^\dagger, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}); \tag{4.85}$$

where the discrete adjoint bilinear form, $a^\dagger(\cdot,\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}_{h,p}^*(\mathcal{T}) \to \mathbb{R}$, is defined by:

$$a^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) = a_L^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) - a_S^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) - a_F^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) + a_{IP}^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) + a_{ext}^\dagger(\hat{\phi}^{\dagger *}, \hat{w}). \tag{4.86}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$\begin{aligned} a_L^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G} \left[\sum_{V_e \in \mathcal{T}} \int_{V_e} \vec{\nabla}\hat{w} \cdot D_e^g \vec{\nabla}\hat{\phi}^{g\dagger *} + \hat{w}\Sigma_{r,e}^g \hat{\phi}^{g\dagger *} \, dV \right], \\ &= \sum_{g=1}^{G} a_L^{g\dagger}(\hat{\phi}^{g\dagger *}, \hat{w}); \end{aligned} \tag{4.87a}$$

$$\begin{aligned} a_S^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G} \left[\sum_{g\prime \neq g}^{G} \sum_{V_e \in \mathcal{T}} \int_{V_e} \hat{w}\Sigma_{s,e}^{g \to g\prime} \hat{\phi}^{g\prime\dagger *} \, dV \right], \\ &= \sum_{g=1}^{G} a_S^{g\dagger}(\hat{\phi}^{\dagger *}, \hat{w}); \end{aligned} \tag{4.87b}$$

$$\begin{aligned} a_F^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G} \left[\sum_{g\prime=1}^{G} \sum_{V_e \in \mathcal{T}} \nu_e^g \Sigma_{f,e}^g \int_{V_e} \hat{w}\chi_e^{g\prime} \hat{\phi}^{g\prime\dagger *} \, dV \right], \\ &= \sum_{g=1}^{G} a_F^{g\dagger}(\hat{\phi}^{\dagger *}, \hat{w}); \end{aligned} \tag{4.87c}$$

$$\begin{aligned} a_{IP}^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G} \Bigg[\sum_{\partial V_f \in \partial V_{IP}} - \int_{\partial V_f} [\![\hat{w}]\!] \{\{D^g \vec{\nabla}\phi^{g\dagger *} \cdot \vec{n}\}\} \, dS - \theta \int_{\partial V_f} \{\{D^g \vec{\nabla}\hat{w} \cdot \vec{n}\}\} [\![\hat{\phi}^{g\dagger *}]\!] \, dS \\ &\quad + \sum_{\partial V_f \in \partial V_{IP}} \mu_f^g \int_{\partial V_f} [\![\hat{w}]\!] [\![\hat{\phi}^{g\dagger *}]\!] \, dS \Bigg], \\ &= \sum_{g=1}^{G} a_{IP}^{g\dagger}(\hat{\phi}^{g\dagger *}, \hat{w}); \end{aligned} \tag{4.87d}$$

$$\begin{aligned} a_{ext}^\dagger(\hat{\phi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G} \left[\sum_{V_e \in \mathcal{T}} \sum_{\partial V_f \in V_e \cap \partial V_V} \frac{1}{2} \int_{\partial V_f} \hat{w}\hat{\phi}^{g\dagger *} \, dS \right], \\ &= \sum_{g=1}^{G} a_{ext}^{g\dagger}(\hat{\phi}^{g\dagger *}, \hat{w}); \end{aligned} \tag{4.87e}$$

and the discrete adjoint linear functional, $b^\dagger(\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \to \mathbb{R}$ is defined as:

$$\begin{aligned} b^\dagger(\hat{w}) &= \sum_{g=1}^{G} \left[\sum_{V_e \in \mathcal{T}} \int_{V_e} \hat{w} q_{x,e}^{g\dagger} \, dV \right], \\ &= \sum_{g=1}^{G} b^{g\dagger}(\hat{w}). \end{aligned} \tag{4.88}$$

For eigenvalue problems, one recalls that the solution is normalised after each outer iteration such that $a_F(\hat{\phi}, \hat{\phi}) = 1$; and the adjoint solution is normalised such that $a_F(\hat{\phi}, \hat{\phi}^\dagger) = 1$; q.v. Section 3.6.2.2.

4.4.1 A Discrete Coercivity Analysis

The performance of the SIP-DG discretisation will depend heavily upon the choice of penalty parameters; viz. they must be large enough to ensure definiteness and invertibility of the matrix equations,

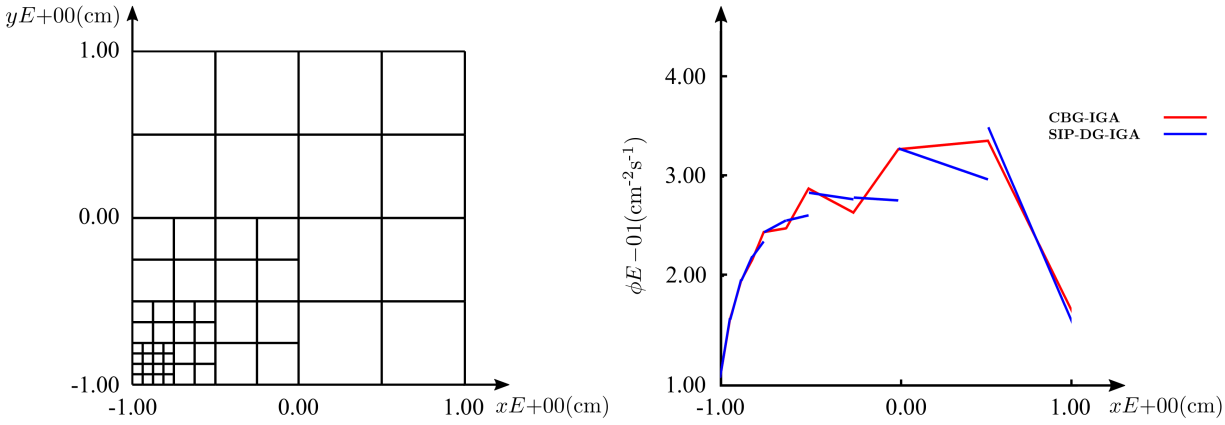


Figure 4.3: Another comparison between the solutions (right) of an SIP-DG-IGA discretisation and its CBG-IGA counterpart along the line at y = -1.00 cm over the irregular mesh (left) of a toy problem; v. Figure 4.1. Across 52 bilinear patches, strict C^0 continuity enforced in the CBG discretisation (73 DoF), whilst this is relaxed in the DG discretisation (208 DoF). (*Vide* the web-based version for reference to colour.)

yet not arbitrarily large that the iterative smoother is slow to converge; and an excessive penalisation at sharp boundary layers may even trigger the Gibbs phenomenon; i.e. the SIP-DG discretisation tends to a CBG one as $\mu_f \to \infty$ [345,347]. Optimal rates of convergence in the H^1-norm and the L_2-norm have been proven for SIP-DG schemes for $\mu_f \propto \mathcal{O}(h^{-1})$ [288,333]. More rigorously, one seeks the lower bound of the *Goldilocks* range of the penalty parameters via an analysis of coercivity following the methodology of Shahbazi [348].

Since the discretised system is only weakly coupled in energy, it is sufficient to consider separately the coercivity of each within-group bilinear form, the left-hand side of Equation 4.82 or Equation 4.83; where a^g_{lhs} is recalled below for some $\hat{u} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, suppressing the spatial dependence for $g = 1, \ldots, G$:

$$a^g_{lhs}(\hat{u},\hat{u}) = a^g_L(\hat{u},\hat{u}) + a^g_{IP}(\hat{u},\hat{u}) + a^g_{ext}(\hat{u},\hat{u}), \tag{4.89a}$$

$$\begin{aligned} &= \sum_{V_e \in \mathcal{T}} \int_{V_e} D^g_e \vec{\nabla}\hat{u} \cdot \vec{\nabla}\hat{u} + \Sigma^g_{r,e}(\hat{u})^2 \, dV - \sum_{\partial V_f \in \partial V_{IP}} (1+\theta) \int_{\partial V_f} [\![\hat{u}]\!] \{\{ D^g \vec{\nabla}\hat{u} \cdot \vec{n} \}\} \, dS \\ &+ \sum_{\partial V_f \in \partial V_{IP}} \mu^g_f \int_{\partial V_f} ([\![\hat{u}]\!])^2 \, dS + \sum_{\partial V_f \in \partial V_V} \frac{1}{2} \int_{\partial V_f} (\hat{u})^2 \, dS. \end{aligned} \tag{4.89b}$$

The within-group bilinear form is self-adjoint, $a^g_{lhs} = a^{g\dagger}_{lhs}$, for $\theta = 1$; and since the discrete variational statements for both the primal problem and the dual problem are formulated on the same mesh and their solutions sought in the same space, $\phi, \phi^\dagger \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, the well-posedness of either depends upon a common coercivity analysis that will yield a single set of energy-dependent penalty parameters.

The only potentially negative term and therefore the only one that may not *contribute* to coercivity is that in the second term of the second line of Equation 4.89b. Whilst the NIP-DG scheme is coercive for any $\mu_f > 0$, a lower bound must be imposed for the IIP and the SIP variations. For reasons already discussed, only the latter will be investigated; however, for completeness, the variable θ is retained.

Sufficient penalisation of each face $\partial V_f \in \partial V_{IP}$, should ensure the coercivity of the bilinear form with respect to the energy-norm, which is defined for some $\hat{u} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ and for $g = 1, \ldots, G$:

$$\|\hat{u}\|^2_{\hat{\epsilon}^g} \coloneqq \sum_{V_e \in \mathcal{T}} \int_{V_e} D^g_e \vec{\nabla}\hat{u} \cdot \vec{\nabla}\hat{u} + \Sigma^g_{r,e}(\hat{u})^2 \, dV + \sum_{\partial V_f \in \partial V_{IP}} \mu^g_f \int_{\partial V_f} (\llbracket \hat{u} \rrbracket)^2 \, dS. \tag{4.90}$$

Most popular in the literature are expressions for the penalty parameters of the general form:

$$\mu_f = \rho \frac{\partial V_f}{V_e} \approx \frac{\rho}{h}, \quad \forall \partial V_f \in V_e \cap \partial V_{IP}; \tag{4.91}$$

whereby some *large enough* user-dependent parameter, ρ, is scaled by the mesh spacing, h; and quite often, only the largest value is retained. More rigorously, Shahbazi establishes a lower bound for the penalty parameters through an analysis of the coercivity of the bilinear form [348]. Warburton and Hesthaven and separately Hillewaert adopted this same methodology and derived closed-form expressions for μ_f for particular straight-sided geometric primitives such as simplices; and quadrilaterals and hexahedra, respectively [345,349]. The current work follows more closely those developments made by Owens et al; viz. agnostic of the element type, the penalty parameters are calculated via generalised eigenvalue problems posed over the discontinuous faces [334].

One recalls Young's inequality for the product of *two* scalars, a and b:

$$ab \leq \frac{\epsilon a^2}{2} + \frac{b^2}{2\epsilon}, \quad \forall \epsilon > 0. \tag{4.92}$$

From the definition of jumps and averages, Equation 4.58a and Equation 4.58b, respectively, and then by letting $a = \llbracket \hat{u} \rrbracket$ and $b = \{\{D^g \vec{\nabla}\hat{u} \cdot \vec{n}\}\}$, the substitution of Equation 4.92 into Equation 4.89b bounds the potentially negative term in the second line over all discontinuous faces, $\forall \partial V_f \in \partial V_{IP}$. Collecting like-terms for $g = 1, \ldots, G$:

$$\begin{aligned} a^g_{lhs}(\hat{u}, \hat{u}) \geq a^g_L(\hat{u}, \hat{u}) &- \sum_{\partial V_f \in \partial V_{IP}} \frac{(1+\theta)}{2\epsilon_f} \int_{\partial V_f} (\{\{D^g \vec{\nabla}\hat{u} \cdot \vec{n}\}\})^2 \, dS \\ &+ \sum_{\partial V_f \in \partial V_{IP}} \left(\mu^g_f - \frac{(1+\theta)\epsilon_f}{2} \right) \int_{\partial V_f} (\llbracket \hat{u} \rrbracket)^2 \, dS + a^g_{ext}(\hat{u}, \hat{u}). \end{aligned} \tag{4.93}$$

By the definition of Equation 4.58b, the term $(\{\{D^g \vec{\nabla}\hat{u}\}\} \cdot \vec{n})^2$ is expanded out, to which Equation 4.92 is applied with $\epsilon = 1$; i.e. $a^2 + b^2 + 2ab \leq 2a^2 + 2b^2$. The result is rewritten as a sum of integrals over patch boundaries, where individual terms are collected onto their respective patches; i.e.:

$$\begin{aligned} a^g_{lhs}(\hat{u}, \hat{u}) \geq a^g_L(\hat{u}, \hat{u}) &- \sum_{V_e \in \mathcal{T}} (1+\theta) \Bigg[\sum_{\partial V_f \in V_e \cap \partial V_{int}} \frac{(D^g_e)^2}{4\epsilon_f} \int_{\partial V_f} (\vec{\nabla}\hat{u} \cdot \vec{n})^2 \, dS \\ &+ \sum_{\partial V_f \in V_e \cap \partial V_D} \frac{(D^g_e)^2}{2\epsilon_f} \int_{\partial V_f} (\vec{\nabla}\hat{u} \cdot \vec{n})^2 \, dS \Bigg] \\ &+ \sum_{\partial V_f \in \partial V_{IP}} \left(\mu^g_f - \frac{(1+\theta)\epsilon_f}{2} \right) \int_{\partial V_f} (\llbracket \hat{u} \rrbracket)^2 \, dS + a^g_{ext}(\hat{u}, \hat{u}). \end{aligned} \tag{4.94}$$

Taking advantage of the equivalence of norms in the chosen function space, there exist discrete trace inequalities that bound from above integral quantities of the restriction of a function, or the normal component of its gradient, as defined over the boundary ∂V_e; i.e. for some function $\hat{u} \in \hat{\mathcal{V}}_{h,p}(V_e)$:

$$\|\gamma_0 \hat{u}\|^2_{L_2(\partial V_e)} \leq C^0_{e,f} \|\hat{u}\|^2_{L_2(V_e)}; \qquad (4.95a) \qquad \|\gamma_1 \hat{u}\|^2_{L_2(\partial V_e)} \leq C^1_{e,f} \|\vec{\nabla} \hat{u}\|^2_{L_2(V_e)}. \qquad (4.95b)$$

For $i = 0, 1$, $C^i_{e,f}$ are the trace inequality constants (TIC) that pertain to the Dirichlet and the Neumann traces in Equation 4.56a and Equation 4.56b, respectively. Those integrals over patch boundaries in Equation 4.94 are substituted for Equation 4.95b; and one seeks the sharpest TIC values to optimise the penalty parameters. Collecting like-terms and comparing against Equation 4.79, the following must be satisfied for $g = 1, \cdots, G$:

$$1 - (1+\theta)\left[\sum_{\partial V_f \in V_e \cap \partial V_{int}} \frac{D^g_e C^1_{e,f}}{4\epsilon_f} + \sum_{\partial V_f \in V_e \cap \partial V_D} \frac{D^g_e C^1_{e,f}}{2\epsilon_f}\right] > 0, \qquad \forall V_e \in \mathcal{T}^g(V); \qquad (4.96a)$$

$$\mu^g_f - \frac{(1+\theta)\epsilon_f}{2} > 0, \qquad \forall \partial V_f \in \partial V_{IP}. \qquad (4.96b)$$

One must choose $\epsilon_f > 0$, such that Equation 4.96a is satisfied over each element. In anisotropic fashion, Drosson et al achieved this by bounding from below each separate term in the sum over the F_e-number of faces of a given element [350]. Furthermore, the maximum necessary penalisation is applied at a given interface to ensure that Equations 4.96 hold over every element that intersects at that interface. Encapsulated below, the penalty parameters can be calculated in a pre-processing step for $g = 1, \cdots, G$ and $\forall \partial V_f \in \partial V_{IP}$:

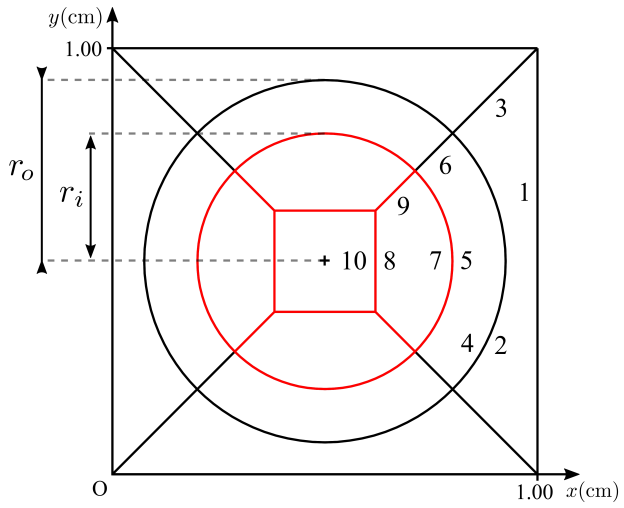
$$\mu^g_f > \begin{cases} \frac{1}{8}(1+\theta)^2 \max\limits_{V_e \in \bigcap\limits_{e'} V_{e'} \ni \partial V_f} D^g_e C^1_{e,f} F_e, & \partial V_f \in \partial V_{int}; \\ \frac{1}{4}(1+\theta)^2 D^g_e C^1_{e,f} N_e, & \partial V_f \in \partial V_D. \end{cases} \qquad (4.97a)$$

4.4.2 The Trace Inequality Constants

In search of the smallest possible penalty parameters, a coercivity analysis has revealed the dependence on a constant, $C^1_{e,f}$, from an inequality that bounds from above the Neumann trace of a function; q.v. Equation 4.95b. The μ_f are proportional to the corresponding TIC and thus, the problem of calculating *optimised* penalty parameters becomes one of minimising these TIC [334].

Warburton and Hesthaven provide closed-form expressions for the TIC of constant-Jacobian simplices [351]. Similarly, Hillewaert provides those for constant-Jacobian quadrilaterals and hexahedra [349]:

$$C^{lit}_{e,f} = (p+1)^2 \frac{A(\partial V_f)}{V(V_e)}; \qquad (4.98)$$



∂V_f	TICs C_f^{lit*}	C_f^{0*}	C_f^{1*}	C_f^0/C_f^1
$7a$	9.00E+00	1.26E+01	5.62E+02	< 0.1
$7b$	9.00E+00	8.21E+00	3.68E+00	~ 2.2
8	9.00E+00	1.39E+01	1.12E+01	~ 1.2
9	9.00E+00	1.41E+01	1.18E+01	~ 1.2
10	9.00E+00	9.00E+00	4.00E+00	~ 2.3

Figure 4.4: For the example of a 2D unit pin-cell grid of 13 non-degenerate bi-quadratic NURBS patches (left), a comparison is made between those TIC, $C_{e,f}^{lit}$, $C_{e,f}^{0}$ and $C_{e,f}^{1}$, as calculated from the literature, Equation 4.98, the mass-matrix method, Equation 4.95a, and the stiffness-matrix method, Equation 4.95b, respectively. The *ten* distinct faces are labelled 1 through 10; and all TIC values are normalised by $\times V(V_e)/A(\partial V_f)$. The distinction between faces $7a$ and $7b$ is such that the former is the the single distinct face of the central (red) circle constructed from a single degenerate NURBS patch; whereas the latter belongs to the set of distinct faces that belong to the central circle as it is presented, constructed from *five* non-degenerate NURBS patches. Those TIC values that are invariant with any change to the radii are presented in the table (right). (*Vide* the web-based version for reference to colour.)

where p denotes the polynomial degree of the element; $A(\partial V_f)$ the area of ∂V_f; and $V(V_e)$ the volume of V_e. Such expressions, however, are derived from a different trace inequality; one that bounds the function itself and not its gradient; and whilst this may yield suitable penalty parameters for straight-sided linear finite elements, they preclude any curved elements, or rather those that possesses a non-linear mapping; viz. NURBS-based IGA. For this reason, Owens et al proposed a stiffness-matrix-based method, derived from the correct trace inequality, to calculate provably sharp TIC for any general element face [278].

First, the function in Equation 4.95b is expanded in a basis of NURBS, such that for $\hat{u} \in \hat{\mathcal{V}}_{h,p}(V_e)$:

$$\vec{\nabla}\hat{u}(\vec{r}) = \sum_{i=1}^{N_e} u_i \vec{\nabla} R_i(\vec{r}) = \underset{(1\times N_e)}{[\underline{u}]^T} \cdot \underset{(N_e\times d)}{\nabla \underline{\underline{R}}}, \qquad \vec{r} \in V_e;$$

and similarly:

$$\vec{\nabla}\hat{u}(\vec{r}) \cdot \vec{n}(\vec{r}) = \sum_{i=1}^{N_e} u_i \vec{\nabla} R_i(\vec{r}) \cdot \vec{n}(\vec{r}) = \underset{(1\times N_e)}{[\underline{u}]^T} \cdot \underset{(N_e\times d)}{\nabla \underline{\underline{R}}} \cdot \underset{(d\times 1)}{\vec{n}}, \qquad \vec{r} \in \partial V_f;$$

where $\underline{u}$ is the column vector of expansion coefficients and $\nabla \underline{\underline{R}}$ is a matrix that stores the gradient of each of the N_e-number of NURBS in each of the d-number of spatial directions as a separate column vector; i.e. $(\nabla \underline{\underline{R}})_{ij} = \partial R_i / \partial x_j \in \mathbb{R}^{N_e \times d}$. Equation 4.95b can thus be rewritten in matrix notation:

$$[\underline{u}]^T \cdot \underline{\underline{S}}_{\partial V_f} \cdot \underline{u} \leq C_{e,f}^1 [\underline{u}]^T \cdot \underline{\underline{S}}_V \cdot \underline{u}. \tag{4.99}$$

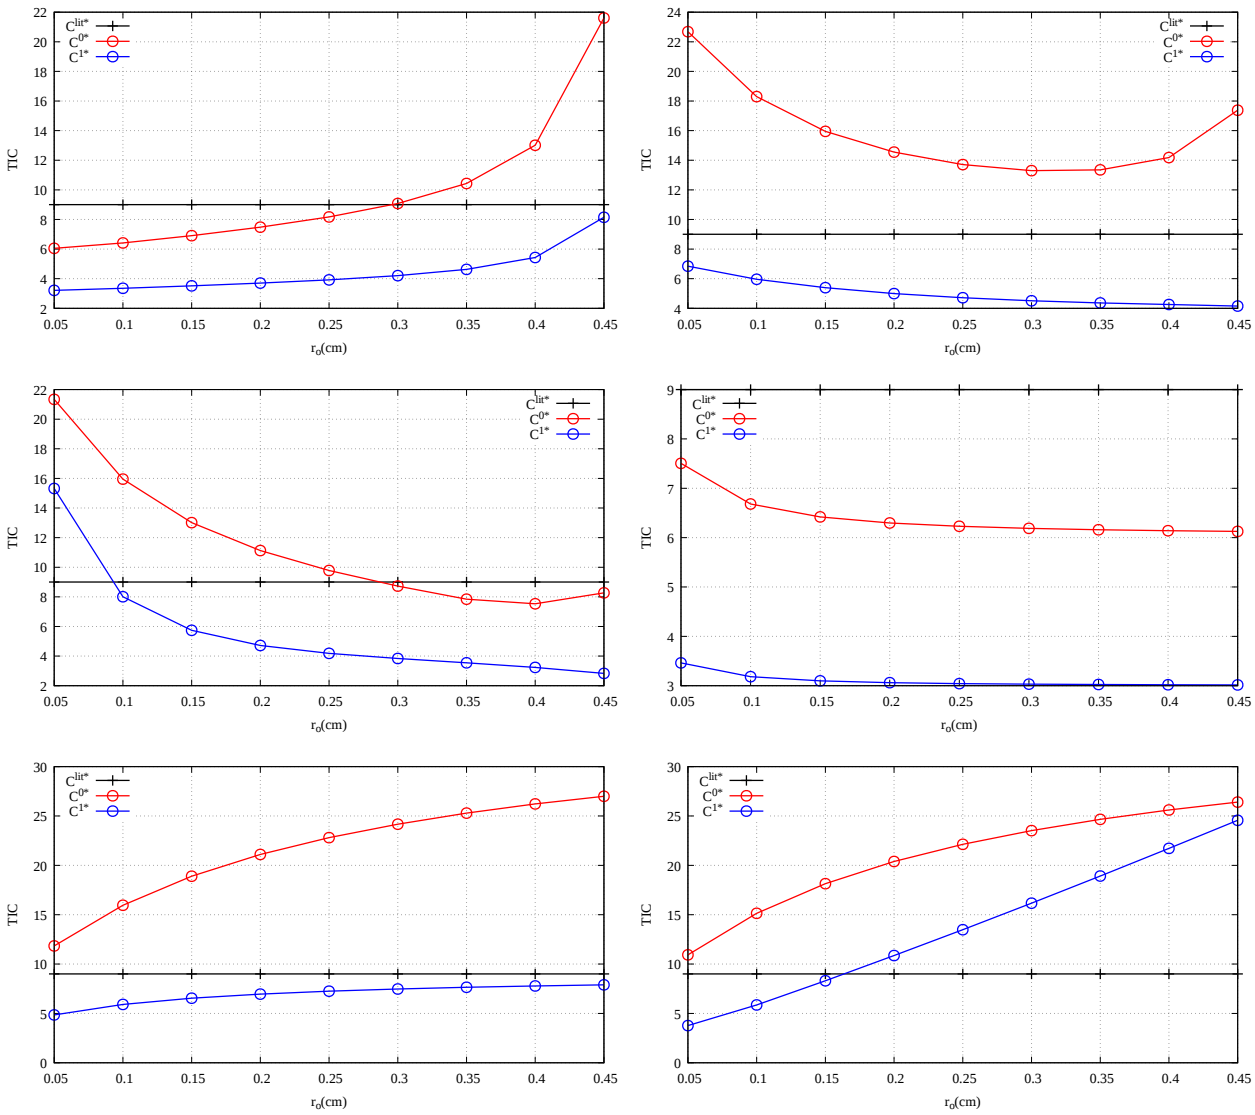


Figure 4.5: A parametric study that compares TIC values for the distinct faces ∂V_1 (top left), ∂V_2 (top right), ∂V_3 (middle left), ∂V_4 (middle right), ∂V_5 (bottom left), and ∂V_6 (bottom right) of a 2D unit pin-cell grid as the outer-radius is varied; v. Figure 4.4. In each instance, the volume patch and the face patches may be defined by r_o, which is varied between 0.05 cm and 0.45 cm. In the case of ∂V_4, ∂V_5 and ∂V_6, r_i is fixed at 0.025 cm. (*Vide* the web-based version for reference to colour.)

The matrix $\underline{\underline{S}}_{\partial V_f}$ is constructed over a given face of a given patch and the matrix $\underline{\underline{S}}_{V_e}$ is constructed over the volume of that same patch; the entries of each are given below, respectively:

$$(\underline{\underline{S}}_{\partial V_f})_{ij} = \left[\int_{\partial V_f} \vec{\nabla} R_i(\vec{r}) \cdot \vec{n}(\vec{r}) \vec{\nabla} R_j(\vec{r}) \cdot \vec{n}(\vec{r}) \, dS \right] \in \mathbb{R}^{N_e \times N_e}; \tag{4.100a}$$

$$(\underline{\underline{S}}_{V_e})_{ij} = \left[\int_{V_e} \vec{\nabla} R_i(\vec{r}) \cdot \vec{\nabla} R_j(\vec{r}) \, dV \right] \in \mathbb{R}^{N_e \times N_e}. \tag{4.100b}$$

Equation 4.99 can be re-posed as a generalised eigenvalue problem that possesses N_e-number of discrete eigenvalues, $\lambda_i \geq 0$ for $i = 1, \ldots, N_e$:

$$\begin{cases} (\underline{\underline{S}}_{\partial V_f} - \lambda \underline{\underline{S}}_V) \cdot \underline{x} = 0, \\ \text{with } [\underline{x}]^T \cdot \underline{\underline{S}}_V \cdot \underline{x} = 1. \end{cases} \tag{4.101}$$

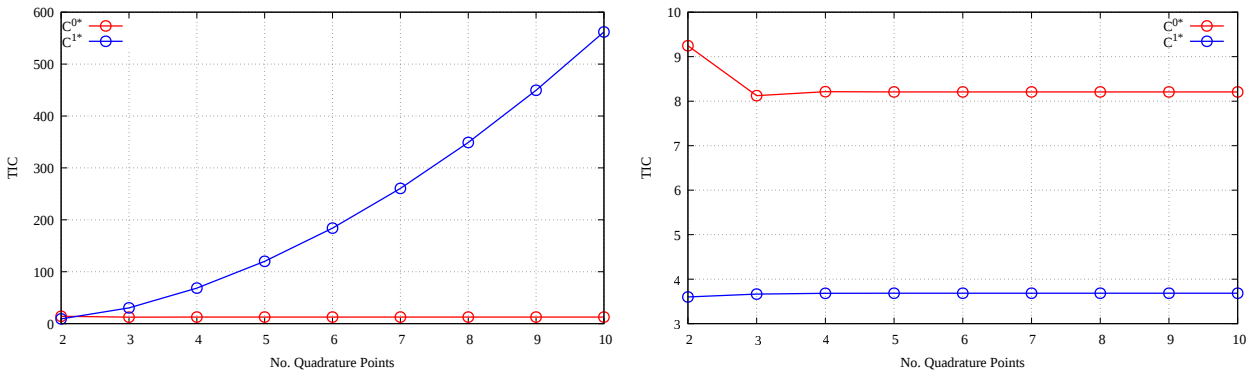


Figure 4.6: A parametric study that compares TIC values for the distinct faces $7a$ (left) and $7b$ (right) of a 2D unit pin-cell grid as the number of QP is varied; v. Figure 4.4. The distinction between faces $7a$ and $7b$ is such that the former is the the single distinct face of the central (red) circle constructed from a single degenerate NURBS patch; whereas the latter belongs to the set of distinct faces that belong to the central circle constructed from *five* non-degenerate NURBS patches. In each instance, the TIC are invariant with any change to the radii. (*Vide* the web-based version for reference to colour.)

The maximum eigenvalue of which corresponds to that sought-after and provably sharp TIC in Equation 4.95b; i.e. $C^1_{e,f} = \max_i \lambda_i$ [334]. That being said, both of the matrices in the above are only symmetric positive semi-definite (SPSD) and do not form a regular pair; accordingly, neither will the generalised eigenvalue problem possess a spectrum of distinctly positive eigenvalues; and there exits a non-trivial intersection of the matrices' null-spaces; i.e. $\text{null}\,(\underline{\underline{S}}_{\partial V_f}) \cap \text{null}\,(\underline{\underline{S}}_{V_e}) \neq \emptyset$. Easily enough, however, this common null-space is identified as that 1D space spanned by the vector that represents the constant function, $\underline{u}_k$, in $\hat{\mathcal{V}}_{h,p}(V_e)$, which, as per the positivity of the right-hand side of Equation 4.95b, corresponds to the degenerate eigenvector associated with the single *zero* eigenvalue. Fortunately, the NURBS-based basis forms a partition of unity, such that the constant function is simply an N_e-dimensional vector of ones; i.e. $\underline{u}_k = [1, \ldots, 1]^T$.

It is possible to remove this singularity and compute the eigenvalues associated with the restriction of the pair of matrices to the complement of the null space [352]. Practically, this can be achieved via the (modified) Gram-Schmidt procedure, which forms an orthonormal set, $\underline{\underline{Z}}$, that spans the subspace of $\mathbb{R}^{N_e}$ by the orthogonal projection of some arbitrary input basis, $\underline{\underline{U}}$, which contains the constant function as the kth basis vector. Denoting $\underline{\underline{Z}}\prime$ as the orthogonal ($N_e \times N_e$-1)-matrix less the kth column, the transformation to the (N_e-1)-dimensional space $\mathbb{R}^{N_e} \setminus \text{span}\,\underline{u}_k$ is performed as:

$$\underline{\underline{S}}_{\partial V_f}\prime = [\underline{\underline{Z}}\prime]^T \cdot \underline{\underline{S}}_{\partial V_f} \cdot \underline{\underline{Z}}\prime; \qquad (4.102a) \qquad\qquad \underline{\underline{S}}_{V_e}\prime = [\underline{\underline{Z}}\prime]^T \cdot \underline{\underline{S}}_{V_e} \cdot \underline{\underline{Z}}\prime; \qquad (4.102b)$$

where the singularity has been orthogonalised out of the function space and the shared kernel removed such that the matrices $\underline{\underline{S}}_{\partial V_f}\prime$ and $\underline{\underline{S}}_{V_e}\prime$ are of dimension one-less than the originals. Although the former is still SPSD, the latter is now SPD and the LAPACK routine `DSYGV` can be employed to solve for the maximum eigenvalue by means of a Cholesky (LU) factorisation [296]. Importantly, the positive eigenvalues of the original generalised eigenvalue problem are preserved such that the new pair of matrices is said to be equivalent to the original pair [352]. In the interest of a wholly-local

AMR strategy, an explicit orthogonalisation of the matrices in Equation 4.102 is not expected to be computationally expensive since an h-refinement forms new discontinuous patches; consequently, the dimensions of any matrices in Equations 4.100 will remain relatively small; i.e. $N_e \sim \mathcal{O}(10^1)$.

It is worth comparing the different TIC obtained from Equation 4.95a and Equation 4.95b, for the example of a 2D unit pin-cell grid, a geometry that is ubiquitous in reactor physics; v. Figure 4.4. The $C^0_{e,f}$ are calculated using a mass-matrix approach; viz. substitution of the stiffness matrices in Equation 4.101 for mass matrices, which do not need to be orthogonalised. The parametric study of TIC values is inspired by that of Owens et al; however, theirs did not allow for a fair comparison since it was performed over a mesh of degenerate NURBS patches [334]. As per the stiffness-matrix method, the derivatives of the volumetric NURBS must be square-integrable at the boundary and thus, the patch must be non-degenerate. Otherwise, $C^1_{e,f}$ are quadrature-dependent, becoming arbitrarily large and increasingly so for an increasingly large quadrature set; q.v. Section 3.3.1. For this reason, the authors concluded that the penalty parameters for the faces of the elements of a pin-cell grid may be underestimated by as much as $\times 50$ when employing the literature expressions, Equation 4.98, or the mass-matrix approach to calculate the TIC.

The geometry for the current analysis is constructed from 13 bi-quadratic non-degenerate NURBS patches; v. Figure 4.4. The central fuel pin is represented by a circle, constructed from *five* patches; the cladding is represented by an annulus, constructed from *four* patches; and so is the surrounding moderator region. Taking into account all lines of symmetry, there are only *four* distinct element types and only those *ten* distinct faces labelled 1 through 10 need to be considered. For each distinct face, $C^0_{e,f}$ and $C^1_{e,f}$ are computed; and to facilitate a further comparison with those TIC, $C^{lit}_{e,f}$, taken from the literature, Equation 4.98, all TIC values are normalised by $\times V(V_e)/A(\partial V_f)$. Numerical integration is performed consistently with a Legendre-Gauss quadrature set of *ten*, $Q = 10$, in each parametric direction over each element.

The moderator region is defined uniquely by the outer-radius, which is varied between 0.05 cm and 0.45 cm; and the results for faces 1, 2 & 3 are presented in the top left, the top right and the middle left, respectively, of Figure 4.5. The shape of the annular region is defined by the ratio of the radii; therefore, the inner-radius was fixed at 0.025 cm and the outer-radius was varied between 0.05 cm and 0.45 cm; and the results for faces 4, 5 & 6 are presented in the middle right, the bottom left and the bottom right, respectively, of Figure 4.5. The shape of the central pin, as represented by *five* non-degenerate patches, is defined by the inner-radius, which is varied between 0.05 cm and 0.45 cm; however, the normalised TIC are invariant with r_i and these results for faces 7*b*, 8, 9 & 10 are presented in the table in Figure 4.4.

For completeness, a comparison is included for those TIC evaluated over the surface of degenerate

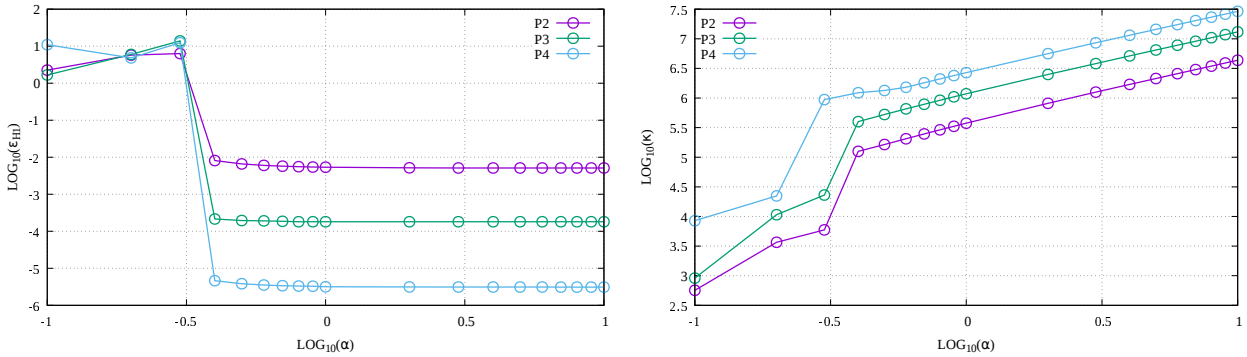


Figure 4.7: A parametric study of the H^1-error (left) and the spectral conditioning (right) of the SIP-DG-IGA discretisation of the 1G NDE for the example of a 2D unit pin-cell grid; qq.v. Table 4.1 and Section 4.5.1.2. The penalty parameters of each face are varied as $\mu_f^* = \alpha \times \mu_f$. (*Vide* the web-based version for reference to colour.)

NURBS patch; viz. the central pin is constructed instead from a single degenerate NURBS patch. The distinction between faces $7a$ and $7b$ is such that the former is the the single distinct face of the central (red) circle constructed from a single degenerate patch; whereas the latter belongs to the set of distinct faces that belong to the central circle constructed from *five* non-degenerate patches. Similarly, those TIC over face $7a$ are invariant with r_i; and one observes that $C_{7a}^0 < 0.01 \times C_{7a}^1$. However, this is deceiving and clearly dependent upon the order of the numerical quadrature set; v. Figure 4.6.

In the consideration of constant-Jacobian elements, both the literature method and the mass-matrix approach yield the same result; or rather, an overestimation of the appropriate TIC value. In the consideration of non-constant-Jacobian elements, again, the mass-matrix approach yields TIC values larger than those of the stiffness-matrix method. This was also generally true of the literature method, with the exception of face 6 for an increasingly large outer-radius. Whilst it has not been demonstrated that use of the mass-matrix approach would jeopardise the coercivity of the bilinear form; yet, when information has not been discarded and the gradient of the function has been taken into account, the stiffness-matrix method allows for sharper bounds on the resulting penalty parameters. Moreover, the proposed approach is robust in its application to any element type.

4.4.3 Optimised Penalty Parameters

For completeness, the effect of varying the penalty parameters for an SIP-DG-IGA discretisation of the 1G NDE is studied for the example of a 2D unit homogeneous pin-cell with reflective BC; qq.v. Table 4.1 and Section 4.5.1.2. The μ_f from Equation 4.97 for each face are varied as $\mu_f^* = \alpha \times \mu_f$, for $\alpha = [0, 10]$, and the effect on the discretisation error and the conditioning of the system is presented in Figure 4.7 for $p = 2, 3, 4$; for each dataset, the mesh is $9 \times$ KI-h-refined globally and uniformly. The was also studied for $9 \times$ KI-1-refinement and the same conditioning effect was observed; however, the

discretisation error was less sensitive to any variation in μ_f because of the higher-order continuity in the basis.

From Figure 4.7 the effect of under-penalisation, apparent by the poor values of the H^1-error; and although not obvious from the scale, the error attains a minimum for the choice of $\mu_f^* = \mu_f$ and thereafter plateaus for $\alpha > 1$; viz. there is no gain in accuracy for super-penalisation. Another obvious effect is that on the conditioning, $\kappa(\underline{\underline{A}})$, of the discretised system; q.v. Equation 1.2; viz. the matrix becomes increasingly ill conditioned for increasingly large μ_f. Taking $\kappa(\underline{\underline{A}})$ as the ratio of the largest to the smallest eigenvalues of $\underline{\underline{A}}$, the spectral condition number is used to gauge the number of iterations that are required to solve the matrix equation. The eigenvalues are computed using the PETSc GMRes solver with the flag `-ksp_monitor_singular_value` [249]. Rather than absolute values, one is more interested in the general trends that varying μ_f has on the conditioning of the system; and whilst the approximation of the largest singular value is often fairly accurate, that of the smallest is only typically so upon convergence; thus, *restarts* are avoided by setting the restart parameter to $10,000$.

The effect of the polynomial degree is also noticeable in Figure 4.7, that is, increasingly better accuracy can be attained for the same number of elements belonging to the mesh for bases constructed from increasingly higher-order polynomials. This only has implications for the asymptotic convergence ranges and assumes that the analysis is performed over a sufficiently graded mesh; and may suggest an instability in a basis of higher-order polynomials over a coarse mesh [312]. Furthermore, the discrete system becomes increasingly more dense as the polynomial degree of the basis increases and it becomes increasingly more expensive to smoothen.

4.5 The Numerical Results

In the application to the multi-group NDE, this section presents some numerical results for the AMR-$h(p)$ strategies developed in Section 3.5 and Section 3.6. Unless otherwise stated, all operations are performed such as to optimise the discretisation error in the H^1-(semi)norm. For a couple of verification test cases and nuclear reactor physics benchmarks, the AMR strategies are compared against global uniform ones. Any data-plots that make reference to the DoF of the discretised system justly account for all unknowns that were solved for in the complete system; viz. $DoF = G \times N$ for G-number of groups solved over an energy-independent mesh of N-number of control-variables.

Both discretisations developed within this chapter have been implemented in a new modern-Fortran code [73]. Adopting a distributed memory computing architecture, one makes use of the Open MPI message passing library; as well as the PETSc library and its capacity for parallelisation of its sparse

linear solvers and its compressed row storage (CRS) data structures for large sparse matrices [249]. One remarks that no sophisticated method of domain decomposition has been employed; rather, elements and their local information are rudely assigned to and held on different MPI parallel processes. All *jobs* are submitted to the general purpose Imperial College HPC cluster computing service (CX1). The within-group matrix equations are converged to machine precision using PETSc's CG smoother, preconditioned by the BoomerAMG package from the HYPRE library; a maximum number of 1E+06 inner iterations is imposed [353]. The scalar neutron flux is converged to a tolerance of $1.00\text{E}-06$ between successive source iterations; and the eigenvalue is converged to a tolerance of $1.00\text{E}-08$ between successive *power* iterations; a maximum number of 2E+03 outer iterations is imposed. An upper limit of *ten* is imposed upon the polynomial degree of the NURBS basis over the coarser current mesh; therefore $p_{max} = 11$ over the reference mesh.

A considerable reduction to the number of outer iterations required to attain convergence is achieved by *bootstrapping* the reference solution from the previous AMR iteration to provide a good initial guess for the computation over the reference mesh in the current i^{th} iteration; i.e. $\phi^{(i)}_{\text{h}/2,p+1} = \prod_{\text{h}/2,p+1}^{(i)} \phi^{(i\text{-}1)}_{\text{h}/2,p+1}$ [246].

4.5.1 The Method of Manufactured Solutions

The MMS is used to investigate the spatial discretisations of the multi-group NDE; cf. an overview by Salari et al [67]. Unlike most multi-dimensional problems, a smooth manufactured solution can be chosen to test against theoretical orders of accuracy [68, 242]:

$$|e|_{H^1} \leq C_1 h^{\min\{p,r\}}; \qquad (4.103a) \qquad \|e\|_{L_2} \leq C_2 h^{\min\{p+1,r\}}; \qquad (4.103b)$$

where $e = \phi - \hat{\phi}$ is the discretisation error; h is the characteristic mesh-spacing, p is the polynomial degree of the underlying NURBS basis; r denotes the regularity index of the true solution; and C_1 and C_2 are constants independent of h and p.

As per the tensor-product structure of the NURBS, Equations 3.7, the polynomial degree in the inequalities above, Equations 4.103, is actually $p = \min_i\{p_i\}$ for $i = \xi, \eta, \zeta$. Unfortunately, the current discretisations could not benefit from high rates of convergence whilst making savings in the description of the coarsest mesh; e.g. at least quartic NURBS must be used to tile the surface of a general spheroid in Cartesian coordinates, whilst linear NURBS would suffice in representing the solid volume as it varies radially from its centre to a point on the surface; however, the numerical method would be reduced to second-order accuracy with respect to Equation 4.103b.

For some manufactured solution, ϕ^g, the (semi-)analytical manufactured source, q^g_x, can be evaluated

for $g = 1, \ldots, G$:

$$q_x^g = (L^g - S^g)\phi. \tag{4.104}$$

The manufactured solution and source are used to study the convergence rates of the global uniform refinement strategy on a simple Cartesian grid and a more challenging pin-cell one; taking into account the implementation of different BC and the numerical quadrature routines. These results are then used as a reference with which to analyse the NRG-AMR-$h(p)$ strategy, which takes into account performance of all mesh-to-mesh interpolation procedures and the capacity of the numerical methods to enforce continuity over irregular meshes that contain *hanging-nodes*.

4.5.1.1 A Cartesian Grid

Material	g	D^g (cm)	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to 1}$ (cm^{-1})
(1) -	1	3.333E−01	1.000E+00	5.000E−01

Table 4.1: The nuclear data for the 2D 1G MMS convergence analysis.

The first verification case consists of a simple Cartesian grid defined over $V = [0, 1]^2$, which comprises of a homogeneous material whose properties are given in Table 4.1. The coarsest mesh is composed of *four* equal-volume squares represented by bilinear NURBS patches. Assuming a *one-speed* energy approximation, the manufactured solution is chosen as a standing wave equation:

$$\begin{cases} \phi^n(\vec{r}) = \cos\left(B_x^n x\right)\cos\left(B_y^n y\right), & \vec{r} \in V; \quad (4.105a) \\ B_x^n, B_y^n := (2n{+}1)\pi, & \forall n \geq 0; \quad (4.105b) \end{cases}$$

where ϕ^n denotes an odd harmonic mode as characterised by the geometric buckling, B_x^n and B_y^n. To allow for the maximum attainable convergence rates, the manufactured solution has been chosen to be smooth, yet it cannot be arbitrarily represented exactly by higher-order NURBS. Moreover, the homogeneous Neumann (reflective) BC is naturally satisfied; it will also be naturally enforced by the discretisations on all external boundaries.

Global Uniform Refinement For the KI-refinement of linear, quadratic and cubic B-splines, the H^1-errors for each of the CBG-IGA and the SIP-DG-IGA spatial discretisations are presented in Figure 4.8 for the manufactured solution chosen as the first, third and fifth harmonics of the standing wave equation; q.v. Equation 4.105. Qualitatively similar plots were obtained for the L_2-error; however, for brevity, these have not been presented. Under refinement, both the continuous and the discontinuous discretisation schemes converge to the analytical solution; and asymptotically, they attain their theoretical orders of accuracy, as presented in Table 4.2.

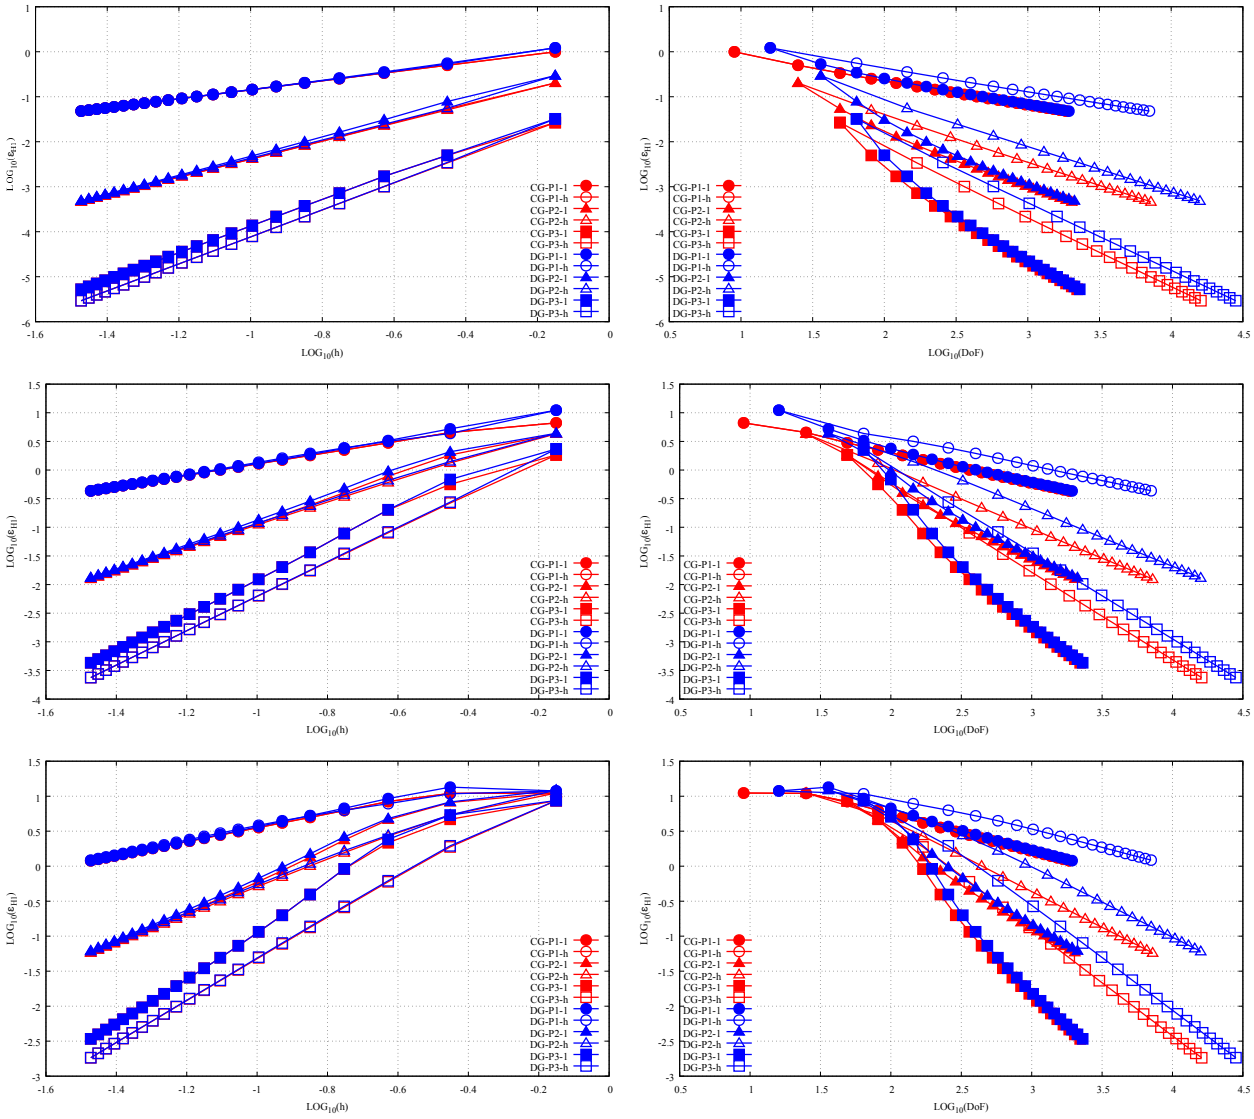


Figure 4.8: The convergence plots, ϵ_{H^1} v h (left) and ϵ_{H^1} v DoF (right), for the verification of the global uniform refinement strategy for the CBG-IGA and the SIP-DG-IGA 1G NDE over a 2D Cartesian grid; the manufactured solution is chosen as per Equation 4.105 for $n = 0$ (top row), $n = 1$ (middle row) and $n = 2$ (bottom row). (*Vide* the web-based version for reference to colour.)

Whilst the mesh spacing, h, does give some insight into how the discretisation error varies under refinement, it does not reveal the true computational effort required; viz. the number of elements in the mesh does not vary under DE-refinement. One must therefore also consider the DoF of the discrete system, which is also a better metric with which to compare non-uniform refinement strategies.

For such a homogeneous problem with a smooth underlying solution, exploiting higher-order continuity in the basis yields more accuracy per DoF than that compared to a similar h-refinement. With respect to the latter, one remarks that DG discretisations are generally more expensive than their continuous counterparts; in either case, h-refinement becomes increasingly more expensive for increasingly large polynomial degrees, such that, for the same accuracy, the relative cost of a DG discretisation decreases in comparison to a continuous one. In regard to the IP-DG discretisation, the DoF may still not reflect the true computational effort since one must also take into consideration the calculation of the TIC,

		CBG-IGA				SIP-DG-IGA			
		KI-1		KI-h		KI-1		KI-h	
	p	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}
	1	1.00	2.00	1.00	2.00	1.00	2.00	1.00	2.00
$(n=0)$	2	2.00	3.00	2.00	3.00	2.05	3.03	2.00	2.99
	3	2.98	3.98	3.00	4.00	2.98	3.98	3.00	4.00
	1	0.99	2.00	0.99	2.00	1.00	2.01	1.02	1.98
$(n=1)$	2	2.01	3.03	2.00	3.00	2.06	3.05	2.00	3.00
	3	3.01	4.03	3.00	4.00	3.01	4.03	3.00	3.99
	1	0.99	2.00	1.00	2.00	1.00	2.01	1.04	1.96
$(n=2)$	2	2.04	3.08	2.00	2.99	2.08	3.09	2.01	3.01
	3	3.06	4.01	3.00	3.99	3.06	4.09	3.01	3.98

Table 4.2: The asymptotic orders of accuracy, Equations 4.103, for the verification of the global uniform refinement strategy for the CBG-IGA and the SIP-DG-IGA 1G NDE over a 2D Cartesian grid; the convergence rates are evaluated as the slope of the last few points from each dataset.

for which it becomes increasingly more expensive to solve the orthogonalised eigenvalue problem for a patch with an increasingly large internal refinement structure; q.v. Section 4.4.2.

A further interesting yet expected observation is that of the prolonged pre-asymptotic range of convergence observed on the left-hand side of Figure 4.8, which becomes increasingly large for increasingly oscillatory solutions. For Galerkin-based methods, unique solvability exists for certain conditions imposed upon the choice the discrete approximation space; viz. this can be understood as a condition of minimum resolution for asymptotic convergence [354].

Local Adaptive Refinement The results of the NRG-AMR-$h(p)$ strategy for chasing the global H^1-error in each of the CBG-IGA and the SIP-DG-IGA discretisations are presented in Figure 4.9 for the manufactured solution chosen as the first, third and fifth harmonics of the standing wave equation; q.v. Equation 4.105. These results are compared against those obtained for the global uniform h-refinement strategy. For an initial coarse mesh of linear B-splines, the NRG-AMR-hp strategy exhibits exponential rates of convergence; such a homogeneous and smooth problem favours p-refinement until p_{max} of the basis is attained, after which only h-refinement is available. For initial coarse meshes of linear and quadratic B-splines, the NRG-AMR-h strategy yields algebraic rates of convergence comparable with those of the global uniform refinement strategy; one recalls that the analysis is actually performed on the reference mesh; viz. *one* global uniform hp-refinement of the current coarser mesh.

In both discretisations, enforcing continuity over irregular meshes does introduce some errors; this is to be expected since a more-refined patch is either constrained or penalised against a less-refined neighbour. Although it appears that the local refinement strategy may not offer much of a compet-

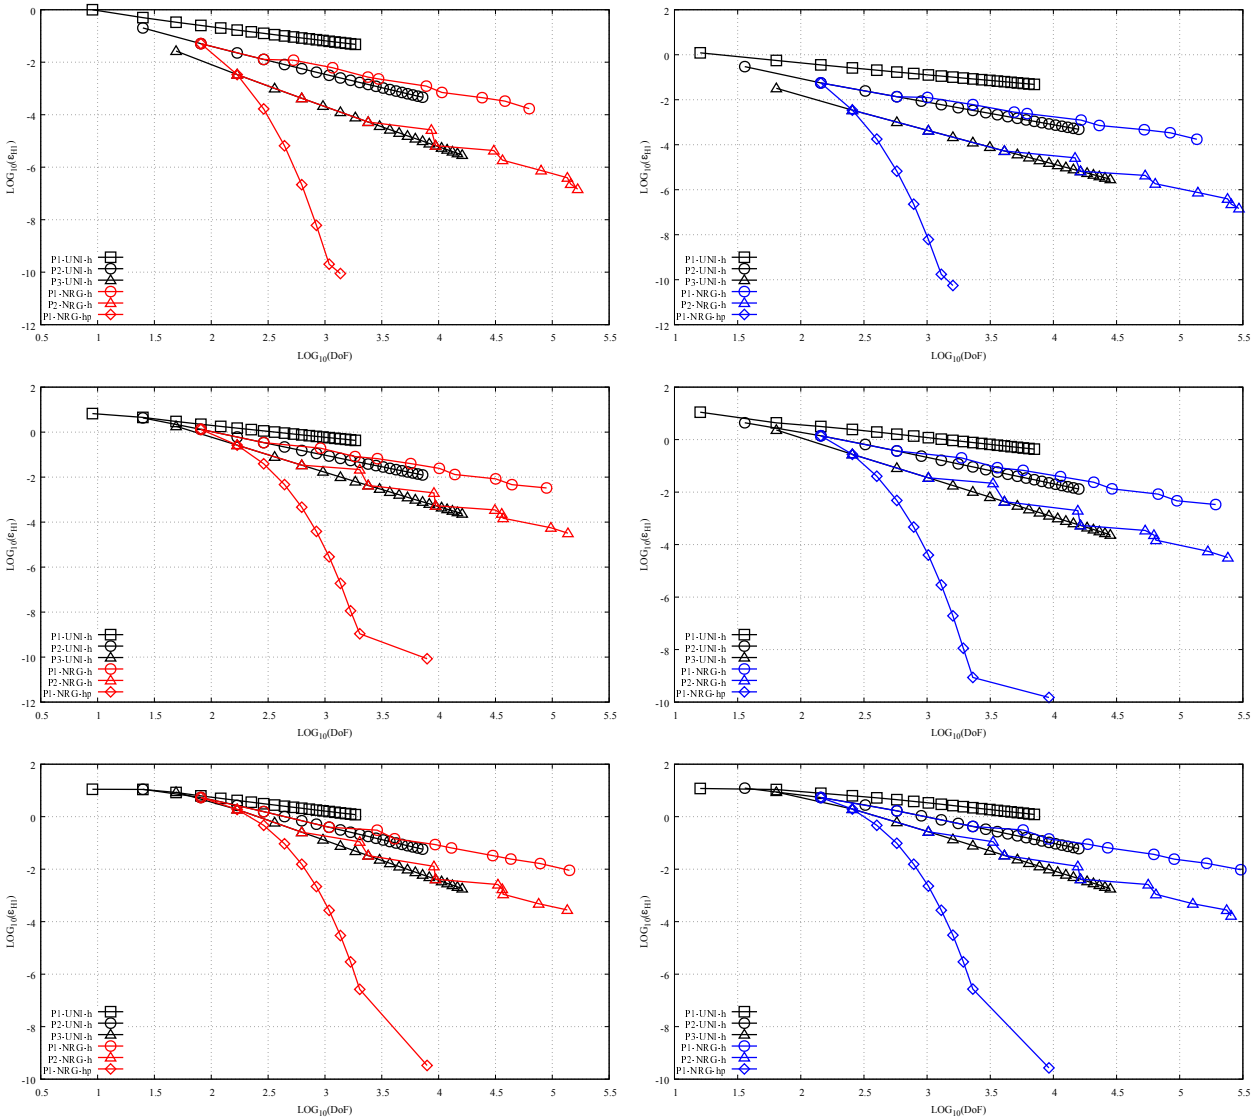


Figure 4.9: The convergence plots, ϵ_{H^1} v DoF, for the verification of the NRG-AMR-$h(p)$ strategy for the CBG-IGA (left) and the SIP-DG-IGA (right) 1G NDE over a 2D Cartesian grid; the manufactured solution is chosen as per Equation 4.105 for $n = 0$ (top row), $n = 1$ (middle row) and $n = 2$ (bottom row). (*Vide* the web-based version for reference to colour.)

itive edge for such a simple problem over a Cartesian grid, the results do verify consistency in the mechanisms of the AMR strategy.

4.5.1.2 A Pin-Cell Grid

Material	g	D^g (cm)	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to1}$ (cm^{-1})	$\Sigma_s^{g\to2}$ (cm^{-1})
(1) -	1	3.333E−01	1.000E+00	5.000E−01	2.500E−01
(1) -	2	3.333E−01	1.000E+00	0.000E+00	5.000E−01

Table 4.3: The nuclear data for the 2D 2G MMS convergence analysis.

The second verification case consists of a more challenging pin-cell grid defined over $V = [0, 1]^2$, which comprises of a homogeneous material whose properties are given in Table 4.3. The coarsest mesh is

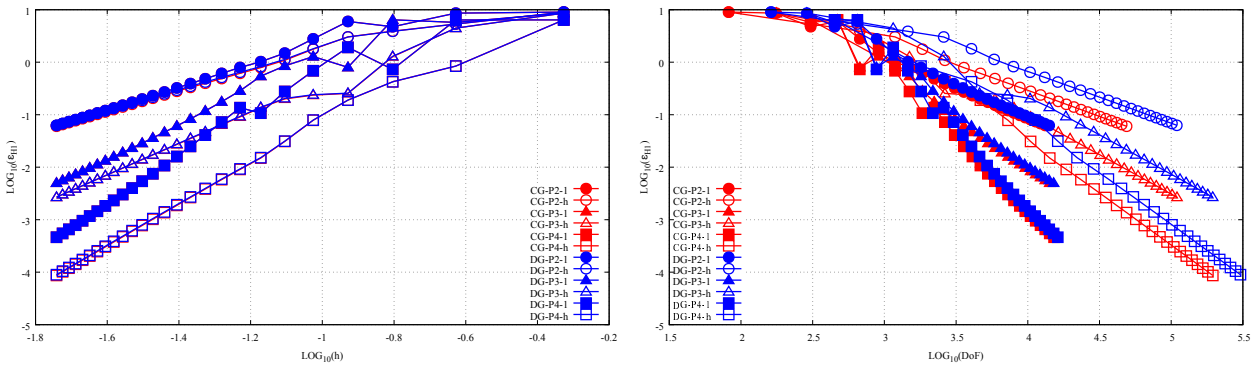


Figure 4.10: The convergence plots, ϵ_{H^1} v h (left) and ϵ_{H^1} v DoF (right), for the verification of the global uniform refinement strategy for the CBG-IGA and the SIP-DG-IGA 2G NDE over a 2D pin-cell grid; the manufactured solutions are chosen as per Equation 4.106 for $g = 1, 2$. (*Vide* the web-based version for reference to colour.)

	CBG-IGA				SIP-DG-IGA			
p	KI-1		KI-h		KI-1		KI-h	
	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}
2	2.06	3.15	1.99	2.98	2.06	3.15	2.00	3.03
3	3.09	4.18	3.00	4.00	3.09	4.18	3.00	3.97
4	4.29	5.40	3.99	4.99	4.29	5.40	3.99	5.02

Table 4.4: The asymptotic orders of accuracy, Equations 4.103, for the verification of the global uniform refinement strategy for the CBG-IGA and the SIP-DG-IGA 2G NDE over a 2D pin-cell grid; the convergence rates are evaluated as the slope of the last few points from each dataset.

composed of *nine* equal-volume non-degenerate bi-quadratic NURBS patches. Assuming a two-group energy approximation, the manufactured solution is chosen as:

$$\begin{cases} \phi^g(\vec{r}) = \sin\left(B_x^g x\right) \sin\left(B_y^g y\right), & \vec{r} \in V; \quad (4.106a) \\ B_x^g, B_y^g := (2g{+}1)\pi, & \text{for } g = 1, 2; \quad (4.106b) \end{cases}$$

where ϕ^g denotes an odd harmonic mode as characterised by the geometric buckling, B_x^g and B_y^g. Again, the manufactured solution has been chosen to be infinitely differentiable. However, now, the homogeneous Dirichlet (zero-flux) BC is naturally satisfied; and this is weakly enforced in the discretisation via penalisation of all external boundaries; q.v. Equation 4.25.

Global Uniform Refinement For the KI-refinement of quadratic, cubic and quartic NURBS, the H^1-errors for each of the CBG-IGA and the SIP-DG-IGA spatial discretisations are presented in Figure 4.10 for the multi-group components of the manufactured solution chosen as the third and fifth harmonics of the standing wave equation; q.v. Equation 4.106. Qualitatively similar plots were obtained for the L_2-error; however, for brevity, these have not been presented. Under refinement, both the continuous and the discontinuous discretisation schemes converge to the analytical solution; and asymptotically, they attain their theoretical orders of accuracy for KI-h-refinement, as presented in

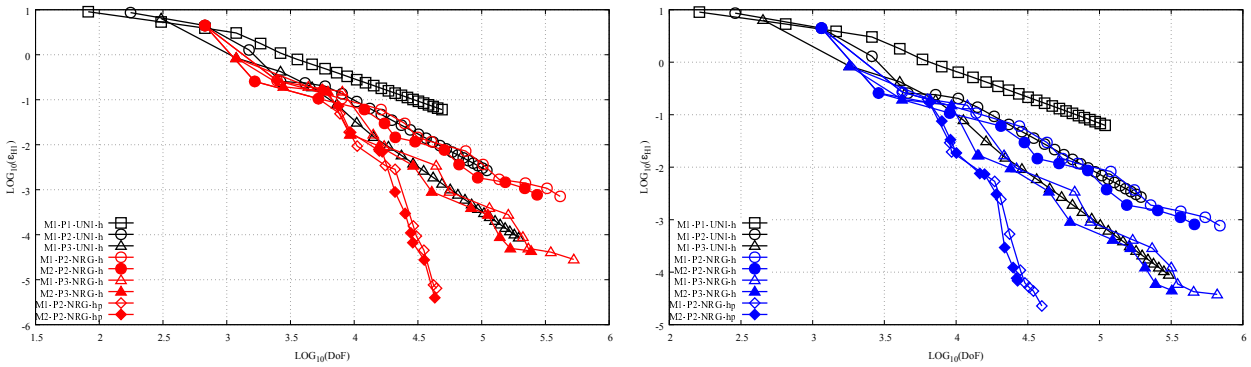


Figure 4.11: The convergence plots, ϵ_{H^1} v DoF, for the verification of the NRG-AMR-$h(p)$ strategy for both the CBG-IGA (left) and the SIP-DG-IGA (right) 2G NDE over a 2D pin-cell grid; the manufactured solutions are chosen as per Equation 4.106 for $g = 1, 2$. (*Vide* the web-based version for reference to colour.)

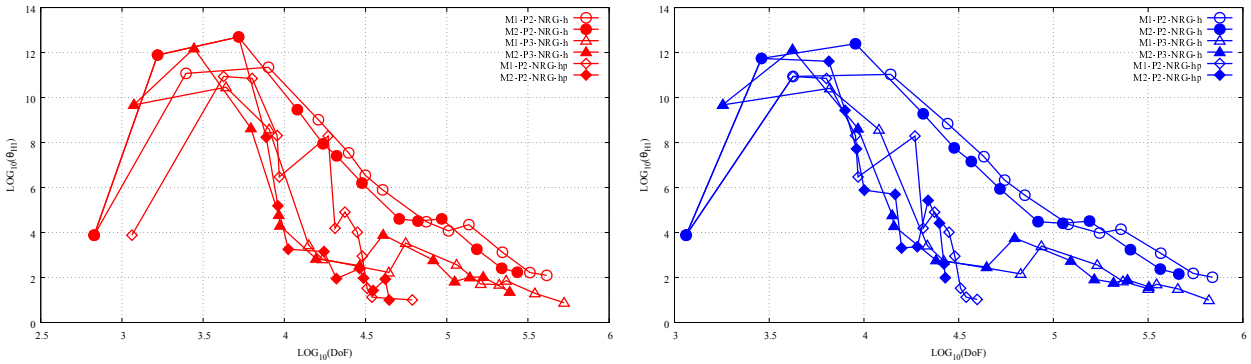


Figure 4.12: The plots of the effectivity index, θ_{H^1} v DoF, for the verification of the NRG-AMR-$h(p)$ strategy for the CBG-IGA (left) and the SIP-DG-IGA (right) 2G NDE over a 2D pin-cell grid; the manufactured solutions are chosen as per Equation 4.106 for $g = 1, 2$. (*Vide* the web-based version for reference to colour.)

Table 4.4. However, for KI-1-refinement, one remarks a degree of super-convergence, which may be explained either by the higher-order continuity in the basis or the weak imposition of the Dirichlet BC. Although not presented, for the same problem with naturally enforced reflective BC and manufactured solutions chosen as per Equation 4.105, one observed the expected rates of convergence.

Local Adaptive Refinement The results of the NRG-AMR-$h(p)$ strategy for chasing the global H^1-error in each of the CBG-IGA and the SIP-DG-IGA discretisations are presented in Figure 4.11 for the multi-group components of the manufactured solution chosen as the third and fifth harmonics of the standing wave equation; q.v. Equation 4.106. These results are compared against those obtained for the global uniform h-refinement strategy. For an initial coarse mesh of quadratic NURBS, the NRG-AMR-hp strategy, again, exhibits faster rates of convergence; although this time the mesh iteration is not dominated by p-refinements. For initial coarse meshes of quadratic and cubic NURBS, the NRG-AMR-h strategy, again, yields algebraic rates of convergence that are comparable with those of the global uniform refinement strategy.

Within a multi-group formalism, it is fitting to test the capacity of the energy-dependent meshes,

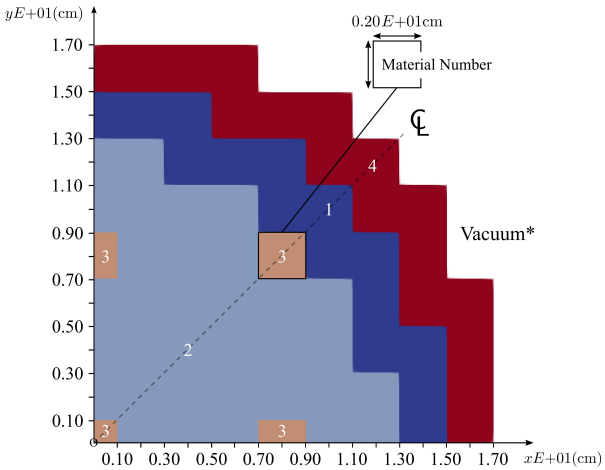


Figure 4.13: The geometry for the 2D 2G IAEA/ANL BSS-11 (LWR) benchmark [355]. The constituent materials are denoted by the numbers in parentheses; viz. fuel-1 (1), fuel-2 (2), fuel-2 & control-rod (3) and reflector (4). (*Vide* the web-based version for reference to colour.)

which, even for such a simple homogeneous problem, offer a competitive edge, yielding more accuracy per DoF when compared to the global uniform h-refinement strategies. The effectivity, θ_{H^1}, of the interpolation-based error measures, as per Equation 3.24, is measured and displayed in Figure 4.12, from which it is observed that θ_{H^1} tends to values slightly larger than unity. This is desirable as one can confidently approximate any remaining error in the discretisation, which can be used to terminate the mesh-iteration procedure. Again, the results do verify consistency in the mechanisms of the AMR strategy over irregular meshes of rational underlying geometries.

4.5.2 The 2D IAEA/ANL BSS-11 (LWR) Benchmark

Material	g	D^g (cm)	Σ_r^{g*} (cm^{-1})	$\Sigma_{s0}^{g\to 2}$ (cm^{-1})	$\nu^g\Sigma_f^g$ (cm^{-1})
(1) Fuel-1	1	1.50000E+00	3.01200E−02	2.00000E−02	0.00000E+00
	2	4.00000E−01	8.00320E−02	0.00000E+00	1.35000E−01
(2) Fuel-2	1	1.50000E+00	3.01200E−02	2.00000E−02	0.00000E+00
	2	4.00000E−01	8.50320E−02	0.00000E+00	1.35000E−01
(3) Fuel-3 & Rod	1	1.50000E+00	3.01200E−02	2.00000E−02	0.00000E+00
	2	4.00000E−01	1.30032E−01	0.00000E+00	1.35000E−01
(4) Reflector	1	2.00000E+00	4.01600E−02	4.00000E−02	0.00000E+00
	2	3.00000E−01	1.00240E−02	0.00000E+00	0.00000E+00

Table 4.5: The nuclear data for the 2D 2G IAEA/ANL BSS-11 (LWR) benchmark [355]. To better approximate the effects of axial leakage, one supplies an axial buckling, $B_z^2 = 8.00000\text{E}-05$; accordingly, the effective macroscopic cross-sections of removal are corrected as $\Sigma_r^{g*} = \Sigma_r^g + D^g B_z^2$.

The 2D IAEA benchmark for the 2G NDE is the BSS-11 quarter-core LWR problem from the Argonne National Laboratory (ANL) Benchmark Problem Book, Supplement 2 [355]. Taken at the mid-plane, $z = 1.90\text{E}+02$ cm, this 2D benchmark is adapted from the full 3D (xyz) one; and thus, an axial buckling, $B_z^2 = 8.00\text{E}-05$, approximates the effects of axial leakage [356]. Accordingly, the *corrected* macroscopic nuclear cross-section data are given in Table 4.5. Materials 1, 2 and 3 represent fuelled

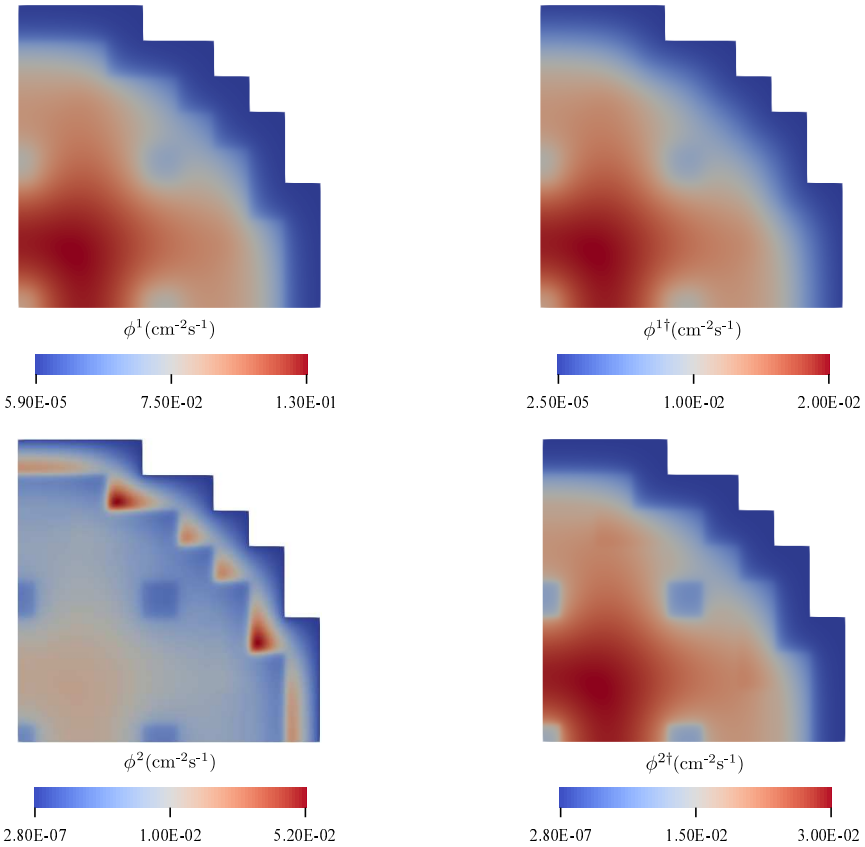


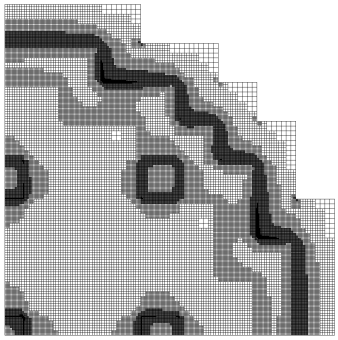
Figure 4.14: The spatial distribution of the multi-group components of the scalar neutron flux (left) and the scalar neutron importance (right) for the 2G NDE for the 2D IAEA/ANL BSS-11 (LWR) benchmark; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

regions, each with slightly different properties to mimic varying levels of enrichment or the presence of control-rods; whereas Material 4 represents a reflector region; v. Figure 4.13. One assumes that there is no incoming current at the external physical boundaries, as modelled by (extrapolated) vacuum BC; and as per the *one-quarter* symmetry in the problem, reflective boundary conditions are prescribed at $x = y = 0.00$ cm. The coarsest Cartesian mesh is constructed from 241 bilinear NURBS patches. A reference solution was obtained over the energy-independent mesh generated after the 11th iteration of the DWR-AMR-*hp* strategy for the CBG-IGA spatial discretisation; viz. $k_{\text{eff}} = 1.0295886369$, which is in good agreement with the literature for those results supplied by FEM and FDM codes [355]. The spatial distribution of the multi-group components of both the scalar neutron flux and the scalar neutron importance are presented in Figure 4.14, for the QoI chosen as the k_{eff}.

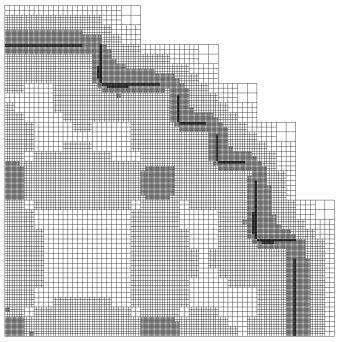
Qualitatively similar meshes were generated for each of the continuous and the discontinuous discretisations for each of the refinement strategies employed; those generated after the 9th AMR iteration for the CBG-IGA AMR-*h* and AMR-*hp* strategies are presented in Figure 4.15 and Figure 4.16, respectively. At the cost of an increasingly ill-conditioned matrices, the *hp*-refinement strategies performed better than the corresponding *h*-refinement ones; fortunately, the computation times for the former were not prohibitive since the chosen algebraic multi-grid preconditioner worked effectively. Further owing to the reduced continuity between the NURBS patches in each discretisation, it is unsurprising



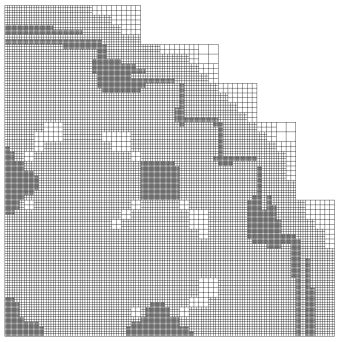
(a) $\mathcal{T}^{(9)}$: 194,393 DoF (×2)



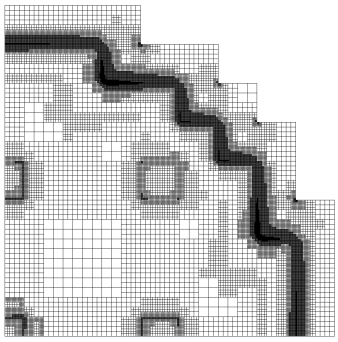
(b) $\mathcal{T}^{(9)}$: 268,033 DoF (×2)



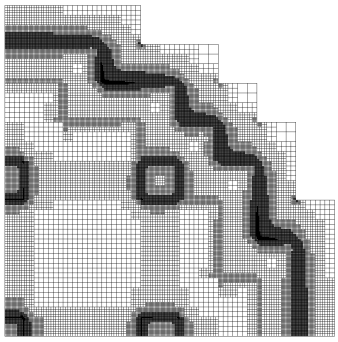
(c) $\mathcal{T}^{1(9)}$: 92,793 DoF



(d) $\mathcal{T}^{1(9)}$: 84,733 DoF



(e) $\mathcal{T}^{2(9)}$: 175,473 DoF

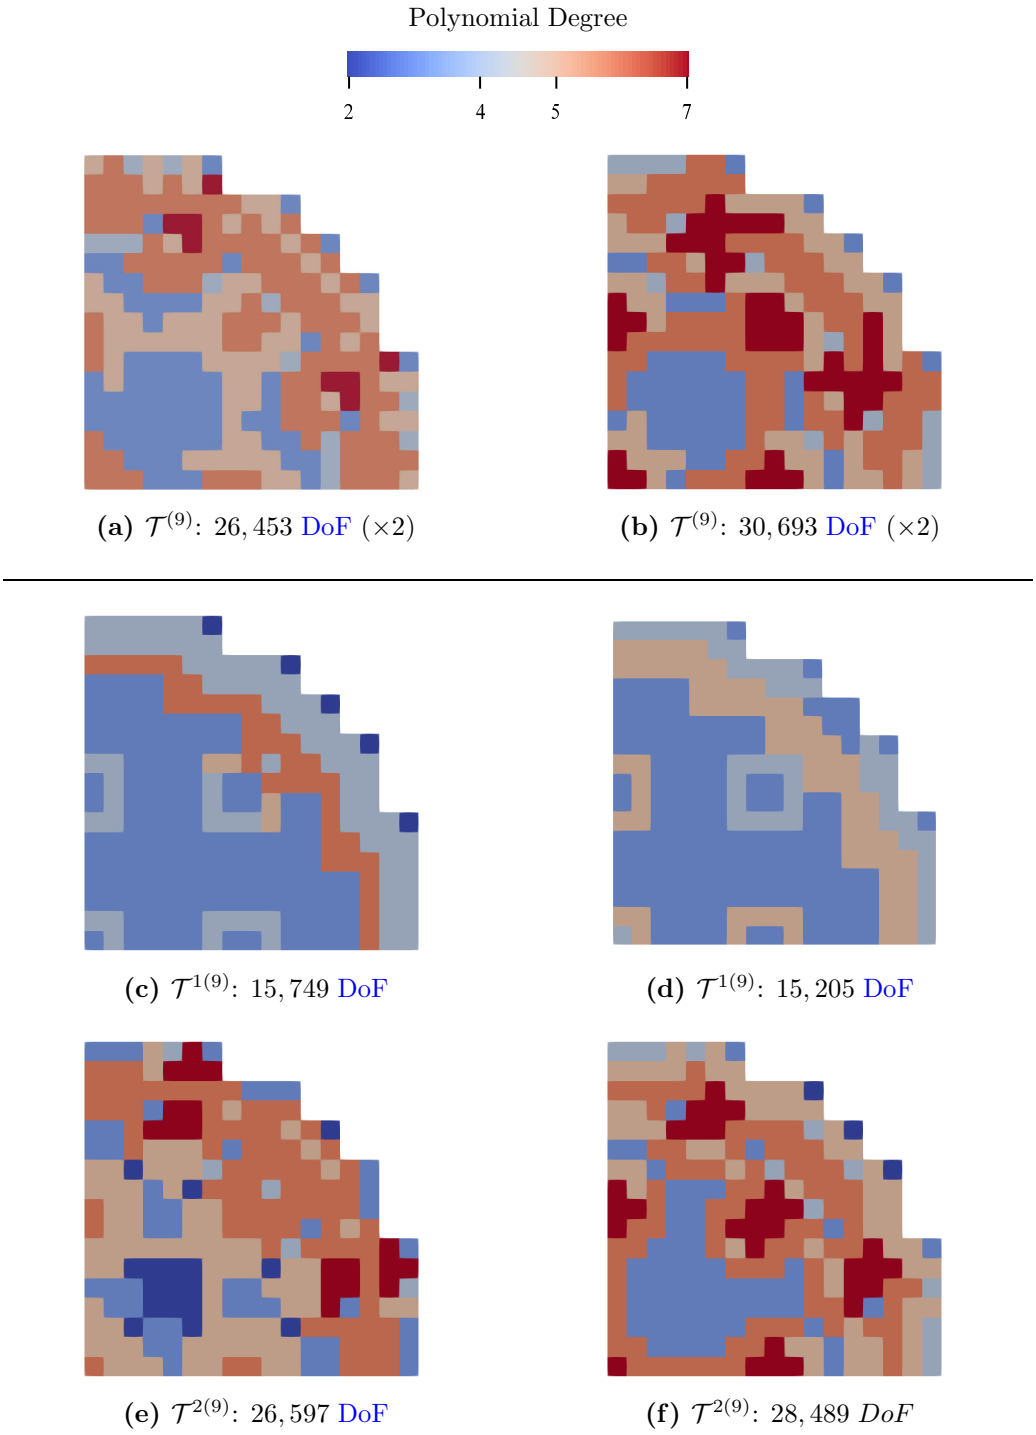


(f) $\mathcal{T}^{2(9)}$: 228,213 DoF

Figure 4.15: The energy-independent meshes (top) and the energy-dependent meshes (middle and bottom) generated at the 9th AMR iteration for the NRG-AMR-h (left) and the DWR-AMR-h (right) strategies for the CBG-IGA 2G NDE for the 2D IAEA/ANL BSS-11 (LWR) benchmark; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

that both discretisations also yield comparable results; v. Figure 4.17. Nevertheless, when compared to their discontinuous counterparts, the CBG-IGA discretisations do appear to gain in accuracy per DoF. Whilst allowing for the imposition of a certain degree of regularity to avoid any multiply constraints, the discretisation developed in Section 4.3 does not appear to introduce any prohibitive sources of errors across the irregular interfaces of this heterogeneous problem.

Both AMR strategies performed better than the global uniform one; yet, this is unsurprising since the k_{eff} is a global quantity and both AMR strategies chase the global error in the primal solution. However, at the cost of solving the dual problem, the DWR-AMR strategies do allow for a better



(a) $\mathcal{T}^{(9)}$: 26,453 DoF (×2) **(b)** $\mathcal{T}^{(9)}$: 30,693 DoF (×2)

(c) $\mathcal{T}^{1(9)}$: 15,749 DoF **(d)** $\mathcal{T}^{1(9)}$: 15,205 DoF

(e) $\mathcal{T}^{2(9)}$: 26,597 DoF **(f)** $\mathcal{T}^{2(9)}$: 28,489 *DoF*

Figure 4.16: The energy-independent meshes (top) and the energy-dependent meshes (middle and bottom) generated at the 9th AMR iteration for the NRG-AMR-*hp* (left) and the DWR-AMR-*hp* (right) strategies for the CBG-IGA 2G NDE for the 2D IAEA/ANL BSS-11 (LWR) benchmark; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

allocation of resources; and better still when employing energy-dependent meshes.

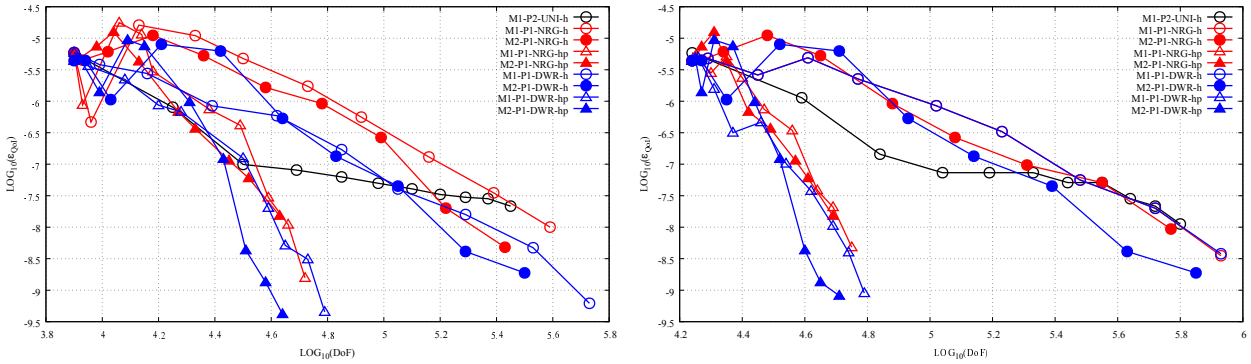


Figure 4.17: The convergence plots, ϵ_{QoI} v DoF, for the global uniform refinement (black), the NRG-AMR (red) and the DWR-AMR (blue) strategies for the CBG-IGA (left) and the SIP-DG-IGA (right) 2G NDE for the 2D IAEA/ANL BSS-11 (LWR) benchmark; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

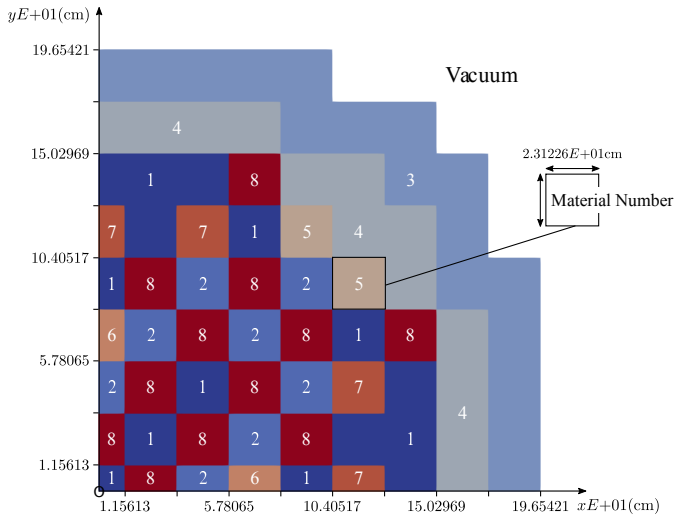


Figure 4.18: The geometry for the 2D 2G BIBLIS (PWR) core benchmark [357]. The constituent materials are denoted by the numbers in parentheses; viz. fuel-1 (1), fuel-2 (2), baffle & reflector (3), fuel-3 (4), fuel-4 (5), fuel-5 (5), fuel-6 (7) and fuel-7 (8). (*Vide* the web-based version for reference to colour.)

4.5.3 The 2D BIBLIS (PWR) Core Benchmark

The *checkerboard* fuel-loading pattern of the 2D BIBLIS (PWR) benchmark for the 2G NDE better represents the simplified core of an *operating* PWR than does the previous benchmark. First proposed by Finemann and Wagner, Herbert provided an updated version to correspond to a *rods-out* configuration; this problem presents *eight* physically distinct regions, including a homogenised baffle-and-reflector region and *seven* smeared fuelled regions of varying degrees of burn up; each element is a square of congruent sides of length 2.31226E+01 cm; v. Figure 4.18 [357, 358]. The corresponding macroscopic nuclear cross-section data are given in Table 4.6; and no axial buckling information is provided. One assumes that there is no incoming current at the external physical boundaries, as modelled by vacuum BC; and as per the *one-quarter* symmetry in the problem, reflective BC are prescribed at $x = y = 0.00$ cm. The coarsest Cartesian mesh is constructed from 257 bilinear NURBS patches. Again, a reference was obtained over the energy-independent mesh generated after the 11th iteration

Material	g	D^g (cm)	Σ_r^g (cm^{-1})	$\Sigma_{s0}^{g\to 2}$ (cm^{-1})	$\nu^g\Sigma_f^g$ (cm^{-1})
(1) -	1	1.43600E+00	2.72582E−02	1.77540E−02	5.87080E−03
	2	3.63500E−01	7.50580E−02	0.00000E+00	9.60670E−02
(2) -	1	1.43660E+00	2.72995E−02	1.76210E−02	6.19080E−03
	2	3.63600E−01	7.84360E−02	0.00000E+00	1.03580E−01
(3) -	1	1.32000E+00	2.57622E−02	2.31060E−02	0.00000E+00
	2	2.77200E−01	7.15960E−02	0.00000E+00	0.00000E+00
(4) -	1	1.43890E+00	2.74640E−02	1.71010E−02	7.45270E−03
	2	3.63800E−01	9.14080E−02	0.00000E+00	1.32360E−01
(5) -	1	1.43810E+00	2.72930E−02	1.72900E−02	6.19080E−03
	2	3.66500E−01	8.48280E−02	0.00000E+00	1.03580E−01
(6) -	1	1.43850E+00	2.73240E−02	1.71920E−02	6.42850E−03
	2	3.66500E−01	8.73140E−02	0.00000E+00	1.09110E−01
(7) -	1	1.43890E+00	2.72900E−02	1.71250E−02	6.19080E−03
	2	3.67900E−01	8.80240E−02	0.00000E+00	1.03580E−01
(8) -	1	1.43930E+00	2.73210E−02	1.70270E−02	6.42850E−03
	2	3.68000E−01	9.05100E−02	0.00000E+00	1.09110E−01

Table 4.6: The nuclear data for the 2D 2G BIBLIS (PWR) core benchmark [357]. The average number of neutrons released per fission is $\nu = 2.47$ for the fission spectrum $\chi = [1.00, 0.00]^T$.

of the DWR-AMR-*hp* strategy for the CBG-IGA spatial discretisation. The spatial distribution of the multi-group components of both the scalar neutron flux and the scalar neutron importance are presented in Figure 4.19. From the literature, $k_{\text{eff}} = 1.0250103$ [358]. However, in this instance, the QoI is chosen as the absorption in Material 5; v. Figure 4.18.

The meshes generated after the 9th AMR iteration for each of the SIP-DG-IGA AMR strategies are presented in Figure 4.21 and Figure 4.21; where, again, qualitatively similar meshes were obtained for the continuous discretisations. The highest rates of convergence were obtained for the *hp*-refinement strategies, with which one benefits from either mesh subdivision or an increasingly large polynomial degree of the basis, as where appropriate.

Due to the loading pattern and the relatively large dimensions of the fuel assemblies, this problem is challenging to solve with low-order spatial discretisations such as linear FEM. One would not expect such difficulties for the proposed AMR strategies since the initial reference mesh is constructed from bi-quadratic NURBS patches. However, the fast and thermal components of the multi-group fluxes do present vastly disparate spatial distributions; viz. the former exhibits a relatively smooth profile whilst the latter does not. Clearly, one single mesh could not suitably allow for an optimum allocation of resources. Accordingly, for the use of energy-dependent meshes, almost twice as many DoF were allocated to the $\mathcal{T}^2$ than to $\mathcal{T}^1$ for the DWR-AMR strategy; whereas this was more than *four* times as many for the NRG-AMR strategy. The latter aims to distribute evenly the error in the solution in a global sense and thus, better approximate those sharp gradients in the thermal flux profile almost everywhere; this can be observed as the *h*-refinement about the material discontinuities in Figure 4.21. However, this does not contribute to any consistent reduction in the error in QoI; v. Figure

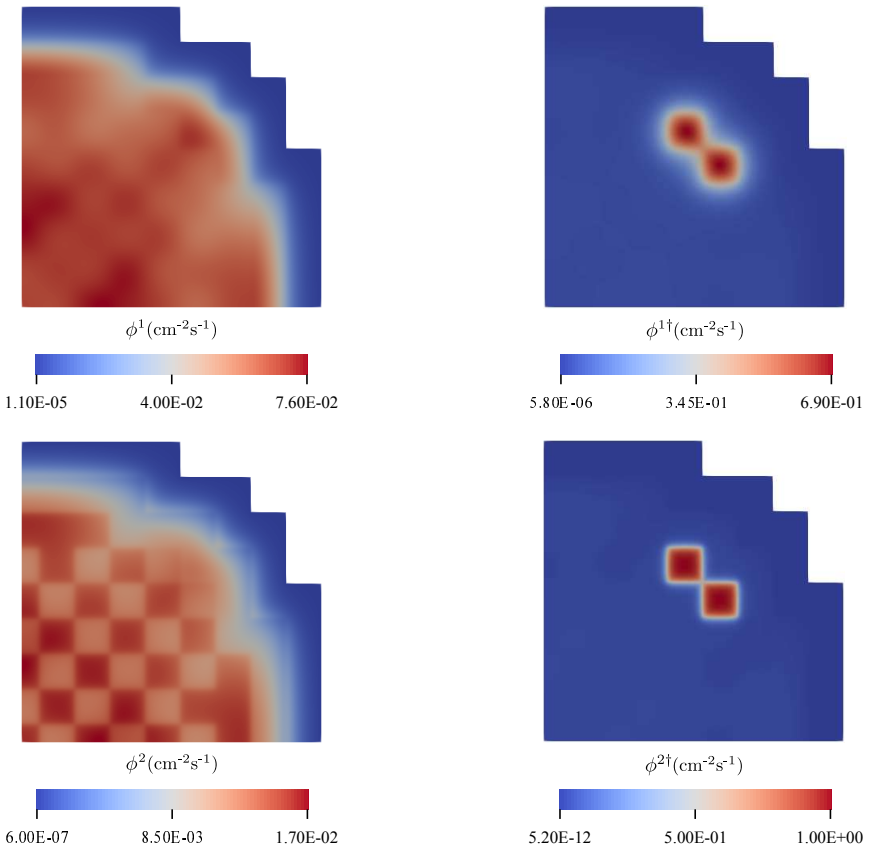


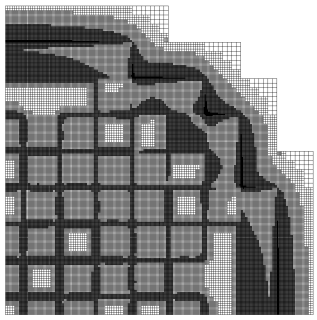
Figure 4.19: The spatial distribution of the multi-group components of the scalar neutron flux (left) and the scalar neutron importance (right) for the 2G NDE for the 2D BIBLIS (PWR) core benchmark; the QoI is the absorption in Material 5; v. Figure 4.18. (*Vide* the web-based version for reference to colour.)

4.22. Accommodating for the need to refine such as to minimise those dominating errors in the thermal neutron flux distribution, the NRG-AMR strategy for the use of an energy-independent mesh outperforms that for the use of energy-dependent ones, since in the case of the former, $\mathcal{T}^1$ benefits from a refinement about Material 5, where it otherwise does not. Conversely, the DWR-AMR strategies consistently yield more accuracy per DoF by expressly refining each mesh about the region of interest.

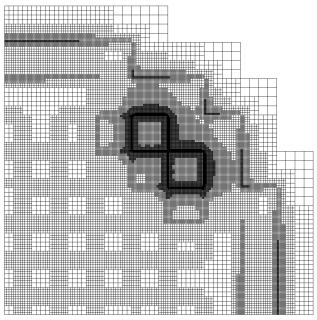
4.6 A Discussion

This chapter saw the application of automised NRG-AMR-$h(p)$ and DWR-AMR-$h(p)$ strategies to separate CBG-IGA and SIP-DG-IGA spatial discretisations of the multi-group NDE.

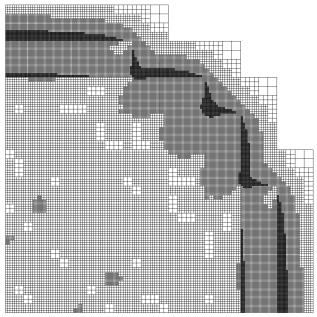
With respect to the continuous discretisation, the potential for higher-order continuity within a NURBS patch was sacrificed to incorporate the wholly-local nature of the AMR algorithms. Where not trivially so, continuity in the solution field across irregular patch interfaces was enforced via constraint equations based upon a master-slave relationship and the conservative interpolation between surface meshes. This was facilitated by the use of open knot-vectors, such that the NURBS basis is interpolatory at the patch boundaries; therefore, any *surfaces* are themselves NURBS objects of one-less



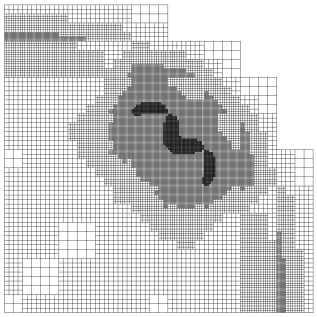
(a) $\mathcal{T}^{(9)}$: $1,204,920$ DoF ($\times 2$)



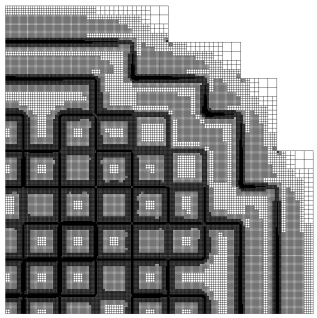
(b) $\mathcal{T}^{(9)}$: $380,556$ DoF ($\times 2$)



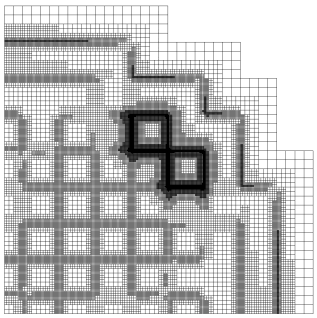
(c) $\mathcal{T}^{1(9)}$: $378,720$ DoF



(d) $\mathcal{T}^{1(9)}$: $196,956$ DoF



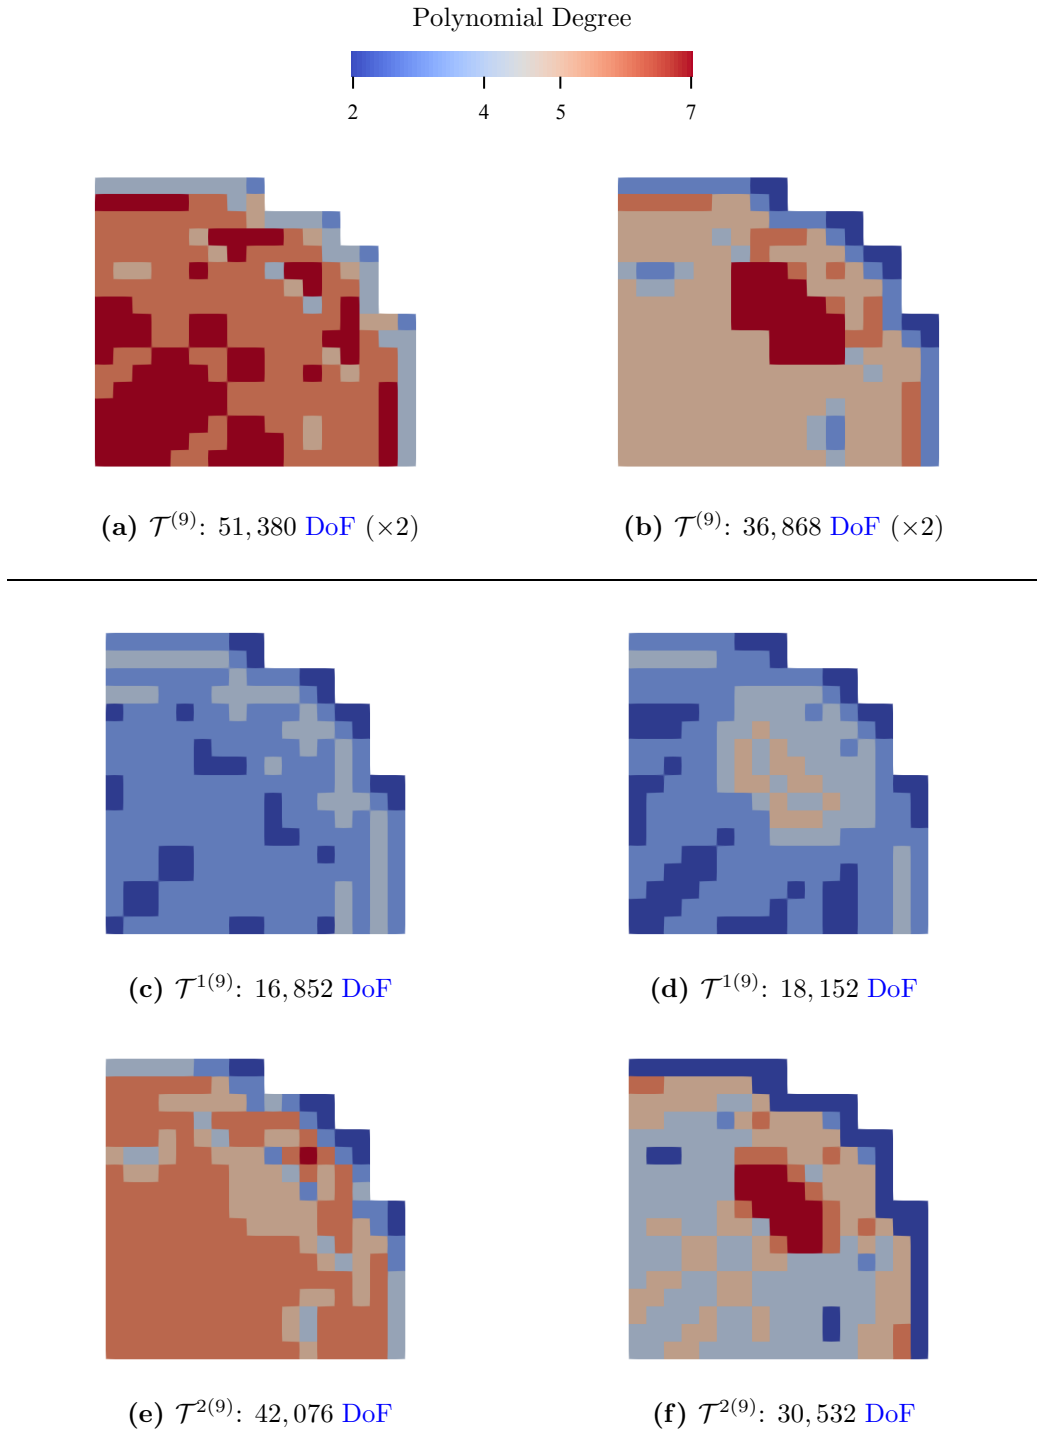
(e) $\mathcal{T}^{2(9)}$: $1,669,752$ DoF



(f) $\mathcal{T}^{2(9)}$: $406,692$ DoF

Figure 4.20: The energy-independent meshes (top) and the energy-dependent meshes (middle and bottom) generated at the 9th AMR iteration for the NRG-AMR-h (left) and the DWR-AMR-h (right) strategies for the SIP-DG-IGA 2G NDE for the 2D BIBLIS (PWR) core benchmark; the QoI is the absorption in Material 5; v. Figure 4.18. (*Vide* the web-based version for reference to colour.)

dimension than the volume. A similar Galerkin projection was presented in Chapter 3 for the conservative interpolation between volume meshes when evaluating the group-to-group source terms over energy-dependent meshes; and similarly so for the calculation of the interpolation-based error measures. For the SIP scheme for the discontinuous discretisation, *optimised* penalty parameters, based upon rigorous coercivity analyses, were derived for general element types to ensure sufficient, yet not excessive, penalisation of the discrete solution field across the spatial mesh(es). Both discretisations result in large, sparse, SPD matrices that benefited from PETSc's capacity for distributed-memory CRS data structures and the parallel implementation of its CG solver. Moreover, the algebraic multi-grid preconditioner appeared to be effective for those problems posed over Cartesian grids.



(a) $\mathcal{T}^{(9)}$: 51,380 DoF (×2)

(b) $\mathcal{T}^{(9)}$: 36,868 DoF (×2)

(c) $\mathcal{T}^{1(9)}$: 16,852 DoF

(d) $\mathcal{T}^{1(9)}$: 18,152 DoF

(e) $\mathcal{T}^{2(9)}$: 42,076 DoF

(f) $\mathcal{T}^{2(9)}$: 30,532 DoF

Figure 4.21: The energy-independent meshes (top) and the energy-dependent meshes (middle and bottom) generated at the 9th AMR iteration for the NRG-AMR-*hp* (left) and the DWR-AMR-*hp* (right) strategies for the SIP-DG-IGA 2G NDE for the 2D BIBLIS (PWR) core benchmark; the QoI is the absorption in Material 5; v. Figure 4.18. (*Vide* the web-based version for reference to colour.)

Whilst mesh irregularities, such as those about hanging-nodes, do violate principles of compatibility and introduce additional discretisation errors, the capacity for local refinement does offer greater practical gains in terms of a better allocation of resources. Allowing for the imposition of a certain degree of regularity to avoid any multiply constraints, the AMR-$h(p)$ algorithms for both of the CBG-IGA and the SIP-DG-IGA discretisations did generate qualitatively similar meshes. Although the latter does incur the cost of at least twice as many DoF as the former, one does appreciate the

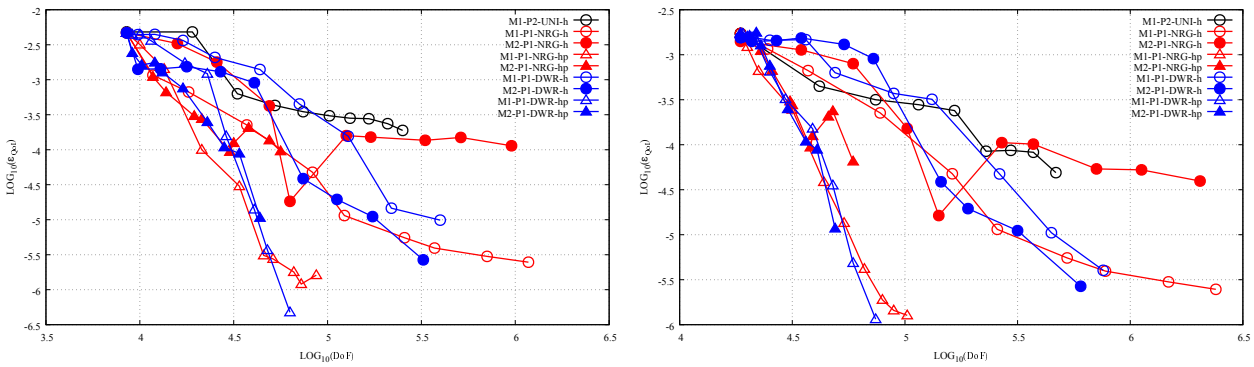


Figure 4.22: The convergence plots, ϵ_{QoI} v DoF, for the global uniform refinement (black) and the NRG-AMR (red) and the DWR-AMR (blue) strategies for the CBG-IGA (left) and the SIP-DG-IGA (right) 2G NDE for the 2D BIBLIS (PWR) core benchmark; the QoI is the absorption in Material 5; v. Figure 4.18. (*Vide* the web-based version for reference to colour.)

flexibility it offers in terms of local refinement and a more compact stencil, as well as its relative ease of implementation. Furthermore, with regard to the mesh-to-mesh interpolation routines, the transfer matrices could be expressed as fully-discontinuous block matrices that were formed and inverted locally. Conversely, to avoid any loss of information for the globally-defined basis functions of the CBG-IGA discretisation, it was necessary to form and inverted globally-assembled transfer matrices.

On an accuracy-per-DoF basis, for global quantities such as the k_{eff}, the NRG-AMR strategies for each discretisation were competitive with the DWR-AMR ones, respectively. However, for localised QoI, the former tended to become less performant than the global uniform strategy for h-refinement; and especially so for the use of energy-dependent meshes, since certain multi-group components were left under-resolved. Conversely, the DWR-AMR strategies consistently and reliably reduced the error in the QoI at each subsequent AMR iteration. For a given problem, it is difficult to compare directly the efficiencies of each of the proposed CBG-IGA and SIP-DG-IGA discretisations using DoF as a metric, since it does not truly account for computational effort attributed to the different stages of analysis; e.g. establishing the constraint-based equations; and calculating the optimised penalty parameters.

Chapter 5

Elliptic Systems: Part 2

This chapter presents the application of automised NRG-AMR-$h(p)$ and DWR-AMR-$h(p)$ strategies to separate CBG-IGA and SIP-DG-IGA spatial discretisations of the multi-group SAAF-S_N equations.

5.1 An Introduction

Whilst diffusion theory does serve as a good approximation to transport theory in some applications, one is more rigorously and more generally interested in modelling for the angular dependence in a neutron population. Since they also constitute elliptic systems and can thus benefit from those techniques developed in the previous chapter, the natural order of progression is to consider those second-order forms of the NTE presented in Section 2.3; namely, the SAAF formulation forms the topic of this current chapter; q.v. Section 2.3.2

Following a multi-group discretisation in energy and a discrete ordinates discretisation in angle, the BVP is reduced to a system weakly coupled in both energy and angle that has yet to be spatially discretised or prescribed BC. Working within *two* different frameworks, the discrete problem is posed as that resulting from either a *conforming* CBG-IGA spatial discretisation or a *non-conforming* DG-IGA spatial discretisation.

Unlike the EOP transport equations, the current formulation models for the full range of the angular neutron flux over the entire unit sphere; with the scalar neutron flux evaluated as per Equation 1.9. Whilst this comes at the cost of twice as many angular unknowns, it does allow for a more direct implementation of the incoming BC, in a manner similar to that for the FONT-S_N equations; and since, for rectangular Cartesian coordinate systems, each discrete angular direction can be solved for *independently* of one another, the SAAF-S_N equations are compatible with the SI procedure; q.v.

Section 2.8 [118]. An obvious shortcoming, however, shared by many second-order forms, is the singular behaviour $1/\Sigma_t \to \infty$ as $\Sigma_t \to 0$; viz. the SAAF-S_N equations are not readily applicable to problems with vacuums without some form of void treatment; and usually this comes at the sacrifice of desirable properties. For example, Wang et al and Schunert et al modified the bilinear form of their SAAF-S_N equations to ensure void-compatibility of the Rattlesnake code within the MOOSE framework; whilst tending towards the standard discretisation of the first-order form of the NTE and admitting discrete adjoint-consistency in the fine mesh limit, $h \to 0$, the modified bilinear form is asymmetric, as might be expected, and the stabilisation depends upon some free-floating parameter, which may require tuning [120, 125]. Morel and McGhee had also previously derived void-compatible SAAF equations; however, their formulation did not constitute a statement of particle conservation [118].

Furthermore, elliptic PIDE govern equilibrium problems and the discretised system is best solved by relaxation methods, whereby the spatial dependence is resolved simultaneously, which may violate theories of causation; qq.v. Section 2.8.1 and Section 2.3.1. Even so, there exist efficient and robust smoothers for SPD matrices, such as the PCG method; and consistent diffusion acceleration schemes can be devised for the discretised SAAF-S_N equations since both second-order elliptic systems can be made subject to the same spatial discretisation; q.v. Chapter 7 [359].

The rest of the chapter is organised as follows: First, the governing equations are presented, the multi-group SAAF-S_N equations, along with their adjoint equations and appropriate BC. This is followed by a CBG-IGA discretisation, where continuity over irregular meshes is enforced in a manner similar to that developed in Section 4.3.1. A separate SIP-DG-IGA discretisation follows, where *optimised* penalty parameters are derived from a coercivity analysis similar to that of Section 4.4.1. Finally, the discretisations are verified and the convergence properties analysed by the MMS; and numerical results are presented for a couple of nuclear benchmarks.

Since the previous chapter laid most of the groundwork for the discretisation of elliptic systems, this current one will be relatively brief in comparison.

5.2 The Self-Adjoint Angular Flux Form of the Multi-Group Neutron Transport Equations

5.2.1 The Direct (Primal) Equations

The direct (primal) multi-group SAAF-$S_N P_L$ equations are posed on a multi-dimensional closed, connected and bounded domain, $V \subset \mathbb{R}^d$, in the real physical space for $d = 2, 3$. The strong form of the

problem is cast as solving $\varphi^{gk}\colon V \to \mathbb{R}$ for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$\begin{cases} -\vec{\Omega}^k \cdot \vec{\nabla} \frac{1}{\Sigma_t^g(\vec{r})} \vec{\Omega}^k \cdot \vec{\nabla} \varphi^{gk}(\vec{r}) + \Sigma_t^g(\vec{r}) \varphi^{gk}(\vec{r}) = q^{gk}(\vec{r}) - \vec{\Omega}^k \cdot \vec{\nabla} \frac{1}{\Sigma_t^g(\vec{r})} q^{gk}(\vec{r}), & \vec{r} \in V; \quad (5.1a) \\ \varphi^{gk}(\vec{r}) = 0, & \vec{r} \in \partial V_V^-; \quad (5.1b) \\ \varphi^{gk}(\vec{r}) = \varphi^{gk\prime}(\vec{r}), & \vec{r} \in \partial V_R^-; \quad (5.1c) \\ \varphi^{gk}(\vec{r}) = \frac{1}{\Sigma_t^g(\vec{r})} \left[-\vec{\Omega}^k \cdot \vec{\nabla} \varphi^{gk}(\vec{r}) + q^{gk}(\vec{r}) \right], & \vec{r} \in \partial V_{ext}^+. \quad (5.1d) \end{cases}$$

The normal vector, $\vec{n}$, can be defined everywhere on the external physical boundary, ∂V_{ext}, over different portions of which, ∂V_{ext}^- and ∂V_{ext}^+, the angular neutron flux, φ^{gk}, may be either *incoming* or *outgoing*, respectively, for a given discrete angular direction of flight, $\vec{\Omega}^k$; where $\partial V_{ext}^{\pm} \coloneqq \{\partial V_f \colon \vec{\Omega}^k \cdot \vec{n} \gtrless 0, \forall \partial V_f \in \partial V_{ext}\}$. From their algebraic derivation, Morel and McGhee suggest that any spurious solutions in the SAAF-S_N equations may be eliminated by requiring the angular neutron flux to satisfy the first-order form of the NTE over all outgoing portions of the physical boundary, ∂V_{ext}^+; q.v. Equation 5.1d [118].

Only Dirichlet-type incoming BC are considered in the above, $\partial V_{ext} \coloneqq \partial V_V \cup \partial V_R$, where ∂V_V and ∂V_R denote the incoming portions of the boundary where vacuum and reflective boundaries conditions are prescribed, respectively; such that $\partial V_V^- \coloneqq \partial V_V \cap \partial V_{ext}^-$ and $\partial V_R^- \coloneqq \partial V_R \cap \partial V_{ext}^-$. One considers only those reflective boundaries that form a perpendicular plane with any of the axes in a rectangular Cartesian coordinate system, such that the specular reflection, $\vec{\Omega}^{k\prime}$, of each discrete ordinate, $\vec{\Omega}^k$, is also a member of the LS or GC angular quadrature sets; where for $\varphi^{gk\prime} = \varphi^g(\vec{\Omega}^{k\prime})$, $|\vec{\Omega}^{k\prime} \cdot \vec{n}| = |\vec{\Omega}^k \cdot \vec{n}|$ and $w_{k\prime} = w_k$.

The defining equations can be recast in block-like group-ordinate operator notation:

$$L\varphi = Qq, \quad (5.2a)$$

$$= Q\left[S\varphi + F\varphi + \frac{1}{4\pi} q_x\right]; \quad (5.2b)$$

with the angular neutron flux, $\varphi = [\varphi^1, \ldots, \varphi^G]^T$, where $\varphi^g = [\varphi^{g1}, \ldots, \varphi^{gK}]^T$ for $g = 1, \ldots, G$; the strength of any extraneous fixed sources, $q_x = [q_x^1, \ldots, q_x^G]^T$; the diagonal-loss operator, L, comprised of within-group losses due to streaming and total collision in each discrete ordinate, L^{gk}, Equation 5.3a; the scatter operator, S, comprised of potentially anisotropic group-to-group scatter terms in each discrete ordinate, S^{gk}, Equation 5.3b; the fission operator, F, comprised of isotropic group-to-group

fission terms in each discrete ordinate, F^{gk}, Equation 5.3c:

$$L^{gk}\varphi \coloneqq -\vec{\Omega}^k \cdot \vec{\nabla}\frac{1}{\Sigma_t^g}\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk} + \Sigma_t^g \varphi^{gk}; \tag{5.3a}$$

$$S^{gk}\varphi \coloneqq \sum_{g\prime=1}^{G}\sum_{l=0}^{L}\sum_{m=-l}^{l}\Sigma_{sl}^{g\prime\to g}Y_l^m\Phi_l^{g\prime m}; \tag{5.3b}$$

$$F^{gk}\varphi \coloneqq \chi^g\sum_{g\prime=1}^{G}\nu^{g\prime}\Sigma_f^{g\prime}Y_0^0\Phi_0^{g\prime 0}. \tag{5.3c}$$

For completeness, a $\mathrm{P_L}$ expansion of the scattering kernel has been included; however, within the scope of this project, one is content in the limit of $L = 0$; likewise, only isotropic fission and extraneous fixed sources are considered. The multi-group components of the spherical harmonic flux moments, Φ_l^{gm}, are evaluated as per Equation 2.44, which is recalled here for $1, \ldots, G$:

$$\Phi_l^{gm}(\vec{r}) \approx \sum_{k=1}^{K} w_k Y_l^m(\vec{\Omega}^k)\varphi^{gk}(\vec{r}). \tag{5.4}$$

The operator, Q, that acts group-ordinate-wise upon each of the source terms on the right-hand side of Equation 5.2 is defined compactly as:

$$Q = 1 + X. \tag{5.5}$$

For the group-ordinate components of the operator X defined for $g = 1, \ldots, G$ and $k = 1, \ldots, K$:

$$X^{gk} \coloneqq -\vec{\Omega}^k \cdot \vec{\nabla}\frac{1}{\Sigma_t^g}, \tag{5.6}$$

defining the operator, A, such that:

$$A\varphi = (L - QS)\varphi, \tag{5.7}$$

one can distinguish fixed-source problems:

$$(A - QF)\varphi = \frac{1}{4\pi}Qq_x, \tag{5.8}$$

from eigenvalue ones:

$$A\varphi = \frac{1}{\mathrm{k_{eff}}}QF\varphi. \tag{5.9}$$

5.2.2 The Adjoint (Dual) Equations

Goal-based (DWR-)AMR relies upon duality arguments; therefore, one introduces the adjoint (dual) problem; q.v. Section 2.2.3. Mathematically, the dual problem should be adjoint-consistent with the primal one; q.v. Equation 2.12. By linearity of the inner-product, adjoint-consistency can be satisfied

in each of the operators defined in Equations 5.3; i.e.: [75]:

$$\sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[\int_V \varphi^{gk\dagger}(\vec{r}) B^{gk} \varphi^{gk}(\vec{r})\, dV \Big] \equiv \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[\int_V \varphi^{gk}(\vec{r}) B^{gk\dagger} \varphi^{gk\dagger}(\vec{r})\, dV \Big]. \tag{5.10}$$

In general, the different possible permutations of the order of discretisation and the adjoint operation do not commute; however, in first establishing the multi-group formalism and performing the spatial discretisation last, it is arbitrary the order in which the discrete ordinates approximation or the adjoint operation are performed, since these two *do* commute; qq.v Section 2.4.1 and Section 2.5.2 [125].

The dual problem is posed on the same domain, $V \in \mathbb{R}^d$, as the primal one; its strong form cast in group-ordinate operator notation as solving $\varphi^{gk\dagger} : V \to \mathbb{R}$ for $g = G, \ldots, 1$ and for $k = 1, \ldots, K$:

$$L^\dagger \varphi^\dagger = Q^\dagger q^\dagger, \tag{5.11a}$$

$$= Q^\dagger \big[S^\dagger \varphi^\dagger + F^\dagger \varphi^\dagger + \frac{1}{4\pi} q_x^\dagger \big]; \tag{5.11b}$$

with the angular neutron importance, $\varphi^\dagger = [\varphi^{1\dagger}, \ldots, \varphi^{G\dagger}]^T$, where $\varphi^{g\dagger} = [\varphi^{g1\dagger}, \ldots, \varphi^{gK\dagger}]^T$ for $g = 1, \ldots, G$; the strength of the adjoint source, $q_x^\dagger = [q_x^{1\dagger}, \ldots, q_x^{G\dagger}]^T$; and the loss, scatter and fission adjoint operators, $L^\dagger$, $S^\dagger$ and $F^\dagger$, respectively:

$$L^{gk\dagger}\varphi^\dagger \coloneqq \text{-}\vec{\Omega}^k \cdot \vec{\nabla} \frac{1}{\Sigma_t^g} \vec{\Omega}^k \cdot \vec{\nabla} \varphi^{gk\dagger} + \Sigma_t^g \varphi^{gk\dagger}; \tag{5.12a}$$

$$S^{gk\dagger}\varphi^\dagger \coloneqq \sum_{g\prime=1}^{G} \sum_{l=0}^{L} \sum_{m=\text{-}l}^{l} \Sigma_{sl}^{g\to g\prime} Y_l^m \Phi_l^{g\prime m\dagger}; \tag{5.12b}$$

$$F^{gk\dagger}\varphi^\dagger \coloneqq \nu^g \Sigma_f^g \sum_{g\prime=1}^{G} \chi^{g\prime} Y_0^0 \Phi_0^{g\prime 0\dagger}; \tag{5.12c}$$

with the multi-group components of the adjoint flux moments, $\Phi_l^{gm\dagger}$, for $1, \ldots, G$:

$$\Phi_l^{gm\dagger}(\vec{r}) = \sum_{k=1} w_k Y_l^m(\vec{\Omega}^k) \varphi^{gk\dagger}(\vec{r}). \tag{5.13}$$

The adjoint operator $Q^\dagger$, that acts group-ordinate-wise upon each of the source terms on the right-hand side of Equation 5.11 is defined compactly as:

$$Q^\dagger = 1 + Y^\dagger; \tag{5.14}$$

with the group-ordinate components of the operator $Y^\dagger$ defined for $g = 1, \ldots, G$ and $k = 1, \ldots, K$:

$$Y^{gk\dagger} \coloneqq +\vec{\Omega}^k \cdot \vec{\nabla} \frac{1}{\Sigma_t^g}. \tag{5.15}$$

The scatter and the fission operators are not self-adjoint; and essentially their adjoint operators are obtained by transposition and the interchanging of arguments of the energy-integral kernels. Furthermore, the *flow* of neutron importance is understood in a reversed sense and thus, one changes

the sign of all odd-parity differential operators [125]. The loss operator is formally self-adjoint, $(L\varphi, \varphi^\dagger) = (\varphi, L\varphi^\dagger)$; however, the streaming term necessitates an appropriate prescription of adjoint BC. Equation 5.11 is multiplied by the angular neutron flux and subtracted from Equation 5.2, which is first multiplied by the angular neutron importance; and by linearity of the inner-product:

$$(L\varphi, \varphi^\dagger) - (Xq, \varphi^\dagger) - (\varphi, L^\dagger\varphi^\dagger) + (\varphi, Y^\dagger q^\dagger) = (q, \varphi^\dagger) - (\varphi, q^\dagger). \tag{5.16}$$

Suppressing any spatial dependence, the inner-product formed with the loss operator is written explicitly in the first line of Equation 5.17, to which the divergence theorem is applied to give the second line, to which the divergence theorem is applied to the second term to give the third line, on which the integrals over the surface and those over the volume are collected:

$$(L\varphi, \varphi^\dagger) = \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[- \int_V \varphi^{gk\dagger}\vec{\Omega}^k \cdot \vec{\nabla}\frac{1}{\Sigma_t^g}\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk}\, dV + \int_V \varphi^{gk\dagger}\Sigma_t^g\varphi^{gk}\, dV \Big], \tag{5.17a}$$

$$= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[- \int_{\partial V} \varphi^{gk\dagger}\frac{\vec{\Omega}^k \cdot \vec{n}}{\Sigma_t^g}\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk}\, dS + \int_V \frac{\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk\dagger}}{\Sigma_g^t}\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk}\, dV + \int_V \varphi^{gk\dagger}\Sigma_t^g\varphi^{gk}\, dV \Big], \tag{5.17b}$$

$$= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[\int_{\partial V} \varphi^{gk}\frac{\vec{\Omega}^k \cdot \vec{n}}{\Sigma_t^g}\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk\dagger}\, dS - \int_{\partial V} \varphi^{gk\dagger}\frac{\vec{\Omega}^k \cdot \vec{n}}{\Sigma_t^g}\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk}\, dS \Big] + (\varphi, L^\dagger\varphi^\dagger). \tag{5.17c}$$

Similar operations are performed on the second and fourth terms on the left-hand side of Equation 5.16, respectively:

$$(Xq, \varphi^\dagger) = \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[- \int_V \varphi^{gk\dagger}\vec{\Omega}^k \cdot \vec{\nabla}\frac{1}{\Sigma_t^g}q^{gk}\, dV \Big], \tag{5.18a}$$

$$= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[- \int_{\partial V} \varphi^{gk\dagger}\frac{\vec{\Omega}^k \cdot \vec{n}}{\Sigma_t^g}q^{gk}\, dS \Big] + (q, X^\dagger\varphi^\dagger); \tag{5.18b}$$

$$(\varphi, Y^\dagger q^\dagger) = \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[\int_V \vec{\Omega}^k \cdot \vec{\nabla}\frac{1}{\Sigma_t^g}q^{gk\dagger}\varphi^{gk}\, dV \Big], \tag{5.19a}$$

$$= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[\int_{\partial V} \varphi^{gk}\frac{\vec{\Omega}^k \cdot \vec{n}}{\Sigma_t^g}q^{gk\dagger}\, dS \Big] + (Y\varphi, q^\dagger). \tag{5.19b}$$

From the above, Equation 5.16 reduces to:

$$\sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[\int_{\partial V} \vec{\Omega}^k \cdot \vec{n} \Big[\varphi^{gk}\frac{1}{\Sigma_t^g}\big[\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk\dagger} + q^{gk\dagger}\big] + \varphi^{gk\dagger}\frac{1}{\Sigma_t^g}\big[-\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk} + q^{gk}\big] \Big]\, dS \Big]$$

$$-(q, X^\dagger\varphi^\dagger) + (Y\varphi, q^\dagger) = (q, \varphi^\dagger) - (\varphi, q^\dagger);$$

where the operators $X^\dagger$ and Y are derived from Equation 5.6 and Equation 5.15, respectively, and the inner-product from Equation 5.10; i.e. for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$X^{gk\dagger} = +\frac{\vec{\Omega}^k \cdot \vec{\nabla}}{\Sigma_t^g}; \qquad (5.20a) \qquad Y^{gk} = -\frac{\vec{\Omega}^k \cdot \vec{\nabla}}{\Sigma_t^g}. \qquad (5.20b)$$

Rearranging the first-order form of the NTE, Equation 2.17, one obtains:

$$\begin{aligned}(Y\varphi, q^\dagger) - (q, X^\dagger\varphi^\dagger) &= (\varphi - q/\Sigma_t, q^\dagger) - (q, \varphi^\dagger - q^\dagger/\Sigma_t),\\ &= (\varphi, q^\dagger) - (q/\Sigma_t, q^\dagger) - (q, \varphi^\dagger) + (q, q^\dagger/\Sigma_t),\\ &= (\varphi, q^\dagger) - (q, \varphi^\dagger),\\ &= 0;\end{aligned}$$

where $1/\Sigma_t$ denotes the group-wise division by the macroscopic cross-section of total collision.

For a set of BC for the primal problem, there is a set for the dual one:

$$\begin{aligned}\sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\int_{\partial V^+} |\vec{\Omega}^k \cdot \vec{n}| \Big[\varphi^{gk} \frac{1}{\Sigma_t^g}\big[\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk\dagger} + q^{gk\dagger}\big] + \varphi^{gk\dagger}\frac{1}{\Sigma_t^g}\big[-\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk} + q^{gk}\big]\Big]\, dS -\\ \int_{\partial V^-} |\vec{\Omega}^k \cdot \vec{n}| \Big[\varphi^{gk} \frac{1}{\Sigma_t^g}\big[\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk\dagger} + q^{gk\dagger}\big] + \varphi^{gk\dagger}\frac{1}{\Sigma_t^g}\big[-\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk} + q^{gk}\big]\Big]\, dS \Bigg] = 0;\end{aligned} \qquad (5.21)$$

where the integral over ∂V_{ext} has been re-expressed as that over *outflow* and *inflow* boundaries, $\partial V_{ext} = \partial V_{ext}^+ \cup \partial V_{ext}^-$, such that for $k = 1, \ldots, K$:

$$\int_{\partial V_{ext}} \vec{\Omega}^k f(\vec{r}) \cdot \vec{n}(\vec{r})\, dS = \int_{\partial V_{ext}^+} |\vec{\Omega}^k \cdot \vec{n}(\vec{r})| f(\vec{r})\, dS - \int_{\partial V_{ext}^-} |\vec{\Omega}^k \cdot \vec{n}(\vec{r})| f(\vec{r})\, dS. \qquad (5.22)$$

The BC for the primal problem are retained from Equations 5.1 and repeated:

$$\begin{cases} \varphi^{gk}(\vec{r}) = 0, & \vec{r} \in \partial V_V^-; \qquad (5.23a)\\ \varphi^{gk}(\vec{r}) = \varphi^{gk\prime}(\vec{r}), & \vec{r} \in \partial V_R^-; \qquad (5.23b)\\ \varphi^{gk}(\vec{r}) = \dfrac{1}{\Sigma_t^g(\vec{r})}\big[-\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk}(\vec{r}) + q^{gk}(\vec{r})\big], & \vec{r} \in \partial V_{ext}^+. \qquad (5.23c)\end{cases}$$

The corresponding set of outgoing adjoint BC are prescribed on the same boundaries as those incoming direct ones, respectively; and the incoming adjoint BC is prescribed on the same portion of the boundary as the outgoing direct one:

$$\begin{cases} \varphi^{gk\dagger}(\vec{r}) = 0, & \vec{r} \in \partial V_V^+; \qquad (5.24a)\\ \varphi^{gk\dagger}(\vec{r}) = \varphi^{gk\prime\dagger}(\vec{r}), & \vec{r} \in \partial V_R^+; \qquad (5.24b)\\ \varphi^{gk\dagger}(\vec{r}) = \dfrac{1}{\Sigma_t^g(\vec{r})}\big[\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk\dagger}(\vec{r}) + q^{gk\dagger}(\vec{r})\big], & \vec{r} \in \partial V_{ext}^-. \qquad (5.24c)\end{cases}$$

The adjoint BC are consistent by substitution into Equation 5.21; and recalling for reflective BC that

$|\vec{\Omega}^{k\prime} \cdot \vec{n}| = |\vec{\Omega}^{k} \cdot \vec{n}|$ and $w_{k\prime} = w_k$; and by a change of dummy variables:

$$\sum_{g=1}^{G} \sum_{k\prime=1}^{K} w_{k\prime} \left[\int_{\partial V_R^+} |\vec{\Omega}^{k\prime} \cdot \vec{n}| \varphi^{gk\prime\dagger} \varphi^{gk} \, dS \right] - \sum_{g=1}^{G} \sum_{k=1}^{K} w_k \left[\int_{\partial V_R^-} |\vec{\Omega}^{k} \cdot \vec{n}| \varphi^{gk\prime} \varphi^{gk\dagger} \, dS \right] = 0. \tag{5.25}$$

In a similar manner to Equation 5.7 for the primal problem, the operator $A^\dagger$ is defined such that:

$$A^\dagger \varphi^\dagger = (L^\dagger - Q^\dagger S^\dagger) \varphi^\dagger; \tag{5.26}$$

and, again, one distinguishes fixed-source problems:

$$(A^\dagger - Q^\dagger F^\dagger) \varphi^\dagger = \frac{1}{4\pi} Q^\dagger q_x^\dagger, \tag{5.27}$$

from eigenvalue ones:

$$A^\dagger \varphi^\dagger = \frac{1}{\mathrm{k_{eff}}} Q^\dagger F^\dagger \varphi^\dagger. \tag{5.28}$$

Under the assumption that the property of (semi-discrete) adjoint-consistency holds, the dual problem possesses the same eigenspectrum as the primal one; and hence, $\mathrm{k}_{\mathrm{eff}}^\dagger = \mathrm{k_{eff}}$; q.v. Equation 4.18.

5.3 A Continuous-Bubnov Galerkin Isogeometric Discretisation

One makes use of the same s-order Sobolev space, $H^s(V)$ introduced in Section 4.3; and the point of departure is the semi-discrete multi-group SAAF-$\mathrm{S_N}$ equations, Equation 5.2, and their adjoint equations, Equation 5.11. Both the primal and the dual problems constitute systems of equations that are solved for the group-ordinate components $\varphi^{gk} : V \to \mathbb{R}$ and $\varphi^{gk\dagger} : V \to \mathbb{R}$, respectively; with BC, such as those in Equations 5.23 and Equations 5.24, respectively, prescribed over separate sections of the external boundary.

Without loss of generality, an emphasis is placed upon manipulating the direct equations; and the same discretisation is applied to the adjoint equations; although one appreciates that discrete adjoint-consistency of the (unmodified) SAAF-$\mathrm{S_N}$ is only attained in the asymptotic limit [125]. The weak problem is formulated by multiplying Equation 5.2 by test functions, $w \in \mathcal{V}(V)$; expanding the angular neutron flux in a basis of trial functions from the same space, $u \in \mathcal{V}(V)$; and then integrating over V. By application of the divergence theorem to each of the streaming terms, the basis functions must only be first-order differentiable; i.e. $\mathcal{V}(V) \subseteq H^1(V)$. Naturally satisfied by the weak form is the homogeneous Dirichlet-type BC, or the vacuum BC, over the inflow portions of the exterior boundary, ∂V_{ext}^-; q.v. Equation 5.23a. The reflective BC is implemented in a weak sense by substitution of Equation 5.23b. BC must also be prescribed over those outflow portions ∂V_{ext}^+; where the solution must satisfy the first-order form of the NTE; q.v. Equation 5.23c.

The domain, V, is decomposed into E-number of non-degenerate, non-overlapping NURBS patches, V_e, the union of which forms the physical mesh; i.e. $\mathcal{T}(V) = \{V_e\}_{e=1}^{E}$ and $V = \bigcup_{e=1}^{E} V_e$. An interface is formed as the intersection of *two* adjoining patches; and thus, an arbitrary patch face belongs to the complete set $\partial V_f \in \partial V = \partial V_{int} \cup \partial V_{ext}$. For the sake of simplicity, one assumes a single energy-independent mesh in the following development; and in any case, all of the angular components of a given group will always be resolved over the same mesh.

The weak form of Equation 5.8 is to find $\varphi \in \mathcal{V}(V)$, such that:

$$a(\varphi, w) = b(w), \quad \forall\, w \in \mathcal{V}(V). \tag{5.29}$$

The bilinear form, $a(\cdot,\cdot)\colon \mathcal{V}(V) \times \mathcal{V}(V) \to \mathbb{R}$, retains its block-like structure and comprises of terms of streaming and total collision, group-to-group scattering and fission and BC:

$$a(\varphi, w) = a_L(\varphi, w) - a_S(\varphi, w) - a_F(\varphi, w) + a_{ext}(\varphi, w). \tag{5.30}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$\begin{aligned} a_L(\varphi, w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e \in \mathcal{T}} \int_{V_e} \vec{\Omega}^k \cdot \vec{\nabla} w \frac{1}{\Sigma_{t,e}^g} \vec{\Omega}^k \cdot \vec{\nabla} \varphi^{gk} + w \Sigma_{t,e}^g \varphi^{gk} \, dV \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_L^{gk}(\varphi^{gk}, w); \end{aligned} \tag{5.31a}$$

$$\begin{aligned} a_S(\varphi, w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{g'=1}^{G} Y_0^0(\vec{\Omega}^k) \sum_{V_e \in \mathcal{T}} \int_{V_e} w \Sigma_{s0,e}^{g' \to g} \Phi_0^{g'0} + \vec{\Omega}^k \cdot \vec{\nabla} w \frac{\Sigma_{s0,e}^{g' \to g}}{\Sigma_{t,e}^g} \Phi_0^{g'0} \, dV \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_S^{gk}(\varphi, w); \end{aligned} \tag{5.31b}$$

$$\begin{aligned} a_F(\varphi, w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{g'=1}^{G} Y_0^0(\vec{\Omega}^k) \sum_{V_e \in \mathcal{T}} \chi_e^g \int_{V_e} w \nu_e^{g'} \Sigma_{f,e}^{g'} \Phi_0^{g'0} + \vec{\Omega}^k \cdot \vec{\nabla} w \frac{\nu_e^{g'} \Sigma_{f,e}^{g'}}{\Sigma_{t,e}^g} \Phi_0^{g'0} \, dV \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_F^{gk}(\varphi, w); \end{aligned} \tag{5.31c}$$

$$\begin{aligned} a_{ext}(\varphi, w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e \in \mathcal{T}} \sum_{\partial V_f \in V_e \cap \partial V_{ext}^+} \int_{\partial V_f} |\vec{\Omega}^k \cdot \vec{n}| w \varphi^{gk} \, dS \\ &\quad - \sum_{\partial V_f \in V_e \cap \partial V_R^-} \int_{\partial V_f} |\vec{\Omega}^k \cdot \vec{n}| w \varphi^{gk\prime} \, dS \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_{ext}^{gk}(\varphi^{gk}, w). \end{aligned} \tag{5.31d}$$

The linear functional, $b(\cdot)\colon \mathcal{V}(V) \to \mathbb{R}$, comprises only of fixed sources:

$$\begin{aligned} b(w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e \in \mathcal{T}} \int_{V_e} w \frac{1}{4\pi} q_{x,e}^{g} + \vec{\Omega}^k \cdot \vec{\nabla} w \frac{1}{4\pi \Sigma_{t,e}^{g}} q_{x,e}^{g} \, dV \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} b^{gk}(w). \end{aligned} \tag{5.32}$$

For multiplying systems, $q_x = 0$, the weak form of Equation 4.6 is to find $\varphi \in \mathcal{V}(V)$, such that:

$$\tilde{a}(\varphi, w) = \frac{1}{\mathrm{k_{eff}}} a_f(\varphi, w), \quad \forall\, w \in \mathcal{V}(V); \tag{5.33}$$

where $\tilde{a}$ is the original bilinear form in Equation 5.30, a, less the fission term; i.e.:

$$\tilde{a} = a + a_F = a_L - a_S + a_{ext}. \tag{5.34}$$

The discrete variational problem is formulated by restricting the function space to a finite-dimensional subset; i.e. $\hat{\mathcal{V}}_{h,p}(V) \subset \mathcal{V}(V)$. As per the isoparametric concept, this space is chosen as that spanned by the same N-number of NURBS that describe the mesh, $\mathcal{T}$; q.v. Equation 4.32. Clearly, $\hat{\mathcal{V}}_{h,p}(V)$ is chosen such that $\hat{\varphi} \to \varphi$ as $N \to \infty$. The weighted-residual statement must be satisfied for each member in the weighting set; and since the space of trial and test functions coincide, this forms the same number of algebraic equations as there are unknowns per energy group; i.e. $DoF^{gk} = \dim \hat{\mathcal{V}}_{h,p}(V) = N$ for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$; and hence, $DoF = \sum_g \sum_k DoF^{gk} = G \times K \times N$. Consequently, for the analysis over a single energy-independent mesh, the discrete problem is posed as a system of $(GKN \times GKN)$-matrix equations.

Similar to the CBG-IGA discretisation of the multi-group NDE, continuity is enforced across the domain in the same strong sense by virtually *gluing* together the NURBS across regular patch interfaces and employing projection-based constraint equations across irregular ones; q.v. Section 4.3.1. Again, the bilinear form, a, possesses the properties of continuity, Equation 4.34, and coercivity, Equation 4.35; and as per the Lax-Milgram Lemma these are sufficient to assert well-posedness [338]. Since $\hat{\mathcal{V}}_{h,p}(V) \subset \mathcal{V}$, the discrete problem inherits this and other properties from the (semi-)analytical weak-problem; and again, error bounds can be established as per Céa's Lemma of quasi-best approximation for conforming and consistent discretisations; q.v. Equation 4.36 [222].

The discrete system is solved via the SI procedure; q.v. Section 2.8. Weakly coupled in energy and angle, the within-group equations are solved sequentially with the most up-to-date available information. For fixed sources, $\hat{\varphi}^{gk} \in \hat{\mathcal{V}}_{h,p}(V)$ is sought for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$a_L^{gk}(\hat{\varphi}^{gk}, \hat{w}) + a_{ext}^{gk+}(\hat{\varphi}^{gk}, \hat{w}) = a_S^{gk}(\hat{\varphi}, \hat{w}) + a_F^{gk}(\hat{\varphi}, \hat{w}) + a_{ext}^{gk-}(\hat{\varphi}^{gk}, \hat{w}) + b^{gk}(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V); \tag{5.35}$$

where $a_{ext} = a_{ext}^{+} - a_{ext}^{-}$ has been decomposed into the outgoing and incoming terms, respectively. The

left-hand side, $a_L^{gk} + a_{ext}^{gk+}$, is SPD and therefore Equation 5.35 can be solved via the PCG method.

Employing the PI procedure to converge $|1/\mathrm{k_{eff}}|$, the eigenvalue problem is cast as finding $\hat{\varphi}^{gk} \in \hat{\mathcal{V}}_{h,p}(V)$ for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$a_L^{gk}(\hat{\varphi}^{gk}, \hat{w}) + a_{ext}^{gk+}(\hat{\varphi}^{gk}, \hat{w}) = a_S^{gk}(\hat{\varphi}, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a_F^{gk}(\hat{\varphi}, \hat{w}) + a_{ext}^{gk\text{-}}(\hat{\varphi}^{gk}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V). \tag{5.36}$$

In kind, the discrete variational statement of the adjoint problem in the presence of fixed adjoint sources, $q_x^\dagger \neq 0$, is cast as finding $\hat{\varphi}^{gk\dagger} \in \hat{\mathcal{V}}_{h,p}(V)$ for $g = G, \ldots, 1$ and for $k = 1, \ldots, K$:

$$\begin{aligned} a_L^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) + a_{ext}^{gk\dagger+}(\hat{\varphi}^{gk\dagger}, \hat{w}) = \\ a_S^{gk\dagger}(\hat{\varphi}^\dagger, \hat{w}) + a_F^{gk\dagger}(\hat{\varphi}^\dagger, \hat{w}) + a_{ext}^{gk\dagger\text{-}}(\hat{\varphi}^{gk\dagger}, \hat{w}) + b^{gk\dagger}(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V). \end{aligned} \tag{5.37}$$

The adjoint eigenvalue problem is cast as finding $\hat{\varphi}^{gk\dagger} \in \hat{\mathcal{V}}_{h,p}(V)$ for $g = G, \ldots, 1$ and for $k = 1, \ldots, K$:

$$\begin{aligned} a_L^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) + a_{ext}^{gk\dagger+}(\hat{\varphi}^{gk\dagger}, \hat{w}) = \\ a_S^{gk\dagger}(\hat{\varphi}^\dagger, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a_F^{gk\dagger}(\hat{\varphi}^\dagger, \hat{w}) + a_{ext}^{gk\dagger\text{-}}(\hat{\varphi}^{gk\dagger}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V); \end{aligned} \tag{5.38}$$

where the adjoint bilinear form, $a^\dagger(\cdot,\cdot)\colon \mathcal{V}(V) \times \mathcal{V}(V) \to \mathbb{R}$, is defined by:

$$a^\dagger(\varphi^\dagger, w) = a_L^\dagger(\varphi^\dagger, w) - a_S^\dagger(\varphi^\dagger, w) - a_F^\dagger(\varphi^\dagger, w) + a_{ext}^\dagger(\varphi^\dagger, w). \tag{5.39}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$\begin{aligned} a_L^\dagger(\varphi^\dagger, w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e\in\mathcal{T}} \int_{V_e} \vec{\Omega}^k \cdot \vec{\nabla} w \frac{1}{\Sigma_{t,e}^g} \vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk\dagger} + w\Sigma_{t,e}^g \varphi^{gk\dagger}\, dV \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_L^{gk\dagger}(\varphi^{gk\dagger}, w); \end{aligned} \tag{5.40a}$$

$$\begin{aligned} a_S^\dagger(\varphi^\dagger, w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{g\prime=1}^{G} Y_0^0(\vec{\Omega}^k) \sum_{V_e\in\mathcal{T}} \int_{V_e} w\Sigma_{s0,e}^{g\to g\prime} \Phi_0^{g\prime 0\dagger} + \vec{\Omega}^k \cdot \vec{\nabla} w \frac{\Sigma_{s0,e}^{g\to g\prime}}{\Sigma_{t,e}^g} \Phi_0^{g\prime 0\dagger}\, dV \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_S^{gk\dagger}(\varphi^\dagger, w); \end{aligned} \tag{5.40b}$$

$$\begin{aligned} a_F^\dagger(\varphi^\dagger, w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{g\prime=1}^{G} Y_0^0(\vec{\Omega}^k) \sum_{V_e\in\mathcal{T}} \nu_e^g \Sigma_{f,e}^g \int_{V_e} w\chi_e^{g\prime} \Phi_0^{g\prime 0\dagger} + \vec{\Omega}^k \cdot \vec{\nabla} w \frac{\chi_e^{g\prime}}{\Sigma_{t,e}^g} \Phi_0^{g\prime 0\dagger}\, dV \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_F^{gk\dagger}(\varphi^\dagger, w); \end{aligned} \tag{5.40c}$$

$$\begin{aligned} a_{ext}^\dagger(\varphi^\dagger, w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e\in\mathcal{T}} \sum_{\partial V_f \in V_e \cap \partial V_{ext}^-} \int_{\partial V_f} |\vec{\Omega}^k \cdot \vec{n}| w\varphi^{gk\dagger}\, dS \\ &\quad - \sum_{\partial V_f \in V_e \cap \partial V_R^+} \int_{\partial V_f} |\vec{\Omega}^k \cdot \vec{n}| w\varphi^{gk\prime\dagger}\, dS \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_{ext}^{gk\dagger}(\varphi^{gk\dagger}, w); \end{aligned} \tag{5.40d}$$

and the adjoint linear functional, $b^\dagger(\cdot)\colon \mathcal{V}(V) \to \mathbb{R}$, is defined as:

$$\begin{aligned} b^\dagger(w) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e\in\mathcal{T}} \int_{V_e} w\frac{1}{4\pi} q_{x,e}^{g\dagger} - \vec{\Omega}^k \cdot \vec{\nabla} w \frac{1}{4\pi\Sigma_{t,e}^g} q_{x,e}^{g\dagger}\, dV \Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} b^{gk\dagger}(w). \end{aligned} \tag{5.41}$$

For eigenvalue problems, the solution is normalised after each outer iteration such that $a_F(\varphi, \varphi) = 1$; and the adjoint solution is normalised such that $a_F(\varphi, \varphi^\dagger) = 1$; q.v. Section 3.6.2.2.

5.4 An Interior Penalty Scheme for a Discontinuous Galerkin Isogeometric Discretisation

Changing tack, one derives now an SIP-DG-IGA discretisation of the multi-group SAAF-S_N equations, Equation 5.2, and their adjoint equations, Equation 5.11, are solved for the group-ordinate components $\varphi^{gk}\colon V \to \mathbb{R}$ and $\varphi^{gk\dagger}\colon V \to \mathbb{R}$, respectively, over the same mesh as that adopted in Section 5.3; i.e. $\mathcal{T}(V) = \{V_e\}_{e=1}^E$ and $V = \bigcup_{e=1}^E V_e$; and $\partial V_f \in \partial V = \partial V_{int} \cup \partial V_{ext}$; where $\partial V_{ext} = \partial V_V \cup \partial V_R$. Within

a discontinuous framework, one exploits the broken Sobolev space, $H^s(\mathcal{T})$; q.v. Equation 4.54. The restriction of values of such functions with support over V_e are defined over ∂V_e by trace operators, $\gamma_i \colon H^s(V_e) \to H^{s\text{-}i\text{-}½}(\partial V_e)$ for $0 \leq i < s$ [222]. Of particular interest is the Dirichlet trace, γ_0, which is well defined in $L_2(\partial V_e)$ for $s > ½$; q.v. Equation 4.56a. One notes that this condition is less strict than that imposed in Section 4.4; where for $s > ⅜$, consistency also required the trace of the normal components of the gradient to be well defined.

Any function $u \in H^s(\mathcal{T})$ is discontinuous across V, yet smooth enough within an element to admit two-valued traces at an interface for $s > ½$; with *jumps* and *averages* defined as per Equation 4.58a and Equation 4.58b, respectively, $\forall \partial V_f \in \partial V_{int}$. Physically, the solution to the primal problem should be sufficiently smooth that any jumps in the angular flux, φ, and any jumps in the normal component of the angular neutron current, $\vec{\Omega}\varphi \cdot \vec{n}$, are vanishing across the interior of the domain in the real physical space; i.e.

$$\left.\begin{aligned} [\![\varphi]\!]_{\partial V_f}; \\ [\![\vec{\Omega}\varphi \cdot \vec{n}]\!]_{\partial V_f} \end{aligned}\right\} = 0, \quad \forall \partial V_f \in \partial V_{int}. \tag{5.42}$$

Likewise, as inherited by the property of adjoint-consistency, the solution to the dual problem, $\varphi^\dagger$, is equally smooth; i.e. $\varphi, \varphi^\dagger \in \mathcal{V}(V) \subset H^{s>½}(V)$. Without loss of generality, an emphasis is placed upon manipulating the direct equations; and the same discretisation is applied to the adjoint equations; although one appreciates that discrete adjoint-consistency of the unmodified SAAF-S_N is only attained in the asymptotic limit [125]. Again, one formulates the weak form by multiplying Equation 5.2 by test functions, $w \in \mathcal{V}(\mathcal{T})$, integrating over each V_e and then taking the sum of these contributions. The angular neutron flux is expanded in a basis of trial functions from the same space, $u \in \mathcal{V}(\mathcal{T}) \subseteq H^s(\mathcal{T})$, which is the s-order broken Sobolev space, Equation 4.54. By application of the divergence theorem to each of the streaming terms, the basis functions must only be first-order differentiable; viz. $H^1(\mathcal{T})$ is an appropriate space.

The proposed spatial discretisation is non-conforming since $\hat{\varphi} \in \mathcal{V}(\mathcal{T}) \not\subset \mathcal{V}(V) \ni \varphi$; and properties of the (semi-)analytical solution cannot be assumed by the discrete one; and it will be necessary to establish directly discrete stability and consistency. The discrete variational problem is posed by restricting the collection of infinite spaces of basis functions to finite-dimensional subsets; i.e. $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \subset \mathcal{V}(\mathcal{T})$. As per the isoparametric concept, these broken spaces are chosen as those spanned by the same N_e-number of NURBS that describe the patches, V_e, that constitute the mesh, $\mathcal{T}$; q.v. Equation 4.61. Clearly, $\hat{\mathcal{V}}_{h,p}(\mathcal{T})$ is chosen such that $\hat{\varphi} \to \varphi$ as $N \to \infty$. The weighted-residual statement must be satisfied for each member in the weighting set; and since the space of trial and test functions coincide, this forms the same number of algebraic equations as there are unknowns per group; i.e. $DoF_e^{gk} = \dim \hat{\mathcal{V}}_{h,p}(V_e) = N_e$ and $DoF^{gk} = \sum_{V_e} DoF_e^{gk} = N$ for $k = 1, \ldots, K$; and

hence, $DoF^g = \sum_k DoF^{gk} = K \times N$ for $g = 1, \ldots, G$; and hence, $DoF = \sum_g DoF^g = G \times K \times N$. Consequently, for the analysis over a single energy-independent mesh, the discrete problem is posed as a system of $(GKN \times GKN)$-matrix equations.

Naturally satisfied by the weak form is the homogeneous Dirichlet-type BC, or the vacuum BC, over the inflow portions of the exterior boundary, ∂V_{ext}^-; q.v. Equation 5.23a. The reflective BC is implemented in a weak sense by substitution of Equation 5.23b. BC must also be prescribed over those outflow portions ∂V_{ext}^+; where the solution must satisfy the first-order form of the NTE; q.v. Equation 5.23c. Equally, this condition is required to hold over those faces to the interior of the mesh, ∂V_{int}.

As it stands, the discrete variational statement of Equation 5.8 is to find $\hat{\varphi} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, such that:

$$a(\hat{\varphi}, \hat{w}) = b(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}). \tag{5.43}$$

The discrete bilinear form, $a(\cdot,\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \to \mathbb{R}$, is block-like structure and comprises of terms of streaming and total collision, group-to-group scattering and fission, leakage across the union of interfaces and outflow portions of the external boundary, $\partial V_{IP} = (\partial V_{ext}^-)^c = \partial V_{int} \cup \partial V_{ext}^+$, and the weak enforcement of incoming BC over ∂V_{ext}^-:

$$a(\hat{\varphi}, \hat{w}) = a_L(\hat{\varphi}, \hat{w}) - a_S(\hat{\varphi}, \hat{w}) - a_F(\hat{\varphi}, \hat{w}) + a_{IP}(\hat{\varphi}, \hat{w}) + a_{ext}(\hat{\varphi}, \hat{w}). \tag{5.44}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$\begin{aligned} a_L(\hat{\varphi},\hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e\in\mathcal{T}} \int_{V_e} \vec{\Omega}^k\cdot\vec{\nabla}\hat{w}\frac{1}{\Sigma_{t,e}^g}\vec{\Omega}^k\cdot\vec{\nabla}\hat{\varphi}^{gk} + \hat{w}\Sigma_{t,e}^g\hat{\varphi}^{gk}\, dV\Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_L^{gk}(\hat{\varphi}^{gk},\hat{w}); \end{aligned} \tag{5.45a}$$

$$\begin{aligned} a_S(\hat{\varphi},\hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{g'=1}^{G} Y_0^0(\vec{\Omega}^k) \sum_{V_e\in\mathcal{T}} \int_{V_e} \hat{w}\Sigma_{s0,e}^{g'\to g}\hat{\Phi}_0^{g'0} + \vec{\Omega}^k\cdot\vec{\nabla}\hat{w}\frac{\Sigma_{s0,e}^{g'\to g}}{\Sigma_{t,e}^g}\hat{\Phi}_0^{g'0}\, dV\Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_S^{gk}(\hat{\varphi},\hat{w}); \end{aligned} \tag{5.45b}$$

$$\begin{aligned} a_F(\hat{\varphi},\hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{g'=1}^{G} Y_0^0(\vec{\Omega}^k) \sum_{V_e\in\mathcal{T}} \chi_e^g \int_{V_e} \hat{w}\nu_e^{g'}\Sigma_{f,e}^{g'}\hat{\Phi}_0^{g'0} + \vec{\Omega}^k\cdot\vec{\nabla}\hat{w}\frac{\nu_e^{g'}\Sigma_{f,e}^{g'}}{\Sigma_{t,e}^g}\hat{\Phi}_0^{g'0}\, dV\Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_F^{gk}(\hat{\varphi},\hat{w}); \end{aligned} \tag{5.45c}$$

$$\begin{aligned} a_{IP}(\hat{\varphi},\hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e\in\mathcal{T}} \sum_{\partial V_f\in V_e\cap\partial V_{IP}} \int_{\partial V_f} \hat{w}\vec{\Omega}^k\hat{\varphi}^{gk}\cdot\vec{n}\, dS\Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{\partial V_f\in\partial V_{IP}} \int_{\partial V_f} [\![\hat{w}\vec{\Omega}^k\hat{\varphi}^{gk}\cdot\vec{n}]\!]\, dS, \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_{IP}^{gk}(\hat{\varphi}^{gk},\hat{w}); \end{aligned} \tag{5.45d}$$

$$\begin{aligned} a_{ext}(\hat{\varphi},\hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e\in\mathcal{T}} \sum_{\partial V_f\in V_e\cap\partial V_R^-} - \int_{\partial V_f} |\vec{\Omega}^k\cdot\vec{n}|\hat{w}\hat{\varphi}^{gk\prime}\, dS\Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} a_{ext}^{gk}(\hat{\varphi}^g,\hat{w}). \end{aligned} \tag{5.45e}$$

The discrete linear functional, $b(\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \to \mathbb{R}$, comprises only of fixed sources:

$$\begin{aligned} b(\hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e\in\mathcal{T}} \int_{V_e} \hat{w}\frac{1}{4\pi}q_{x,e}^g + \vec{\Omega}^k\cdot\vec{\nabla}\hat{w}\frac{1}{4\pi\Sigma_{t,e}^g}q_{x,e}^g\, dV\Bigg], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} b^{gk}(\hat{w}). \end{aligned} \tag{5.46}$$

For multiplying systems, $q_x = 0$, the weak form of Equation 5.9 is to find $\hat{\varphi} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, such that:

$$\tilde{a}(\hat{\varphi},\hat{w}) = \frac{1}{\mathrm{k_{eff}}} a_f(\hat{\varphi},\hat{w}), \quad \forall\ \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}); \tag{5.47}$$

where $\tilde{a}$ is the original bilinear form in Equation 5.44, a, less the fission term; i.e.:

$$\tilde{a} = a + a_F = a_L - a_S + a_{IP} + a_{ext}. \tag{5.48}$$

For those reasons already broached, the properties of neutron conservation and symmetry will be

preserved in devising an interior penalty scheme for the discretised bilinear form.

Consistency The discretisation is consistent if the true solution satisfies the discrete system of equations; however, one cannot simply substitute φ into Equation 5.43 because the bilinear form, a, has been defined on $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}_{h,p}(\mathcal{T})$; and similarly for $\tilde{a}$ in Equation 5.47. Therefore, one must assume that $\varphi \in \mathcal{V}^*(V) \subset \mathcal{V}(V)$; and second, that $\hat{\varphi}^* \in \hat{\mathcal{V}}^*_{h,p}(\mathcal{T}) = \mathcal{V}^*(V) + \hat{\mathcal{V}}_{h,p}(\mathcal{T})$; such that the bilinear form can be extended to be bounded on $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}^*_{hp}(\mathcal{T})$. The space $\hat{\mathcal{V}}^*_{h,p}(\mathcal{T})$ is equipped with the norm $\|\cdot\|_{\hat{e}^*}$; q.v. Equations 4.68.

From the definition of jumps, Equation 4.58a, the first line of Equation 5.45d expressed as a sum of integrals over element boundaries can be re-expressed in the second line as a sum of integrals over interfaces and outflow portions of the external boundary, which can then be developed further using the definition of averages, Equation 4.58b; i.e. suppressing the spatial dependence for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$\begin{aligned} &\sum_{\partial V_f \in \partial V_{IP}} \int_{\partial V_f} [\![\hat{w} \vec{\Omega}^k \hat{\varphi}^{gk*} \cdot \vec{n}]\!] \, dS \\ &= \sum_{\partial V_f \in \partial V_{int}} \int_{\partial V_f} [\![\hat{w} \vec{\Omega}^k \hat{\varphi}^{gk*} \cdot \vec{n}]\!] \, dS + \sum_{V_e \in \mathcal{T}} \sum_{\partial V_f \in V_e \cap \partial V^+_{ext}} \int_{\partial V_f} \hat{w} \vec{\Omega}^k \hat{\varphi}^{gk*} \cdot \vec{n} \, dS, \end{aligned} \tag{5.49a}$$

$$= \sum_{\partial V_f \in \partial V_{int}} \int_{\partial V_f} \{\!\{ \hat{w} \}\!\} \cancelto{0}{[\![\vec{\Omega}^k \hat{\varphi}^{gk*} \cdot \vec{n}]\!]} \, dS + \sum_{\partial V_f \in \partial V_{IP}} \int_{\partial V_f} [\![\hat{w}]\!] \{\!\{ \vec{\Omega}^k \hat{\varphi}^{gk*} \cdot \vec{n} \}\!\} \, dS. \tag{5.49b}$$

The strong form, Equation 5.1a, and the BC, Equations 5.23, are retrieved by substitution of φ into the discrete weak-problem, Equation 5.43; and by application of the divergence theorem to the first term in Equation 5.45a, the double-differential term is recouped. This is followed by the re-expression of a_{IP}, Equation 5.45d, in terms of Equation 5.49b; and then accounting for the regularity in the true solution, Equations 5.42. Thus, the discrete problem is consistent provided that the numerical fluxes are consistent; and exploiting continuity of the normal component of $\vec{\Omega}^k$ across interfaces, the following modification can be made for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$\implies a^{gk}_{IP}(\hat{\varphi}^{gk*}, \hat{w}) = \sum_{\partial V_f \in \partial V_{IP}} \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n} [\![\hat{w}]\!] \{\!\{ \hat{\varphi}^{gk*} \}\!\} \, dS. \tag{5.50}$$

Galerkin's property of orthogonality stems from consistency such that the error in the numerical solution is orthogonal to the chosen subspace since $\mathcal{V}^*(V) \subset \hat{\mathcal{V}}^*_{h,p}(\mathcal{T})$; q.v. Equation 4.71.

A strong definition of consistency has been enforced, yet integral quantities will be evaluated via numerical quadrature; q.v. Section 3.2.3. Consequently, one expects consistency, asymptotically. Equally important in upholding the properties of local conservation and consistency is an ability to evaluate sufficiently any boundary integrals. The use of p+1 QP may no longer adequate because the

integrands are polynomial constructs of a combined maximal polynomial degree of $p+p-1+p=3p-1$; q.v. Equation 5.50. Where each component of the normal vector, $\vec{n}$, is itself a linear combination of derivatives of the same functions; viz. polynomial constructs of degree p-1. Using Legendre-Gauss quadrature rules, the boundary integral should be evaluated sufficiently with Q-number of QP in each of the (d-1)-directions over the canonical representation of the boundary in the parent space; i.e.:

$$2N_q - 1 \geq 3p - 1; \tag{5.51a}$$

$$\implies N_q = \begin{cases} \dfrac{3p}{2}, & \text{for } p\text{-even}; \\ \dfrac{3p+1}{2}, & \text{for } p\text{-odd}. \end{cases} \tag{5.51b}$$

In the interest of AMR, asymptotic consistency and the conservation of neutrons must hold over irregular meshes. To satisfy Equations 5.42, one employs quadrature rules that adequately sample the bases either side of an interface, $\forall \partial V_f \in \partial V_{int}$; q.v. Equation 4.73. For those irregularities induced by h-refinement and the subsequent integration over element boundaries that contain such hanging-nodes, Equations 5.51 would still be inadequate because the integrand will not generally be a polynomial. Nonetheless, the integrand over partial boundaries remains a continuous function over smooth NURBS objects and it is possible to gain in accuracy by supplying additional QP.

Symmetry Foreshadowing the call upon duality arguments and the desire to employ the conjugate gradient method, the following is appended to a_{IP} for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$\sum_{\partial V_f \in \partial V_{int}} \theta \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n} \{\{\hat{w}\}\} [\![\hat{\varphi}^{gk*}]\!] \, dS = 0 \quad \forall \partial V_f \in \partial V_{int}; \tag{5.52}$$

or rather, redundantly, symmetry is reclaimed by letting $\theta = 1$ in the following modification to a_{IP}^{gk}, Equation 5.50, for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$\implies a_{IP}^{gk}(\hat{\varphi}^{gk*}, \hat{w}) = a_{IP_c}^{gk}(\hat{\varphi}^{gk*}, \hat{w}) + a_{IP_s}^{gk}(\hat{\varphi}^{gk*}, \hat{w}), \tag{5.53a}$$

$$\begin{aligned} &= \sum_{\partial V_f \in \partial V_{IP}} \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n} [\![\hat{w}]\!] \{\{\hat{\varphi}^{gk*}\}\} \, dS \\ &+ \sum_{\partial V_f \in \partial V_{int}} \theta \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n} \{\{\hat{w}\}\} [\![\hat{\varphi}^{gk*}]\!] \, dS. \end{aligned} \tag{5.53b}$$

Consistency is maintained as $[\![\varphi]\!] = 0$ across the interior of the domain.

With the exception of a_{IP}, Equations 5.45 are given as a sum over elements or boundaries whose local contributions are the product of basis functions with local support over the same domain. Differently, Equation 5.53 comprises of a sum over interfaces, whose local contributions are the product of basis functions that may not necessarily have local support over the same element, such that a_{IP} can be block-partitioned; q.v. Equation 4.76 [237]. Again, savings can be made by considering that the

symmetry-part of the block is simply the transpose of the consistency-part.

Penalty As per the Lax-Milgram Lemma, the properties of continuity, Equation 4.78, and more specifically coercivity, Equation 4.79, are sufficient conditions to assert well-posedness [338]. The latter is often simply guaranteed by the nature of the problem; however, owing to the arbitrariness of the normal vector over an interface, the sign-indefinite consistency and the symmetry terms in Equation 5.53 are potentially negative terms. In general, it is not possible to uphold a statement of coercivity without some stabilisation mechanism; q.v. Section 4.4 [330].

Discontinuities in the numerical approximation result from the non-conformity in seeking a discrete solution in a broken function space. A final modification is thus made to a_{IP}, such that appended to the bilinear form is a residual-dependent term that *contributes* to the coercivity by penalising the error, or rather quashing any jumps $[\![\hat{\varphi}]\!]$ across the set of discontinuous patch interfaces; q.v. Equation 5.54. Stability of the discretisation is enhanced and invertibility is ensured under refinement for some suitable selection of penalty parameters, $\mu_f^g > 0$, for $g = 1, \dots, G$ and for $k = 1, \dots, K$:

$$\begin{aligned} \implies a_{IP}^{gk}(\hat{\varphi}^{gk*}, \hat{w}) = & \sum_{\partial V_f \in \partial V_{ext}^+} \int_{\partial V_f} \hat{w} \vec{\Omega}^k \hat{\varphi}^{gk*} \cdot \vec{n} \, dS \\ & + \sum_{\partial V_f \in \partial V_{int}} \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n} [\![\hat{w}]\!] \{\{\hat{\varphi}^{gk*}\}\} \, dS + \theta \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n} \{\{\hat{w}\}\} [\![\hat{\varphi}^{gk*}]\!] \, dS \\ & + \sum_{\partial V_f \in \partial V_{int}} \mu_f^g \int_{\partial V_f} [\![\vec{\Omega}^k \hat{w} \cdot \vec{n}]\!] [\![\vec{\Omega}^k \hat{\varphi}^{gk*} \cdot \vec{n}]\!] \, dS. \end{aligned} \tag{5.54}$$

One notes the particular choice to penalise the angular neutron current will ease the derivation of the penalty parameters in Section 5.4.1. Since Dirichlet BC are satisfied *naturally*, they need not be enforced by penalisation. Again, the properties of symmetry and consistency have been maintained; and error bounds can be established as per Strang's Second Lemma of quasi-best approximation for non-conforming and consistent discretisations; q.v. Equation 4.81 [237].

The discrete system is solved via the SI procedure; q.v. Section 2.8. Weakly coupled in energy and angle, the within-group equations are solved sequentially with the most up-to-date available information. For fixed sources, $\hat{\varphi}^{gk} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ is sought for $g = 1, \dots, G$ and for $k = 1, \dots, K$:

$$a_L^{gk}(\hat{\varphi}^{gk}, \hat{w}) + a_{IP}^{gk}(\hat{\varphi}^{gk}, \hat{w}) = a_S^{gk}(\hat{\varphi}, \hat{w}) + a_F^{gk}(\hat{\varphi}, \hat{w}) - a_{ext}^{gk}(\hat{\varphi}^{gk}, \hat{w}) + b^{gk}(\hat{w}), \quad \forall \, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}). \tag{5.55}$$

The left-hand side, $a_L^{gk} + a_{IP}^{gk}$, is SPD for $\theta = 1$ and the above can be solved by the PCG method.

Employing the PI procedure to converge $|1/\mathrm{k_{eff}}|$, the eigenvalue problem is cast as finding $\hat{\varphi}^{gk} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$

for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$a_L^{gk}(\hat{\varphi}^{gk}, \hat{w}) + a_{IP}^{gk}(\hat{\varphi}^{gk}, \hat{w}) = a_S^{gk}(\hat{\varphi}, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a_F^{gk}(\hat{\varphi}, \hat{w}) - a_{ext}^{gk}(\hat{\varphi}^{gk}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}). \tag{5.56}$$

In kind, the discrete variational statement of the adjoint problem in the presence of fixed adjoint sources, $q_x^\dagger \neq 0$, is cast as finding $\hat{\varphi}^{gk\dagger} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ for $g = G, \ldots, 1$ and for $k = 1, \ldots, K$:

$$\begin{aligned} a_L^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) + a_{IP}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) = \\ a_S^{gk\dagger}(\hat{\varphi}^\dagger, \hat{w}) + a_F^{gk\dagger}(\hat{\varphi}^\dagger, \hat{w}) - a_{ext}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) + b^{gk\dagger}(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}). \end{aligned} \tag{5.57}$$

The adjoint eigenvalue problem is cast as finding $\hat{\varphi}^{gk\dagger} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ for $g = G, \ldots, 1$ and for $k = 1, \ldots, K$:

$$\begin{aligned} a_L^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) + a_{IP}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) = \\ a_S^{gk\dagger}(\hat{\varphi}^\dagger, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a_F^{gk\dagger}(\hat{\varphi}^\dagger, \hat{w}) - a_{ext}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}); \end{aligned} \tag{5.58}$$

where the discrete adjoint bilinear form, $a^\dagger(\cdot,\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}_{h,p}^*(\mathcal{T}) \to \mathbb{R}$, is defined by:

$$a^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) = a_L^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) - a_S^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) - a_F^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) + a_{IP}^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) + a_{ext}^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}). \tag{5.59}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$
\begin{aligned}
a_L^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e \in \mathcal{T}} \int_{V_e} \vec{\Omega}^k \cdot \vec{\nabla}\hat{w} \frac{1}{\Sigma_{t,e}^g} \vec{\Omega}^k \cdot \vec{\nabla}\hat{\varphi}^{gk\dagger *} + \hat{w}\Sigma_{t,e}^g \hat{\varphi}^{gk\dagger *} \, dV \Bigg], \\
&= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k a_L^{gk\dagger}(\hat{\varphi}^{gk\dagger *}, \hat{w});
\end{aligned} \tag{5.60a}
$$

$$
\begin{aligned}
a_S^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{g\prime=1}^{G} Y_0^0(\vec{\Omega}^k) \sum_{V_e \in \mathcal{T}} \int_{V_e} \hat{w}\Sigma_{s0,e}^{g\to g\prime} \hat{\Phi}_0^{g\prime 0\dagger *} - \vec{\Omega}^k \cdot \vec{\nabla}\hat{w} \frac{\Sigma_{s0,e}^{g\to g\prime}}{\Sigma_{t,e}^g} \hat{\Phi}_0^{g\prime 0\dagger *} \, dV \Bigg], \\
&= \sum_{g=1}^{G}\sum_{k=1}^{K} a_S^{gk\dagger}(\hat{\varphi}^{\dagger *}, \hat{w});
\end{aligned} \tag{5.60b}
$$

$$
\begin{aligned}
a_F^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{g\prime=1}^{G} Y_0^0(\vec{\Omega}^k) \sum_{V_e \in \mathcal{T}} \nu_e^g \Sigma_{f,e}^g \int_{V_e} \hat{w}\chi_e^{g\prime} \hat{\Phi}_0^{g\prime 0\dagger *} - \vec{\Omega}^k \cdot \vec{\nabla}\hat{w} \frac{\chi_e^{g\prime}}{\Sigma_{t,e}^g} \hat{\Phi}_0^{g\prime 0\dagger *} \, dV \Bigg], \\
&= \sum_{g=1}^{G}\sum_{k=1}^{K} a_F^{gk\dagger}(\hat{\varphi}^{\dagger *}, \hat{w});
\end{aligned} \tag{5.60c}
$$

$$
\begin{aligned}
a_{IP}^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{\partial V_f \in \partial V_{ext}^-} \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n}\hat{w}\hat{\varphi}^{gk\dagger *} \, dS \\
&+ \sum_{\partial V_f \in \partial V_{int}} \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n} [\![\hat{w}]\!] \{\!\{\hat{\varphi}^{gk\dagger *}\}\!\} \, dS + \theta \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n} \{\!\{\hat{w}\}\!\} [\![\hat{\varphi}^{gk\dagger *}]\!] \, dS \\
&+ \sum_{\partial V_f \in \partial V_{int}} \mu_f^g \int_{\partial V_f} [\![\vec{\Omega}^k \hat{w} \cdot \vec{n}]\!] [\![\vec{\Omega}^k \hat{\varphi}^{gk\dagger *} \cdot \vec{n}]\!] \, dS \Bigg], \\
&= \sum_{g=1}^{G}\sum_{k=1}^{K} a_{IP}^{gk\dagger}(\hat{\varphi}^{gk\dagger *}, \hat{w});
\end{aligned} \tag{5.60d}
$$

$$
\begin{aligned}
a_{ext}^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e \in \mathcal{T}} \sum_{\partial V_f \in V_e \cap \partial V_R^+} - \int_{\partial V_f} |\vec{\Omega}^k \cdot \vec{n}| \hat{w}\hat{\varphi}^{gk\prime\dagger *} \, dS \Bigg], \\
&= \sum_{g=1}^{G}\sum_{k=1}^{K} a_{ext}^{gk\dagger}(\hat{\varphi}^{g\dagger *}, \hat{w});
\end{aligned} \tag{5.60e}
$$

and the discrete adjoint linear functional, $b^\dagger(\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \to \mathbb{R}$ is defined as:

$$
\begin{aligned}
b^\dagger(\hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Bigg[\sum_{V_e \in \mathcal{T}} \int_{V_e} \hat{w} \frac{1}{4\pi} q_{x,e}^{g\dagger} - \vec{\Omega}^k \cdot \vec{\nabla}\hat{w} \frac{1}{4\pi\Sigma_{t,e}^g} q_{x,e}^{g\dagger} \, dV \Bigg], \\
&= \sum_{g=1}^{G}\sum_{k=1}^{K} b^{gk\dagger}(\hat{w}).
\end{aligned} \tag{5.61}
$$

For eigenvalue problems, the solution is normalised after each outer iteration such that $a_F(\varphi, \varphi) = 1$; and the adjoint solution is normalised such that $a_F(\varphi, \varphi^\dagger) = 1$; q.v. Section 3.6.2.2.

5.4.1 A Discrete Coercivity Analysis

The performance of the SIP-DG discretisation will depend heavily upon the choice of penalty parameters; viz. they must be large enough to ensure definiteness and invertibility of the matrix equations, yet not arbitrarily large; q.v. Section 4.4.3. Suffice it to say, under-penalisation is just as unwelcome as over-penalisation; one therefore seeks the lower bound of the *Goldilocks* range of penalty parameters via a coercivity analysis, based upon the methodology as presented by Shahbazi [348].

Since the discretised system is only weakly coupled in energy and angle, it is sufficient to consider separately the coercivity of each group-ordinate bilinear form, the left-hand side of Equation 5.55 and Equation 5.56; where $a^{gk}_{lhs}(\hat{u},\hat{u})$ is recalled below for some $\hat{u} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, suppressing the spatial dependence for $g = 1,\dots,G$ and for $k = 1,\dots,K$:

$$a^{gk}_{lhs}(\hat{u},\hat{u}) = a^{gk}_{L}(\hat{u},\hat{u}) + a^{gk}_{IP}(\hat{u},\hat{u}), \tag{5.62a}$$

$$\begin{aligned} &= \sum_{V_e\in\mathcal{T}}\int_{V_e}\frac{1}{\Sigma^{g}_{t,e}}(\vec{\Omega}^{k}\cdot\vec{\nabla}\hat{u})^2 + \Sigma^{g}_{t,e}(\hat{u})^2\,dV + \sum_{\partial V_f\in\partial V^{+}_{ext}}\int_{\partial V_f}|\vec{\Omega}^{k}\cdot\vec{n}|(\hat{u})^2\,dS \\ &- \sum_{\partial V_f\in\partial V_{int}}(1+\theta)\int_{\partial V_f}[\![\vec{\Omega}^{k}\hat{u}\cdot\vec{n}]\!]\{\{\hat{u}\}\}\,dS + \sum_{\partial V_f\in\partial V_{int}}\mu^{g}_{f}\int_{\partial V_f}([\![\vec{\Omega}^{k}\hat{u}\cdot\vec{n}]\!])^2\,dS. \end{aligned} \tag{5.62b}$$

The group-ordinate bilinear form is self-adjoint, $a^{gk}_{lhs} = a^{gk\dagger}_{lhs}$, for $\theta = 1$; and since the discrete variational statements for both the primal problem and the dual problem are formulated on the same mesh and their solutions sought in the same space, $\varphi,\varphi^{\dagger} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, the well-posedness of either depends upon a common coercivity analysis that will yield a single set of energy-dependent penalty parameters.

As per the arbitrariness of the normal vector defined over an interface, the only potentially negative term and thus, the only one that may not *contribute* towards the coercivity is that in the second line of Equation 5.62b. Clearly, the NIP-DG scheme is coercive for any choice of $\mu_f > 0$; whereas a lower bound must be imposed on the IIP and the SIP variations. For reasons already discussed, only the latter is investigated; however, for completeness, the variable θ is retained.

Sufficient penalisation of each face $\partial V_f \in \partial V_{IP}$, ensures coercivity of the bilinear form with respect to the energy-norm, which is defined for some $\hat{u} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ and for $g = 1,\dots,G$ and for $k = 1,\dots,K$:

$$\begin{aligned} \|\hat{u}\|^2_{\hat{e}^{gk}} \coloneqq &\sum_{V_e\in\mathcal{T}}\int_{V_e}\frac{1}{\Sigma^{g}_{t,e}}(\vec{\Omega}^{k}\cdot\vec{\nabla}\hat{u})^2 + \Sigma^{g}_{t,e}(\hat{u})^2\,dV + \sum_{\partial V_f\in\partial V^{+}_{ext}}\int_{\partial V_f}|\vec{\Omega}^{k}\cdot\vec{n}|(\hat{u})^2\,dS \\ &+ \sum_{\partial V_f\in\partial V_{int}}\mu^{g}_{f}\int_{\partial V_f}([\![\vec{\Omega}^{k}\hat{u}\cdot\vec{n}]\!])^2\,dS. \end{aligned} \tag{5.63}$$

A lower bound is established for μ_f through a coercivity analysis of the bilinear form; q.v. Section 4.4.1 [348]. Following the developments made by Owens et al, the TIC are calculated as generalised eigenvalue problems, agnostic of the underlying element type; q.v. Section 4.4.2 [334]. However, the

trace inequality employed here is that for the Dirichlet trace, Equation 4.56a.

From the definition of jumps, Equation 4.58a, and averages, Equation 4.58b; and letting $a = [\![\vec{\Omega}^k \hat{u} \cdot \vec{n}]\!]$ and $b = \{\{\hat{u}\}\}$ in Young's inequality, Equation 4.92; and the substitution thereof into Equation 5.62b bounds the potentially negative term in the second line over all discontinuous faces, $\forall \partial V_f \in \partial V_{int}$; and collecting like-terms for $g = 1, \dots, G$ and for $k = 1, \dots, K$:

$$\begin{aligned} a_{lhs}^{gk}(\hat{u}, \hat{u}) \geq a_L^{gk}(\hat{u}, \hat{u}) &+ \sum_{\partial V_f \in \partial V_{ext}^+} \int_{\partial V_f} |\vec{\Omega}^k \cdot \vec{n}| (\hat{u})^2 \, dS \\ &- \sum_{\partial V_f \in \partial V_{int}} \frac{(1+\theta)}{2\epsilon_f} \int_{\partial V_f} (\{\{\hat{u}\}\})^2 \, dS \\ &+ \sum_{\partial V_f \in \partial V_{int}} \left(\mu_f^g - \frac{(1+\theta)\epsilon_f}{2} \right) \int_{\partial V_f} ([\![\vec{\Omega}^k \hat{u} \cdot \vec{n}]\!])^2 \, dS. \end{aligned} \tag{5.64}$$

By the definition provided in Equation 4.58b, the term $(\{\{\hat{u}\}\})^2$ is expanded out, to which Equation 4.92 is applied with $\epsilon = 1$; i.e. $a^2 + b^2 + 2ab \leq 2a^2 + 2b^2$. The result is rewritten as a sum of integrals over patch boundaries, where individual terms are collected onto their respective patches; i.e.:

$$\begin{aligned} a_{lhs}^{gk}(\hat{u}, \hat{u}) \geq a_L^{gk}(\hat{u}, \hat{u}) &+ \sum_{\partial V_f \in \partial V_{ext}^+} \int_{\partial V_f} |\vec{\Omega}^k \cdot \vec{n}| (\hat{u})^2 \, dS \\ &- \sum_{V_e \in \mathcal{T}} (1+\theta) \left[\sum_{\partial V_f \in V_e \cap \partial V_{int}} \frac{1}{4\epsilon_f} \int_{\partial V_f} (\hat{u})^2 \, dS \right] \\ &+ \sum_{\partial V_f \in \partial V_{int}} \left(\mu_f^g - \frac{(1+\theta)\epsilon_f}{2} \right) \int_{\partial V_f} ([\![\vec{\Omega}^k \hat{u} \cdot \vec{n}]\!])^2 \, dS. \end{aligned} \tag{5.65}$$

Taking advantage of the equivalence of norms in the chosen function space, there exists a discrete trace inequality that bounds from above, up to some constant, the integral quantity of the restriction of a function, as defined over the boundary ∂V_e; i.e. for some function $\hat{u} \in \hat{\mathcal{V}}_{h,p}(V_e)$:

$$\|\gamma_0 \hat{u}\|_{L_2(\partial V_e)}^2 \leq C_{e,f}^0 \|\hat{u}\|_{L_2(V_e)}^2; \tag{5.66}$$

where $C_{e,f}^0$ is the TIC that pertains to the Dirichlet trace in Equation 4.56a, for which those integrals over patch boundaries in second line of Equation 5.65 are substituted; and one seeks the sharpest TIC values in order to optimise the penalty parameters. Collecting like-terms and comparing against Equation 4.79, the following conditions must be satisfied for $g = 1, \cdots, G$ and for $k = 1, \cdots, K$:

$$1 - (1+\theta) \left[\sum_{\partial V_f \in V_e \cap \partial V_{int}} \frac{C_{e,f}^0}{4\Sigma_{t,e}^g \epsilon_f} \right] > 0, \qquad \forall V_e \in \mathcal{T}^g(V); \tag{5.67a}$$

$$\mu_f^g - \frac{(1+\theta)\epsilon_f}{2} > 0, \qquad \forall \partial V_f \in \partial V_{int}; \tag{5.67b}$$

One must choose $\epsilon_f > 0$, such that Equation 5.67a is satisfied over each element. This is achieved by separately bounding from below each term in the sum over the F_e-number of faces of a given element;

and to ensure that Equations 5.67 hold, the maximum penalisation is applied across an interface [350]. Encapsulated below, the group-dependent penalty parameters may be calculated in a pre-processing step for $g = 1, \cdots, G$:

$$\mu_f^g > \frac{1}{8}(1+\theta)^2 \max_{V_e \in \bigcap_{e\prime} V_{e\prime} \ni \partial V_f} \frac{C_{e,f}^0 F_e}{\Sigma_{t,e}^g}, \quad \forall \partial V_f \in \partial V_{int}. \tag{5.68a}$$

Similar parametric studies to those presented in Section 4.4.3 were performed on the SIP-DG-IGA discretisation of the *one-speed* SAAF-S_N equations; since qualitatively similar results were obtained, they are not presented here for brevity; viz. to the detriment of the conditioning of the matrix equations, there was no significant gain in accuracy for super-penalisation; and the discretisation was not convergent for any gross under-penalisation.

5.5 The Numerical Results

In the application to the multi-group SAAF-S_N equations, this section presents some numerical results for the AMR-$h(p)$ strategies developed in Section 3.5 and Section 3.6. Unless otherwise stated, all operations are performed such as to optimise the H^1-error. For a couple of verification test cases and nuclear reactor physics benchmarks, the AMR strategies are compared against global uniform ones. Any data-plots that make reference to the DoF of the discretised system justly account for all unknowns that were solved for in the complete system; viz. $DoF = G \times K \times N$ for G-number of groups and K-number of ordinates solved over an energy-independent mesh of N-number of control-variables.

Both discretisations developed within this chapter have been implemented in the same new modern-Fortran code as that presented in Section 4.5; and thus, the same libraries, tolerances and parameters are employed here [73]. Again, all *jobs* are submitted to the general purpose Imperial College HPC cluster computing service (CX1); and at each iteration of the AMR procedure, the reference solution is *bootstrapped* from the previous iteration to serve as an initial guess over the current reference mesh.

5.5.1 The Method of Manufactured Solutions

For some manufactured solution, φ^{gk}, the (semi-)analytical anisotropic manufactured source, q_x^{gk}, can be evaluated for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$Q^{gk} q_x = (L^{gk} - Q^{gk} S^{gk})\varphi; \tag{5.69}$$

where Q is the operator that acts group-ordinate-wise as defined in Equation 5.5; and q_x is the left-hand side of the FONT-S_N equations; i.e. for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$q_x^{gk} = \vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk} + \Sigma_t^g \varphi^{gk} - \sum_{g\prime=1}^{G} \Sigma_{s0}^{g\prime \to g} Y_0^0(\vec{\Omega}^k)\Phi_0^{g\prime 0}. \tag{5.70}$$

The manufactured solution and source are used to study the convergence rates of the global uniform refinement strategy on a simple Cartesian grid and a more challenging pin-cell one; taking into account the implementation of different BC and the numerical quadrature routines. These results are then used as a reference with which to analyse the NRG-AMR-$h(p)$ strategy, which takes into account performance of all mesh-to-mesh interpolation procedures and the capacity of the numerical methods to enforce continuity over irregular meshes that contain *hanging-nodes*.

5.5.1.1 A Cartesian Grid

The first verification case consists of a simple Cartesian grid defined over $V = [0, 1]^2$, which comprises of a homogeneous material whose properties are given in Table 4.1. The coarsest mesh is composed of *four* equal-volume squares represented by bilinear NURBS patches. Assuming a *one-speed* energy approximation and the S_4 LS angular quadrature set, the manufactured solution is chosen as a modified standing wave equation for $k = 1, \ldots, K$:

$$\begin{cases} \varphi^{nk}(\vec{r}) = \sin\left(B_x^n x\right)\sin\left(B_y^n y\right)[1 + \vec{\Omega}^k \cdot \vec{\Omega}^k], & \vec{r} \in V; \quad (5.71a) \\ B_x^n, B_y^n \coloneqq (2n{+}1)\pi, & \forall n \geq 0; \quad (5.71b) \end{cases}$$

with:

$$\Phi_0^0(\vec{r}) \approx \sum_{k=1}^{K} w_k Y_0^0(\vec{\Omega}^k)\sin\left(B_x^n x\right)\sin\left(B_y^n y\right)[1 + \vec{\Omega}^k \cdot \vec{\Omega}^k], \tag{5.72}$$

$$\approx \sin\left(B_x^n x\right)\sin\left(B_y^n y\right)\sum_{k=1}^{K} w_k Y_0^0(\vec{\Omega}^k)[1 + \vec{\Omega}^k \cdot \vec{\Omega}^k]; \tag{5.73}$$

where φ^{nk} denotes an odd harmonic mode as characterised by the geometric buckling, B_x^n and B_y^n. To allow for the maximum attainable convergence rates, the manufactured solution has been chosen to be smooth, yet it cannot be arbitrarily represented exactly by higher-order NURBS. Moreover, the homogeneous Dirichlet (vacuum) BC is naturally satisfied; it will also be naturally enforced by the discretisations on all external boundaries. One emphasises that the manufactured solution is chosen for the discrete ordinates equations rather than the full NTE; it is thus used to isolate those errors that result from the spatial discretisation [68].

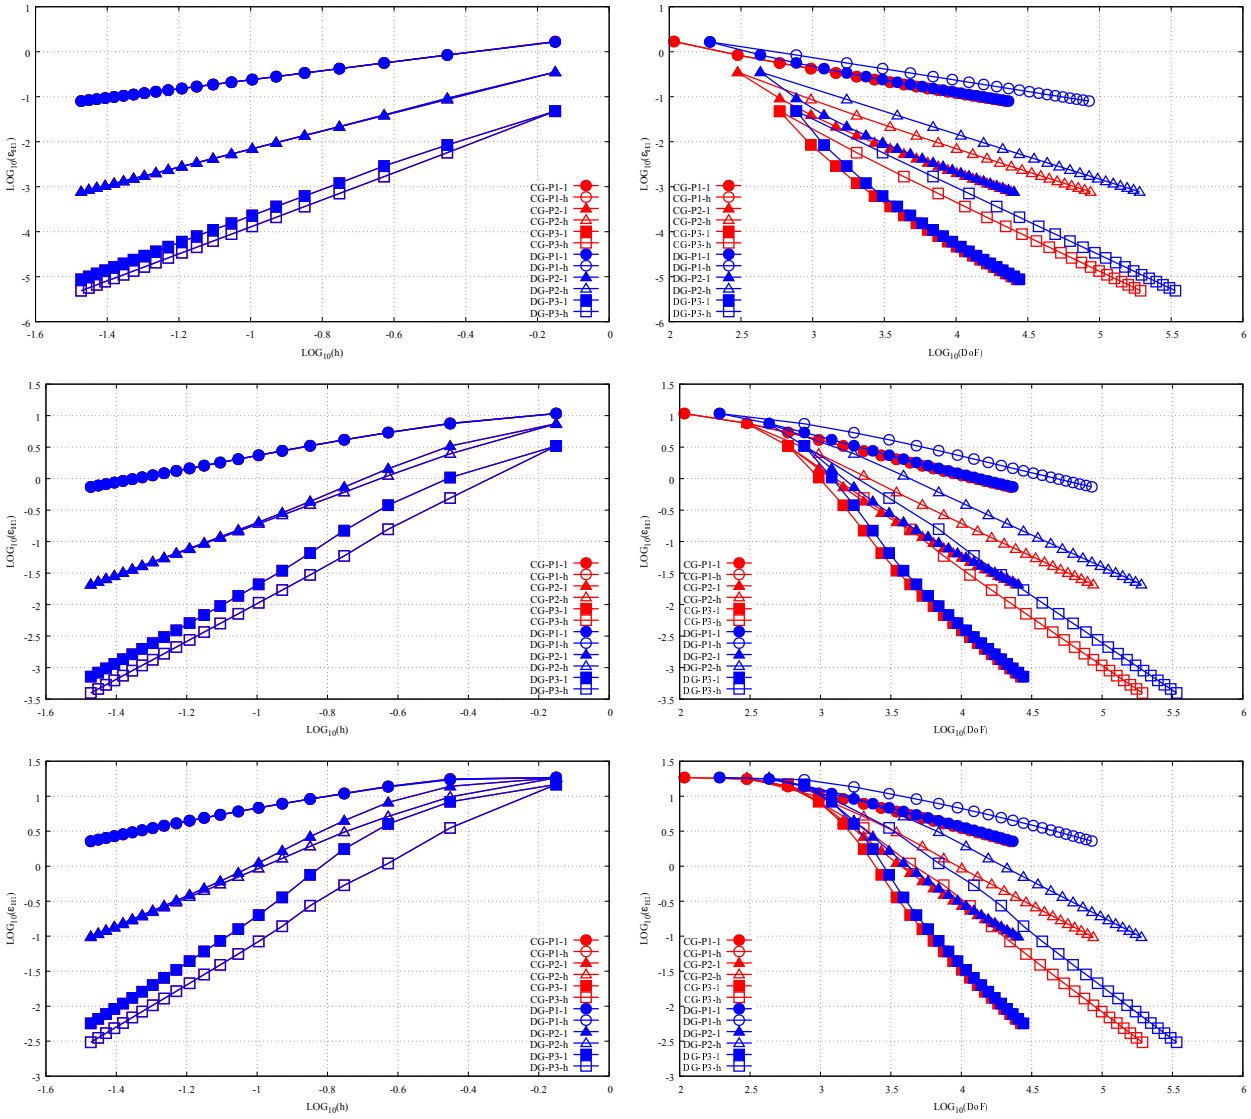


Figure 5.1: The convergence plots, ϵ_{H^1} v h (left) and ϵ_{H^1} v DoF (right), for the verification of the global uniform refinement strategy for the CBG-IGA and the SIP-DG-IGA discretisations of the 1G SAAF-S_4 equations over a 2D Cartesian grid; the manufactured solution is chosen as per Equation 4.105 for $n = 0$ (top row), $n = 1$ (middle row) and $n = 2$ (bottom row). (*Vide* the web-based version for reference to colour.)

Global Uniform Refinement For the KI-refinement of linear, quadratic and cubic B-splines, the H^1-errors for each of the CBG-IGA and the SIP-DG-IGA spatial discretisations are presented in Figure 5.1 for the manufactured solution chosen as the first, third and fifth harmonics of a modified standing wave equation; q.v. Equation 5.71. Qualitatively similar plots were obtained for the L_2-error; however, for brevity, these have not been presented. Under refinement, both the continuous and the discontinuous discretisation schemes converge to the analytical solution; and asymptotically, they attain their theoretical orders of accuracy, as presented in Table 5.1. These results are qualitatively similar to those obtained for the verification of the discretisation of the NDE and thus, for brevity, the discussion is not repeated here; q.v. Section 4.5.1.1.

		CBG-IGA				SIP-DG-IGA			
		KI-1		KI-h		KI-1		KI-h	
	p	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}
	1	1.00	1.99	1.00	1.99	1.00	2.00	1.00	2.00
$(n=0)$	2	2.00	3.01	2.00	3.00	2.00	3.01	2.00	3.00
	3	2.98	3.98	3.00	4.00	2.98	3.99	3.00	4.00
	1	1.03	1.98	1.03	1.99	1.04	1.99	1.04	1.99
$(n=1)$	2	2.02	3.02	2.00	3.06	2.02	3.02	2.00	3.05
	3	3.01	4.03	3.00	4.01	3.01	4.03	3.00	4.01
	1	1.09	1.97	1.09	1.97	1.09	1.98	1.09	1.98
$(n=2)$	2	2.07	3.04	2.03	3.03	2.07	3.04	2.02	3.03
	3	3.06	4.09	3.00	4.03	3.06	4.09	3.00	4.03

Table 5.1: The asymptotic orders of accuracy, Equations 4.103, for the verification of the global uniform refinement strategy for the CBG-IGA and the SIP-DG-IGA discretisations of the 1G SAAF-S_4 equations over a 2D Cartesian grid; the convergence rates are evaluated as the slope of the last few points from each dataset.

Local Adaptive Refinement The results of the H^1-error-driven NRG-AMR-$h(p)$ strategy for each of the CBG-IGA and the SIP-DG-IGA discretisations are presented in Figure 5.2 for the manufactured solution chosen as the first, third and fifth harmonics of a modified standing wave equation; q.v. Equation 5.71. These results are compared against those obtained for the global uniform h-refinement strategy. Since the findings are qualitatively similar to those obtained for the verification of the discretisation of the NDE, the discussion is not repeated here for brevity; q.v. Section 4.5.1.1.

5.5.1.2 A Pincell-like Grid

The second verification case consists of a more challenging pin-cell grid defined over $V = [0,1]^2$, which comprises of a homogeneous material whose properties are given in Table 4.3. The coarsest mesh is composed of *nine* equal-volume non-degenerate bi-quadratic NURBS patches. Assuming a two-group energy approximation and the S_4 LS angular quadrature set, the manufactured solution is chosen for $k = 1, \ldots, K$:

$$\begin{cases} \varphi^{gk}(\vec{r}) = \cos\left(B_x^g x\right)\cos\left(B_y^g y\right)[1 + \vec{\Omega}^k \cdot \vec{\Omega}^k], & \vec{r} \in V; \quad (5.74a) \\ B_x^g, B_y^g \coloneqq (2g{+}1)\pi, & \text{for } g = 1,2; \quad (5.74b) \end{cases}$$

with:

$$\Phi_l^{gm}(\vec{r}) \approx \sum_{k=1}^{K} w_k Y_l^m(\vec{\Omega}^k)\cos\left(B_x^{gk} x\right)\cos\left(B_y^{gk} y\right)[1 + \vec{\Omega}^k \cdot \vec{\Omega}^k], \quad (5.75)$$

$$\approx \cos\left(B_x^{gk} x\right)\cos\left(B_y^{gk} y\right)\sum_{k=1}^{K} w_k Y_l^m(\vec{\Omega}^k)[1 + \vec{\Omega}^k \cdot \vec{\Omega}^k]; \quad (5.76)$$

where φ^{gk} denotes an odd harmonic mode as characterised by the geometric buckling, B_x^g and B_y^g. Again, the manufactured solution has been chosen to be infinitely differentiable. However, now, the

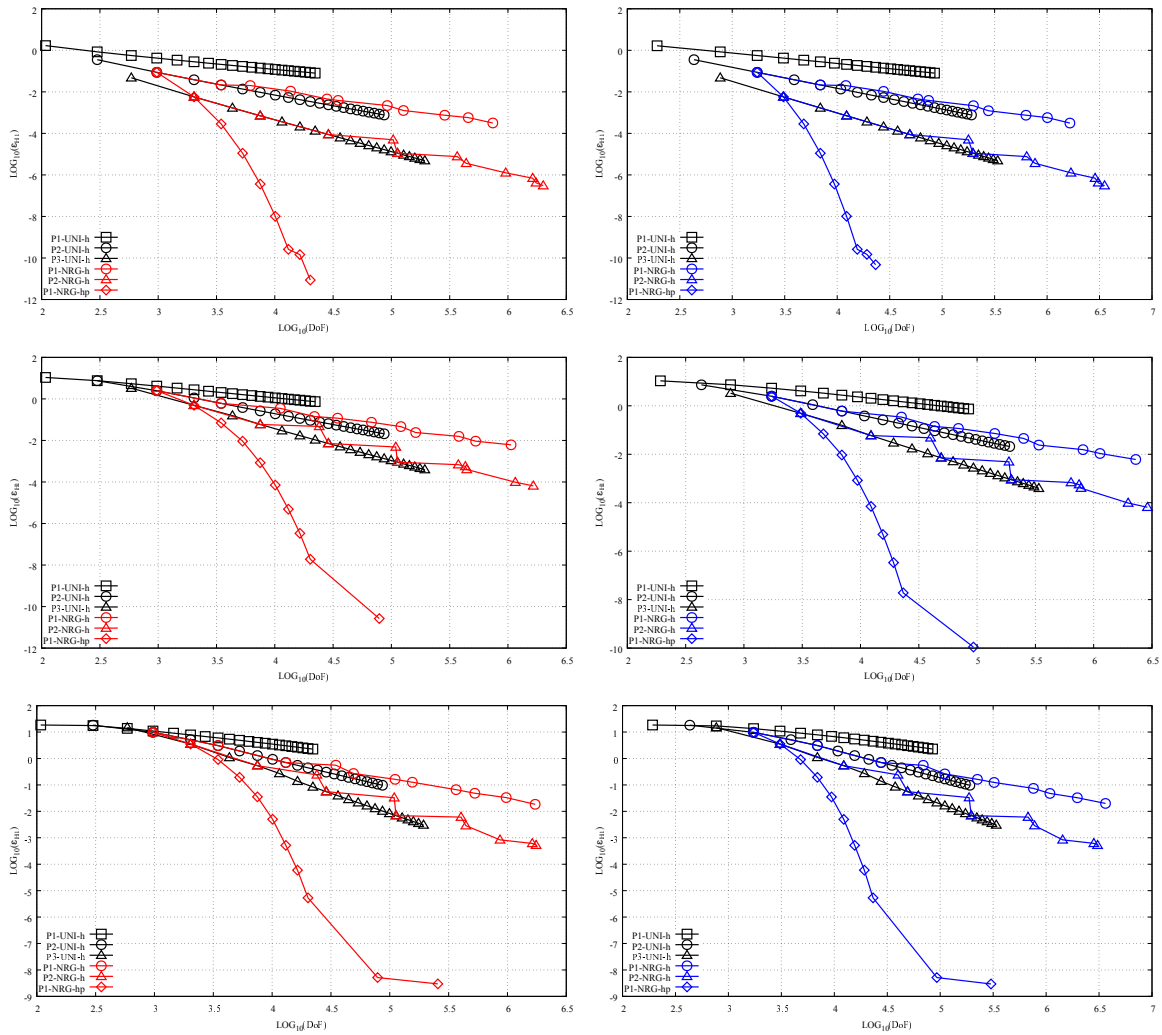


Figure 5.2: The convergence plots, ϵ_{H^1} v DoF, for the verification of the NRG-AMR-$h(p)$ strategy for the CBG-IGA (left) and the SIP-DG-IGA (right) discretisations of the 1G SAAF-S_4 equations over a 2D Cartesian grid the manufactured solution is chosen as per Equation 5.71 for $n = 0$ (top row), $n = 1$ (middle row) and $n = 2$ (bottom row). (*Vide* the web-based version for reference to colour.)

non-homogeneous Dirichlet reflective BC is naturally satisfied; and this is weakly enforced in the discretisation; q.v. Equation 5.23b.

Global Uniform Refinement For the KI-refinement of quadratic, cubic and quartic NURBS, the H^1-errors for each of the CBG-IGA and the SIP-DG-IGA spatial discretisations are presented in Figure 5.3 for the multi-group components of the manufactured solution chosen as the third and fifth harmonics of a modified standing wave equation; q.v. Equation 5.74. Qualitatively similar plots were obtained for the L_2-error; however, for brevity, these have not been presented. Under refinement, both the continuous and the discontinuous discretisation schemes converge to the analytical solution; and asymptotically, they attain their theoretical orders of accuracy for KI-h-refinement, as presented in Table 5.2. Again, however, for KI-1-refinement, one remarks a degree of super-convergence, which may be explained by the higher-order continuity in the basis; q.v. Section 4.5.1.2.

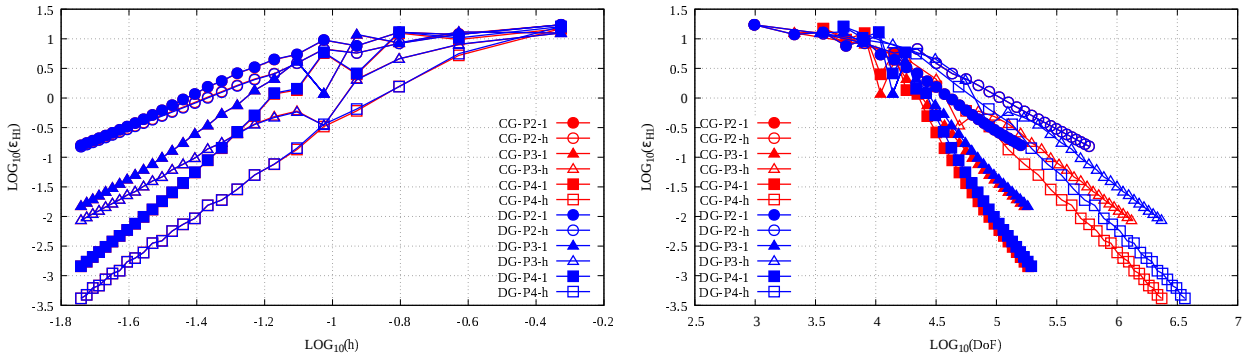


Figure 5.3: The convergence plots, ϵ_{H^1} v h (left) and ϵ_{H^1} v DoF (right), for the verification of the global uniform refinement strategy for the CBG-IGA and the SIP-DG-IGA 2G SAAF-S_4 equations over a 2D pin-cell grid; the manufactured solutions are chosen as per Equation 5.74 for $g = 1, 2$. (*Vide* the web-based version for reference to colour.)

p	CBG-IGA				SIP-DG-IGA			
	KI-1		KI-h		KI-1		KI-h	
	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}
2	2.24	3.54	2.04	3.03	2.24	3.54	2.03	3.01
3	3.30	4.43	3.06	4.01	3.31	4.43	3.02	4.00
4	4.46	5.41	4.04	5.07	4.48	5.36	4.04	5.04

Table 5.2: The asymptotic orders of accuracy, Equations 4.103, for the verification of the global uniform refinement strategy for the CBG-IGA and the SIP-DG-IGA discretisations of the 2G SAAF-S_4 equations over a 2D pin-cell grid; the convergence rates are evaluated as the slope of the last few points from each dataset.

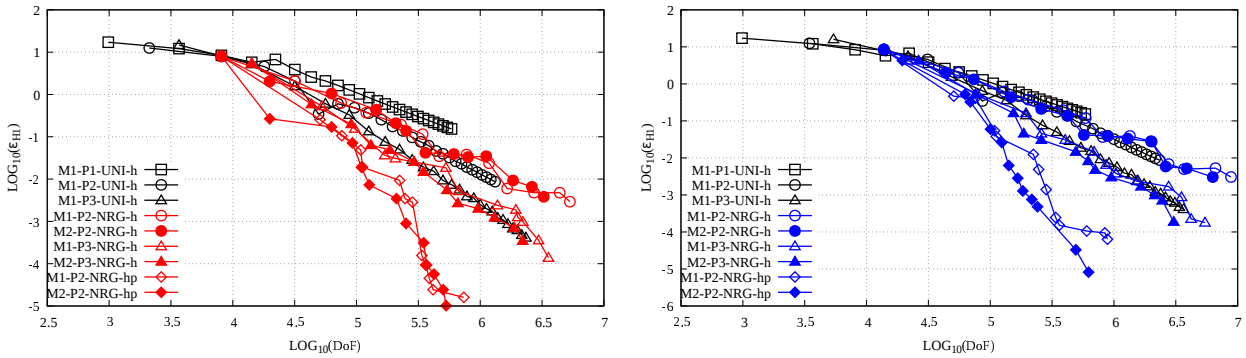


Figure 5.4: The convergence plots, ϵ_{H^1} v DoF, for the verification of the NRG-AMR-$h(p)$ strategy for both the CBG-IGA (left) and the SIP-DG-IGA (right) 2G SAAF-S_4 equations over a 2D pin-cell grid; the manufactured solutions are chosen as per Equation 5.74 for $g = 1, 2$. (*Vide* the web-based version for reference to colour.)

Local Adaptive Refinement The results of the NRG-AMR-$h(p)$ strategy for chasing the global H^1-error in each of the CBG-IGA and the SIP-DG-IGA discretisations are presented in Figure 5.4 for the multi-group components of the manufactured solution chosen as the third and fifth harmonics of the standing wave equation; q.v. Equation 5.74. The effectivity, θ_{H^1}, of the interpolation-based error measures, as per Equation 3.24, is measured and displayed in Figure 5.5. The capacity for energy-dependent meshes is also tested. These results are qualitatively similar to those obtained for the verification of the discretisation of the NDE and thus, for brevity, the discussion is not repeated here; q.v. Section 4.5.1.2. Nevertheless, one emphasises that any savings made in the refinement

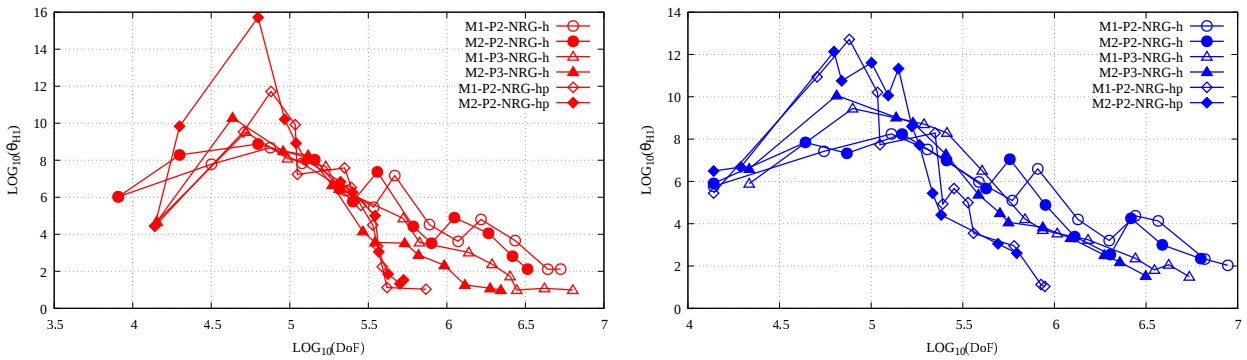


Figure 5.5: The plots of the effectivity index, θ_{H^1} v DoF, for the verification of the NRG-AMR-$h(p)$ strategy for the CBG-IGA (left) and the SIP-DG-IGA (right) 2G SAAF-S_4 equations over a 2D pin-cell grid; the manufactured solutions are chosen as per Equation 5.74 for $g = 1, 2$. (*Vide* the web-based version for reference to colour.)

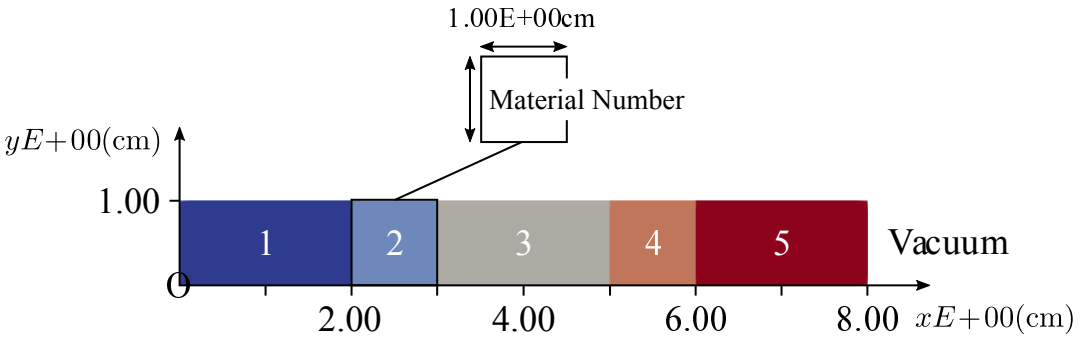


Figure 5.6: The geometry for Reed's 1D 1G problem, which has been adapted from the original problem and presented in 2D. The constituent materials are denoted by the numbers in parentheses; viz. fuel (1), can (2), void (3) and moderator with a source (4) and without (5). (*Vide* the web-based version for reference to colour.)

of a given mesh for a given energy group actually correspond to K-number of savings, since each group-ordinate-wise system of equations is resolved over that same mesh.

5.5.2 Reed's 1D Problem

Material	g	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to1}$ (cm^{-1})	q_x^g (cm^{-3}.s^{-1})
(1) Fuel	1	5.000E+01	0.000E+00	5.000E+01
(2) Can	1	5.000E+00	0.000E+00	0.000E+00
(3) Void	1	0.000E+00	0.000E+00	0.000E+00
(4) Moderator	1	1.000E+00	9.000E−01	1.000E+00
(5) Moderator	1	1.000E+00	9.000E−01	0.000E+00

Table 5.3: The nuclear data for Reed's 1D 1G problem [234].

Reed's 1D problem for the *one-speed* NTE-S_8 is posed over a heterogeneous slab that presents *four* regions of physically distinct materials, including a fuel, a can, a void and a moderator; v. Figure 5.6. The corresponding macroscopic nuclear cross-section data are given in Table 4.5 [234]. The strengths of *two* different extraneous sources of neutrons are assigned between 1.00 cm $< x <$ 1.00 cm and 5.00 cm $< x <$ 6.00 cm. One prescribes reflective and vacuum BC at $x = 0.00$ cm and $x = 8.0$ cm,

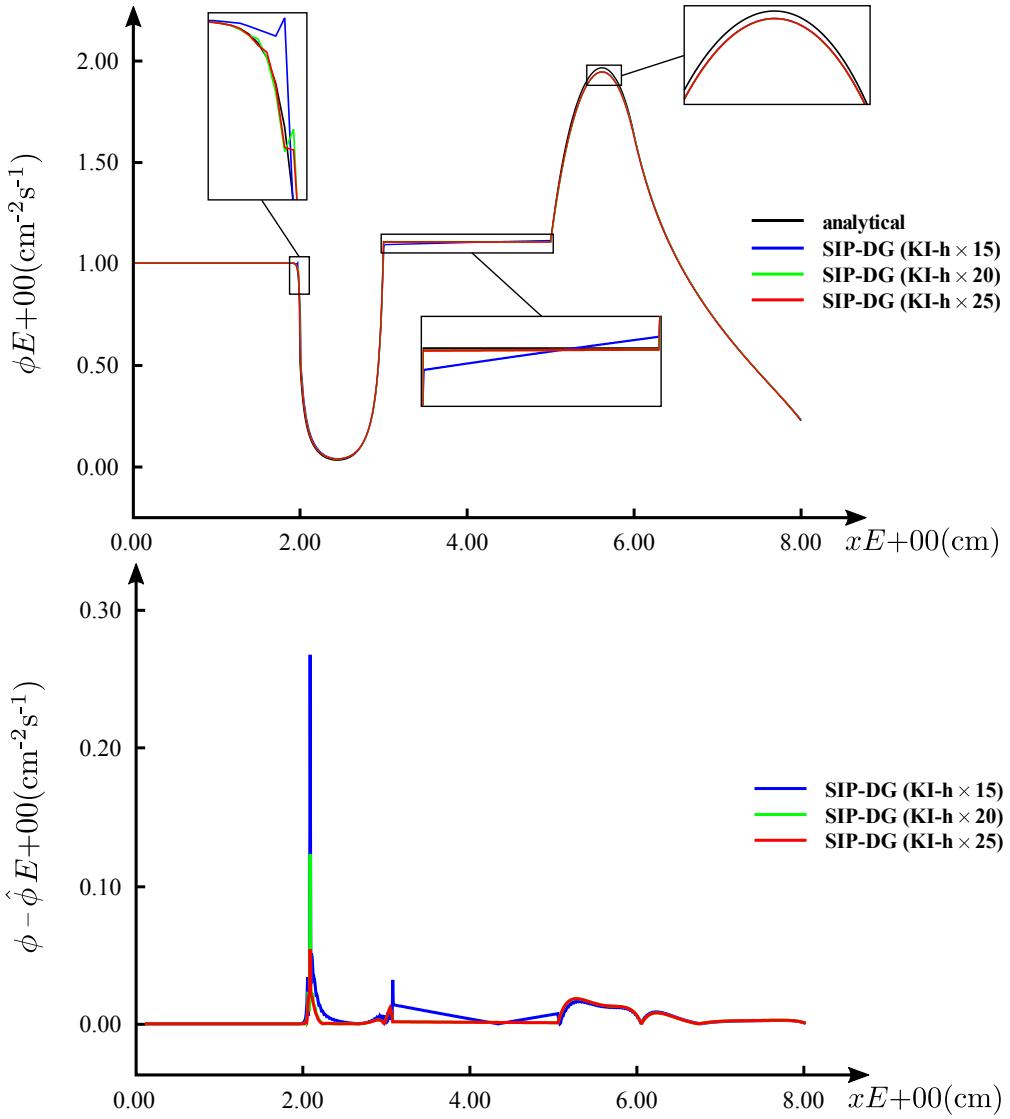


Figure 5.7: The numerical solutions (top) to the SIP-DG-IGA 1G SAAF-S_8 equations and the corresponding discretisation errors (bottom), $\phi - \hat{\phi}$, for different levels of global uniform h-refinement for Reed's 1D problem; the unresolved boundary layer, $x = 2.0$ cm, the void region, $3.0 < x < 5.0$ cm, and the peak, $x \approx 6.0$ cm, are magnified. (*Vide* the web-based version for reference to colour.)

respectively; and since this problem has been adapted into 2D, reflective BC are also prescribed at $y = 0.00$ cm and $y = 1.00$ cm. The coarsest Cartesian mesh is constructed from *eight* bilinear NURBS patches. As was determined by Warsa with the help of MAPLE, a symbolic algebra program, there does exist an analytical solution to the mono-group component of scalar neutron flux for Reed's 1D problem; v. Figure 5.7 [184].

Whilst the general focus of the current research is that of developing AMR-$h(p)$ schemes, this problem was selected to investigate further the convergence properties of the SIP-DG-IGA SAAF-S_N equations; more specifically, to determine the advantages, if any, for allowing discontinuities in the solution over a highly heterogeneous configuration. As such, only global uniform h-refinement is performed with additional knots inserted solely in the ξ/x-direction.

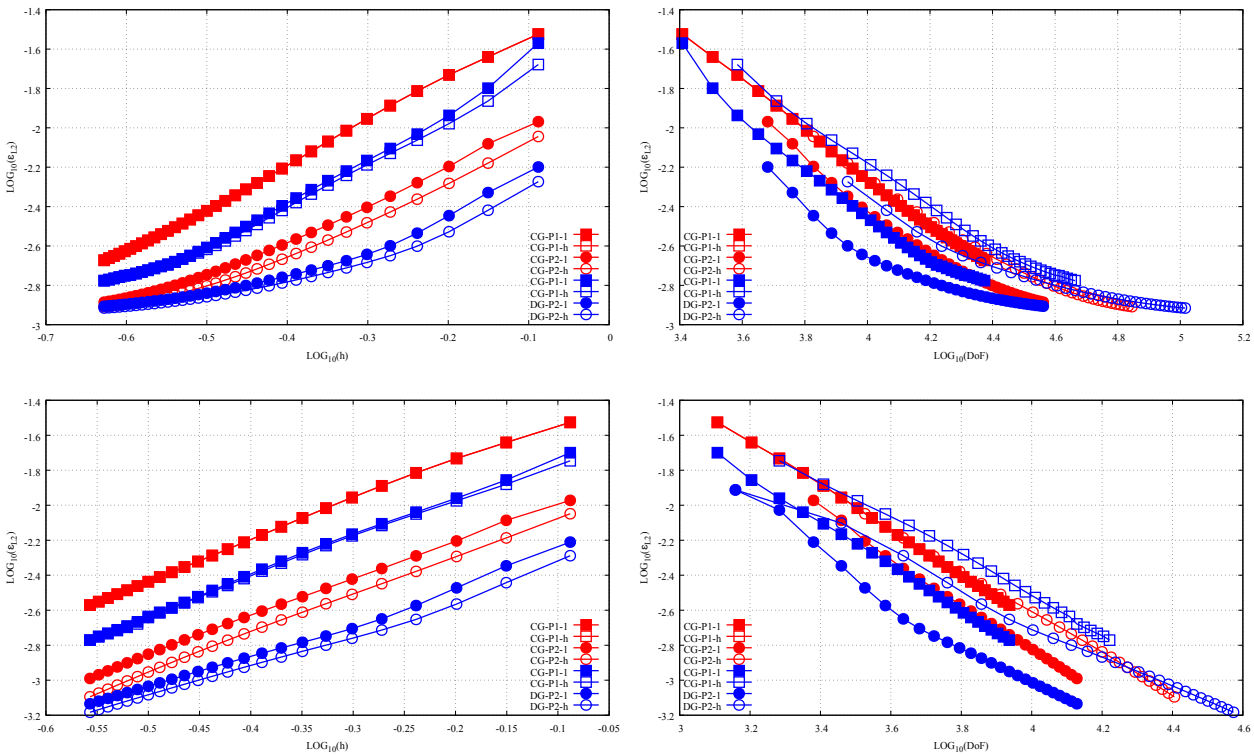


Figure 5.8: The convergence plots, ϵ_{L_2} v h (left) and ϵ_{L_2} v DoF (right), for the global uniform refinement strategy for the CBG-IGA and the SIP-DG-IGA 1G SAAF-S_8 equations for Reed's 1D problem; over the entire domain (top), $0 \leq x \leq 8$ cm, and over the can-and-void region (bottom), $2 \leq x \leq 5$ cm. (*Vide* the web-based version for reference to colour.)

To model the void region, one employs a low-density macroscopic cross-section of total collision, $\Sigma_{t,3} \approx \tau = 1.00\text{E}-06$ cm^{-1}; and to retain the SPD properties of the discrete linear systems, no other attempt is made to incorporate any sophisticated void treatment [120,125]. One recalls that the penalty parameters for the SIP-DG-IGA discretisation vary as $\mu_f \propto 1/\Sigma_t$; q.v. Section 5.4. Therefore, to avoid any over-penalisation of coarser meshes over the void region, one introduces the following modification $1/\Sigma_t^* = \min\{1/\tau, 1/h\}$, such that $1/\Sigma_t^* \to \infty$ as $\tau, h \to 0$.

For various degrees of h-refinement, the scalar neutron flux profiles for the SIP-DG-IGA SAAF-S_8 equations are presented in Figure 5.7 along with their respective discretisation error distributions. The corresponding convergence plots are presented in Figure 5.8, including for KI-1-refinement; and these are compared against similar CBG-IGA discretisations. Clearly, it is advantageous to allow for discontinuities in the solution about the material interfaces; and especially so for the KI-1-refinement strategy for coarser meshes over the void region where the SIP-DG-IGA discretisation becomes competitive on a per DoF basis.

Since the analytical solution presents discontinuous first-order derivatives, one does not expect to attain full rates of convergence. In fact, the errors attributed to the unresolved boundary layer at $x = 2.0$ cm are vanishing under refinement. However, modelling the void region with low-density cross-sections and no scatter implies that there must be some absorption of neutrons; and thus, owing

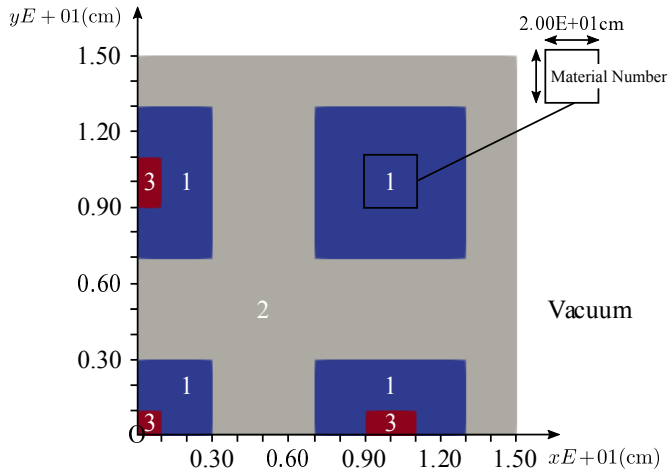


Figure 5.9: The geometry for Ackroyd's 2D 1G small reactor problem [105]. The constituent materials are denoted by the numbers in parentheses; viz. fuel (1), moderator (2) and control-rod (3). (*Vide* the web-based version for reference to colour.)

to a lack of neutrons travelling in the positive x-direction, the result is an underestimation of the peak at $x \approx 6.0$ cm; where one notes that the corresponding discretisation error appears to be constant under refinement. As such, the numerical solutions do not converge to the analytical one.

5.5.3 Ackroyd's 2D Small Reactor Problem

Material	g	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to 1}$ (cm^{-1})	q_x^g (cm^{-3}.s^{-1})
(1) Fuel	1	1.51461E−01	8.93030E−02	1.48083E−02
(2) Moderator	1	1.48934E−01	9.93430E−02	0.00000E+00
(3) Control-Rod	1	1.50000E+00	0.00000E+00	0.00000E+00

Table 5.4: The nuclear data for Ackroyd's 2D 1G small reactor problem [105].

Embedded in a moderator and enclosed by a bare surface, Ackroyd's 2D small reactor problem consists of *four* fuel clusters with *three* control-rods inserted; as represented by Materials 2, 1 and 3, respectively; v. Figure 5.9 [105]. The corresponding macroscopic nuclear cross-section data are provided in Table 5.4. Modelling for only a single thermal energy group, one assumes a uniformly dispersed and isotropic slowing-down source across the fuelled regions. As per the *one-quarter* symmetry in the problem, reflective BC are prescribed at $x = y = 0.00$ cm. The coarsest Cartesian mesh is constructed from 225 bilinear NURBS patches. $\mathrm{P_{N\text{-}1}}$ reference solutions were obtained with the combined GEM-EVENT codes, an independent FEM mesh generator and multi-group even-parity neutron transport solver, respectively, over a mesh of 104,000 bi-quadratic elements [216]. The spatial distribution of the mono-group component of both the scalar neutron flux and the scalar neutron importance for the S_2, S_4 and S_8 approximations with GC angular quadrature sets are presented in Figure 5.10, for the QoI chosen as the absorption in Material 1 in the top right-hand corner of Figure 5.9.

The AMR strategies presented in Section 3.6 were devised to minimise iteratively those errors due to

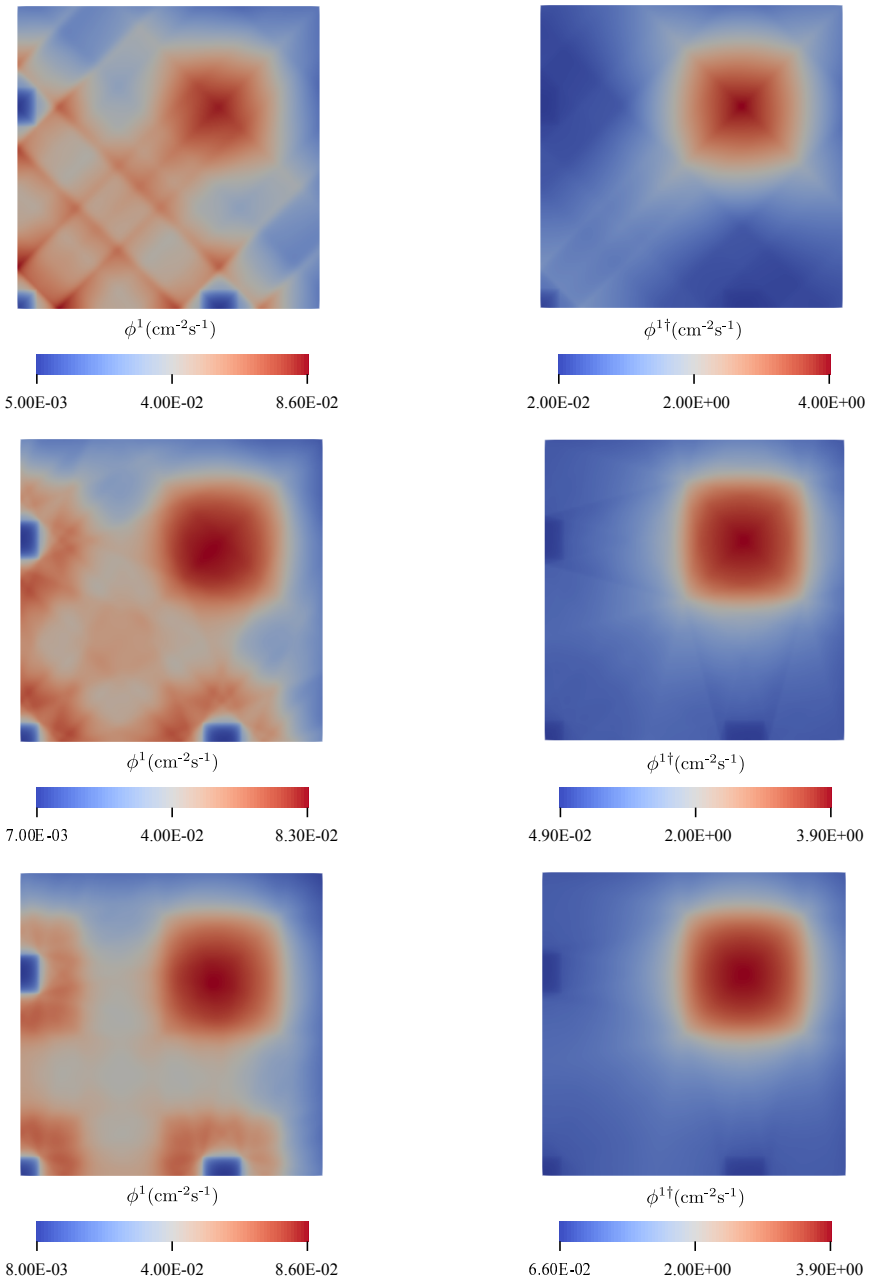
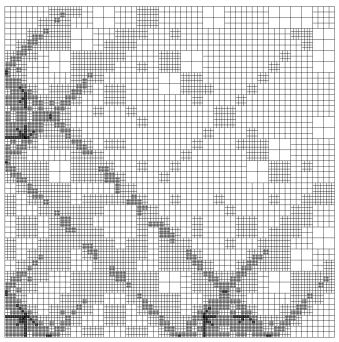


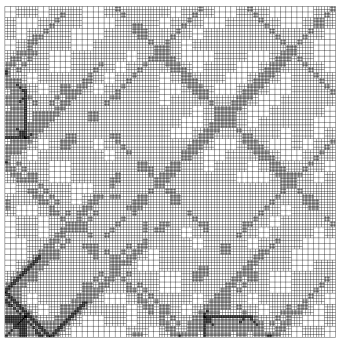
Figure 5.10: The spatial distribution of the mono-group component of the scalar neutron flux (left) and the scalar neutron importance (right) for the 1G SAAF-S_N equations for Ackroyd's 2D small reactor problem; for S_2 (top), S_4 (middle) and S_8 (bottom) GC angular quadrature sets; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 5.9. (*Vide* the web-based version for reference to colour.)

the spatial discretisation. However, these can be strongly coupled to the angular discretisation; and especially so for dominating ray effects [231]. The meshes generated after the 9th AMR iteration for each of the CBG-IGA AMR strategies for an increasingly large order of the GC angular quadrature set, $N = \{2, 4, 8\}$, are presented in Figure 5.11 and Figure 5.12.

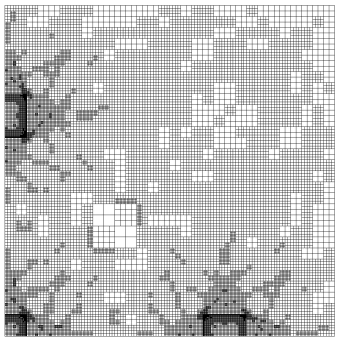
Whilst much less apparent for the SAAF-S_8 equations, the ray effects are clearly visible in the refinement of the mesh for low orders of N; unsurprisingly, this suggests that accurate solutions demand good resolution along these rays. Therefore, the refinement focusses along the angular directions as chosen as the discrete ordinates from the angular quadrature set.



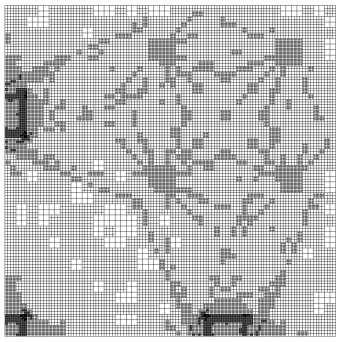
(a) $\mathcal{T}^{(9)}$ for S_2: $89,461$ DoF



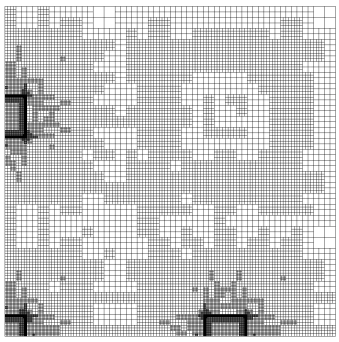
(b) $\mathcal{T}^{(9)}$ for S_2: $106,681$ DoF



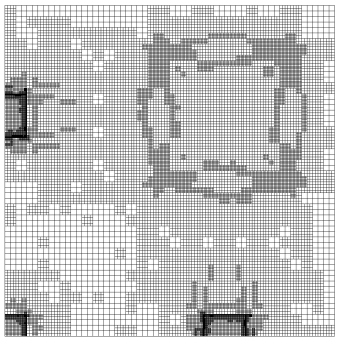
(c) $\mathcal{T}^{(9)}$ for S_4: $106,817$ DoF



(d) $\mathcal{T}^{(9)}$ for S_4: $115,533$ DoF



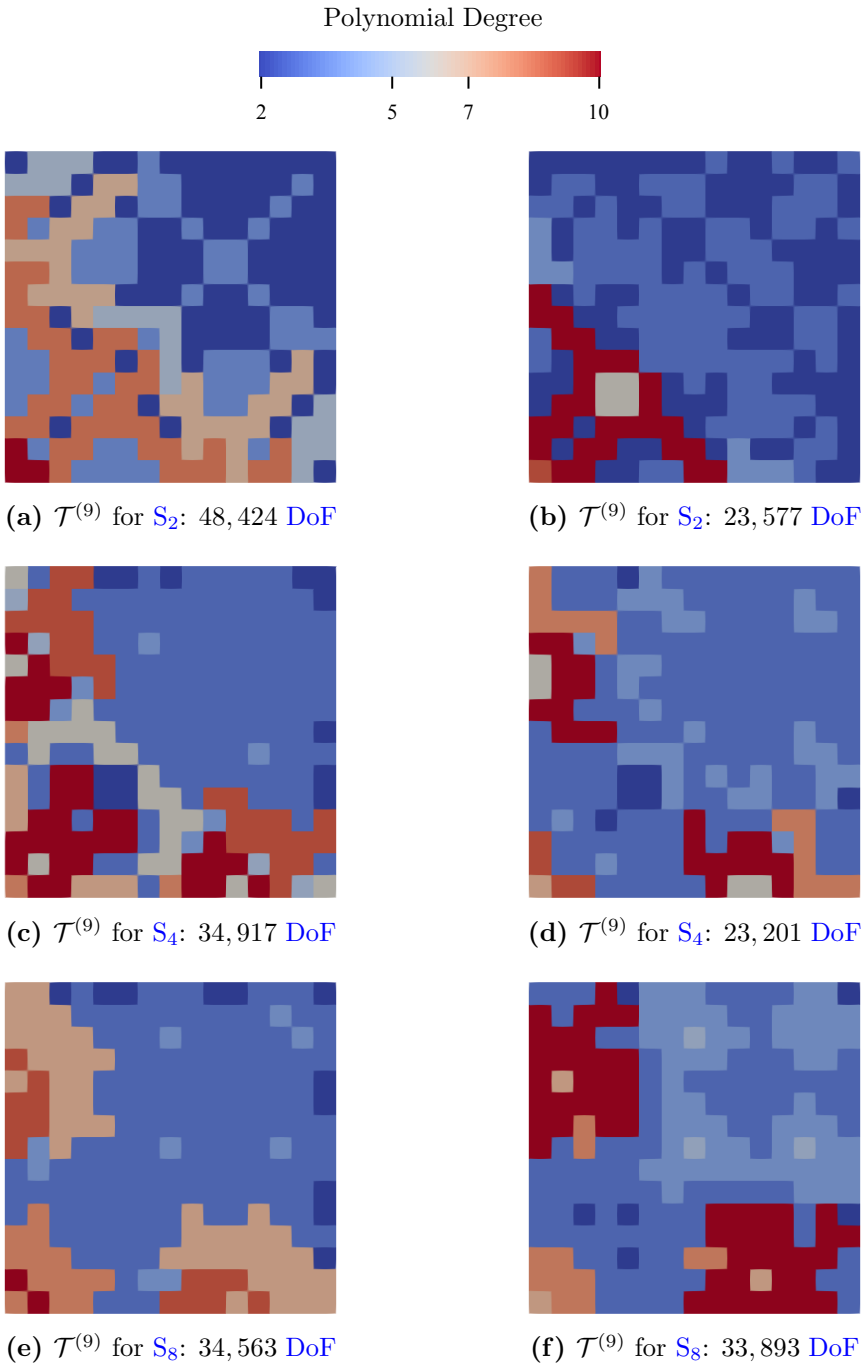
(e) $\mathcal{T}^{(9)}$ for S_8: $92,337$ DoF



(f) $\mathcal{T}^{(9)}$ for S_8: $102,097$ DoF

Figure 5.11: The energy-independent meshes generated at the 9th AMR iteration for the NRG-AMR-h (left) and the DWR-AMR-h (right) strategies for the CBG-IGA 1G SAAF-S_2 (top), SAAF-S_4 (middle) and SAAF-S_8 equations for Ackroyd's 2D small reactor problem; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 5.9. (*Vide* the web-based version for reference to colour.)

For the QoI, the NRG-AMR-h strategy tends to be the least performant, since the refinement concentrates about the sharp gradients in the solution about the control-rods. The DWR-AMR-h strategy is slightly more competitive with the uniform h-refinement one, since the refinement also focusses about the fuelled region of interest. The AMR-hp strategies deliver the best accuracy per DoF; and specifically, the DWR-AMR-hp is the most performant; v. Figure 5.13. However, the individual group-ordinate-wise banded matrices do become increasingly ill conditioned for increasingly large polynomial degrees of the NURBS basis; and with each AMR iteration, increasingly more computational effort is required to converge the within-group equations since the algebraic multi-grid preconditioner appears to become increasingly less effective.



(a) $\mathcal{T}^{(9)}$ for S_2: $48,424$ DoF **(b)** $\mathcal{T}^{(9)}$ for S_2: $23,577$ DoF

(c) $\mathcal{T}^{(9)}$ for S_4: $34,917$ DoF **(d)** $\mathcal{T}^{(9)}$ for S_4: $23,201$ DoF

(e) $\mathcal{T}^{(9)}$ for S_8: $34,563$ DoF **(f)** $\mathcal{T}^{(9)}$ for S_8: $33,893$ DoF

Figure 5.12: The energy-independent meshes generated at the 9th AMR iteration for the NRG-AMR-*hp* (left) and the DWR-AMR-*hp* (right) strategies for the CBG-IGA 1G SAAF-S_2 (top), SAAF-S_4 (middle) and SAAF-S_8 equations for Ackroyd's 2D small reactor problem; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 5.9. (*Vide* the web-based version for reference to colour.)

Both the CBG-IGA and the SIP-DG-IGA discretisations do generate qualitatively similar meshes for similar DoF; yet, the errors introduced by the constraint equations in the former appear to be more significant than those introduced in the latter. Moreover, one recalls that DG discretisations generally offer more flexibility in terms of local mesh refinement.

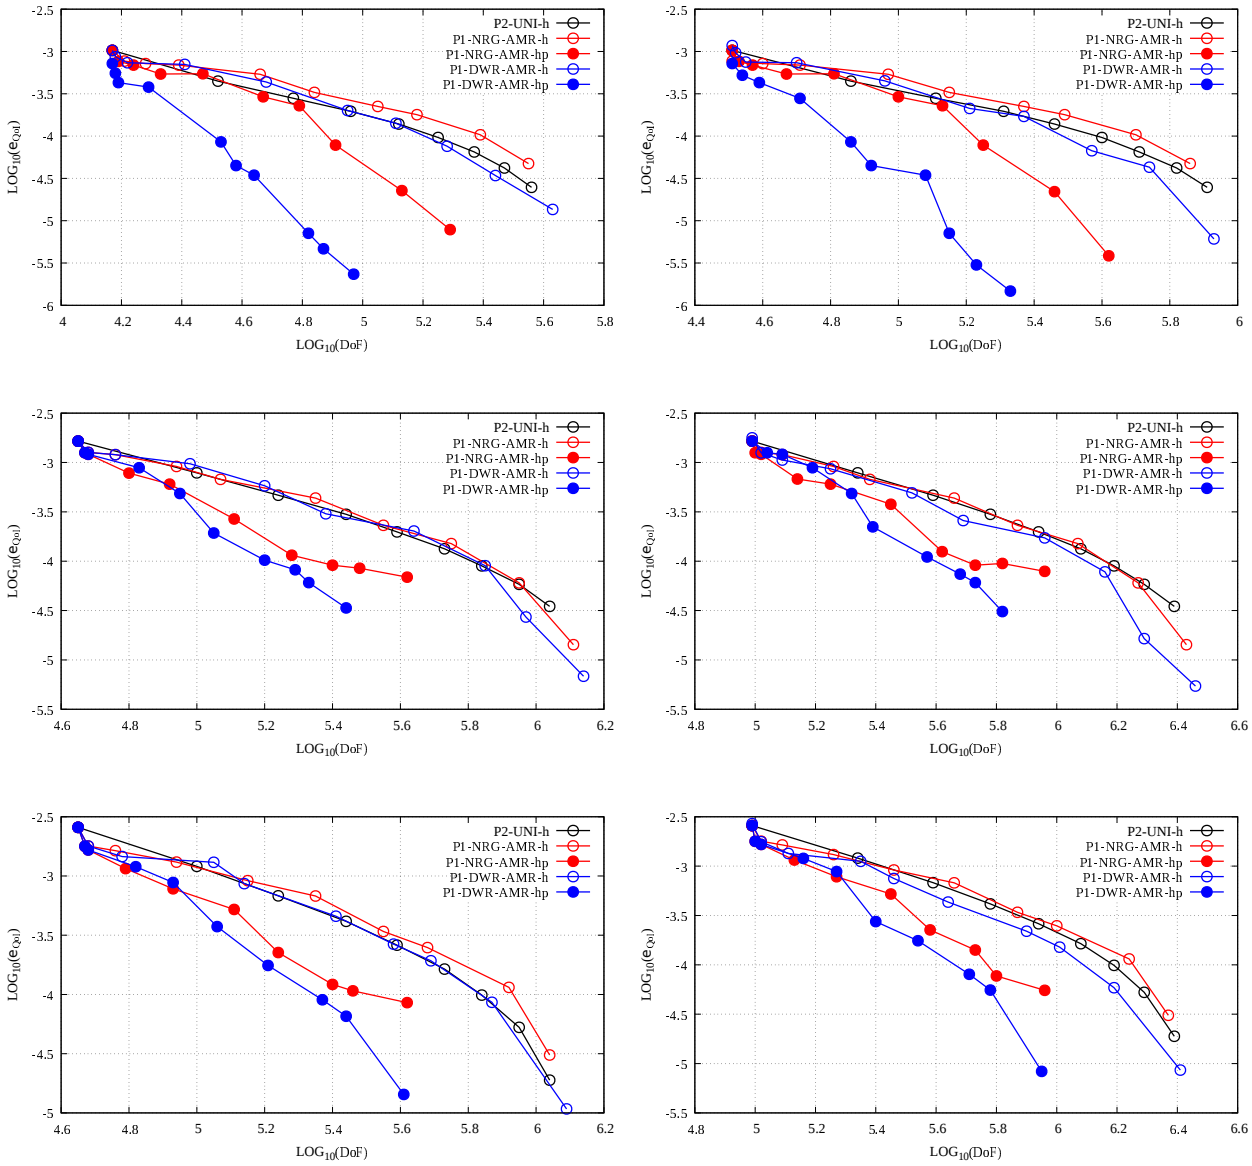


Figure 5.13: The convergence plots, ϵ_{QoI} v DoF, for the global uniform refinement (black) and the NRG-AMR (red) and the DWR-AMR (blue) strategies for the CBG-IGA (left) and the SIP-DG-IGA (right) 1G SAAF-S_2 (top), SAAF-S_4 (middle) and SAAF-S_8 equations for Ackroyd's 2D small reactor problem; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 5.9. (*Vide* the web-based version for reference to colour.)

5.5.4 Hong and Cho's 2D BWR Problem

Hong and Cho's problem is a four-by-four BWR lattice with *two* adjacent burnable Gadolinium poison pins [360]. This problem presents *four* physically distinct regions, including 14 homogenised regular fuel pins (3*wt.*% UO_2) and *two* poisoned ones (3*wt.*% Gd_2O_3), each pin encased in a cladding of Zr-2 and moderated by water, as represented by Materials 1, 2, 3 and 4, respectively; v. Figure 5.14. For the two-group condensation of the 172-group WLUP nuclear data library at the delimitation of 4 eV, this problem was first modelled by DRAGON, a neutronics lattice code [361]. These few-group macroscopic cross-section data are reproduced in Table 5.5 [362]. To model an infinite lattice, reflective BC are specified at every external physical boundary. The coarsest mesh is constructed from 208 bi-

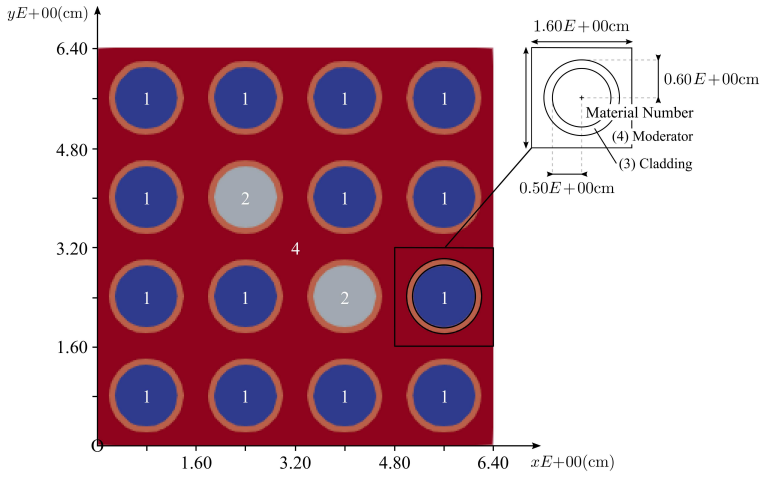


Figure 5.14: The geometry for Hong and Cho's 2D 2G BWR problem with burnable Gadolinium poison pins [360]. The constituent materials are denoted by the numbers in parentheses; viz. Uranium oxide (1), Uranium & Gadolinium Oxides (2), cladding (3) and moderator (4). (*Vide* the web-based version for reference to colour.)

Material	g	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to1}$ (cm^{-1})	$\Sigma_{s0}^{g\to2}$ (cm^{-1})	Σ_f^g (cm^{-1})	$\nu^g\Sigma_f^g$ (cm^{-1})
(1) UO_2	1	3.62022E−01	3.33748E−01	6.64881E−04	7.22964E−03	1.86278E−02
	2	5.72155E−01	0.00000E+00	3.80898E−01	1.41126E−01	3.44137E−01
(2) UO_2 & Gd_2O_2	1	3.71785E−01	3.38096E−01	6.92807E−04	6.97904E−03	1.79336E−02
	2	1.75000E+00	0.00000E+00	3.83204E−01	6.47524E−02	1.57929E−01
(3) Cladding	1	2.74144E−01	2.72377E−01	1.90838E−04	0.00000E+00	0.00000E+00
	2	2.80890E−01	0.00000E+00	2.77230E−01	0.00000E+00	0.00000E+00
(4) Moderator	1	6.40711E−01	6.07382E−01	3.31316E−02	0.00000E+00	0.00000E+00
	2	1.69131E+00	0.00000E+00	1.68428E+00	0.00000E+00	0.00000E+00

Table 5.5: The nuclear data for Hong and Cho's 2D 2G BWR problem with burnable gadolinium poison pins [360]. The fission spectrum is $\chi = [1.00, 0.00]^T$.

quadratic NURBS patches. A reference solution is obtained from the DRAGON code, as provided by Yang and Satvat; viz. $k_{eff} = 0.986561$ [362]. The spatial distribution of the multi-group components of both the scalar neutron flux and the scalar neutron importance for the SAAF-S_6 equations with the LS angular quadrature set are presented in Figure 5.15, for the QoI chosen as the k_{eff}.

The meshes generated after the 9th AMR iteration for the SIP-DG-IGA AMR-h and AMR-hp strategies are presented in Figure 5.16 and Figure 5.17, respectively. For the use of energy-independent meshes, the AMR-$h(p)$ appear to be competitive with each other, respectively; although, generally, the NRG-AMR strategies demand significantly more DoF than the DWR-AMR ones; viz. at each AMR iteration the meshes generated by the former are described by *one-and-a-half* to *two* times as many DoF as those by the latter. However, these additional DoF do not appear to be allocated in a manner that more accurately approximates the QoI; despite the fact that the k_{eff} is a global quantity. This is obvious for the use of energy-dependent meshes, where the NRG-AMR strategies generate unnecessarily refined meshes for the relatively smooth profile of ϕ^1. Conversely, the more pronounced variations in the spatial distribution of $\phi^{2\dagger}$ in Figure 5.15 suggest that it would be more important to resolve those

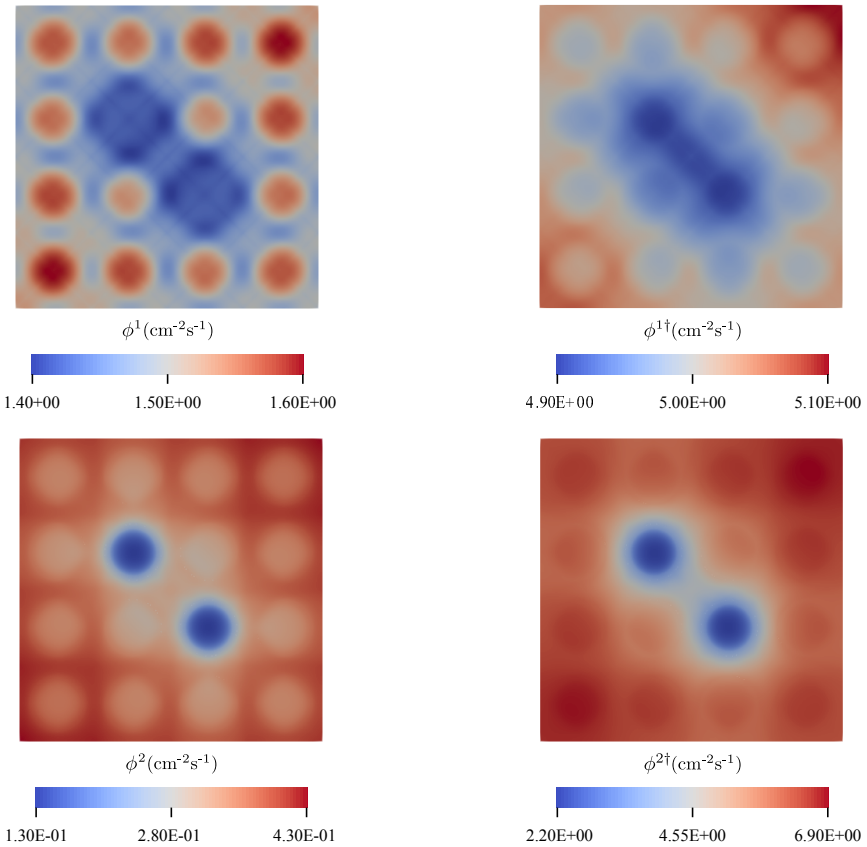


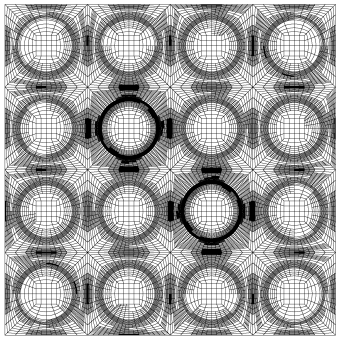
Figure 5.15: The spatial distribution of the multi-group components of the scalar neutron flux (left) and the scalar neutron importance (right) for the 2G SAAF-S_6 equations for Hong and Cho's 2D BWR problem with burnable Gadolinium poison pins; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

sharp gradients about the poison pins in $\mathcal{T}^2$; this is reflected in those meshes generated by the DWR-AMR strategies, where relatively few resources are allocated to $\mathcal{T}^1$. Accordingly, the DWR-AMR-h strategy is able to achieve at least an order-of-magnitude more accuracy than a similar NRG-AMR-h strategy and the uniform h-refinement strategy for the same DoF.

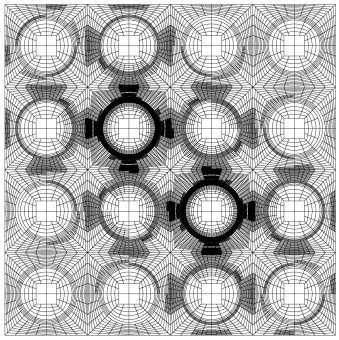
Again, from the convergence plots in Figure 5.18, the DWR-AMR strategies consistently reduce the error in the QoI at each subsequent AMR iteration; and unsurprisingly, the best rates of convergence were obtained for the DWR-AMR-hp strategies with energy-dependent meshes.

5.6 A Discussion

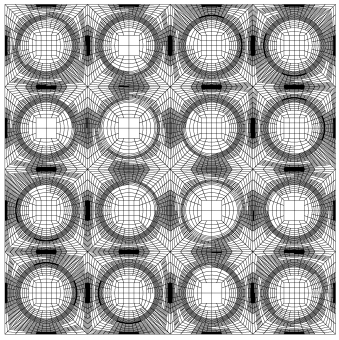
This chapter saw the application of automised NRG-AMR-$h(p)$ and DWR-AMR-$h(p)$ strategies to separate CBG-IGA and SIP-DG-IGA spatial discretisations of the multi-group SAAF-S_N equations. These discretisation methods are similar to those presented in Chapter 4. Although not investigated in great depth, it was observed that one could gain in accuracy per DoF by allowing for a discontinuous solution field over configurations that present voids and a high-degree of heterogeneity; and especially so in the coarse-mesh limit.



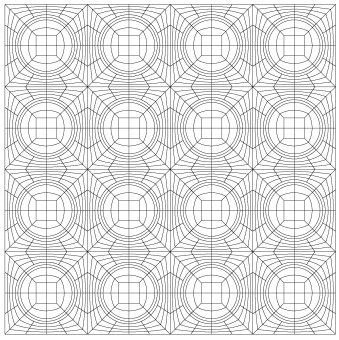
(a) $\mathcal{T}^{(9)}$: $498,240$ DoF $(\times 2)$



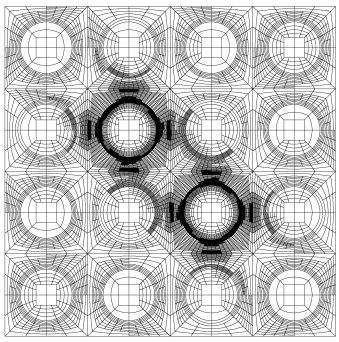
(b) $\mathcal{T}^{(9)}$: $357,248$ DoF $(\times 2)$



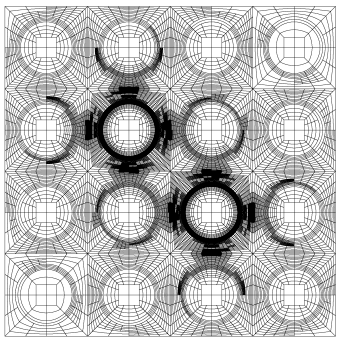
(c) $\mathcal{T}^{1(9)}$: $423,232$ DoF



(d) $\mathcal{T}^{1(9)}$: $25,600$ DoF



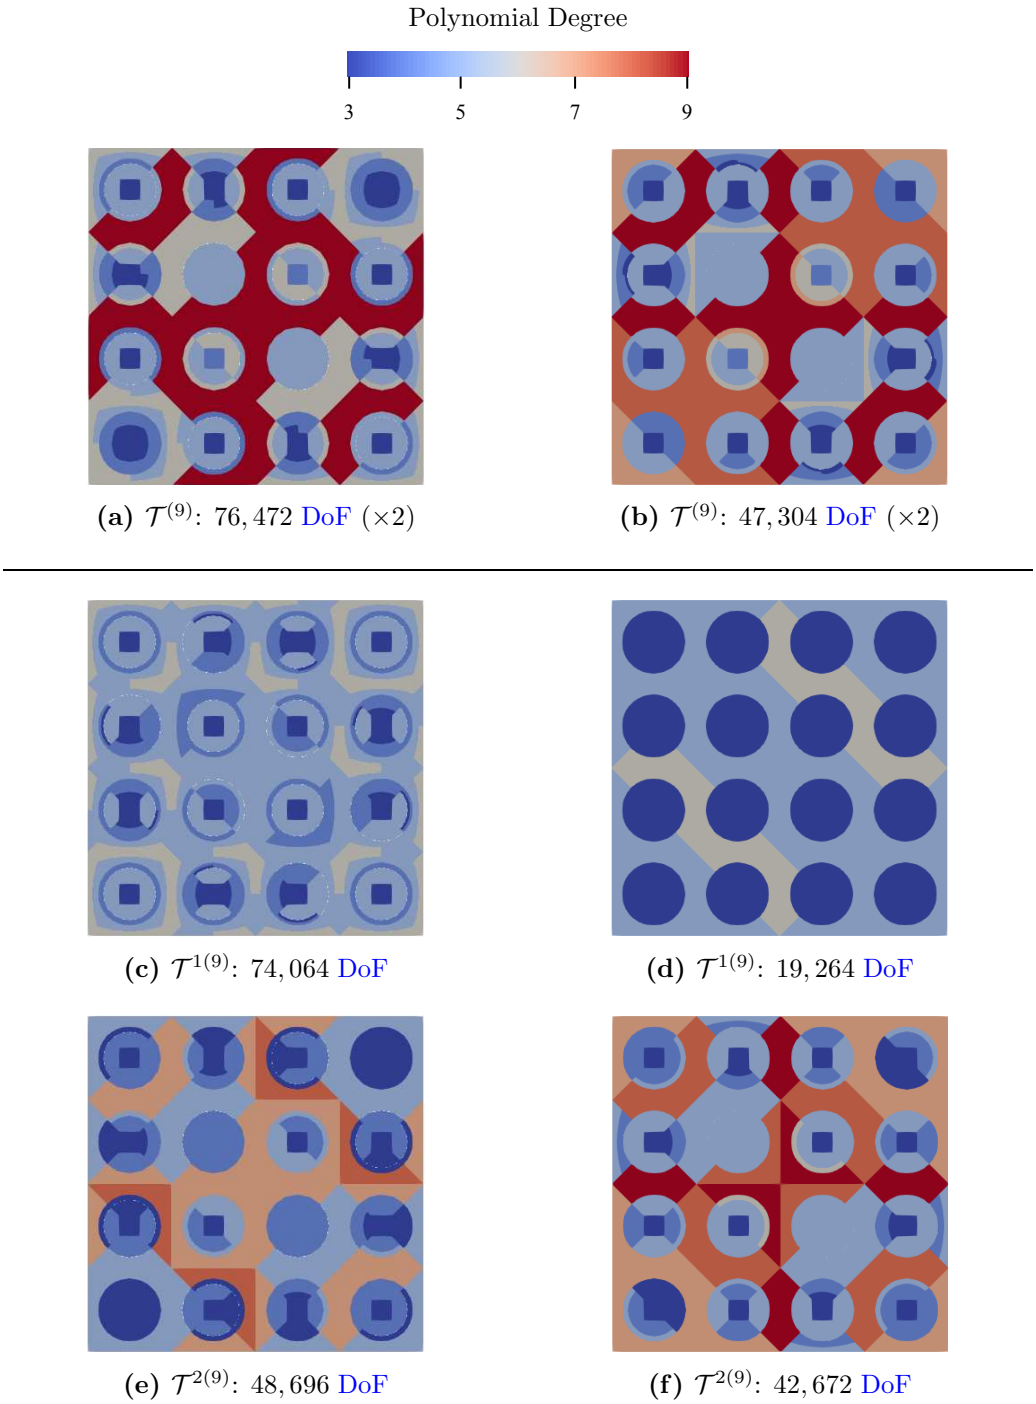
(e) $\mathcal{T}^{2(9)}$: $262,400$ DoF



(f) $\mathcal{T}^{2(9)}$: $421,632$ DoF

Figure 5.16: The energy-independent meshes (top) and the energy-dependent meshes (middle and bottom) generated at the 9th AMR iteration for the NRG-AMR-h (left) and the DWR-AMR-h (right) strategies for the SIP-DG-IGA 2G SAAF-S_6 equations for Hong and Cho's 2D BWR problem with burnable Gadolinium poison pins; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

Again, allowing for the imposition of a certain degree of regularity to avoid any multiply constraints, the AMR-$h(p)$ algorithms for both of the CBG-IGA and the SIP-DG-IGA discretisations did generate qualitatively similar meshes for similar DoF. In each case, mesh refinement tended to be dominated by the ray effects for low-order angular quadrature sets. One expects this behaviour, since any under-resolution of these features would only serve to mask the true behaviour of the solution to the S_N equations; cf. Figure 2.2. Importantly, the spatial discretisation schemes developed in this chapter would be applicable to any effective ray-effect mitigation techniques that maintain the same underlying structure of the S_N method.



(a) $\mathcal{T}^{(9)}$: 76,472 DoF (×2) **(b)** $\mathcal{T}^{(9)}$: 47,304 DoF (×2)

(c) $\mathcal{T}^{1(9)}$: 74,064 DoF **(d)** $\mathcal{T}^{1(9)}$: 19,264 DoF

(e) $\mathcal{T}^{2(9)}$: 48,696 DoF **(f)** $\mathcal{T}^{2(9)}$: 42,672 DoF

Figure 5.17: The energy-independent meshes (top) and the energy-dependent meshes (middle and bottom) generated at the 9th AMR iteration for the NRG-AMR-*hp* (left) and the DWR-AMR-*hp* (right) strategies for the SIP-DG-IGA 2G SAAF-S_6 equations for Hong and Cho's 2D BWR problem with burnable Gadolinium poison pins; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

For localised QoI, the DWR-AMR strategies consistently and reliably reduced the error in the QoI at each subsequent AMR iteration; whereas the NRG-AMR algorithms would do so only sporadically. Even for global quantities, such as the k_{eff}, the latter tended to become less performant than the global uniform strategy for h-refinement; and especially so for the use of energy-dependent meshes, since certain multi-group components were left under-resolved.

Whilst the (h)p-refinement strategies did exhibit exponential rates of convergence, it was observed,

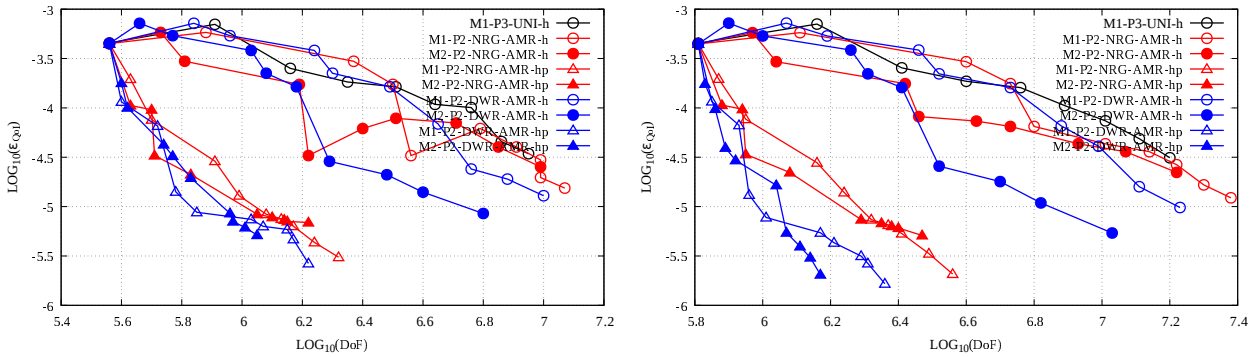


Figure 5.18: The convergence plots, ϵ_{QoI} v DoF, for the uniform refinement (black), the NRG-AMR (red) and the DWR-AMR (blue) strategies for the CBG-IGA (left) and the SIP-DG-IGA (right) 2G SAAF-S_6 equations for Hong and Cho's 2D BWR problem with burnable Gadolinium poison pins; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

even for the local refinement of Cartesian grids, that the resulting banded matrices became increasingly ill conditioned for increasingly large polynomial degrees of the NURBS basis; and although it was not presented, this did demand increasingly more computational time and effort to invert the group-ordinate-wise banded matrices with PETSc's parallel implementation of the CG method for sparse matrices; and especially so, since it would appear that the algebraic multi-grid preconditioner became increasingly less effective.

Chapter 6

Hyperbolic Systems

This chapter presents the application of automised NRG-AMR-$h(p)$ and DWR-AMR-$h(p)$ strategies to a sweep-based DG-IGA spatial discretisation of the multi-group FONT-S_N equations.

6.1 An Introduction

Galerkin-based discretisations of second-order forms of the neutron transport equation generally yield SPD systems that can benefit from efficient smoothers; q.v. Chapter 5. Additionally, it should be straightforward to devise consistent diffusion acceleration schemes since both the high and low-order approximations can be subject to the same spatial discretisation [40]. Even so, memory requirements may become prohibitive for large problems in multiple dimensions; viz. banded matrices must be stored group-ordinate-wise. Furthermore, problems of causality may be encountered and it is difficult to model for the streaming of neutrons through voids, where the discretised equations are undefined; q.v. Section 2.3 [111]. One may therefore opt to treat the FONT equation.

Unfortunately, CBG discretisations of the FONT-S_N equations are unstable for heterogeneous problems due to a lack of local conservation [207, 242]; q.v. Section 6.3. Allowing for geometric flexibility, local resolution and the propagation of information in specific directions, such wave-dominated problems are more amenable to upwinding schemes for DG discretisations [238]. The coupling between elements is achieved solely through integrals over the interfaces and weakly enforcing continuity in such a manner is equivalent to treating the jump conditions as natural interface conditions [194, 207, 236]. Thus, alternatively, this could be presented as a least-squares penalisation for a particular choice of penalty parameters [363]. However, such a discretisation becomes unstable for arbitrarily large penalty parameters [347]. Whilst no restriction is placed upon the smoothness between elements, locality does come at the cost of significantly more DoF than that of a similar CBG discretisation.

Choosing the numerical fluxes as per an upwinding scheme allows for the spatial dependence to be resolved locally, element by element. In a given direction of neutron travel, the locally-formulated problem depends only upon those values upwind; and any dependencies between elements can be expressed as a *digraph*, which is clearly ordinate-dependent; q.v. Section 6.4.1. If acyclic, the elements can be ordered such that globally connected model is strictly block lower-triangular, with each *block* corresponding to the *unknowns* of a single element, which can be inverted successively, obviating the assembly of any large banded matrices [345]. For this reason, marching schemes for DG-FEM discretisations have been applied primarily to meshes of straight-sided simplices; q.v. Section 2.7.3.1 [69, 236, 240, 364]. Although the locally formulated matrix equations are generally non-SPD and preclude the use of PCG methods, they are usually small enough to be solved by direct methods.

Otherwise, should the graph contain cycles, the discretised equations could not be solved *exactly* in *one* full transport sweep; q.v. Section 2.8.1. In the general case of unstructured meshes for re-entrant geometries, the normal vector over a non-planar face is not constant and the neutrons travelling in a given direction can be perceived as both *incoming* and *outgoing*, resulting in a mutual dependency between those elements that share that face; e.g. the curved element faces in a NURBS description of a pin-cell present such dependencies; v. Figure 6.2. Cycles can be resolved by defining an overall direction of neutron travel based upon the value of a face-averaged normal vector [236]. However, this would not be appropriate within an IGA framework, as it detracts from the goal of performing the analysis over an exact representation of the underlying geometry. The first application of an upwinding scheme to a DG-IGA discretisation of hyperbolic systems was investigated by Owens et al, who demonstrated degraded performance for the use of face-averaged normals [247]. Thus, one must ensure that the numerical integration over the separate *inflow* and *outflow* portions of an element face are performed with sufficient accuracy such as not to affect the convergence of the numerical method [278]. Howbeit, strongly coupled components (SCC) can still be induced in the digraph of a mesh that does not contain re-entrant elements; e.g. most general 2D quadrilateral meshes and 3D hexahedral and tetrahedral meshes. Should it not be possible to form a definite block lower-triangular ordering of the unknowns, the cycles must be broken somehow and the transport sweep performed using the best available information for the upwinded fluxes; viz. any unresolved dependencies are simply lagged within the inner-iteration procedure [236, 365].

The rest of the chapter is organised as follows: First, the governing equations are presented, the multi-group FONT-S_N equations, along with their physical adjoint equations and appropriate BC. This is followed by an upwinding scheme for a DG-IGA discretisation, for which any cycles must be broken and a sweeping order established before solving the system in a *matrix-free* fashion. Finally, the discretisation is verified and the convergence properties analysed by the MMS; and numerical results are presented for a couple of nuclear benchmarks.

6.2 The First-Order Form of the Multi-Group Neutron Transport Equations

6.2.1 The Direct (Primal) Equations

The direct (primal) FONT-$S_N P_L$ equations are posed on a multi-dimensional open, connected and bounded domain, $V \subset \mathbb{R}^d$ for $d = 2, 3$. The strong form of the problem is cast as solving $\varphi^{gk} \colon V \to \mathbb{R}$ for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$\begin{cases} \vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk}(\vec{r}) + \Sigma_t^g(\vec{r})\varphi^{gk}(\vec{r}) = q^{gk}(\vec{r}), & \vec{r} \in V; \quad (6.1a) \\ \varphi^{gk}(\vec{r}) = 0, & \vec{r} \in \partial V_V^-; \quad (6.1b) \\ \varphi^{gk}(\vec{r}) = \varphi^{gk\prime}(\vec{r}), & \vec{r} \in \partial V_R^-. \quad (6.1c) \end{cases}$$

The normal vector, $\vec{n}$, can be defined everywhere over ∂V_{ext}, over different portions of which, ∂V_{ext}^-, the angular neutron flux, φ^{gk}, may be *incoming* for a given $\vec{\Omega}^k$; where $\partial V_{ext}^{\pm} \coloneqq \{\partial V_f \colon \vec{\Omega}^k \cdot \vec{n} \gtrless 0, \forall \partial V_f \in \partial V_{ext}\}$. Only Dirichlet-type incoming BC have been considered in the above, $\partial V_{ext} \coloneqq \partial V_V \cup \partial V_R$, where ∂V_V and ∂V_R denote the incoming portions of the boundary where vacuum and reflective BC are prescribed, respectively; i.e. $\partial V_V^- \coloneqq \partial V_V \cap \partial V_{ext}^-$ and $\partial V_R^- \coloneqq \partial V_R \cap \partial V_{ext}^-$. One considers only those reflective boundaries that form a perpendicular plane with any of the axes in a rectangular Cartesian coordinate system, such that the specular reflection, $\vec{\Omega}^{k\prime}$, of each discrete ordinate, $\vec{\Omega}^k$, is also a member of the LS or GC angular quadrature sets; where for $\varphi^{gk\prime} = \varphi^g(\vec{\Omega}^{k\prime})$, $|\vec{\Omega}^{k\prime} \cdot \vec{n}| = |\vec{\Omega}^k \cdot \vec{n}|$ and $w_{k\prime} = w_k$.

The defining equations can be recast in block-like group-ordinate operator notation:

$$\begin{aligned} L\varphi &= q, && (6.2a) \\ &= S\varphi + F\varphi + \frac{1}{4\pi} q_x; && (6.2b) \end{aligned}$$

with the angular neutron flux, $\varphi = [\varphi^1, \ldots, \varphi^G]^T$, where $\varphi^g = [\varphi^{g1}, \ldots, \varphi^{gK}]^T$ for $g = 1, \ldots, G$; the strength of any extraneous isotropic fixed sources, $q_x = [q_x^1, \ldots, q_x^G]^T$; the diagonal-loss operator, L, comprised of within-group losses due to streaming and total collision in each discrete ordinate, L^{gk}, Equation 6.3a; the scatter operator, S, comprised of potentially anisotropic group-to-group scatter terms in each discrete ordinate, S^{gk}, Equation 6.3b; and the fission operator, F, comprised of isotropic

group-to-group fission terms in each discrete ordinate, F^{gk}, Equation 6.3c.

$$L^{gk}\varphi \coloneqq \vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk} + \Sigma_t^g \varphi^{gk}; \tag{6.3a}$$

$$S^{gk}\varphi \coloneqq \sum_{g\prime=1}^{G} \sum_{l=0}^{L} \sum_{m=-l}^{l} \Sigma_{sl}^{g\prime\to g} Y_l^m \Phi_l^{g\prime m}; \tag{6.3b}$$

$$F^{gk}\varphi \coloneqq \chi^g \sum_{g\prime}^{G} \nu^{g\prime} \Sigma_f^{g\prime} Y_0^0 \Phi_0^{g\prime 0}. \tag{6.3c}$$

For completeness, a $\mathrm{P_L}$ expansion of the scattering kernel has been included; however, within the scope of this project, one is content in the limit of $L = 0$; likewise, only isotropic fission and extraneous fixed sources are considered. The multi-group components of the spherical harmonic flux moments, Φ_l^{gm}, are evaluated as per Equation 2.44, which is recalled here for $1, \ldots, G$:

$$\Phi_l^{gm}(\vec{r}) \approx \sum_{k=1}^{K} w_k Y_l^m(\vec{\Omega}^k) \varphi^{gk}(\vec{r}). \tag{6.4}$$

Defining the operator, A, such that:

$$A\varphi = (L - S)\varphi, \tag{6.5}$$

one can distinguish fixed-source problems:

$$(A - F)\varphi = \frac{1}{4\pi} q_x, \tag{6.6}$$

from eigenvalue ones:

$$A\varphi = \frac{1}{\mathrm{k_{eff}}} F\varphi. \tag{6.7}$$

6.2.2 The Adjoint (Dual) Equations

Goal-based (DWR-)AMR relies upon duality arguments; therefore, one introduces the adjoint (dual) problem; q.v. Section 2.2.3. Mathematically, the dual problem should be adjoint-consistent with the primal one; q.v. Equation 2.12. By linearity of the inner-product, adjoint-consistency can be satisfied in each of the operators defined in Equations 6.3; q.v. Equation 5.10 [75]. In general, the different possible permutations of the order of discretisation and the adjoint operation do not commute; however, in first establishing the multi-group formalism, it is arbitrary the order in which the discrete ordinates approximation, the spatial discretisation or the adjoint operation are performed, since these three do commute; qq.v Section 2.4.1 and Section 2.5.2 [70, 125].

The dual problem is posed on the same domain, $V \in \mathbb{R}^d$, as the primal one; its strong form cast in

group-ordinate operator notation as solving $\varphi^{gk\dagger}: V \to \mathbb{R}$ for $g = G, \ldots, 1$ and for $k = 1, \ldots, K$:

$$L^\dagger \varphi^\dagger = q^\dagger, \tag{6.8a}$$

$$= S^\dagger \varphi^\dagger + F^\dagger \varphi^\dagger + \frac{1}{4\pi} q_x^\dagger; \tag{6.8b}$$

with the angular neutron importance, $\varphi^\dagger = [\varphi^{1\dagger}, \ldots, \varphi^{G\dagger}]^T$, where $\varphi^{g\dagger} = [\varphi^{g1\dagger}, \ldots, \varphi^{gK\dagger}]^T$ for $g = 1, \ldots, G$; the strength of the adjoint source, $q_x^\dagger = [q_x^{1\dagger}, \ldots, q_x^{G\dagger}]^T$; and the loss, scatter and fission adjoint operators, $L^\dagger$, $S^\dagger$ and $F^\dagger$, respectively:

$$L^{gk\dagger}\varphi^\dagger := -\vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk\dagger} + \Sigma_t^g \varphi^{gk\dagger}; \tag{6.9a}$$

$$S^{gk\dagger}\varphi^\dagger := \sum_{g\prime=1}^{G} \sum_{l=0}^{L} \sum_{m=-l}^{l} \Sigma_{sl}^{g\to g\prime} Y_l^m \Phi_l^{g\prime m\dagger}; \tag{6.9b}$$

$$F^{gk\dagger}\varphi^\dagger := \nu^g \Sigma_f^g \sum_{g\prime}^{G} \chi^{g\prime} Y_0^0 \Phi_0^{g\prime 0\dagger}; \tag{6.9c}$$

with the multi-group components of the adjoint flux moments, $\Phi_l^{gm\dagger}$, for $1, \ldots, G$:

$$\Phi_l^{gm\dagger}(\vec{r}) = \sum_{k=1}^{K} w_k Y_l^m(\vec{\Omega}^k) \varphi^{gk\dagger}(\vec{r}). \tag{6.10}$$

The scatter and the fission operators are not self-adjoint; and essentially their adjoint operators are obtained by transposition and the interchanging of arguments of the energy-integral kernels. Furthermore, the *flow* of neutron importance is understood in a reversed sense and thus, one changes the sign of all odd-parity differential operators [125]. The loss operator is formally self-adjoint, $(L\varphi, \varphi^\dagger) = (\varphi, L\varphi^\dagger)$; however, the streaming term necessitates an appropriate prescription of adjoint BC. Equation 6.8 is multiplied by the angular neutron flux and subtracted from Equation 6.2, which is first multiplied by the angular neutron importance; and by linearity of the inner-product:

$$(L\varphi, \varphi^\dagger) - (\varphi, L^\dagger \varphi^\dagger) = (q, \varphi^\dagger) - (\varphi, q^\dagger). \tag{6.11}$$

Suppressing the spatial dependence, the inner-product formed with the loss operator is written explicitly in the first line of Equation 6.12, to which the divergence theorem is applied to give the second line, on which the integrals over the volume are collected:

$$(L\varphi, \varphi^\dagger) = \sum_{g=1}^{G} \sum_{k=1}^{K} w_k \Big[\int_V \varphi^{gk\dagger}(\vec{r}) \vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk}(\vec{r})\, dV + \int_V \varphi^{gk\dagger}(\vec{r}) \Sigma_t^g(\vec{r}) \varphi^{gk}(\vec{r})\, dV \Big], \tag{6.12a}$$

$$= \sum_{g=1}^{G} \sum_{k=1}^{K} w_k \Big[\int_{\partial V} \varphi^{gk\dagger} \vec{\Omega}^k \cdot \vec{n} \varphi^{gk}\, dS \Big] + (\varphi, L^\dagger \varphi^\dagger). \tag{6.12b}$$

For a set of BC for the primal problem, there is a set for the dual one:

$$\sum_{g=1}^{G}\sum_{k=1}^{K} w_k \Big[\int_{\partial V^+} |\vec{\Omega}^k \cdot \vec{n}| \varphi^{gk\dagger} \varphi^{gk} \, dS - \int_{\partial V^-} |\vec{\Omega}^k \cdot \vec{n}| \varphi^{gk\dagger} \varphi^{gk} \, dS \Big] = 0; \tag{6.13}$$

where the integral over ∂V_{ext} has been re-expressed as that over *outflow* and *inflow* boundaries, $\partial V_{ext} = \partial V_{ext}^{+} \cup \partial V_{ext}^{-}$, such that for $k = 1, \ldots, K$:

$$\int_{\partial V_{ext}} \vec{\Omega}^k f(\vec{r}) \cdot \vec{n}(\vec{r}) \, dS = \int_{\partial V_{ext}^{+}} |\vec{\Omega}^k \cdot \vec{n}(\vec{r})| f(\vec{r}) \, dS - \int_{\partial V_{ext}^{-}} |\vec{\Omega}^k \cdot \vec{n}(\vec{r})| f(\vec{r}) \, dS. \tag{6.14}$$

The BC for the primal problem are retained from Equations 6.1 and repeated:

$$\begin{cases} \varphi^{gk}(\vec{r}) = 0, & \vec{r} \in \partial V_V^-; \quad (6.15a) \\ \varphi^{gk}(\vec{r}) = \varphi^{gk\prime}(\vec{r}), & \vec{r} \in \partial V_R^-. \quad (6.15b) \end{cases}$$

The corresponding set of outgoing adjoint BC are prescribed on the same boundaries as those incoming direct ones, respectively:

$$\begin{cases} \varphi^{gk\dagger}(\vec{r}) = 0, & \vec{r} \in \partial V_V^+; \quad (6.16a) \\ \varphi^{gk\dagger}(\vec{r}) = \varphi^{gk\prime\dagger}(\vec{r}), & \vec{r} \in \partial V_R^+. \quad (6.16b) \end{cases}$$

The adjoint BC are consistent by substitution into Equation 6.13; and recalling for reflective BC that $|\vec{\Omega}^{k\prime} \cdot \vec{n}| = |\vec{\Omega}^k \cdot \vec{n}|$ and $w_{k\prime} = w_k$; and by a change of dummy variables; q.v. Equation 5.25.

In a similar manner to Equation 6.5 for the primal problem, the operator $A^\dagger$ is defined such that:

$$A^\dagger \varphi^\dagger = (L^\dagger - S^\dagger)\varphi^\dagger; \tag{6.17}$$

and again, one distinguishes fixed-source problems:

$$(A^\dagger - F^\dagger)\varphi^\dagger = \frac{1}{4\pi} q_x^\dagger, \tag{6.18}$$

from eigenvalue ones:

$$A^\dagger \varphi^\dagger = \frac{1}{\mathrm{k_{eff}}} F^\dagger \varphi^\dagger. \tag{6.19}$$

Under the assumption that the property of (semi-discrete) adjoint-consistency holds, the dual problem possesses the same eigenspectrum as the primal one; and hence, $\mathrm{k}_{\mathrm{eff}}^\dagger = \mathrm{k_{eff}}$; q.v. Equation 4.18.

6.3 A Continuous Bubnov-Galerkin Isogeometric Discretisation

A CBG-IGA discretisation of the multi-group FONT-S_N equations is not suitable for problems that present strong material heterogeneity; for brevity, this is demonstrated informally for the specific

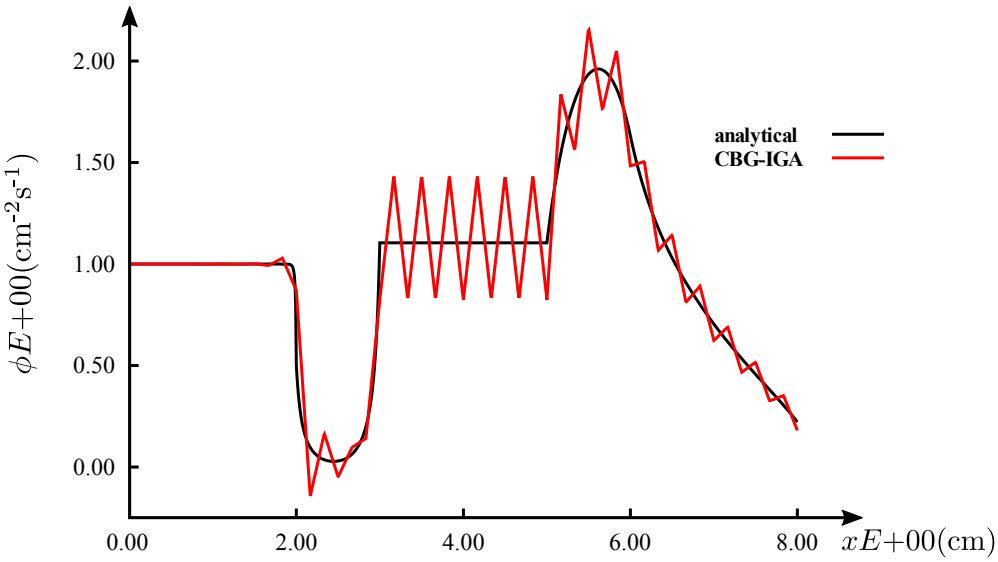


Figure 6.1: An example of the spurious oscillations exhibited by the CBG-IGA 1G FONT-S_8 equations for the example of Reed's 1D problem; modelled by 288 bi-linear NURBS patches, or 350 DoF per ordinate; and solved by direct methods; q.v. Section 5.5.2 [234]. (*Vide* the web-based version for reference to colour.)

example of Reed's 1D 1G problem; q.v. Section 5.5.2 [234].

One makes use of the same s-order Sobolev space, $H^s(V)$ introduced in Section 4.3 and the weak problem is formulated by multiplying Equation 6.2 by test functions, $w \in \mathcal{V}(V)$; expanding φ in a basis of trial functions, $u \in \mathcal{V}(V)$; integrating over V, which is followed by an application of the divergence theorem to the streaming term, such that $\mathcal{V}(V) \subseteq H^1(V)$. The vacuum BC are naturally satisfied; whilst the reflective BC are implemented by substitution of Equation 6.15b. The domain, V, is decomposed into a regular Cartesian grid of 288 bilinear NURBS patches; viz. 350 DoF per ordinate. The global integral statement is assembled and general solvers must be employed because of the size and the asymmetric off-diagonal structure of the matrix equations; e.g. the GMRes method. As demonstrated by the violent oscillations in the numerical approximation about the analytical solution in Figure 6.1, the discretisation becomes unstable after a few *transport sweeps*; and there were no significant gains in accuracy for increasing the polynomial degree of the basis.

The discrete ordinates discretisation allows for a natural decoupling of those *forward* moving neutrons from those *backward* moving ones; however, a strong enforcement of continuity across V and an insufficient resolution of rapidly varying fluxes about boundary layers introduces spurious numerical artefacts and large errors propagate through the voided region. The unsuitability of CBG discretisations of the FONT-S_N equations has already been discussed and thus, for brevity, it is not repeated here; q.v. Section 2.7.3.2.

6.4 An Upwinding Scheme for a Discontinuous Galerkin Isogeometric Discretisation

The multi-group FONT-S_N equations, Equation 6.2, and their adjoint equations, Equation 6.8, are solved for the $\varphi^{gk}\colon V \to \mathbb{R}$ and $\varphi^{gk\dagger}\colon V \to \mathbb{R}$, respectively, over the physical mesh, $\mathcal{T}$, formed as the decomposition of the domain V into non-degenerate and non-overlapping NURBS patches, V_e, such that $\mathcal{T}(V) = \{V_e\}_{e=1}^{E}$ and $V = \bigcup_{e=1}^{E} V_e$; and an arbitrary face $\partial V_f \in \partial V = V_{int} \cup \partial V_{ext}$; where $\partial V_{ext} = \partial V_V \cup \partial V_R$; and $\partial V^{\pm} := \{\partial V_f\colon \vec{\Omega}^k \cdot \vec{n} \gtrless 0, \forall \partial V_f \in \partial V\}$. Again, one exploits the broken Sobolev space, $H^s(\mathcal{T})$; q.v. Equation 4.54. The restriction of values of such functions over V_e are defined over ∂V_e by trace operators, $\gamma_i\colon H^s(V_e) \to H^{s\text{-}i\text{-}½}(\partial V_e)$ for $0 \leq i < s$ [222]. Of particular interest is the Dirichlet trace, γ_0, which is well defined in $L_2(\partial V_e)$ for $s > ½$; q.v. Equation 4.56a.

Any function $u \in H^s(\mathcal{T})$ is discontinuous across V, yet smooth enough within an element to admit two-valued traces at an interface for $s > ½$; with *jumps* and *averages* defined as per Equation 4.58a and Equation 4.58b, respectively, $\forall \partial V_f \in \partial V_{int}$. Physically, the solution to the primal problem should be sufficiently smooth that $[\![\varphi]\!]$ and $[\![\vec{\Omega}\varphi \cdot \vec{n}]\!]$ are vanishing to the interior of $\mathcal{T}$; q.v. Equations 5.42. Likewise, as inherited by the property of adjoint-consistency, the solution to the dual problem, $\varphi^{\dagger}$, is equally smooth; i.e. $\varphi, \varphi^{\dagger} \in \mathcal{V}(V) \subset H^{s>½}(V)$. Without loss of generality, an emphasis is placed upon manipulating the direct equations; and the same discretisation is applied to the adjoint equations.

The weak form is obtained by multiplying Equation 6.2 by test functions, $w \in \mathcal{V}(\mathcal{T})$, integrating over each V_e and then taking the sum of these contributions. The angular neutron flux is expanded in a basis of trial functions, $u \in \mathcal{V}(\mathcal{T}) \subseteq H^s(\mathcal{T})$. By application of the divergence theorem to the streaming term, the basis functions must only be first-order differentiable; viz. $H^1(\mathcal{T})$ is an appropriate space. Again, the proposed spatial discretisation is non-conforming since $\hat{\varphi} \in \mathcal{V}(\mathcal{T}) \not\subset \mathcal{V}(V) \ni \varphi$; and properties of the (semi-)analytical solution cannot be assumed by the discrete one [237]. The discrete variational problem is posed by restricting the collection of infinite spaces of basis functions to finite-dimensional subsets; i.e. $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \subset \mathcal{V}(\mathcal{T})$. As per the isoparametric concept, these broken spaces are chosen as those spanned by the same N_e-number of NURBS that describe the patches, V_e, that constitute the mesh, $\mathcal{T}$; q.v. Equation 4.61. Clearly, $\hat{\mathcal{V}}_{h,p}(\mathcal{T})$ is chosen such that $\hat{\varphi} \to \varphi$ as $N \to \infty$. The weighted-residual statement must be satisfied for each member in the weighting set; and since the space of trial and test functions coincide, this forms the same number of algebraic equations as there are unknowns per group; i.e. $DoF_e^{gk} = \dim \hat{\mathcal{V}}_{h,p}(V_e) = N_e$ and $DoF^{gk} = \sum_{V_e} DoF_e^{gk} = N$ for $k = 1, \ldots, K$; and hence, $DoF^{g} = \sum_{k} DoF^{gk} = K \times N$ for $g = 1, \ldots, G$; and hence, $DoF = \sum_{g} DoF^{g} = G \times K \times N$. Consequently, for the analysis over a single energy-independent mesh, the discrete problem is posed as a $(GKN \times GKN)$-matrix system of equations.

Naturally satisfied by the weak form is the homogeneous Dirichlet-type BC, or the vacuum BC, over ∂V_{ext}^{-}; q.v. Equation 6.15a. The reflective BC is implemented in a weak sense by substitution of Equation 6.15b.

As it stands, the discrete variational statement of Equation 6.6 is to find $\hat{\varphi} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, such that:

$$a(\hat{\varphi}, \hat{w}) = b(\hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}). \tag{6.20}$$

The discrete bilinear form, $a(\cdot,\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \to \mathbb{R}$, is block-like structure and comprises of terms of streaming and total collision, group-to-group scattering and fission, leakage across the union of interfaces and outflow portions of the external boundary, $(\partial V_{ext}^{-})^{c} = \partial V \backslash \partial V_{ext}^{-}$, and the weak enforcement of incoming BC over ∂V_{ext}^{-}:

$$a(\hat{\varphi}, \hat{w}) = a_L(\hat{\varphi}, \hat{w}) - a_S(\hat{\varphi}, \hat{w}) - a_F(\hat{\varphi}, \hat{w}) + a_{int}(\hat{\varphi}, \hat{w}) + a_{ext}(\hat{\varphi}, \hat{w}). \tag{6.21}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$\begin{aligned} a_L(\hat{\varphi}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{V_e \in \mathcal{T}} \left[- \int_{V_e} \vec{\Omega}^k \cdot \vec{\nabla} \hat{w} \hat{\varphi}^{gk} + \hat{w} \Sigma_{t,e}^{g} \hat{\varphi}^{gk} \, dV \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} \sum_{V_e \in \mathcal{T}} a_{L,e}^{gk}(\hat{\varphi}^{gk}, \hat{w}); \end{aligned} \tag{6.22a}$$

$$\begin{aligned} a_S(\hat{\varphi}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{g\prime=1}^{G} \sum_{V_e \in \mathcal{T}} \left[Y_0^0(\vec{\Omega}^k) \int_{V_e} \hat{w} \Sigma_{s0,e}^{g\prime \to g} \hat{\Phi}_0^{g\prime 0} \, dV \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} \sum_{V_e \in \mathcal{T}} a_{S,e}^{gk}(\hat{\varphi}, \hat{w}); \end{aligned} \tag{6.22b}$$

$$\begin{aligned} a_F(\hat{\varphi}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{g\prime=1}^{G} \sum_{V_e \in \mathcal{T}} \left[Y_0^0(\vec{\Omega}^k) \chi_e^g \int_{V_e} \hat{w} \nu_e^{g\prime} \Sigma_{f,e}^{g\prime} \hat{\Phi}_0^{g\prime 0} \, dV \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} \sum_{V_e \in \mathcal{T}} a_{F,e}^{gk}(\hat{\varphi}, \hat{w}); \end{aligned} \tag{6.22c}$$

$$\begin{aligned} a_{int}(\hat{\varphi}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{V_e \in \mathcal{T}} \left[\sum_{\partial V_f \in V_e \cap (\partial V_{ext}^{-})^{c}} \int_{\partial V_f} \hat{w} \vec{\Omega}^k \hat{\varphi}^{gk} \cdot \vec{n} \, dS \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} \sum_{V_e \in \mathcal{T}} a_{int,e}^{gk}(\hat{\varphi}^{gk}, \hat{w}); \end{aligned} \tag{6.22d}$$

$$\begin{aligned} a_{ext}(\hat{\varphi}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{V_e \in \mathcal{T}} \left[\sum_{\partial V_f \in V_e \cap \partial V_R^{-}} - \int_{\partial V_f} |\vec{\Omega}^k \cdot \vec{n}| \hat{w} \hat{\varphi}^{gk\prime} \, dS \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K} \sum_{V_e \in \mathcal{T}} a_{ext,e}^{gk}(\hat{\varphi}^{g}, \hat{w}). \end{aligned} \tag{6.22e}$$

The discrete linear functional, $b(\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \to \mathbb{R}$, comprises only of fixed sources:

$$\begin{aligned} b(\hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{V_e\in\mathcal{T}} \int_{V_e} \left[\hat{w}\frac{1}{4\pi} q_{x,e}^{g}\, dV\right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K}\sum_{V_e\in\mathcal{T}} \int_{V_e} b_e^{gk}(\hat{w}). \end{aligned} \tag{6.23}$$

For multiplying systems, $q_x = 0$, the weak form of Equation 6.7 is to find $\hat{\varphi} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$, such that:

$$\tilde{a}(\hat{\varphi}, \hat{w}) = \frac{1}{\mathrm{k_{eff}}} a_f(\hat{\varphi}, \hat{w}), \quad \forall\ \hat{w} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T}); \tag{6.24}$$

where $\tilde{a}$ is the original bilinear form in Equation 6.21, a, less the fission term; i.e.:

$$\tilde{a} = a + a_F = a_L - a_S + a_{int} + a_{ext}. \tag{6.25}$$

Following those same arguments for the IP-DG-IGA discretisation of the SAAF-$\mathrm{S_N}$ equations, consistency of the discretisation cannot be established by simple substitution of the true solution into Equation 6.20; and more rigorously, the bilinear form should be extended to be bounded on $\hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}^{*}_{hp}(\mathcal{T})$; q.v. Section 5.4. Nonetheless, one assumes that the discrete problem is consistent provided that the numerical fluxes are consistent, from which stems Galerkin's property of orthogonality, Equation 4.71. Alternatively, the current discretisation could be presented as an interior penalty scheme by re-expressing the bilinear form a_{int}, Equation 6.22d, in terms of jumps across interfaces, as per Equation 5.45d; rearranging as per Equation 5.49b; and accounting for the smoothness of the true solution as per Equations 5.42; and for the particular choice of penalising the interface jumps [363]:

$$\begin{aligned} a_{int}^{gk}(\hat{\varphi}^{gk*}, \hat{w}) = &\sum_{\partial V_f \in \partial V_{ext}^{+}} \int_{\partial V_f} \hat{w}\vec{\Omega}^{k}\hat{\varphi}^{gk*}\cdot\vec{n}\, dS + \sum_{\partial V_f\in\partial V_{int}} \int_{\partial V_f} \vec{\Omega}^{k}\cdot\vec{n}[\![\hat{w}]\!]\{\{\hat{\varphi}^{gk*}\}\}\, dS \\ &+ \sum_{\partial V_f\in\partial V_{int}} \frac{1}{2}\int_{\partial V_f} |\vec{\Omega}^{k}\cdot\vec{n}|[\![\hat{w}]\!][\![\hat{\varphi}^{gk*}]\!]\, dS. \end{aligned} \tag{6.26}$$

Either way, one expects consistency asymptotically, since the integral quantities are evaluated using numerical quadrature; q.v. Section 3.2.3; and the minimum quadrature rules, Equations 5.51, will be employed over any boundaries. Furthermore, minimal discrete stability is ensured since the bilinear form is L_2-coercive and thus, the discrete problem admits a unique solution [237].

The discrete system is solved via the SI procedure; q.v. Section 2.8. Weakly coupled in energy and angle, the within-group equations are solved sequentially with the most up-to-date available information; however, different from the foregoing chapters, hyperbolic systems govern propagation problems and are best solved via marching schemes; q.v. Section 2.8.1. That is, for fixed sources,

$\hat{\varphi}^{gk} \in \hat{\mathcal{V}}_{h,p}(V_e)$ is sought for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$ and $\forall V_e \in \mathcal{T}$:

$$\begin{aligned} a^{gk}_{L,e}(\hat{\varphi}^{gk}, \hat{w}) + a^{gk}_{\downarrow,e}(\hat{\varphi}^{gk}, \hat{w}) = a^{gk}_{S,e}(\hat{\varphi}, \hat{w}) + a^{gk}_{F,e}(\hat{\varphi}, \hat{w}) + b^{gk}_{e}(\hat{w}) \\ + a^{gk}_{\uparrow,e}(\hat{\varphi}^{gk}, \hat{w}) - a^{gk}_{ext,e}(\hat{\varphi}^{gk}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V_e). \end{aligned} \tag{6.27}$$

Employing the PI procedure to converge $|1/\mathrm{k_{eff}}|$, the eigenvalue problem is cast as finding $\hat{\varphi}^{gk} \in \hat{\mathcal{V}}_{h,p}(V_e)$ for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$ and $\forall V_e \in \mathcal{T}$:

$$\begin{aligned} a^{gk}_{L,e}(\hat{\varphi}^{gk}, \hat{w}) + a^{gk}_{\downarrow,e}(\hat{\varphi}^{gk}, \hat{w}) = a^{gk}_{S,e}(\hat{\varphi}, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a^{gk}_{F,e}(\hat{\varphi}, \hat{w}) \\ + a^{gk}_{\uparrow,e}(\hat{\varphi}^{gk}, \hat{w}) - a^{gk}_{ext,e}(\hat{\varphi}^{gk}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V_e). \end{aligned} \tag{6.28}$$

The local formulation of the problem significantly reduces the in-core memory requirements, since the globally connected model is not assembled; and although the left-hand side, $a^{gk}_{L,e} + a^{gk}_{\downarrow,e}$, is not SPD, it is usually small enough to invert by direct methods; viz. $\dim \hat{\mathcal{V}}_{h,p}(V_e) = (p+1)^d$ for $d = 2, 3$, for a wholly-local AMR strategy. Otherwise, for a complex internal refinement structure of a given patch, the local system becomes large and it becomes more efficient to employ iterative methods [68, 366].

The coupling between elements is achieved through the integrals over element boundaries in Equation 6.22d, which are separated into *outflow*, ∂V_e^+, and *inflow* portions, ∂V_e^-, such that $\partial V_e^\pm \coloneqq \{\partial V_f : \vec{\Omega}^k \cdot \vec{n} \gtrless 0, \forall \partial V_f \in \partial V_e\}$; and $\partial V_e^+ \cap \partial V_e^- = \emptyset$ and $\partial V_e^+ \cup \partial V_e^- = \partial V_e$; and $a_{int} = a_\downarrow - a_\uparrow$; i.e. for a given V_e and for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$ [236]:

$$\begin{aligned} \sum_{\partial V_f \in \partial V_e} \int_{\partial V_f} \hat{w} \vec{\Omega}^k \hat{\varphi}^{gk*} \cdot \vec{n}\, dS = \sum_{\partial V_f^+ \in \partial V_e} \int_{\partial V_f^+} |\vec{\Omega}^k \cdot \vec{n}| \hat{w} \hat{\varphi}^{gk*}\, dS \\ - \sum_{\partial V_f^- \in \partial V_e} \int_{\partial V_f^-} |\vec{\Omega}^k \cdot \vec{n}| \hat{w} \hat{\varphi}^{gk\uparrow *}\, dS; \end{aligned} \tag{6.29a}$$

$$\implies a^{gk}_{int,e}(\hat{\varphi}^{gk*}, \hat{w}) = a^{gk}_{\downarrow,e}(\hat{\varphi}^{gk*}, \hat{w}) - a^{gk}_{\uparrow,e}(\hat{\varphi}^{gk*}, \hat{w}). \tag{6.29b}$$

Following the convention for upwinding schemes for DG-FEM, the test functions are taken from within V_e and the trial functions are taken as those *upwind* with respect to the direction of neutron travel, $\vec{\Omega}^k$; viz. $\hat{\varphi}^k$ is taken from within V_e for the outflow portions; whilst the angular neutron flux is taken as that from those neighbouring elements that lie upwind for the inflow portions, as denoted as $\hat{\varphi}^\uparrow$; e.g. for $\mathcal{T} = \{V_1, V_2\}$ and $\partial V_{int} = V_1 \cap V_2$, over which $\vec{n}$ points from V_1 to V_2, then $\hat{\varphi}^{k\uparrow} = \hat{\varphi}^k|_{V_1}$ for $\vec{\Omega} \cdot \vec{n} > 0$, otherwise $\hat{\varphi}^{k\uparrow} = \hat{\varphi}^k|_{V_2}$ for $\vec{\Omega}^k \cdot \vec{n} < 0$. One cannot assume that the restriction of the $\hat{\varphi}^{k\uparrow}$ to ∂V_{int} can be exactly represented in the basis of the element being upwinded; and thus, in the most general case, the upwinded flux must be obtained as a projection from the function space of the neighbouring element; viz. this interpolation between *surfaces* meshes can be performed in a similar manner to that presented in Section 4.3.1; e.g. for $\vec{\Omega}^k \cdot \vec{n} < 0$, $\prod^1_{h,p} \hat{\varphi}^{k\uparrow}|_{V_2} : \hat{\mathcal{V}}^2(\partial V_{int}) \to \hat{\mathcal{V}}^1(\partial V_{int})$ for the restriction of the local approximation spaces to the interface, $\hat{\mathcal{V}}^e(\partial V_{int})$ for $e = 1, 2$, as per

Equation 4.45. In order to uphold the conservation of neutrons, one employs quadrature rules that sufficiently sample the bases either side of the interface; q.v. Equations 5.51 subject to Equation 4.73. Should the inflow portion of an element's boundary lie on the physical external boundary, the upwinded flux is determined by the associated BC, Equations 6.15. In the case of regular meshes and in the absence of hanging-nodes, the upwinded flux could be represented arbitrarily in either of the bases in the restriction to the interface; and thus, the upwinding could be performed locally without any knowledge of the neighbouring elements.

For the geometries envisaged within this work and for a large enough angular quadrature set, an arbitrary face ∂V_f of a NURBS patch with curved boundaries may contain both inflow and outflow portions; v. Figure 6.2. For such rational geometries, it would not be appropriate to assume a constant averaged value of the normal vector; and since $\vec{n}(\vec{r})$ for $\vec{r} \in \partial V_f$, the normal components of the direction of neutron travel must be evaluated at each QP [247]. The integral over partial inflow or outflow boundaries, $\partial V_f^{\pm}$, is treated simply as an integral over the entire boundary, ∂V_f; however, the integrand is made to vanish for $\vec{\Omega}^k \cdot \vec{n} \gtrless 0$; i.e.

$$\int_{\partial V_f^{\pm}} \vec{\Omega}^k \cdot \vec{n}(\vec{r}) f(\vec{r})\, dS = \int_{\partial V_f} \vec{\Omega}^k \cdot \vec{n}(\vec{r}) g(\vec{r})\, dS, \quad \text{for } g(\vec{r}) \coloneqq \begin{cases} f(\vec{r}), & \text{for } \vec{\Omega}^k \cdot \vec{n}(\vec{r}) \gtrless 0; \\ 0, & \text{for } \vec{\Omega}^k \cdot \vec{n}(\vec{r}) \gtrless 0. \end{cases} \tag{6.30}$$

The same arguments made for the integration over those partial boundaries with hanging-nodes induced by h-refinement can be repeated here for the integration over $\partial V_f^{\pm}$; viz. the integrand remains a continuous function over a smooth NURBS object and greater accuracy can be gained by supplying more QP than required by Equations 5.51 [70].

In kind, the discrete variational statement of the adjoint problem in the presence of fixed adjoint sources, $q_x^{\dagger} \neq 0$, is cast as finding $\hat{\varphi}^{gk\dagger} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ for $g = G, \ldots, 1$ and for $k = 1, \ldots, K$ and $\forall V_e \in \mathcal{T}$:

$$\begin{aligned} a_{L,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) + a_{\uparrow,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) = a_{S,e}^{gk\dagger}(\hat{\varphi}^{\dagger}, \hat{w}) + a_{F,e}^{gk\dagger}(\hat{\varphi}^{\dagger}, \hat{w}) + b_e^{gk\dagger}(\hat{w}) \\ + a_{\downarrow,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) - a_{ext,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V_e). \end{aligned} \tag{6.31}$$

The adjoint eigenvalue problem is cast as finding $\hat{\varphi}^{gk\dagger} \in \hat{\mathcal{V}}_{h,p}(\mathcal{T})$ for $g = G, \ldots, 1$ and for $k = 1, \ldots, K$ and $\forall V_e \in \mathcal{T}$:

$$\begin{aligned} a_{L,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) + a_{\uparrow,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}) = a_{S,e}^{gk\dagger}(\hat{\varphi}^{\dagger}, \hat{w}) + \frac{1}{\mathrm{k_{eff}}} a_{F,e}^{gk\dagger}(\hat{\varphi}^{\dagger}, \hat{w}) \\ + a_{\downarrow,e}^{gk}(\hat{\varphi}^{gk\dagger}, \hat{w}) - a_{ext,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger}, \hat{w}), \quad \forall\, \hat{w} \in \hat{\mathcal{V}}_{h,p}(V_e); \end{aligned} \tag{6.32}$$

where the discrete adjoint bilinear form, $a^{\dagger}(\cdot,\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \times \hat{\mathcal{V}}_{h,p}^{*}(\mathcal{T}) \to \mathbb{R}$, is defined by:

$$a(\hat{\varphi}^{\dagger *}, \hat{w}) = a_L^{\dagger}(\hat{\varphi}^{\dagger *}, \hat{w}) - a_S^{\dagger}(\hat{\varphi}^{\dagger *}, \hat{w}) - a_F^{\dagger}(\hat{\varphi}^{\dagger *}, \hat{w}) + a_{int}^{\dagger}(\hat{\varphi}^{\dagger *}, \hat{w}) + a_{ext}^{\dagger}(\hat{\varphi}^{\dagger *}, \hat{w}). \tag{6.33}$$

In turn, each of these is defined below, respectively, suppressing any spatial dependence:

$$\begin{aligned} a_L^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{V_e \in \mathcal{T}} \left[\int_{V_e} \vec{\Omega}^k \cdot \vec{\nabla}\hat{w}\hat{\varphi}^{gk\dagger *} + \hat{w}\Sigma_{t,e}^g \hat{\varphi}^{gk\dagger *}\, dV \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K}\sum_{V_e \in \mathcal{T}} a_{L,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger *}, \hat{w}); \end{aligned} \tag{6.34a}$$

$$\begin{aligned} a_S^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{g\prime=1}^{G} \sum_{V_e \in \mathcal{T}} \left[Y_0^0(\vec{\Omega}^k) \int_{V_e} \hat{w}\Sigma_{s0,e}^{g\rightarrow g\prime} \hat{\Phi}_0^{g\prime 0}\, dV \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K}\sum_{V_e \in \mathcal{T}} a_{S,e}^{gk\dagger}(\hat{\varphi}^{\dagger *}, \hat{w}); \end{aligned} \tag{6.34b}$$

$$\begin{aligned} a_F^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{g\prime=1}^{G} \sum_{V_e \in \mathcal{T}} \left[Y_0^0(\vec{\Omega}^k) \nu_e^g \Sigma_{f,e}^g \int_{V_e} \hat{w}\chi_e^{g\prime} \hat{\Phi}_0^{g\prime 0}\, dV \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K}\sum_{V_e \in \mathcal{T}} a_{F,e}^{gk\dagger}(\hat{\varphi}^{\dagger *}, \hat{w}); \end{aligned} \tag{6.34c}$$

$$\begin{aligned} a_{int}^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{V_e \in \mathcal{T}} \left[\sum_{\partial V_f \in V_e \cap (\partial V_{ext}^+)^c} - \int_{\partial V_f} \hat{w}\vec{\Omega}^k \hat{\varphi}^{gk\dagger *} \cdot \vec{n}\, dS \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K}\sum_{V_e \in \mathcal{T}} a_{int,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger *}, \hat{w}); \end{aligned} \tag{6.34d}$$

$$\begin{aligned} a_{ext}^\dagger(\hat{\varphi}^{\dagger *}, \hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{V_e \in \mathcal{T}} \left[\sum_{\partial V_f \in V_e \cap \partial V_R^+} - \int_{\partial V_f} |\vec{\Omega}^k \cdot \vec{n}| \hat{w}\hat{\varphi}^{gk\prime\dagger *}\, dS \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K}\sum_{V_e \in \mathcal{T}} a_{ext,e}^{gk\dagger}(\hat{\varphi}^{g\dagger *}, \hat{w}); \end{aligned} \tag{6.34e}$$

and the discrete adjoint linear functional, $b^\dagger(\cdot)\colon \hat{\mathcal{V}}_{h,p}(\mathcal{T}) \rightarrow \mathbb{R}$ is defined as:

$$\begin{aligned} b^\dagger(\hat{w}) &= \sum_{g=1}^{G}\sum_{k=1}^{K} w_k \sum_{V_e \in \mathcal{T}} \int_{V_e} \left[\hat{w}\frac{1}{4\pi} q_{x,e}^{g\dagger}\, dV \right], \\ &= \sum_{g=1}^{G}\sum_{k=1}^{K}\sum_{V_e \in \mathcal{T}} \int_{V_e} b_e^{gk\dagger}(\hat{w}). \end{aligned} \tag{6.35}$$

For eigenvalue problems, the solution is normalised after each outer iteration such that $a_F(\varphi, \varphi) = 1$; and the adjoint solution is normalised such that $a_F(\varphi, \varphi^\dagger) = 1$; q.v. Section 3.6.2.2.

Again, the coupling between elements is achieved through the integrals over element boundaries in Equation 6.34d, which are separated in a manner not dissimilar to that of Equations 6.29, such that $a_{int}^\dagger = a_\uparrow^\dagger - a_\downarrow^\dagger$; i.e. for a given V_e and for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$\sum_{\partial V_f \in \partial V_e} - \int_{\partial V_f} \hat{w}\vec{\Omega}^k \hat{\varphi}^{gk\dagger *} \cdot \vec{n} dS = \sum_{\partial V_f^- \in \partial V_e} \int_{\partial V_f^-} |\vec{\Omega}^k \cdot \vec{n}| \hat{w}\hat{\varphi}^{gk\dagger *} dS - \sum_{\partial V_f^+ \in \partial V_e} \int_{\partial V_f^+} |\vec{\Omega}^k \cdot \vec{n}| \hat{w}\hat{\varphi}^{gk\dagger\downarrow *} dS; \tag{6.36a}$$

$$\implies a_{int,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger *}, \hat{w}) = a_{\uparrow,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger *}, \hat{w}) - a_{\downarrow,e}^{gk\dagger}(\hat{\varphi}^{gk\dagger *}, \hat{w}). \tag{6.36b}$$

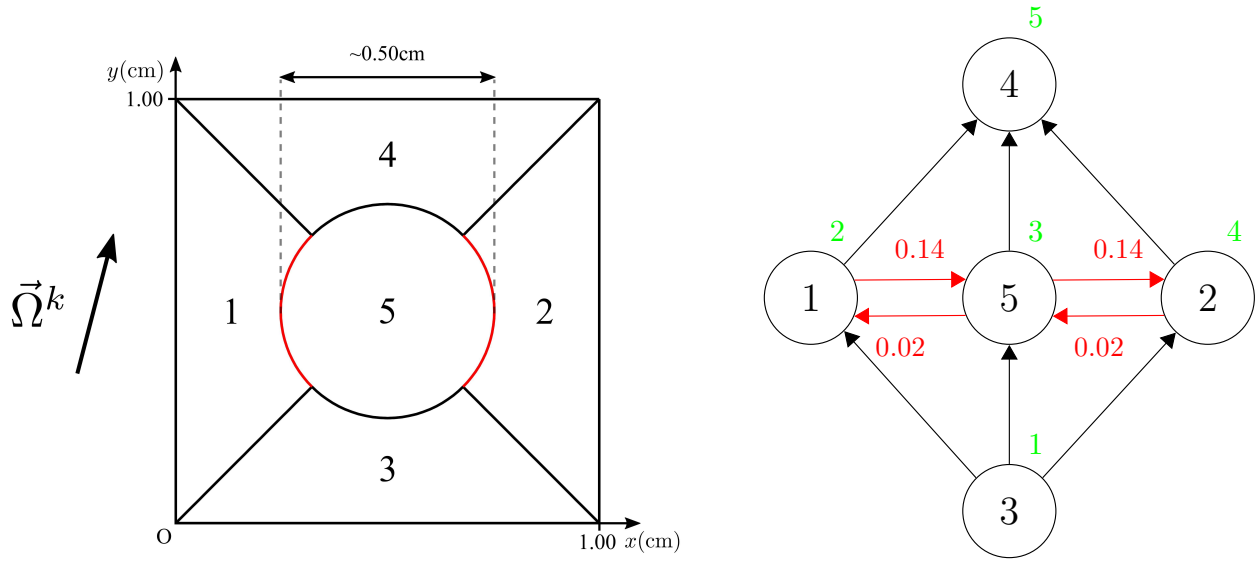


Figure 6.2: A schematic of the cyclic digraph (right) that corresponds to the direction of neutron flight $\vec{\Omega}^k = (0.33, 0.88)$, taken from an S_4 LS angular quadrature set, for the example of a 2D pin-cell constructed from *five* re-entrant NURBS patches (left), which constitute the vertices, v_i, of the digraph and are numbered in black. The directed edges, (v_i, v_j), appear as arrows on the digraph; those highlighted in red form cycles over those re-entrant boundaries also highlighted in red; and their associated weights of *outgoingness* are given as per Equation 6.37. The cycles within the subgraph induced by $\{v_1, v_5, v_2\}$ are *resolved* by systematically deleting the *weaker* edges; viz. (v_5, v_1) and (v_2, v_5); and the final marching order is indicated in green; viz. $\{v_3, v_1, v_5, v_2, v_4\}$. (*Vide* the web-based version for reference to colour.)

Since neutron importance propagates in the opposite direction to that in which neutron travels, the trial functions are taken as those downwind; viz. $\hat{\varphi}^{k\dagger}$ is taken from within V_e for the inflow portions, $\vec{\Omega}^k \cdot \vec{n} < 0$; whilst the angular neutron importance is taken as that from those neighbouring elements that lie downwind for the outflow portions, $\vec{\Omega}^k \cdot \vec{n} > 0$, as denoted as $\hat{\varphi}^{k\dagger\downarrow}$. Should the outflow portion of an element's boundary lie on the physical external boundary, the downwinded adjoint flux is determined by the associated BC; q.v. Equations 6.16.

6.4.1 Graph Theory

As per an upwinding scheme, the weak form of the problem is formulated locally over the elements, $\forall V_e \in \mathcal{T}(V)$. The solution to the FONT-S_N equations ressembles that of a propagating wavefront and thus, one must devise a systematic method to establish, *a priori*, the order in which the mesh must be *swept* in a given direction of neutron travel, $\vec{\Omega}^k$. To this end, borrowing the terminology from Gibbons, one defines the directed graph, or *digraph*, $\mathcal{G}^k$, associated with $\vec{\Omega}^k$ over $\mathcal{T}$; as characterised by a set of *vertices*, $\mathcal{V} = \{v_i\}_{i=1}^{N_v}$, interconnected by a set of *edges*, $\mathcal{E}^k = \{e_i^k\}_{i=1}^{N_e}$; i.e. $\mathcal{G}^k = (\mathcal{V}, \mathcal{E}^k)$, where the total number of vertices and edges are denoted by $N_v = \dim \mathcal{V}$ and $N_e = \dim \mathcal{E}^k$, respectively [367]. More generally, in the case of energy-dependent meshes, the digraph will also depend upon the energy group index, g; however, for the sake of simplicity, this discussion is restricted to the case of a single energy-independent mesh, $\mathcal{T}$.

Each vertex represents a NURBS patch of the same subscript, all of which are assumed to be connected, at least in a weak sense, such that there exist no isolated vertices. A *directed edge* exists between *two* end-points if a neutron can travel in $\vec{\Omega}^k$ between the two; viz. the order of the vertices is significant and a directed edge, $e^k = (v_i, v_j)$, is said to be incident from v_i and incident to v_j; and, obviously, the patches that correspond to v_i and v_j must be adjacent to one another. Moreover, for an interface, ∂V_{ij}, formed between *two* patches, V_i and V_j, as represented in the associated digraph by the vertices v_i and v_j, respectively, a directed edge exists for $\vec{\Omega}^k \cdot \vec{n} > 0$ at any point over ∂V_{ij}, with the normal vector, $\vec{n}$, pointing from V_i towards V_j. To quantify the *outgoingness*, one assigns the weights, $w^k(e^k)$, as per Equation 6.30, to each directed edge for $k = 1, \ldots, K$:

$$w^k(e^k) \coloneqq \int_{\partial V_{ij}^+} \vec{\Omega}^k \cdot \vec{n}(\vec{r})\, dS, \quad \forall e^k \in \mathcal{E}^k. \tag{6.37}$$

For the example of a 2D pin-cell, a schematic of the digraph associated to a given $\vec{\Omega}^k$ is presented in Figure 6.2. A *path* is a sequence of alternating vertices and edges that follows the direction of the edges; the path is *cyclic* if it begins and ends with the same vertex. In the same vein, *two* vertices are said to be *strongly connected* should there exist a directed path from one to the other; and, *vice versa*. Otherwise, *two* vertices are said to be *weakly connected* if they are connected only in an undirected sense; e.g. for the same 2D pin-cell in Figure 6.2, but for $\vec{\Omega}^k = (0.33, 0.33)$, the vertices v_1 and v_3 of the associated digraph would only be weakly connected since $\vec{\Omega}^k$ runs parallel to ∂V_{13}.

Strong connections partition the vertices into *subgraphs* that constitute the SCC, $\mathcal{G}_i^k$, of the digraph; viz. $\mathcal{G}_i^k \subseteq \mathcal{G}^k$ and $\mathcal{G}^k = \{\mathcal{G}_i^k\}_{i=1}^G$; e.g. for the same 2D pin-cell in Figure 6.2, $G = 3$ for the partitioning of the vertices into the subsets $\{v_3\}$, $\{v_1, v_5, v_2\}$ and $\{v_4\}$. The equations formulated over these SCC correspond to the *blocks* in the globally connected model, between which a topological acyclic directed path always exists; e.g. $\{v_3\} \rightarrow \{v_1, v_5, v_2\} \rightarrow \{v_4\}$. Based upon a depth-first search (DFS), the condensation of a digraph into its SCC is unique and can be obtained by Kosaraju's algorithm [368]. For a digraph described by an *adjacency matrix*, this algorithm becomes *complex* and inefficient since it requires $\mathcal{O}((N_v)^2)$ computational steps; otherwise, described by *adjacency lists*, this algorithm can be performed in linear time, $\mathcal{O}(N_v + N_e)$ [367].

6.4.1.1 Breaking Cycles

In the presence of cycles, it becomes unclear the order in which the elements should be swept; and since an acyclic directed path always exists for the condensation of the digraph, any dependencies must be resolved only within each SCC. This could be achieved either by constructing explicitly and then solving simultaneously the corresponding block of matrix equations; or by determining a sweeping order of the vertices within an SCC and then iterating on the angular neutron flux. Owens et al

and *eidem* deemed neither of these approaches to be computationally efficient, since full convergence would generally be slow to attain; and especially so for any SCC formed from a large number of elements [231,247]. Rather, the same authors determined an *optimal* marching order within each SCC by breaking the cycles therewithin; and thus, each element is visited only once per transport sweep in a given direction of neutron travel [231, 247]. The same approach is adopted here.

Restricting the attention to a given subgraph, $\mathcal{G}_i^k$, induced by the vertices of a given SCC, the sweeping order therewithin is determined by systematically deleting the *weakest* edges until an acyclic directed path is obtained. The strength of a directed edge corresponds to its *outgoingness*, as defined in Equation 6.37; viz. the smaller the weight, the weaker the connection between the *two* end-points in the direction $\vec{\Omega}^k$. For example, for the digraph in Figure 6.2, the cycles within the SCC induced by the vertices $\{v_1, v_5, v_2\}$ are *resolved* by deleting the edges (v_5, v_1) and (v_2, v_5); and the final marching order is obtained as $\{v_3, v_1, v_5, v_2, v_4\}$.

Such an algorithm for edge deletion always guarantees a valid marching order for any mesh; however, the modified transport sweep will not exactly invert the global problem and any unresolved dependencies are simply *lagged* within the inner-iteration procedure [236,365]. In other words, the incoming current for a given patch depends only upon that from the neighbouring patches that lie upwind; e.g. for the digraph presented in Figure 6.2, V_1 depends upon V_3 and V_5, yet, only V_3 will contribute any *new* information, since it comes before V_1 in the sweeping order; whilst V_5 can only contribute old information, since it comes later in the sweeping order.

With regard to the interpolation-based nature of the AMR strategy, an inherent bonus of defining the more-refined reference mesh, $\mathcal{T}_{ref} = \mathcal{T}_{h/2,p+1}$, as that which results from *one* uniform *hp*-refinement of the coarser current mesh is that some regularity is implicitly enforced over the mesh where the numerical analysis is performed; viz. some of the dependencies present in the digraph are naturally resolved, or at least weakened, under *h*-refinement; and thus, one expects to dedicate increasingly less effort to the breaking of cycles as the mesh is refined for any of the problems considered in this work.

Finally, the adjoint equations are subject to the same spatial discretisation; yet, the information propagates in an opposite sense and thus, the neutron importance is solved for as per a *downwinding* scheme; or rather, the sweeping order is obtained by simply reversing that established for the primal problem; e.g. the directed edges in digraph presented in Figure 6.2 are reversed for the corresponding dual problem and the sweeping order in $\vec{\Omega}^k = (0.33, 0.88)$ is taken as $\{v_4, v_2, v_5, v_1, v_3\}$.

6.5 The Numerical Results

In the application to the multi-group FONT-S_N equations, this section presents some numerical results for the AMR-$h(p)$ strategies developed in Section 3.5 and Section 3.6. Unless otherwise stated, all operations are performed such as to optimise the L_2-error. For a couple of verification test cases and nuclear reactor physics benchmarks, the AMR strategies are compared against global uniform ones. Any data-plots that make reference to the DoF of the discretised system justly account for all unknowns that were solved for in the complete system; viz. $DoF = G \times K \times N$ for G-number of groups and K-number of ordinates solved over an energy-independent mesh of N-number of control-variables.

The discretisation developed within this chapter has been implemented in the same new modern-Fortran code as that presented in the foregoing chapters; and thus, the same libraries, tolerances and parameters are employed here [73]. However, now, there is no assembly or inversion of any large sparse matrices; and rather, one solves many more smaller systems of algebraic equations. Owing to the decoupling of angles as per the discrete ordinates approximation and by the nature of the marching scheme, a different approach to *parallelisation* is adopted; viz. for rectangular Cartesian coordinate systems, it is trivial to parallelise the discrete problem over the discrete ordinates.

Owens et al provide an extensive investigation of adaptive h-refinement strategies for DG-IGA discretisations of the multi-group FONT-S_N equations in 2D [231, 247, 278]. Thus, this section serves to verify the spatial discretisation and the AMR-$h(p)$ strategies developed and implemented within the numerical code. Again, adopting a distributed memory computing architecture, one makes use of the Open MPI message passing library; in addition to the LAPACK library and its capacity for solving dense linear algebra problems; viz. the efficient double precision routines `DGETRF` to compute the LU decomposition of a general matrix; and `DGETRI` to compute the inverse of a matrix using its LU decomposition [296]. One remarks that no sophisticated method of domain decomposition has been employed; rather, the discrete ordinates are rudely assigned to and their corresponding angular neutron flux solutions held on different MPI parallel processes.

Again, all *jobs* are submitted to the general purpose Imperial College HPC cluster computing service (CX1); and at each iteration of the AMR procedure, the reference solution is *bootstrapped* from the previous iteration to serve as an initial guess over the current reference mesh.

6.5.1 The Method of Manufactured Solutions

For some manufactured solution, φ^{gk}, the (semi-)analytical anisotropic manufactured source, q_x^{gk}, can be evaluated for $g = 1, \ldots, G$ and for $k = 1, \ldots, K$:

$$q_x^{gk} = \vec{\Omega}^k \cdot \vec{\nabla}\varphi^{gk} + \Sigma_t^g \varphi^{gk} - \sum_{g\prime=1}^{G} \Sigma_{s0}^{g\prime \to g} Y_0^0(\vec{\Omega}^k)\Phi_0^{g\prime 0}. \tag{6.38}$$

The manufactured solution and source are used to study the convergence rates of the global uniform refinement strategy on a simple Cartesian grid and a more challenging pin-cell one; taking into account the implementation of different BC and the numerical quadrature routines. These results are then used as a reference with which to analyse the NRG-AMR-$h(p)$ strategy, which takes into account performance of all mesh-to-mesh interpolation procedures and the capacity of the numerical methods to enforce continuity over irregular meshes that contain *hanging-nodes*.

The solution to the FONT-S_N equations is not typically a smooth function in angle nor space. Modelling for the propagation of such discontinuities tends to degrade the accuracy of the numerical method, since a certain degree of regularity, r, is demanded of the true solution; q.v. Equations 4.103. For a smooth underlying solution, $r \to \infty$, *second-order* accuracy can be demonstrated as per Equation 4.103b for $p = 1$; yet this is somewhat reduced as $r \approx 3/2$ for piecewise constant cross-section data; and $r \approx 1/2$ in the presence of voids and highly absorbing materials [369, 370].

6.5.1.1 A Cartesian Grid

The first verification case consists of a simple Cartesian grid defined over $V = [0, 1]^2$, which comprises of a homogeneous material whose properties are given in Table 4.1. The coarsest mesh is composed of *four* equal-volume squares represented by bilinear NURBS patches. Assuming a *one-speed* energy approximation and an S_4 LS angular quadrature set, the manufactured solution, φ, is chosen as that in Equation 5.71; with Φ evaluated as per Equation 5.72.

To allow for the maximum attainable convergence rates, the manufactured solution has been chosen to be smooth, yet it cannot be arbitrarily represented exactly by higher-order NURBS. Moreover, the homogeneous Dirichlet (vacuum) BC is naturally satisfied; it will also be naturally enforced by the discretisations on all external boundaries. One emphasises that the manufactured solution is chosen for the discrete ordinates equations rather than the full NTE; it is thus used to isolate those errors that result from the spatial discretisation [68].

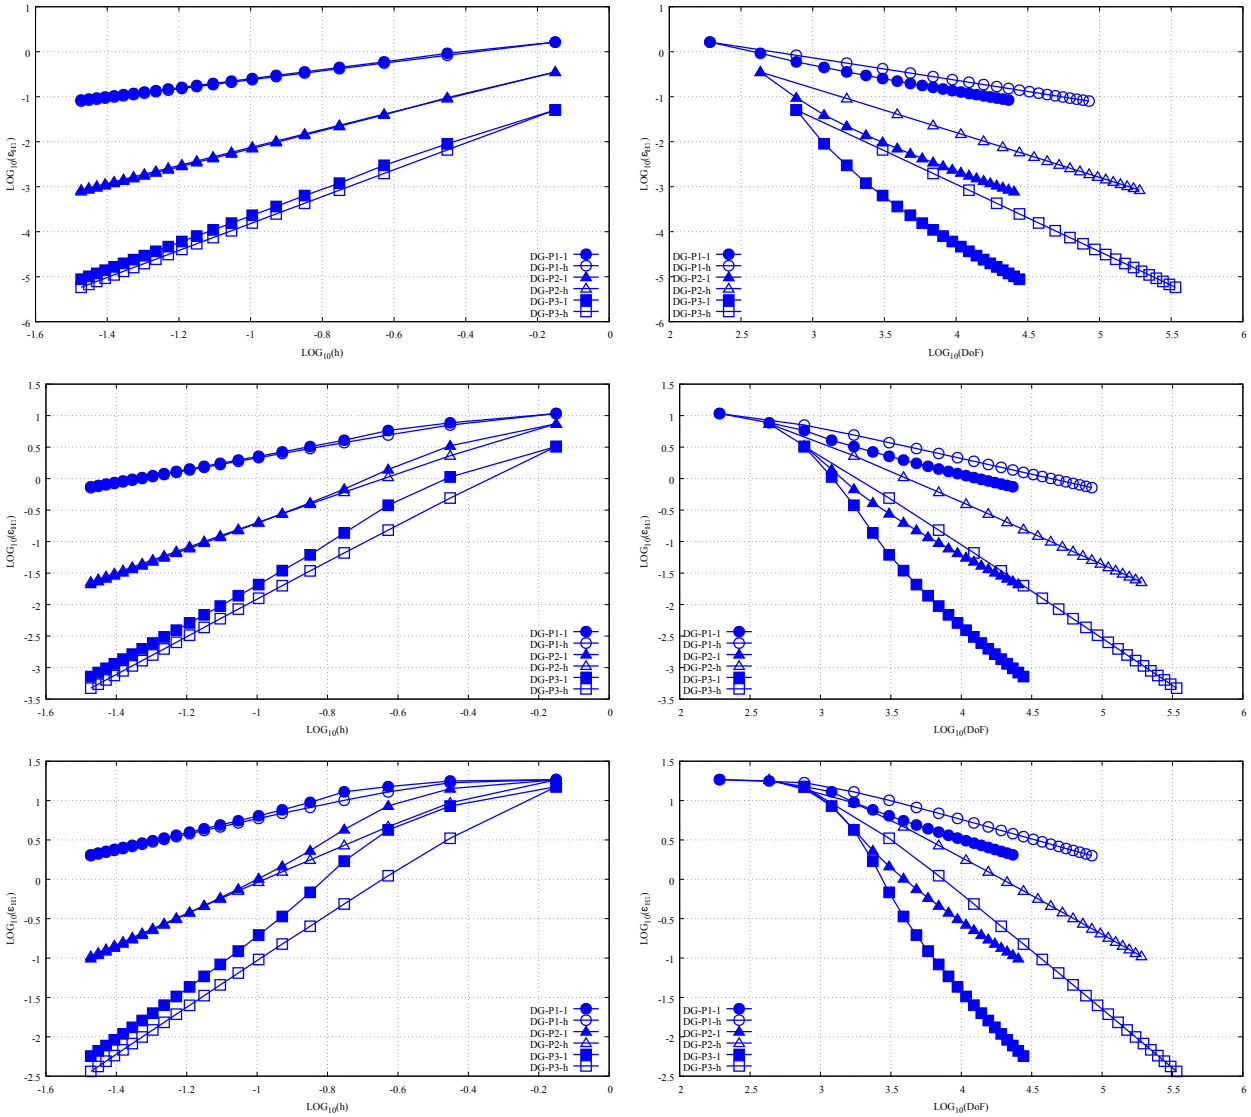


Figure 6.3: The convergence plots, ϵ_{H^1} v h (left) and ϵ_{H^1} v DoF (right), for the verification of the global uniform refinement strategy for the sweep-based DG-IGA discretisation of the 1G FONT-S_4 equations over a 2D Cartesian grid; the manufactured solution is chosen as per Equation 4.105 for $n = 0$ (top row), $n = 1$ (middle row) and $n = 2$ (bottom row). (*Vide* the web-based version for reference to colour.)

Global Uniform Refinement For the KI-refinement of linear, quadratic and cubic B-splines, the H^1-errors for the sweep-based DG-IGA spatial discretisation are presented in Figure 6.3 for the manufactured solution chosen as the first, third and fifth harmonics of a modified standing wave equation; q.v. Equation 5.71. Qualitatively similar plots were obtained for the L_2-error; however, for brevity, these have not been presented. For such a simple Cartesian 2D grid, a unique acyclic sweeping order exists; and for the prescription of vacuum BC, the full system can be resolved in *one* full transport sweep. Under refinement, the discontinuous discretisation scheme converges to the analytical solution; and asymptotically, attains theoretical orders of accuracy, as presented in Table 6.1. For such a homogeneous problem with a smooth underlying solution, exploiting higher-order continuity in the basis yields more accuracy per DoF than that compared to a similar h-refinement.

One acknowledges that for such sweep-based solution methods, neither the mesh spacing nor the DoF

		DG-IGA			
		KI-1		KI-h	
	p	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}
	1	1.01	2.02	1.00	2.00
$(n=0)$	2	2.04	3.00	2.00	2.99
	3	2.99	3.99	3.00	3.98
	1	0.99	2.00	0.99	2.00
$(n=1)$	2	2.02	3.02	2.00	3.00
	3	2.99	4.03	3.00	4.00
	1	1.01	2.03	1.00	2.01
$(n=2)$	2	2.00	3.01	1.99	3.00
	3	3.04	4.05	2.99	4.00

Table 6.1: The asymptotic orders of accuracy, Equations 4.103, for the verification of the global uniform refinement strategy for the sweep-based DG-IGA discretisation of the 1G FONT-S_4 equations over a 2D Cartesian grid; the convergence rates are evaluated as the slope of the last few points from each dataset.

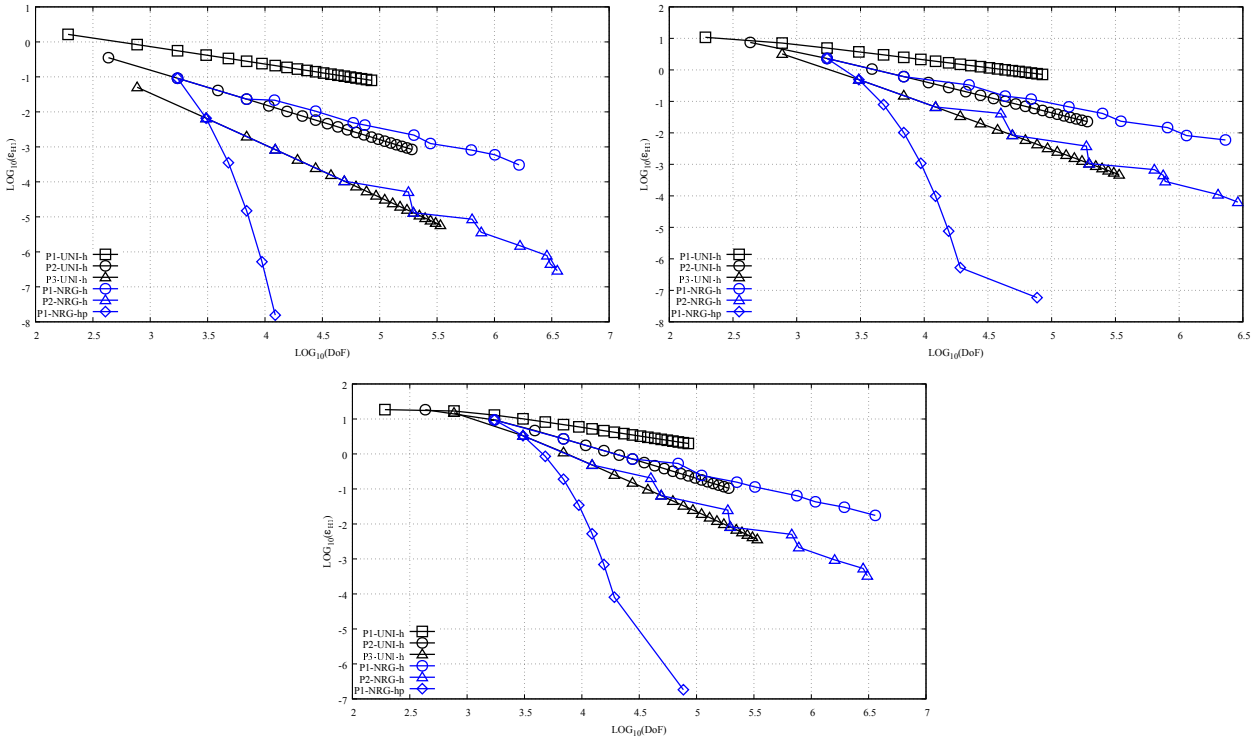


Figure 6.4: The convergence plots, ϵ_{H^1} v DoF, for the verification of the NRG-AMR-$h(p)$ strategy for sweep-based DG-IGA discretisation of the 1G FONT-S_4 equations over a 2D Cartesian grid; the manufactured solution is chosen as per Equation 5.71 for $n=0$ (top left), $n=1$ (top right) and $n=2$ (bottom). (*Vide* the web-based version for reference to colour.)

are particularly useful metrics since the computational effort varies linearly with the number of spatial cells; viz. a more-refined mesh decomposed into twice as many elements, *ceteris paribus*, would require twice the computational effort to resolve [371]. Nevertheless, in the interest of AMR-$h(p)$ strategies, not every element belonging to the same mesh may be constructed from a basis of NURBS of the same polynomial degree; and therefore, it is still useful to compare total DoF for the system.

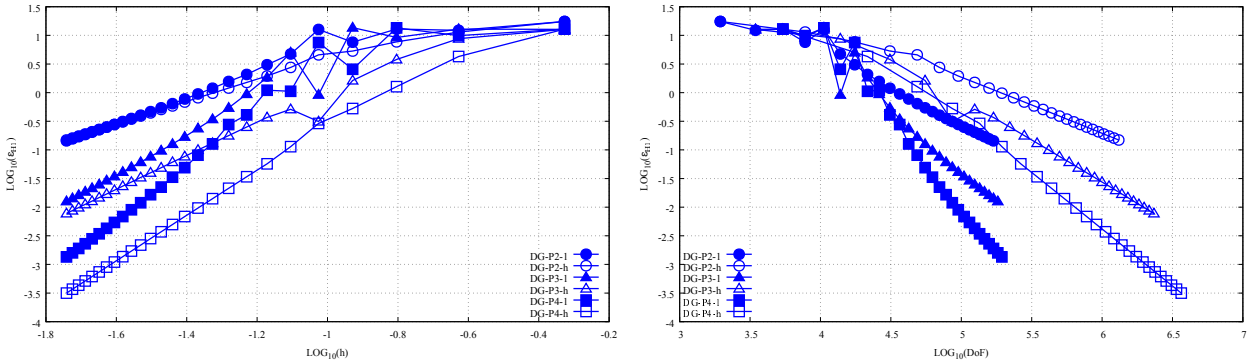


Figure 6.5: The convergence plots, ϵ_{H^1} v h (left) and ϵ_{H^1} v DoF (right), for the verification of the global uniform refinement strategy for the sweep-based DG-IGA discretisation of the 2G FONT-S_4 equations over a 2D pin-cell grid; the manufactured solutions are chosen as per Equation 5.74 for $g = 1, 2$. (*Vide* the web-based version for reference to colour.)

Local Adaptive Refinement The results of the NRG-AMR-$h(p)$ strategy for chasing the global H^1-error in the sweep-based DG-IGA discretisation are presented in Figure 6.4 for the manufactured solution chosen as the first, third and fifth harmonics of a modified standing wave equation; q.v. Equation 5.71. These results are compared against those obtained for the global uniform h-refinement strategy. The findings are qualitatively similar to those obtained for the verification of the discretisation of the NDE and thus, for brevity, the discussion is not repeated here; q.v. Section 4.5.1.1.

6.5.1.2 A Pincell-like Grid

The second verification case consists of a more challenging pin-cell grid defined over $V = [0, 1]^2$, which comprises of a homogeneous material whose properties are given in Table 4.3. The coarsest mesh is composed of *nine* equal-volume non-degenerate bi-quadratic NURBS patches. Assuming a two-group energy approximation and an S_4 LS angular quadrature set, the manufactured solution, φ, is chosen as that in Equation 5.74; with Φ evaluated as per Equation 5.75.

	DG-IGA			
p	KI-1		KI-h	
	ϵ_{H^1}	ϵ_{L_2}	ϵ_{H^1}	ϵ_{L_2}
2	2.11	3.17	1.99	3.07
3	3.14	4.28	3.02	4.04
4	4.22	5.24	3.99	5.05

Table 6.2: The asymptotic orders of accuracy, Equations 4.103, for the verification of the global uniform refinement strategy for the sweep-based DG-IGA discretisation of the 2G FONT-S_4 equations over a 2D pin-cell grid; the convergence rates are evaluated as the slope of the last few points from each dataset.

Global Uniform Refinement For the KI-refinement of quadratic, cubic and quartic NURBS, the H^1-errors for the sweep-based DG-IGA spatial discretisation are presented in Figure 6.5 for the multi-

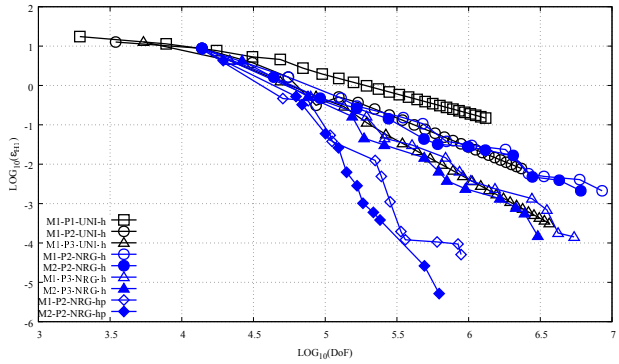


Figure 6.6: The convergence plots, ϵ_{H^1} v DoF, for the verification of the NRG-AMR-$h(p)$ strategy for sweep-based DG-IGA 2G FONT-S_4 equations over a 2D pin-cell grid; the manufactured solutions are chosen as per Equation 5.74 for $g = 1, 2$. (*Vide* the web-based version for reference to colour.)

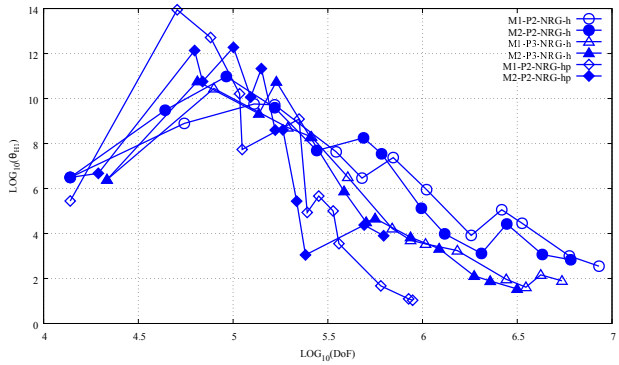


Figure 6.7: The plots of the effectivity index, θ_{H^1} v DoF, for the verification of the NRG-AMR-$h(p)$ strategy for the sweep-based DG-IGA 2G FONT-S_4 equations over a 2D pin-cell grid; the manufactured solutions are chosen as per Equation 5.74 for $g = 1, 2$. (*Vide* the web-based version for reference to colour.)

group components of the manufactured solution chosen as the third and fifth harmonics of a modified standing wave equation; q.v. Equation 5.74. Qualitatively similar plots were obtained for the L_2-error; however, for brevity, these have not been presented. Under refinement, the sweep-based discontinuous discretisation scheme converges to the analytical solution; and asymptotically, attains the theoretical orders of accuracy for KI-h-refinement, as presented in Table 6.2. Again, however, for KI-1-refinement, one remarks a degree of super-convergence, which may be explained by the higher-order continuity in the basis; q.v. Section 4.5.1.2.

These results do verify consistency of the sweep-based DG-IGA discretisation over meshes of curved elements with reentrant surfaces; and thus, one is confident that the algorithm for determining a *once-through* sweeping order from the corresponding cyclic digraphs is robust and that a sufficient number of QP are employed over partial inflow and outflow boundaries in the upwinding scheme.

Local Adaptive Refinement The results of the NRG-AMR-$h(p)$ strategy for chasing the global H^1-error in the sweep-based DG-IGA discretisation are presented in Figure 6.6 for the multi-group components of the manufactured solution chosen as the third and fifth harmonics of the standing

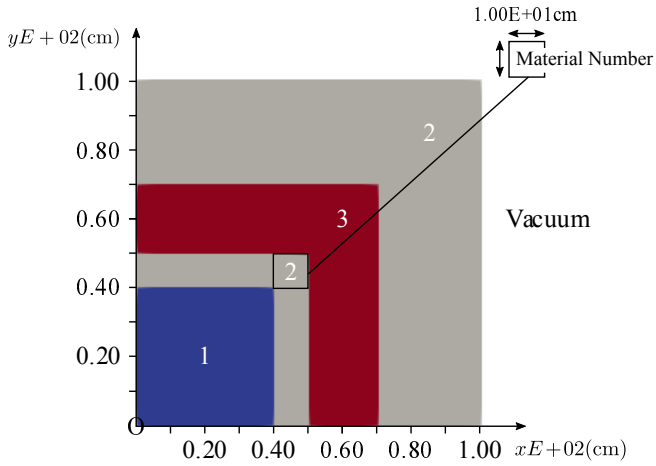


Figure 6.8: The geometry for Azmy's 2D 1G radiation shielding problem. The constituent materials are denoted by the numbers given in parentheses, accordingly: source (1), water (2); and iron (3) [255]. (For an interpretation of any reference made to colour, the lector is referred to the web-based version of this work.)

wave equation; q.v. Equation 5.74. The effectivity, θ_{H^1}, of the interpolation-based error measure, as per Equation 3.24, is measured and displayed in Figure 6.7. The capacity for energy-dependent meshes is also tested. These results are qualitatively similar to those obtained for the verification of the discretisation of the NDE and thus, for brevity the discussion is not repeated here; q.v. Section 4.5.1.2. Nevertheless, one emphasises that any savings made in the refinement of a given mesh for a given energy group actually correspond to K-number of savings, since each group-ordinate-wise system of equations is resolved over that same mesh.

6.5.2 Azmy's 2D Radiation Shielding Problem

Material	Group (g)	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to 1}$ (cm^{-1})	q_x^g (cm^{-3}.s^{-1})
(1) Source	1	1.000E+00	5.000E-01	1.000E+00
(2) Water	1	1.000E-01	1.000E-02	0.000E+00
(3) Iron	1	3.000E-01	1.000E-01	0.000E+00

Table 6.3: The nuclear data for Azmy's 2D 1G radiation shielding problem [372].

To avoid excessively slow rates of convergence, Azmy's problem is an adaptation of Khalil's 2D steel-and-water radiation shielding problem for a smaller average scattering ratio, $c = \Sigma_{s0}^{g\to g}/\Sigma_t^g \leq 1$ [372, 373]. This problem presents *three* physically distinct regions, including a homogenised source region and *two* shield regions of water and iron, as represented by Materials 1, 2 and 3, respectively; v. Figure 6.8. The corresponding macroscopic nuclear cross-section data are provided in Table 6.3. One assumes that an isotropic source is uniformly dispersed across Material 1 and that the surfaces of the physical boundary are bare; moreover, as per the *one-quarter* symmetry in the problem, reflective BC are prescribed at $x = y = 0.00$ cm. The coarsest Cartesian mesh is constructed from 100 bilinear NURBS patches. $\mathrm{P_{N\text{-}1}}$ reference solutions were obtained with the combined GEM-EVENT codes, an independent FEM mesh generator and multi-group even-parity neutron transport solver, respec-

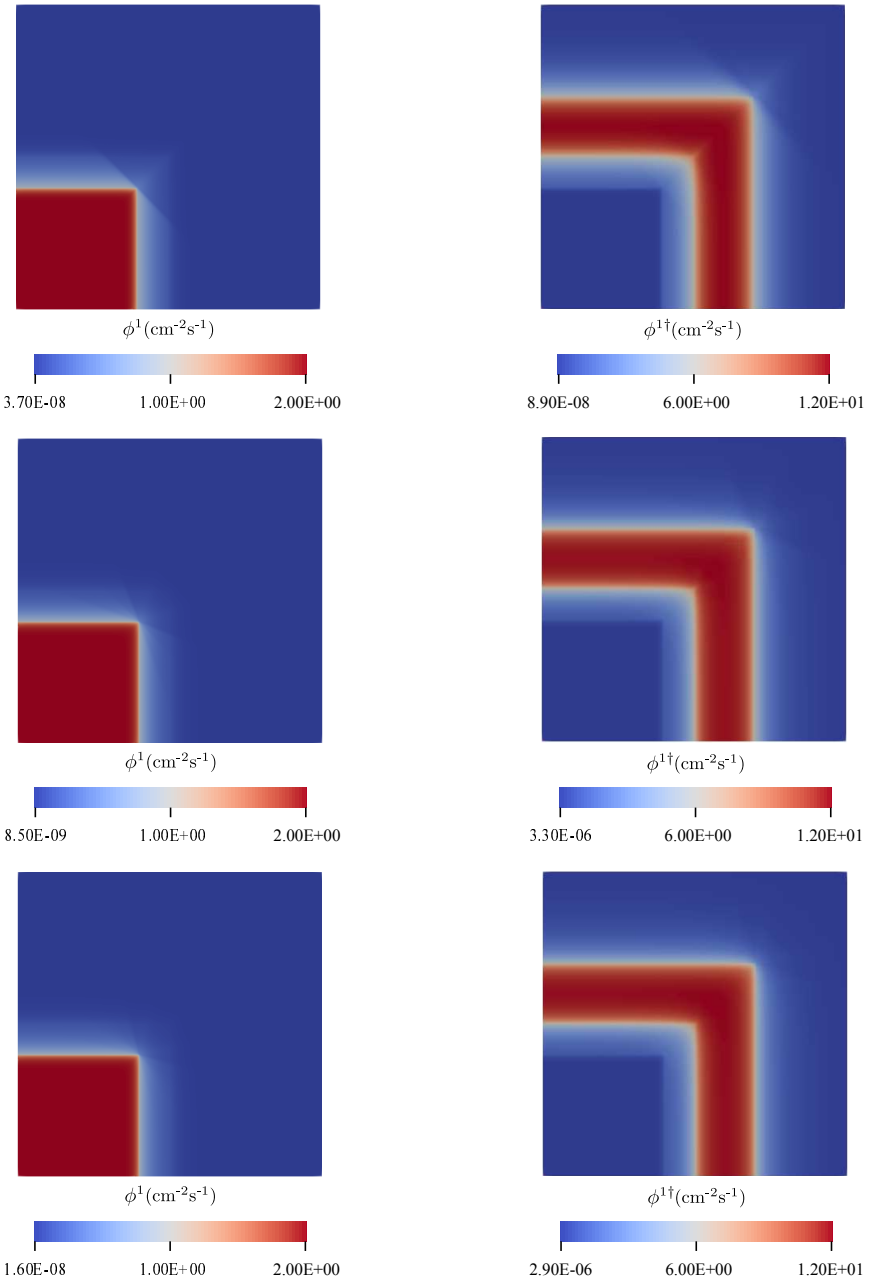
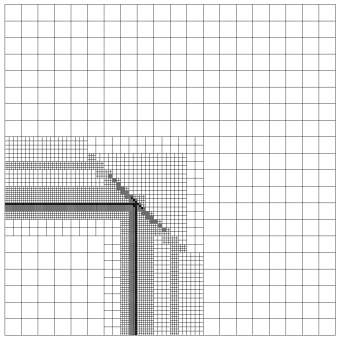


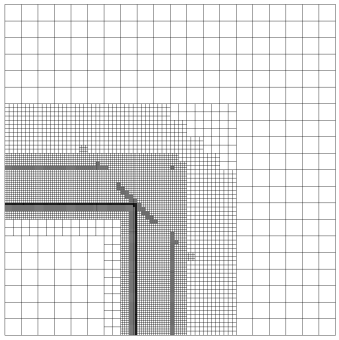
Figure 6.9: The spatial distribution of the mono-group component of the scalar neutron flux (left) and the scalar neutron importance (right) for the 1G FONT-S_N equations for Azmy's 2D radiation shielding problem; for the S_2 (top), S_4 (middle) and S_8 (bottom) LS angular quadrature sets; the QoI is the absorption in Material 3; v. Figure 6.8. (*Vide* the web-based version for reference to colour.)

tively, over a mesh of 104,000 bi-quadratic elements [216]. The spatial distribution of the mono-group component of both the scalar neutron flux and the scalar neutron importance for the S_2, S_4 and S_8 approximations with the LS angular quadrature sets are presented in Figure 6.9, for the QoI chosen as the absorption in Material 3; v. Figure 6.8.

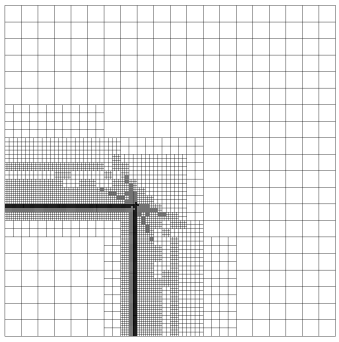
From the meshes generated after the 9th AMR iteration for each of the sweep-based DG-IGA AMR strategies for an increasingly large order of the LS angular quadrature set, $N = \{2, 4, 8\}$, as presented in Figure 6.10 and Figure 6.11, one observes the coupling between the spatial discretisation and the angular one; specifically, the effect of varying N on the AMR strategies, albeit less markedly so than



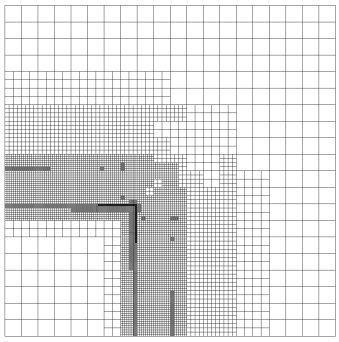
(a) $\mathcal{T}^{(9)}$ for S_2: $62,568$ DoF



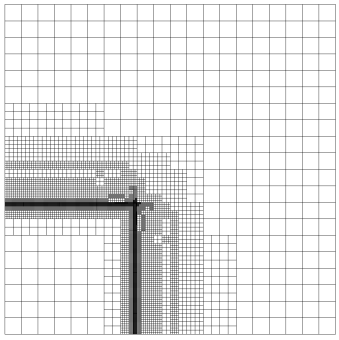
(b) $\mathcal{T}^{(9)}$ for S_2: $88,812$ DoF



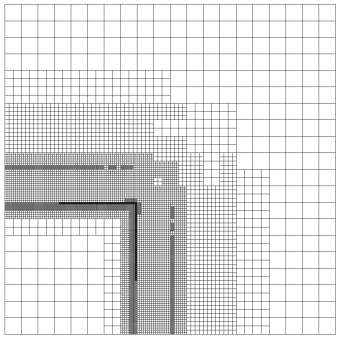
(c) $\mathcal{T}^{(9)}$ for S_4: $81,252$ DoF



(d) $\mathcal{T}^{(9)}$ for S_4: $70,992$ DoF



(e) $\mathcal{T}^{(9)}$ for S_8: $90,432$ DoF

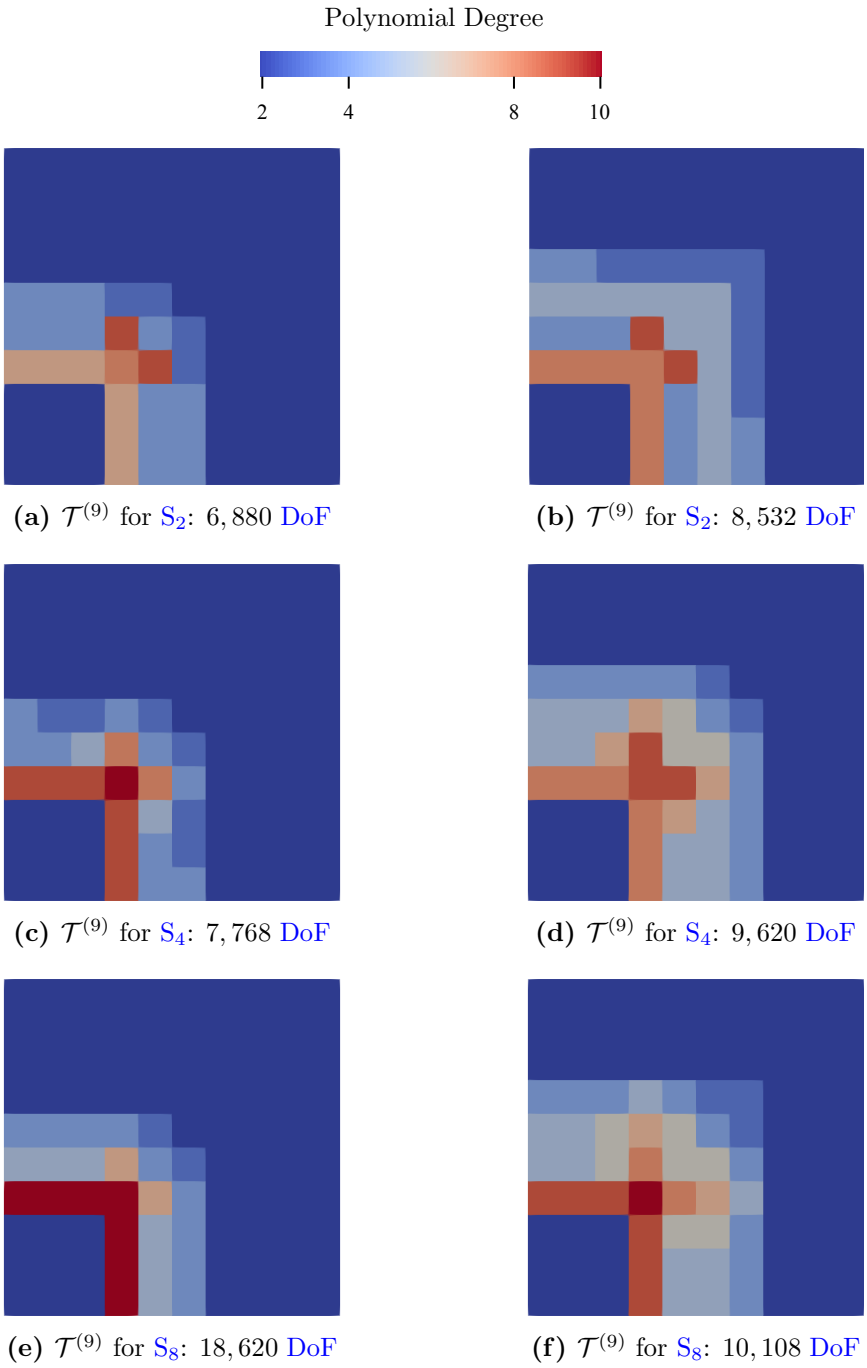


(f) $\mathcal{T}^{(9)}$ for S_8: $82,116$ DoF

Figure 6.10: The energy-independent meshes generated at the 9th AMR iteration for the NRG-AMR-h (left) and the DWR-AMR-h (right) strategies for the sweep-based DG-IGA 1G FONT-S_2 (top), FONT-S_4 (middle) and FONT-S_8 equations for Azmy's 2D radiation shielding problem; the QoI is the absorption in Material 3; v. Figure 6.8. (*Vide* the web-based version for reference to colour.)

for Ackroyd's 2D small reactor problem; q.v. Section 5.5.3.

Both AMR-h strategies perform better than the uniform h-refinement strategies; v. Figure 6.12. However, the NRG-AMR strategies focus on resolving the sharp gradient in the solution at the interface between Material 1 and Material 2; and thus they appear to yield only limited additional accuracy in the QoI before plateauing; whereas the DWR-AMR strategies expressly refine the region of interest and thus, consistently reduce the error in the goal at each AMR iteration.



(a) $\mathcal{T}^{(9)}$ for S_2: $6,880$ DoF **(b)** $\mathcal{T}^{(9)}$ for S_2: $8,532$ DoF

(c) $\mathcal{T}^{(9)}$ for S_4: $7,768$ DoF **(d)** $\mathcal{T}^{(9)}$ for S_4: $9,620$ DoF

(e) $\mathcal{T}^{(9)}$ for S_8: $18,620$ DoF **(f)** $\mathcal{T}^{(9)}$ for S_8: $10,108$ DoF

Figure 6.11: The energy-independent meshes generated at the 9th AMR iteration for the NRG-AMR-*hp* (left) and the DWR-AMR-*hp* (right) strategies for the sweep-based DG-IGA 1G FONT-S_2 (top), FONT-S_4 (middle) and FONT-S_8 equations for Azmy's 2D radiation shielding problem; the QoI is the absorption in Material 3; v. Figure 6.8. (*Vide* the web-based version for reference to colour.)

6.5.3 A 2D CANDU-6 Annular Cell Test Case

This 2D CANDU-6 annular cell test case forms the first part of the DRAGON test case TCWU06, which is representative of a typical CANDU-6 fuel bundle [361]. This problem presents *seven* physically distinct regions, including a central coolant Zr-1 tube of *four* concentric rings of naturally enriched UO_2 fuel pins, each one encased in a cladding of Zr-3, and the remaining space filled by heavy water (D_2O), all of which is enclosed in a Zr-2 tube filled with Helium and placed in an outer-D_2O

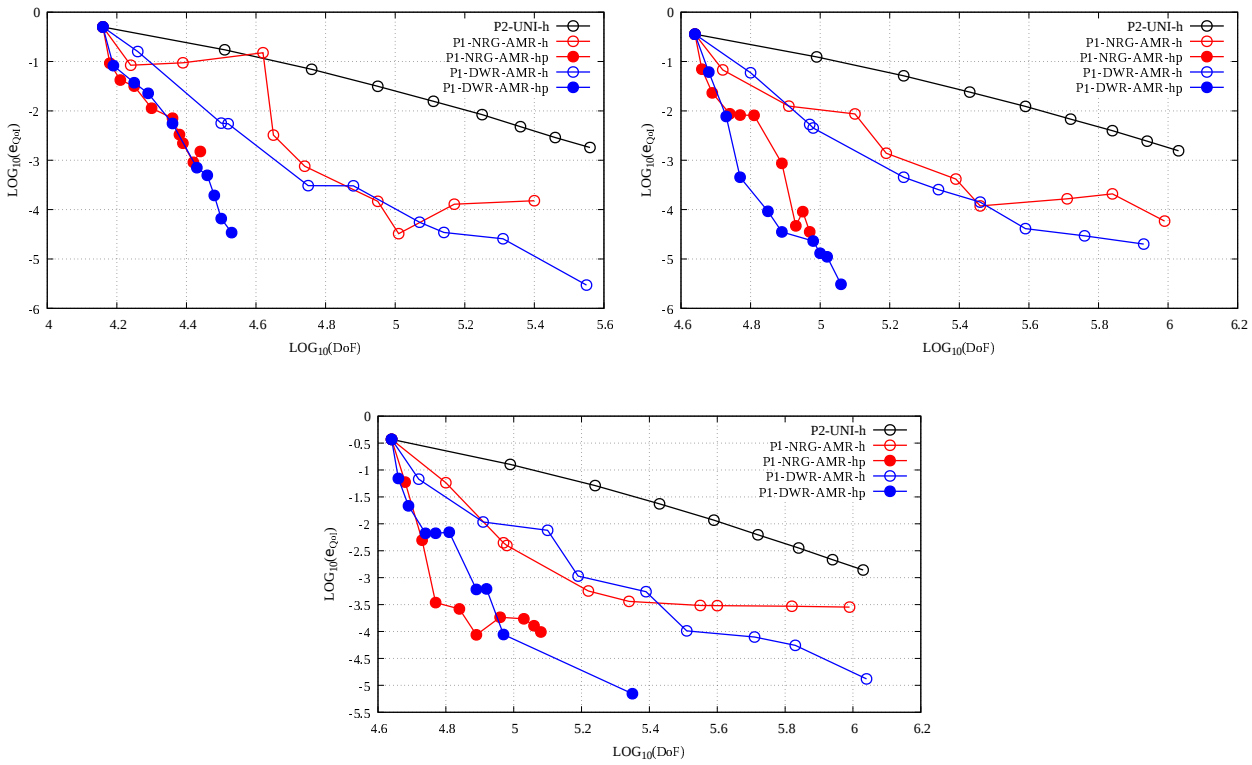


Figure 6.12: The convergence plots, ϵ_{QoI} v DoF, for the global uniform refinement (black) and the NRG-AMR (red) and the DWR-AMR (blue) strategies for the sweep-based DG-IGA 1G FONT-S_2 (top left), FONT-S_4 (top right) and FONT-S_8 (bottom) equations for Azmy's 2D radiation shielding problem; the QoI is the absorption in Material 3; v. Figure 6.8. (*Vide* the web-based version for reference to colour.)

square moderator region of congruent sides of length 2.8575E+01 cm; v. Figure 6.13. Again, for the two-group condensation of the 172-group WLUP nuclear data library at the delimitation of 4 eV, this problem was modelled by the DRAGON code [361]. These few-group macroscopic cross-section data are reproduced in Table 6.4 [362]. To model an infinite lattice, reflective BC are specified at every external physical boundary; and as per the *one-quarter* symmetry in the problem, reflective BC are also prescribed at $x = y = 0.00$ cm. The coarsest mesh is constructed from 259 bi-quadratic NURBS patches. A reference solution is obtained from the DRAGON code, as provided by Yang and Satvat; viz. $k_{eff} = 1.112837$ [362]. The spatial distribution of the multi-group components of both the scalar neutron flux and the scalar neutron importance for the FONT-S_6 equations with a GC angular quadrature set are presented in Figure 6.14, for the QoI chosen as the k_{eff}.

The meshes generated after the 9th AMR iteration for each of the sweep-based DG-IGA AMR-*h* and AMR-*hp* strategies are presented in Figure 6.15 and Figure 6.16, respectively. From the convergence plots in Figure 6.17, both of the AMR-$h(p)$ strategies perform better than the uniform-*h* refinement one; and for the use of energy-independent meshes, they allocate a similar number of DoF to the final mesh; yet, the NRG-AMR strategies appear only to reduce sporadically the error in the QoI, whereas the DWR-AMR strategies consistently reduce this error with each subsequent AMR iteration.

For the use of energy-dependent meshes, the DWR-AMR-$h(p)$ strategies benefit from a better alloca-

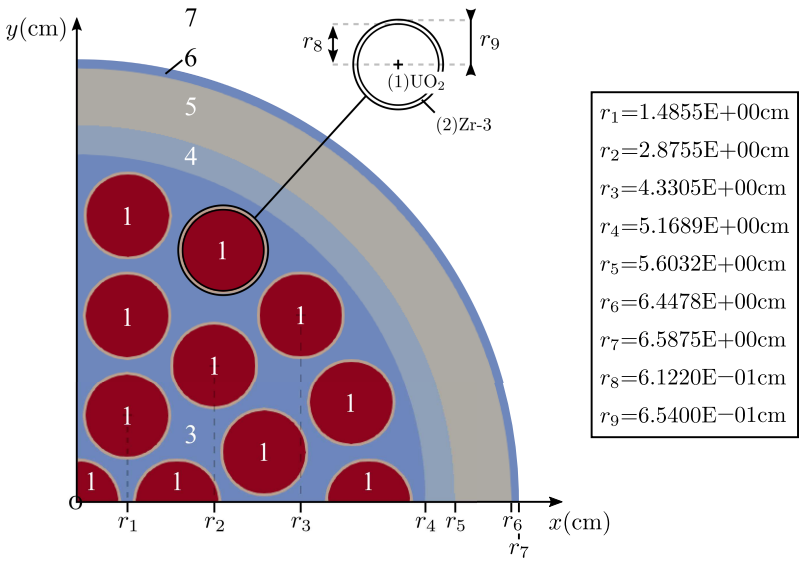


Figure 6.13: The geometry for a 2D 2G CANDU-6 annular cell test case based upon the first part of the DRAGON test case TCWU06; the constituent materials are denoted by the numbers in parentheses: Uranium oxide (1), Zirconium-3 (2), inner-Deutrium (3), Zirconium-1 (4), Helium (5), Zirconium-2 (6) and outer-Deutrium (7) [361]. (*Vide* the web-based version for reference to colour.)

Material	g	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to1}$ (cm^{-1})	$\Sigma_{s0}^{g\to2}$ (cm^{-1})	$\nu^g\Sigma_f^g$ (cm^{-1})
(1) UO_2	1	3.73334E−01	3.55414E−01	5.65522E−04	9.45470E−03
	2	5.05007E−01	0.00000E+00	4.00902E−01	1.36271E−01
(2) Zr-3	1	2.83817E−01	2.82283E−01	1.58650E−04	0.00000E+00
	2	2.80823E−01	0.00000E+00	2.75008E−01	0.00000E+00
(3) D_2O (inner)	1	2.36057E−01	2.32571E−01	3.47825E−03	0.00000E+00
	2	3.46547E−01	0.00000E+00	3.46449E−01	0.00000E+00
(4) Zr-1	1	3.02340E−01	2.99881E−01	2.33735E−04	0.00000E+00
	2	2.85758E−01	0.00000E+00	2.79187E−01	0.00000E+00
(5) He	1	3.05524E−04	3.02747E−04	2.77636E−06	0.00000E+00
	2	1.71235E−04	0.00000E+00	1.71235E−04	0.00000E+00
(6) Zr-2	1	3.04423E−01	3.02281E−01	2.62590E−04	0.00000E+00
	2	2.81233E−01	0.00000E+00	2.74710E−01	0.00000E+00
(7) D_2O (outer)	1	3.33045E−01	3.20134E−01	1.29081E−02	0.00000E+00
	2	4.55438E−01	0.00000E+00	4.55395E−01	0.00000E+00

Table 6.4: The nuclear data for a 2D 2G CANDU-6 annular cell test case based upon the first part of the DRAGON test case TCWU06 [361]. The macroscopic cross-sections for fission in Material 1 are Σ_f^g = 3.51468E−03 cm^{-1} and Σ_f^g = 5.58899E−02 cm^{-1}; and the fission spectrum is $\chi = [1.00, 0.00]^T$.

tion of resources; conversely, this is not the case for the NRG-AMR-$h(p)$ strategies. Indeed, both of the AMR-$h(p)$ strategies generate qualitatively similar $\mathcal{T}^1$ because pronounced variations in the spatial distribution of the fast-component of the scalar neutron flux are exhibited in those regions that correspond to the largest absolute values of ϕ^1. Accordingly, the corresponding elements are flagged for further refinement since they present larger errors than those elements in which pronounced variations in the spatial distribution of the thermal-component of the scalar neutron flux are exhibited because the smaller absolute values of ϕ^2 incur smaller errors by comparison. As a result, for the use of energy-dependent meshes, $\mathcal{T}^2$ is left under-refined by the NRG-AMR strategies; viz. approximately

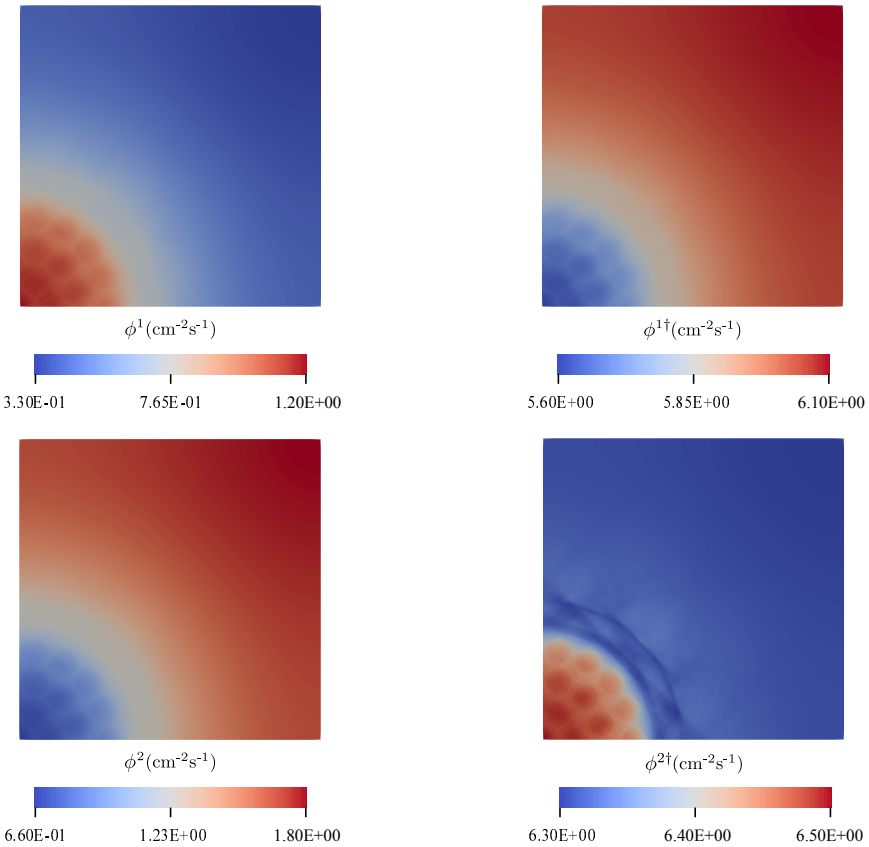


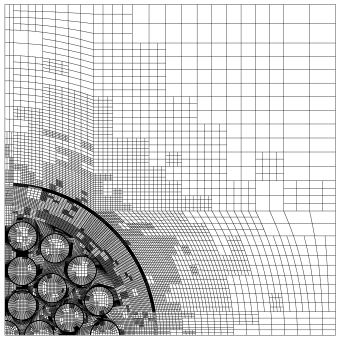
Figure 6.14: The spatial distribution of the multi-group components of the scalar neutron flux (left) and the scalar neutron importance (right) for the 2G FONT-S_6 equations for a 2D CANDU-6 annular cell test case based upon the first part of the DRAGON test case TCWU06 [361]. (*Vide* the web-based version for reference to colour.)

half as many DoF are allocated to $\mathcal{T}^2$ as are allocated by the complementary DWR-AMR strategies.

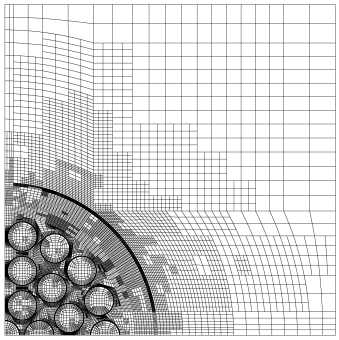
6.6 A Discussion

This chapter saw the application of automised NRG-AMR-$h(p)$ and DWR-AMR-$h(p)$ strategies to a sweep-based DG-IGA spatial discretisation of the multi-group FONT-S_N equations. Due to the nature of hyperbolic systems, CBG discretisations of the FONT-S_N equations are unstable for heterogenous configurations; rather, they are more amenable to upwinding schemes that *follow the neutron*.

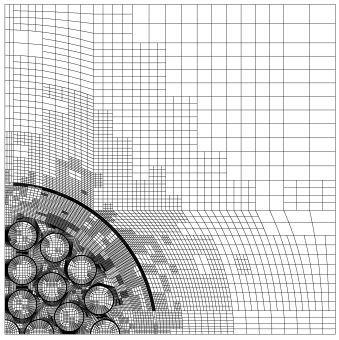
As per an upwinding scheme, the spatial dependence is resolved locally, element by element, since the locally-formulated depends only upon those values upwind. The coupling between elements is achieved solely through integrals over their interfaces; and continuity is weakly enforced in such a manner that it is not dissimilar to a least-squares penalisation for a particular choice of penalty parameter. For a given group-ordinate system, any dependencies between the elements belonging to the same mesh can be expressed as a digraph. In the absence of cycles, the elements can be ordered as a globally-assembled block lower-triangular matrix, which can be inverted locally and thus, obviating the need for



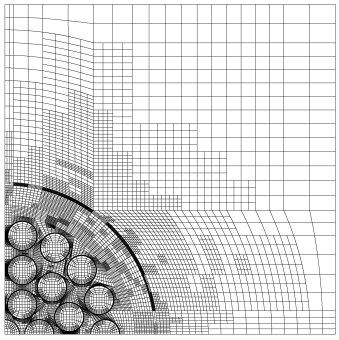
(a) $\mathcal{T}^{(9)}$: 262,336 DoF (×2)



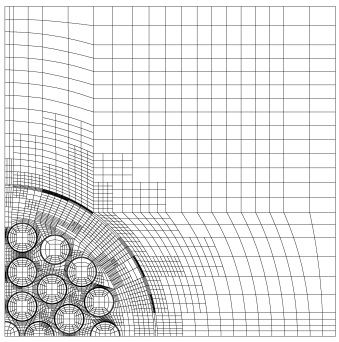
(b) $\mathcal{T}^{(9)}$: 264,016 DoF (×2)



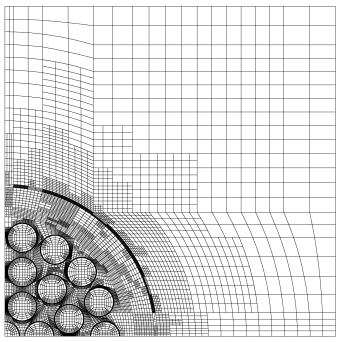
(c) $\mathcal{T}^{1(9)}$: 185,536 DoF



(d) $\mathcal{T}^{1(9)}$: 125,248 DoF



(e) $\mathcal{T}^{2(9)}$: 54,208 DoF

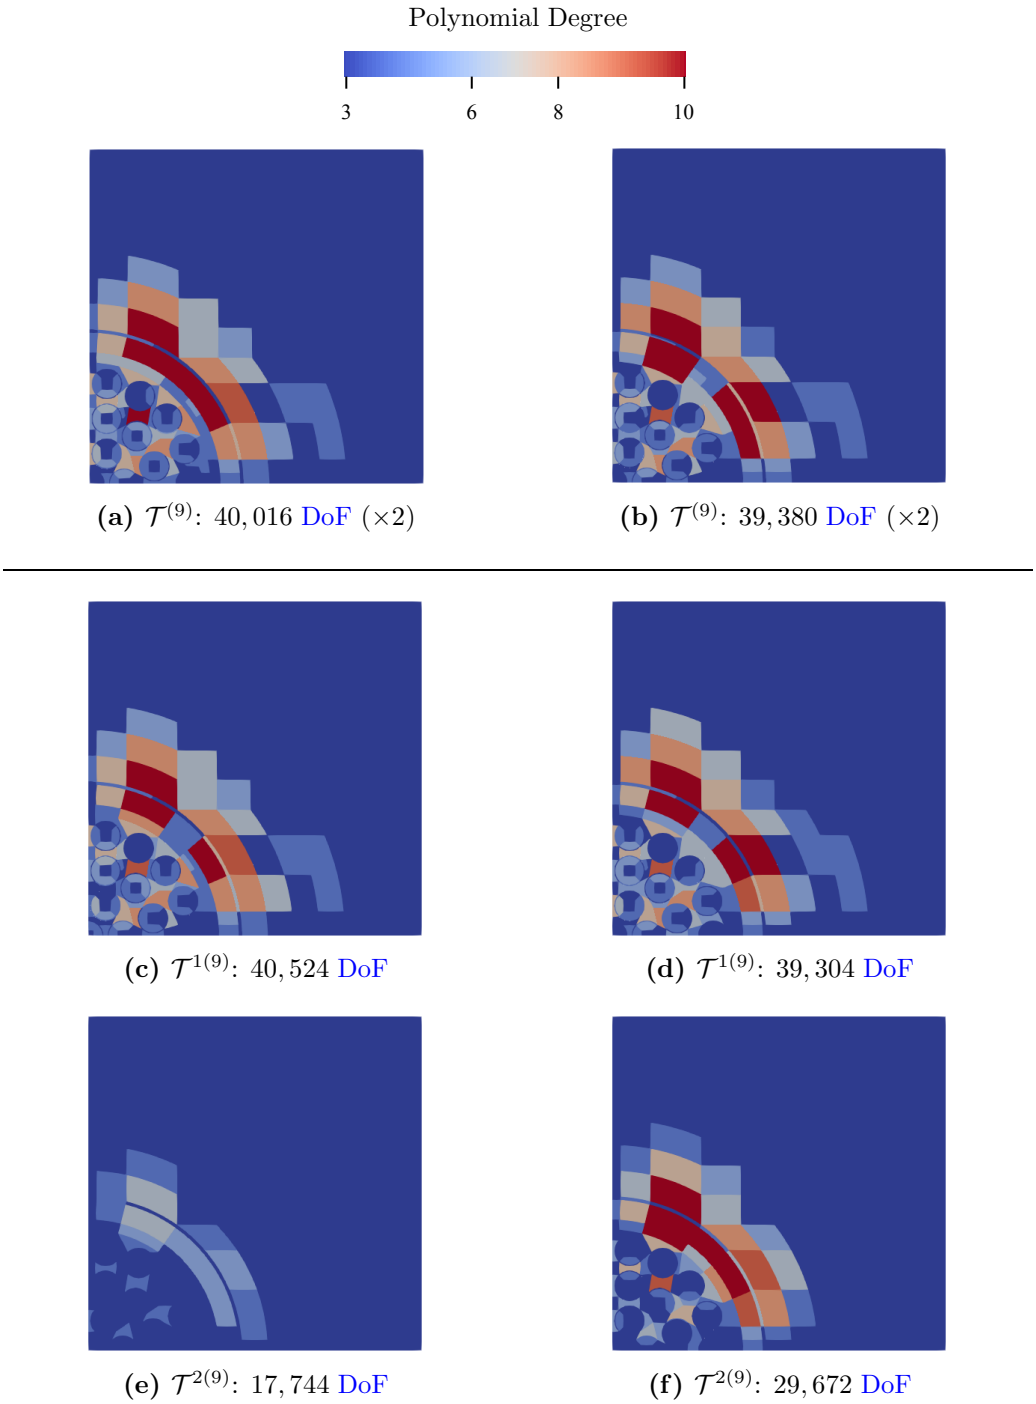


(f) $\mathcal{T}^{2(9)}$: 118,336 DoF

Figure 6.15: The energy-independent meshes (top) and the energy-dependent meshes (middle and bottom) generated at the 9th AMR iteration for the NRG-AMR-h (left) and the DWR-AMR-h (right) strategies for the sweep-based DG-IGA 2G FONT-S_6 equations for a 2D CANDU-6 annular cell test case; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

any global assembly. Otherwise, the dependencies of the SCC must be resolved in order to determine a unique sweeping order; e.g. for curved elements with reentrant boundaries. A systematic approach of deleting the weakest *outgoing* edges was adopted; any any unresolved dependencies were simply lagged within the inner-iteration procedure. As a result, for a given angular direction, the elements in a given transport sweep are visited only once.

Again, as expected, mesh refinement tended to be dominated by the ray effects for low-order angular quadrature sets; cf. Figure 2.2. For localised QoI, the DWR-AMR strategies consistently and reliably reduced the error in the QoI at each subsequent AMR iteration; whereas the NRG-AMR algorithms



(a) $\mathcal{T}^{(9)}$: $40,016$ DoF $(\times 2)$ **(b)** $\mathcal{T}^{(9)}$: $39,380$ DoF $(\times 2)$

(c) $\mathcal{T}^{1(9)}$: $40,524$ DoF **(d)** $\mathcal{T}^{1(9)}$: $39,304$ DoF

(e) $\mathcal{T}^{2(9)}$: $17,744$ DoF **(f)** $\mathcal{T}^{2(9)}$: $29,672$ DoF

Figure 6.16: The energy-independent meshes (top) and the energy-dependent meshes (middle and bottom) generated at the 9th AMR iteration for the NRG-AMR-*hp* (left) and the DWR-AMR-*hp* (right) strategies for the sweep-based DG-IGA 2G FONT-S_6 equations for a 2D CANDU-6 annular cell test case; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

would do so only sporadically. Even for global quantities, such as the k_{eff}, the latter tended to become increasingly less performant; and especially so for the use of energy-dependent meshes, since certain multi-group components were left under-resolved. Nevertheless, for the test cases considered, each of the AMR-*h* strategies was always more performant than the global uniform one. Again, the AMR-*hp* strategies yielded the best accuracy per DoF; yet, when compared to similar algorithms for the SAAF-S_N equations, there was little additional computational effort demanded at each subsequent

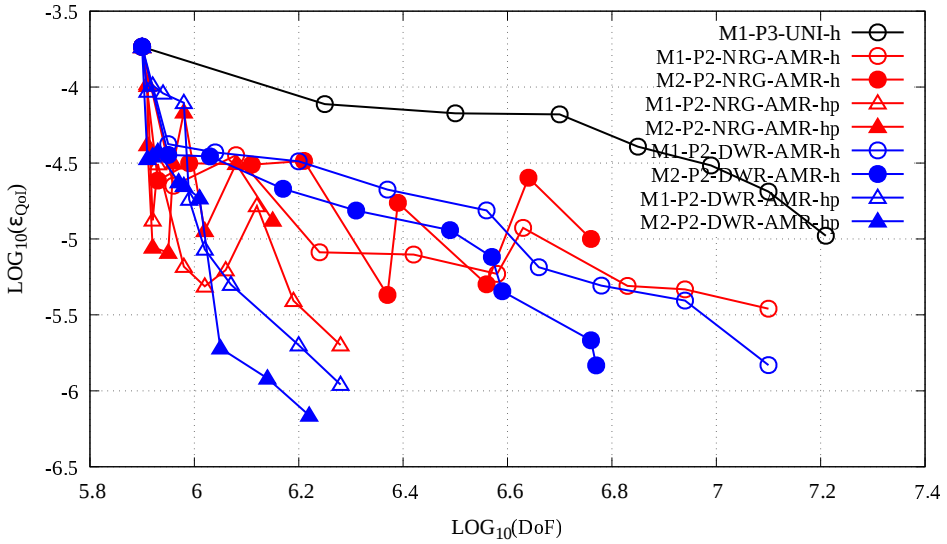


Figure 6.17: The convergence plots, ϵ_{QoI} v DoF, for the uniform refinement (black), the NRG-AMR (red) and the DWR-AMR (blue) strategies for the sweep-based DG-IGA 2G FONT-S_6 equations for a 2D CANDU-6 annular cell test case; the QoI is the k_{eff}. (*Vide* the web-based version for reference to colour.)

AMR iteration, for there was no dramatic increase in the number of elements to sweep through.

Chapter 7

Acceleration Techniques

In the application to those IGA spatial discretisations of the first and second-order forms of the NTE presented in the foregoing chapters, this current one introduces some acceleration techniques; namely, the GMRes method and diffusion *preconditioning* schemes.

7.1 An Introduction

The previous chapters saw the discretisation of different forms of the NTE with respect to each of their independent variables; in each case, the result was a system of linear algebraic equations in the form of Equation 1.2. Due to their size, direct solution methods are impractical and thus, iterative ones were employed. It is worth (re-)emphasising the distinction made between the general *transport solver* that resolves the discretised system in a global sense and the specific approach adopted in the transport sweep to invert the individual group-ordinate matrix equations; qq.v. Section 2.8 and Section 2.8.1, respectively. Therefore, to the title of this chapter, one adds the caveat that, regardless of the form of the discretised field equations, any *acceleration* is understood as that with respect to the SI procedure; viz. any speed up of the convergence of the numerical solution that may precipitate from either the algebraic manipulation or the *preconditioning* of the global system.

The SI procedure is the simplest of stationary, or fixed-point, iterative methods for linear systems; viz. a new iteration of the neutron flux is obtained from the previous one by the application of a linear operator constructed from the actions performed in a transport sweep. This is equivalent to a Richardson iteration scheme; q.v. Section 7.2.1.1 [371]. For non-multiplying systems, a neutron will undergo a finite number of collisions in its lifetime; and thus, the SI procedure is guaranteed to converge. For each of the within-group equations, the spectral radius, ρ, of the iteration matrix determines the convergence behaviour of the inner-iteration procedure, which correlates directly to

the number of collisions a neutron is likely to experience. By means of a Fourier analysis, this can be quantified in the analytical limit as the within-group scattering ratio, $c = \Sigma_{s0}^{g\to g}/\Sigma_t^g \leq 1$ [40]. This holds for both the first-order and the second-order forms of the NTE, since they are just different expressions of the same underlying physics. More generally, in the case of heterogeneous problems, the observed spectral radius is bounded from above by the largest of the within-group scattering ratios and bounded from below by the smallest; i.e. $c_{min} \leq \rho \leq c_{max}$ [374]. Thus, the rate of convergence is rapid for those optically *thin* and *leaky* systems that present a significant amount of absorption. Conversely, the SI procedure is inefficient for those optically *thick* and highly diffusive systems where within-group scattering is important and particle histories are long [55].

The largest absolute eigenvalue of the SI procedure occurs for those error modes that have weak spatial and angular dependence; viz. those spurious error modes with strong spatial and angular variation are effectively quashed by requiring the difference in the flux between successive iterations to become vanishing. Howbeit, standard tests for convergence may be inappropriate, since slowly converging methods are susceptible to problems of *false convergence*; q.v. Section 7.2.1.2 [244,245]. Accordingly, the focus of much research has been that of unconditionally effective and efficient convergence acceleration schemes in multiple dimensions; cf. the review by Adams and Larsen [40]. The techniques considered here may be applied equally to accelerate the inner and the outer iterations of a transport code; yet, within the scope of this project, one considers only those applied to the within-group equations; and thus, one is content in the limit of the *one-speed* approximation and isotropic scattering for fixed-source problems. Indeed, a multi-group formulation could benefit from any accelerated convergence of the self-scatter sources; however, it becomes increasingly more important to accelerate the group-to-group sources in the outer iteration for an increasingly fine group structure, since the within-group scatter decreases as the probability for *out-scatter* increases [80].

Where transport theory is slow to converge, diffusion theory can serve as a good low-order approximation; q.v. Section 2.6 [55]. The problem could be simplified by solving the NTE only where necessary and the NDE elsewhere [178]. A better approach would be to employ such a low-order approximation as a relatively *cheap* means of estimating the error across a transport sweep and then using it to yield a corrective term to be added to the (intermediate) iterate of the high-order transport approximation. Such is the concept introduced by Kopp for synthetic acceleration schemes, for which *one* full iteration consists of at least *two* steps [375]:

1. *One* fixed-source transport sweep of the high-order (transport) approximation; followed by
2. An additive correction from solving the low-order (diffusion) approximation.

To be effective, the added cost of solving the low-order approximation at each iteration should significantly reduce the total number thereof; and should the low-order approximation take on the form of the

NDE, the scheme is justly dubbed diffusion synthetic acceleration (DSA) [183,376]. Gelbard and Hageman derived an upper bound of $0.2247c$, for $c \in [0,1]$, for the spectral radius of the (semi-)analytical DSA scheme for the FONT-S_N equations over 2D Cartesian grids; however, this was observed only for spatially discretised systems in the fine mesh limit [377]. Thus, DSA-SI has the potential to suppress effectively those error modes that have weak spatial and angular dependence, as well as those strongly dependent ones. Reed observed that this same DSA scheme was divergent for optically thick systems, although, when convergent, it was *more performant* than Chebyshev acceleration; a necessary condition for convergence was expressed as a lower bound for the diffusion coefficient, which translates as a cell thickness of less than ~ 1 MFP [244]. For standard symmetric S_N quadrature sets for $N > 2$, the angular discretisation has only a slight effect upon the convergence properties of a DSA scheme, since its effectiveness depends upon the discrete diffusion acceleration equations, which are rotationally invariant. Conversely, however, the spatial discretisation does introduce more severe complications. Alcouffe determined that rapidly convergent DSA schemes for any cell thickness could be devised for a *consistent* spatial discretisation of both the low-order and the high-order approximations; or rather, the discrete NDE should be derived directly from the discrete transport ones, such that the form of the former is equivalent to that of the latter for a weakly anisotropic neutron flux distribution [376]. Whilst this presents difficulties for the FONT equation, it does not for those second-order forms, since both the low-order and high-order approximations would constitute elliptic systems that could enjoy the same spatial discretisation; qq.v. Chapter 4 and Chapter 5, respectively.

In truth, Alcouffe's DSA scheme for the DD-S_N equations was only *partially consistent*; yet, it was robust enough to outperform the coarse-mesh rebalance method; especially for optically thick systems dominated by scattering [183]. Fully-consistent discretisations of the low-order acceleration equations can be derived from the exact P_1 equations, for which the resulting DSA scheme is unconditionally effective for homogeneous problems of any mesh thickness [378,379]. However, this becomes difficult to achieve for general problems in multiple dimensions and computationally inefficient to solve; besides, consistency alone is not sufficient to guarantee stability. Thus, there is the motivation for partially-consistent DSA schemes, such that any loss of effectiveness due to a lack of consistency may be compensated by gains in computational efficiency and overall simplicity. Moreover, for their Fourier analysis of a fully-consistent DSA scheme for the FONT-S_N equations over an unstructured mesh of tetrahedral elements, Warsa et al observed degraded effectiveness in the presence of severe material heterogeneity; and markedly so for $c \to 1$ in at least *one* region [380]. This appears to be an inherent shortcoming of the DSA scheme, one that can be observed by numerical experiments in the analytical limit of fine meshes, irrespective of the consistency of the scheme; q.v. Section 7.4.1 [381].

Where the SI procedure is equivalent to a Richardson iteration, the use of DSA is mathematically equivalent to a *diffusion preconditioned* Richardson iteration; q.v. Section 7.3.1 [382]. A precondi-

tioner is applied with the aim of reducing the spectral radius of the iteration scheme; and importantly, the converged solution must satisfy the original matrix equation, irregardless of the low-order *preconditioner* [105]. The SI procedure, however, is a rather basic one-step iterative method; and instead, one could expedite convergence by employing a more sophisticated class of iterative solvers, such as the Krylov subspace methods or projection methods; e.g. the PCG method and the GMRes method [383,384]. In the general case, the discretisation of the NTE considered here do not yield symmetric linear systems of equations and thus, preclude the use of PCG methods. A rather attractive feature of the GMRes method, however, is that it can be implemented with relative ease by *wrapping* it around any existing code written for the SI procedure; the Krylov subspace is built from the action of inverting the *leakage-and-collision* operator in a transport sweep [385]. Since these algebraic iterative methods also have matrix representations, they too can be preconditioned [386]. Indeed, partially-consistent DSA schemes can serve as effective preconditioners for the GMRes method; and especially so for those problems that present strong material heterogeneity [380].

The rest of the chapter is organised as follows: First, the SI procedure is re-expressed as Richardson iteration in standard operator notation and then the GMRes method is presented as an alternative more-powerful transport solver, each of which is preconditioned by partially-consistent DSA schemes; viz. the same approximation space spanned by NURBS is employed in the spatial discretisation of both the high-order transport approximation and the low-order diffusion one over matching meshes. Whilst the full derivation is omitted for brevity, the acceleration of the sweep-based DG-IGA FONT-S_N equations is based upon the MIP-DG-IGA DSA scheme proposed by Wang and Ragusa; which will also be employed to accelerate the SIP-DG-IGA SAAF-S_N equations [387]. Differently, a CBG-IGA DSA scheme is applied to the CBG-IGA SAAF-S_N equations; one recalls that this is a special case of a DG discretisation for globally continuous functions with vanishing jumps. Finally, numerical results are presented for a couple of nuclear benchmarks to test the effectiveness of the acceleration techniques for the different discretisations of the different forms of the NTE for the analysis performed over exact representations of rational geometries and heterogeneous configurations; importantly, one demonstrates compatibility with the DWR-AMR-$h(p)$ strategy developed in Section 3.6.2.

7.2 Iterative Transport Solvers

The general strategy to solve the discretised NTE of the form of Equation 1.2 is based upon an operator splitting; viz. the transport operator is split as per Equation 6.5:

$$\underline{\underline{A}} = \underline{\underline{L}} - \underline{\underline{S}} \cdot \underline{\underline{\Omega}}. \tag{7.1}$$

The matrix operators $\underline{\underline{L}}$ and $\underline{\underline{S}} \cdot \underline{\underline{\Omega}}$ are the discrete approximations of the (semi-)analytical leakage-and-collision and scattering operators; qq.v. Equation 6.22a and Equation 6.22b, respectively; the latter accounts for both the mapping of angular fluxes onto scalar moments, $\underline{\Phi} = \underline{\underline{\Omega}} \cdot \underline{\varphi}$, and, *vice-versa*; i.e:

$$\underline{\underline{A}} \cdot \underline{\varphi} = \underline{b}; \tag{7.2a}$$

$$\implies \underline{\underline{L}} \cdot \underline{\varphi} = \underline{\underline{S}} \cdot \underline{\underline{\Omega}} \cdot \underline{\varphi} + \underline{b}; \tag{7.2b}$$

$$\implies \underline{\varphi} = [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}} \cdot \underline{\varphi} + [\underline{\underline{L}}]^{-1} \cdot \underline{b}; \tag{7.2c}$$

$$\implies (\underline{\underline{I}} - [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}}) \cdot \underline{\varphi} = [\underline{\underline{L}}]^{-1} \cdot \underline{b}; \tag{7.2d}$$

$$\implies [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{A}} \cdot \underline{\varphi} = [\underline{\underline{L}}]^{-1} \cdot \underline{b}; \tag{7.2e}$$

where $\underline{\underline{I}}$ represents an identity matrix of appropriate dimensions; $[\underline{\underline{L}}]^{-1} \cdot \underline{b}$ represents the uncollided flux for the vector $\underline{b}$ that represents the discrete approximation of the extraneous fixed sources of neutrons; q.v. Equation 6.23; and $\underline{\underline{M}} \coloneqq [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}} = \underline{\underline{I}} - [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{A}}$ represents the iteration matrix. By linearity of the problem, the discrete system has been recast as a (left-)preconditioned matrix equation, where the action of applying the inverse of the non-singular implicit part, $[\underline{\underline{L}}]^{-1}$, is exactly that performed by a physical transport sweep; q.v. Section 2.8.1. Without loss of generality, an emphasis is placed upon manipulating the linear FONT-S_N equations, albeit for the manner in which the action of $[\underline{\underline{L}}]^{-1}$ is performed, the same solution methodologies will be applied to the SAAF-S_N equations.

7.2.1 The Source-Iteration Procedure

In the foregoing chapters, only the most basic of iterative methods had been employed to resolve the neutron flux distribution in space, angle and energy; q.v. Section 2.8. Since the focus of this chapter is 1G fixed-source problems, the SI procedure collapses into a single (inner) iteration on the self-scattering source. A fixed linear operator, $[\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}}$, is applied to the best available information for the neutron flux in order to obtain a new estimate thereof; the process is repeated until convergence is achieved with respect to some numerical tolerance. In essence, the SI procedure is a fixed-point iteration method equivalent to (one-step) Richardson iteration.

7.2.1.1 Richardson Iteration

A discrete representation of the true solution, $\underline{\varphi}$, exactly satisfies Equations 7.2; yet, practically and iteratively one seeks an approximation, $\underline{\hat{\varphi}}$, by numerical methods, where the superscript in parentheses, (i), will denote that at the ith iteration; viz. $\underline{\hat{\varphi}}^{(i)}$ is the numerical solution after i-number of iterations. The imbalance of particle conservation between the gains and the losses in the system due to any

inaccuracy in $\underline{\hat{\varphi}}$ can be quantified at the ith-iterate by the *residual*, $\underline{r}^{(i)}$; i.e. for $i \geq 0$:

$$\underline{r}^{(i)} = \underline{b} - \underline{\underline{A}} \cdot \underline{\hat{\varphi}}^{(i)}. \tag{7.3}$$

From Equation 7.2c, for $i \geq 0$, the Richardson iteration for the (left) preconditioned system is:

$$\begin{aligned}
\underline{\hat{\varphi}}^{(i+1)} &= [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}} \cdot \underline{\hat{\varphi}}^{(i)} + [\underline{\underline{L}}]^{-1} \cdot \underline{b}, && \text{(7.4a)} \\
&= [\underline{\underline{L}}]^{-1} \cdot (\underline{\underline{L}} - \underline{\underline{A}}) \cdot \underline{\hat{\varphi}}^{(i)} + [\underline{\underline{L}}]^{-1} \cdot \underline{b}, && \text{(7.4b)} \\
&= (\underline{\underline{I}} - [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{A}}) \cdot \underline{\hat{\varphi}}^{(i)} + [\underline{\underline{L}}]^{-1} \cdot \underline{b}, && \text{(7.4c)} \\
&= \underline{\hat{\varphi}}^{(i)} + [\underline{\underline{L}}]^{-1} \cdot (\underline{b} - \underline{\underline{A}} \cdot \underline{\hat{\varphi}}^{(i)}), && \text{(7.4d)} \\
&= \underline{\hat{\varphi}}^{(i)} + [\underline{\underline{L}}]^{-1} \cdot \underline{r}^{(i)}. && \text{(7.4e)}
\end{aligned}$$

At each iteration one solves a linear system with a matrix, $\underline{\underline{L}}$, which is readily invertible and computationally cheaper to effect than that of the full transport operator, $\underline{\underline{A}}$. Indeed, convergence of the iteration scheme could be attained in *one* single iteration, if it were preconditioned instead by $[\underline{\underline{A}}]^{-1}$; however, this would come at the cost of a direct solver.

Defining the discretisation error at the ith-iterate as $\underline{e}^{(i)} = \underline{\varphi} - \underline{\hat{\varphi}}^{(i)}$, the Richardson iteration can be re-expressed in terms of errors by subtracting Equation 7.4a from Equation 7.2c; i.e. for $i \geq 0$:

$$\begin{aligned}
\underline{\varphi} - \underline{\hat{\varphi}}^{(i+1)} &= [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}} \cdot \underline{\varphi} + [\underline{\underline{L}}]^{-1} \cdot \underline{b} - [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}} \cdot \underline{\hat{\varphi}}^{(i)} - [\underline{\underline{L}}]^{-1} \cdot \underline{b}, && \text{(7.5a)} \\
&= \underline{\underline{M}} \cdot (\underline{\varphi} - \underline{\hat{\varphi}}^{(i)}); && \text{(7.5b)} \\
\implies \underline{e}^{(i+1)} &= \underline{\underline{M}} \cdot \underline{e}^{(i)}, && \text{(7.5c)} \\
&= \prod_{\hat{i}=0}^{i} [\underline{\underline{M}}] \cdot \underline{e}^{(0)}, && \text{(7.5d)} \\
&= \left[\underline{\underline{M}}\right]^{i+1} \cdot \underline{e}^{(0)}; && \text{(7.5e)}
\end{aligned}$$

where $\underline{e}^{(0)}$ is the initial error vector for some initial guess, $\underline{\hat{\varphi}}^{(0)}$. Clearly, any error in the uncollided flux does not vary between iterations. Conversely, however, increasingly more collided fluxes are computed with each transport sweep and essentially, for some initial guess, the iteration matrix is the scattering operator being acted upon, a sufficient number of times, by $[\underline{\underline{L}}]^{-1}$ until convergence is attained.

A necessary and sufficient condition for convergence is that the spectral radius, or the largest absolute eigenvalue, of the iteration matrix, $\rho(\underline{\underline{M}})$, be less than unity; i.e. $\lim_{i \to \infty} \underline{e}^{(i)} \to \underline{0}$ for $\rho(\underline{\underline{M}}) < 1$. Generally, it is computationally expensive to compute the spectral radius; yet, $\rho(\underline{\underline{M}}) \leq \|\underline{\underline{M}}\|$ for every natural matrix norm, $\|\cdot\|$; thus, for any $\underline{\hat{\varphi}}^{(0)}$ and for the regular splitting of the transport operator as per Equation 7.1, the iteration scheme will converge provided that the iteration matrix is square, $\underline{\underline{I}} - \underline{\underline{M}}$ is non-singular and $\rho(\underline{\underline{M}}) \leq \|\underline{\underline{M}}\| < 1$. From Equation 7.5c, the spectral radius can be approximated

as the asymptotic reduction in error between consecutive iterations:

$$\rho(\underline{\underline{M}}) \approx \lim_{i\to\infty} \frac{\|\underline{e}^{(i+1)}\|}{\|\underline{e}^{(i)}\|}; \tag{7.6}$$

where $\|\cdot\|$ is any well-defined norm; e.g. the L_2-norm. To ensure that this estimate of the spectral radius approaches its asymptotic value, one must allow for a sufficient number of iterations to quash those high-frequency error modes that may introduce spurious numerical artefacts [388]. Since the true solution may not be known, the spectral radius is more readily approximated as the numerical spectral radius (NSR), denoted as $\hat{\rho}$. Substitution of Equation 7.3 into Equation 7.2a yields, for $i \geq 0$:

$$\underline{\underline{A}} \cdot \underline{e}^{(i)} = \underline{r}^{(i)}; \tag{7.7}$$

whence Equation 7.6 becomes, for $i > 0$:

$$\hat{\rho}(\underline{\underline{M}}) \approx \lim_{i\to\infty} \frac{\|\underline{r}^{(i)}\|}{\|\underline{r}^{(i\text{-}1)}\|}, \tag{7.8a}$$

$$\approx \lim_{i\to\infty} \frac{\|\underline{\hat{\varphi}}^{(i+1)} - \underline{\hat{\varphi}}^{(i)}\|}{\|\underline{\hat{\varphi}}^{(i)} - \underline{\hat{\varphi}}^{(i\text{-}1)}\|}; \tag{7.8b}$$

where Equation 7.4e expresses $\underline{r}$ in terms of $\underline{\hat{\varphi}}$ at consecutive iterations; and the spectral radius can be well estimated after a sufficient number of iterations.

7.2.1.2 False Convergence

When employing iterative methods the discretised system is never solved exactly; instead, one obtains an approximate solution, $\underline{\hat{\varphi}}$, to a degree of accuracy as prescribed by some user-defined convergence criterion. From Equation 7.5e one can determine the following error estimate for $i \geq 0$:

$$\|\underline{e}^{(i+1)}\| \leq \|\underline{\underline{M}}\|^{i+1} \|\underline{e}^{(0)}\|, \tag{7.9a}$$

$$\lesssim (\hat{\rho})^{i+1} \|\underline{e}^{(0)}\|. \tag{7.9b}$$

The SI procedure is rapidly convergent for $\hat{\rho} \ll 1$; yet, it becomes arbitrarily slow for scattering-dominated problems, where $\hat{\rho} \approx c \to 1$. A second error estimate can be determined from Equation 7.5c for $i \geq 0$ [40]:

$$\underline{e}^{(i+1)} = \underline{\underline{M}} \cdot (\underline{\varphi} - \underline{\hat{\varphi}}^{(i)}), \tag{7.10a}$$

$$= \underline{\underline{M}} \cdot \underline{e}^{(i+1)} + \underline{\underline{M}} \cdot (\underline{\hat{\varphi}}^{(i+1)} - \underline{\hat{\varphi}}^{(i)}); \tag{7.10b}$$

$$\implies \underline{e}^{(i+1)} = [\underline{\underline{I}} - \underline{\underline{M}}]^{\text{-}1} \cdot \underline{\underline{M}} \cdot (\underline{\hat{\varphi}}^{(i+1)} - \underline{\hat{\varphi}}^{(i)}); \tag{7.10c}$$

$$\implies \|\underline{e}^{(i+1)}\| \lesssim \lim_{i\to\infty} \frac{\hat{\rho}}{1-\hat{\rho}} \|\underline{\hat{\varphi}}^{(i+1)} - \underline{\hat{\varphi}}^{(i)}\|. \tag{7.10d}$$

A general test for convergence is that the difference in $\underline{\hat{\varphi}}$ between successive (inner) iterations be less than some numerical tolerance; τ_i; i.e. $\|\underline{\hat{\varphi}}^{(i)} - \underline{\hat{\varphi}}^{(i\text{-}1)}\| \leq \tau_i$; q.v. Section 2.8; i.e.:

$$\|\underline{e}^{(i)}\| \lesssim \frac{\hat{\rho}\,\tau_i}{1-\hat{\rho}};$$

whereby the error in the final iterate of a slowly converging SI procedure can become much greater than that prescribed by τ_i as $\hat{\rho} \to 1$.

Slow convergence is an undesirable attribute of any iterative scheme since it implies inefficiency and proposes problems of false convergence. The former can be abated by the different acceleration techniques implemented in the remainder of this chapter; whilst the latter can be remedied by the imposition of a stricter convergence criterion of $\|\underline{\hat{\varphi}}^{(i)} - \underline{\hat{\varphi}}^{(i\text{-}1)}\| \leq \tau_i(1-\rho)$; viz. $\|\underline{e}^{(i)}\| \to \tau_i$ in the final iterate as $\rho \to 1$ for a good estimate of the spectral radius; q.v. Equation 7.8.

7.2.2 Krylov Subspace Iterations

For fixed-source problems, the convergence properties of the SI procedure are determined exclusively by the spectral radius of the linear system, as determined by the scattering ratio; thus, the SI procedure is not well suited to problems with optically thick and diffusive regions. Moreover, such a passive one-step iterative method retains only the information from the previous iterate and discards the rest.

Alternatively, the Krylov subspace methods constitute a more sophisticated class of solver, whose convergence properties generally depend upon the clustering of the eigenvalue spectrum of the discretised system; and thus, even in the presence of scattering-dominated regions, they can exhibit rapid convergence; viz. the GMRes method is fast to converge when the eigenvalues are distributed away from the origin [380, 389]. These are projection methods for solving linear systems, Equation 7.2a, such that for some initial guess, $\hat{\varphi}^{(0)}$, an N-dimensional approximate solution, $\hat{\varphi}$, is sought in an affine subspace, $\hat{\varphi}^{(0)} + \hat{\mathcal{K}}_M$, by imposition of the Petrov-Galerkin orthogonality condition, $b - A\hat{\varphi} \perp \hat{\mathcal{L}}_M$; where the subspace of approximations, $\hat{\mathcal{K}}_M$, and the subspace of constraints, $\hat{\mathcal{L}}_M$, are each M-dimensional Krylov subspaces of $\mathbb{R}^N$; i.e. [384]:

$$\hat{\mathcal{K}}_M\big(A, r^{(0)}\big) = \text{span}\left\{r^{(0)}, Ar^{(0)}, \ldots, A^{M\text{-}1}r^{(0)}\right\}; \tag{7.11}$$

where $r^{(0)}$ is the initial residual, as defined by Equation 7.3 for $i = 0$. Naturally, the dimension of Krylov subspace, M, is limited by that of the discrete linear operator system, N; in practice, the former is always much smaller than the latter. Furthermore, the specific set of constraints imposed in building the approximation has important implications upon the iterative technique and the algorithms employed; choosing $\hat{\mathcal{L}}_M = \hat{\mathcal{K}}_M$ yields the *orthogonal* projection techniques; e.g. in the case of the full

orthogonalisation method (FOM), the constraints are simply the Galerkin orthogonality condition; q.v. Equation 4.71. The choice of *two* different subspaces yields the *oblique* projection techniques; e.g. the choice of $\hat{\mathcal{L}}_M = A\hat{\mathcal{K}}_M$ gives rise to the GMRes method by Saad and Schulz [390].

The approximate solution can be re-written as some correction to an initial guess, $\hat{\varphi} = \hat{\varphi}^{(0)} + \delta$, such that the orthogonality conditions can be re-expressed as $r^{(0)} - A\delta \perp \hat{\mathcal{L}}_M$; and the approximate solution may therefore be sought in the general projection step:

$$\hat{\varphi} = \hat{\varphi}^{(0)} + \delta, \qquad \delta \in \hat{\mathcal{K}}_M; \tag{7.12a}$$

$$\left(r^{(0)} - A\delta, w\right) = 0, \qquad \forall w \in \hat{\mathcal{L}}_M. \tag{7.12b}$$

Letting $\underline{\underline{V}}_M = [\underline{v}_1, \ldots, \underline{v}_M]$ and $\underline{\underline{W}}_M = [\underline{w}_1, \ldots, \underline{w}_M]$ be $(N \times M)$-matrices whose column-vectors form bases of $\hat{\mathcal{K}}_M$ and $\hat{\mathcal{L}}_M$, respectively, the projection may be represented in matrix form by first re-writing the correction to the initial guess as $\underline{\underline{V}}_M \cdot \underline{y}$, such that the orthogonality condition can then be re-expressed as a discrete system of linear equations for $\underline{y}$:

$$\begin{aligned} [\underline{\underline{W}}_M]^T \cdot \underline{\underline{A}} \cdot \underline{\underline{V}}_M \cdot \underline{y} &= [\underline{\underline{W}}_M]^T \cdot \underline{r}^{(0)}, \\ \underline{\underline{B}} \cdot \underline{y} &= [\underline{\underline{W}}_M]^T \cdot \underline{r}^{(0)}. \end{aligned} \tag{7.13a}$$

For a non-singular projected matrix $\underline{\underline{B}} \coloneqq [\underline{\underline{W}}_M]^T \cdot \underline{\underline{A}} \cdot \underline{\underline{V}}$, the approximate solution may be sought as:

$$\underline{\hat{\varphi}} = \underline{\hat{\varphi}}^{(0)} + \underline{\underline{V}}_M \cdot \underline{y}, \tag{7.14a}$$

$$= \underline{\hat{\varphi}}^{(0)} + \underline{\underline{V}}_M \cdot [\underline{\underline{B}}]^{-1} \cdot [\underline{\underline{W}}_M]^T \cdot \underline{r}^{(0)}. \tag{7.14b}$$

Most iteration techniques employ a succession of such projections, building a new pair of Krylov subspaces of increasingly large dimension at each step and taking $\underline{\hat{\varphi}}^{(0)}$ as the approximation from the previous projection.

7.2.2.1 The GMRes Method

In anticipation of the additional storage and computational overhead envisaged by a more sophisticated projection method, the order of the discrete linear system of equations, as presented in Equations 7.2, can be reduced significantly by re-expressing the NTE in terms of $\underline{\hat{\Phi}}$; i.e. by pre-multiplying Equation

7.2c by the matrix operator $\underline{\underline{\Omega}}$:

$$\underline{\underline{\Omega}} \cdot \left[\underline{\varphi} = [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}} \cdot \underline{\varphi} + [\underline{\underline{L}}]^{-1} \cdot \underline{b}\right]; \tag{7.15a}$$

$$\implies \underline{\underline{\Omega}} \cdot \left[(\underline{\underline{I}} - [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}}) \cdot \underline{\varphi} = [\underline{\underline{L}}]^{-1} \cdot \underline{b}\right], \tag{7.15b}$$

$$\left(\underline{\underline{I}} - \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}}\right) \cdot \hat{\underline{\Phi}} = \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{b}, \tag{7.15c}$$

$$\tilde{\underline{\underline{A}}} \cdot \hat{\underline{\Phi}} = \tilde{\underline{b}}. \tag{7.15d}$$

The system of equations is not generally SPD; nonetheless, the above requires only the storage of vectors large enough to hold the scalar information rather than each of the discrete angular fluxes. For those solvers based upon stationary Richardson iteration, convergence of the scalar neutron flux for fixed-source problems can be effectively accelerated by minimising the QoI over a Krylov subspace constructed as a linear combination of *all* previous iterations by *wrapping* the GMRes algorithm around those routines that perform the action of a physical transport sweep [371, 380, 385, 391, 392].

Arnoldi's Method One recalls that only a set of linearly independent vectors that span a given vector space may also form a basis thereof. To ensure linear independence of those vectors that span the Krylov subspace, Equation 7.11, one employs Arnoldi's method to find a sequence of orthonormal *Arnoldi vectors*, $\underline{v}_1^{\perp}, \underline{v}_2^{\perp}, \ldots, \underline{v}_M^{\perp}$, that form a basis of $\hat{\mathcal{K}}_M$ [393].

An existing transport solver based upon the SI procedure can be retrofitted such that Arnoldi's method wraps around those parts of the numerical code that perform the physical transport sweeps. In *one* complete transport sweep, a new flux guess is computed from an old one in an affine way and the action of applying $[\underline{\underline{L}}]^{-1}$ does not require the operator itself to be known or stored as a matrix; i.e. at the jth iteration, supplying $\hat{\underline{\Phi}} = \underline{v}_j$ to the right-hand side of Equation 7.15a and letting $\underline{b} = 0$:

$$\hat{\underline{\Phi}} = \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \hat{\underline{\Phi}} + \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \cancelto{0}{\underline{b}}; \tag{7.16a}$$

$$\implies \hat{\underline{\Phi}} = \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{v}_j. \tag{7.16b}$$

Now, the matrix-vector multiplication $\tilde{\underline{\underline{A}}} \cdot \underline{v}_j$ can be computed as the difference in the Arnoldi vector supplied across a transport sweep:

$$\underline{v}_j - \hat{\underline{\Phi}} = \underline{v}_j - \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{v}_j, \tag{7.17a}$$

$$= \left(\underline{\underline{I}} - \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}}\right) \cdot \underline{v}_j, \tag{7.17b}$$

$$= \tilde{\underline{\underline{A}}} \cdot \underline{v}_j. \tag{7.17c}$$

At the jth iteration, starting with $\underline{v}_1^{\perp} = \underline{r}^{(0)}/\beta$ with $\beta = \|\underline{r}^{(0)}\|_{L_2}$, such that $\|\underline{v}_1^{\perp}\|_{L_2} = 1$, the vector $\underline{w}_j = \tilde{\underline{\underline{A}}} \cdot \underline{v}_j^{\perp}$ can now be evaluated as $\underline{w}_j = \underline{v}_j^{\perp} - \hat{\underline{\Phi}}$, which is then orthogonalised against all other

previous Arnoldi vectors by the modified Gram-Schmidt procedure or the Householder algorithm [390]. Obviously, for non-singular matrices, the next Arnoldi vector could not be generated for $\|\underline{w}_j\| = 0$. Yet, in such a rare occurence, the projection onto the Krylov subspace would be exact and the residual vanishing; and thus, such a breakdown would be considered *lucky*.

Also produced is the upper-Hessenberg $(M{+}1 \times M)$-matrix, $\underline{\underline{\bar{H}}}_M$; where $\underline{\underline{H}}_M$ is the $(M \times M)$-matrix obtained by deleting the last row of $\underline{\underline{\bar{H}}}_M$, or rather the orthogonal projection of $\underline{\underline{\tilde{A}}}$ onto $\hat{\mathcal{K}}_M$, such that:

$$\underline{\underline{\tilde{A}}} \cdot \underline{\underline{V}}^{\perp}_M = \underline{\underline{V}}^{\perp}_M \cdot \underline{\underline{H}}_M + \underline{w}_M \cdot [\underline{e}_M]^T, \tag{7.18a}$$

$$= \underline{\underline{V}}^{\perp}_{M+1} \cdot \underline{\underline{\bar{H}}}_M; \tag{7.18b}$$

where $\underline{e}_i$ represents the ith-column of the $(N \times N)$ identity matrix, $\underline{\underline{I}}$.

The problem of minimising the residual over the Krylov subspace can now be expressed as that of computing Equation 7.14a for the M-dimensional vector $\underline{y}$ that minimises the residual such that:

$$\underline{r} = \tilde{\underline{b}} - \underline{\underline{\tilde{A}}} \cdot \hat{\underline{\Phi}}, \tag{7.19a}$$

$$\approx \tilde{\underline{b}} - \underline{\underline{\tilde{A}}} \cdot \left[\hat{\underline{\Phi}}^{(0)} - \underline{\underline{V}}^{\perp}_M \cdot \underline{y}\right], \tag{7.19b}$$

$$\approx \underline{r}^{(0)} - \underline{\underline{\tilde{A}}} \cdot \underline{\underline{V}}^{\perp}_M \cdot \underline{y}, \tag{7.19c}$$

$$\approx \underline{r}^{(0)} - \underline{\underline{V}}^{\perp}_{M+1} \cdot \underline{\underline{\bar{H}}}_M \cdot \underline{y}, \tag{7.19d}$$

$$\approx \beta \underline{v}^{\perp}_1 - \underline{\underline{V}}^{\perp}_{M+1} \cdot \underline{\underline{\bar{H}}}_M \cdot \underline{y}, \tag{7.19e}$$

$$\approx \beta \underline{\underline{V}}^{\perp}_{M+1} \cdot \underline{e}_1 - \underline{\underline{V}}^{\perp}_{M+1} \cdot \underline{\underline{\bar{H}}}_M \cdot \underline{y}, \tag{7.19f}$$

$$\approx \underline{\underline{V}}^{\perp}_{M+1} \cdot \left[\beta \underline{e}_1 - \underline{\underline{\bar{H}}}_M \cdot \underline{y}\right]; \tag{7.19g}$$

where Equation 7.18b is substituted into the above to give Equation 7.19d; and the next line is obtained from the definition of the first Arnoldi vector, $\underline{v}^{\perp}_1$; and from the definition of the vector $\underline{e}_1$, such that $\underline{v}^{\perp}_1 = \underline{\underline{V}}^{\perp}_M \cdot \underline{e}_1 = \underline{\underline{V}}^{\perp}_{M+1} \cdot \underline{e}_1$, follows Equation 7.19f.

Since the column-vectors of $\underline{\underline{V}}^{\perp}_{M+1}$ are orthonormal, the problem of finding $\hat{\underline{\Phi}}$ that minimises the residual can instead be posed as finding the minimiser, $\underline{y}$, to a different functional; i.e.:

$$\min_{\hat{\underline{\Phi}} \in \hat{\underline{\Phi}}^{(0)} + \hat{\mathcal{K}}_M} \|\tilde{\underline{b}} - \underline{\underline{\tilde{A}}} \cdot \hat{\underline{\Phi}}\|_{L_2} = \min_{\underline{y} \in \hat{\mathcal{K}}_M} \|\beta \underline{e}_1 - \underline{\underline{\bar{H}}}_M \cdot \underline{y}\|_{L_2}. \tag{7.20}$$

Most commonly, the $(M{+}1 \times M)$-least-squares problem on the right-hand side of the above is solved by first transforming the upper-Hessenberg matrix into an upper-triangular one by unitary plane rotations; the resulting system can then be resolved by backward substitution [390].

The GMRes method is guaranteed to converge in N-number of iterations; however, general conver-

gence estimates are difficult to obtain for non-stationary iteration methods. With each iteration, the dimension of the Krylov subspace increases by *one* until the norm of the (relative) residual is reduced to some tolerance; obviously, the storage of an additional N-dimensional vectors becomes increasingly expensive. Thus, practically, the GMRes method is periodically restarted every M^{th} iteration, as per the M-restarted GMRes method or the GMRes(M) method, with the most current approximate solution serving as the next initial guess; the general algorithm is presented below:

Algorithm 1: The restarted GMRes(M) method

Compute $\underline{r}^{(0)} = \tilde{\underline{b}} - \underline{\underline{\tilde{A}}} \cdot \hat{\underline{\Phi}}^{(0)}$ and $\beta = \|\underline{r}^{(0)}\|_{L_2}$;
Set $\underline{v}_1^{\perp} \coloneqq \underline{r}^{(0)}/\beta$;
for $j = 1, \ldots, M$ **do**
    Arnoldi's method
        for $i = 1, \ldots, j$ **do**
            Generate $(\underline{\underline{\bar{H}}}_M)_{ij}$;
        Generate $(\underline{\underline{\bar{H}}}_M)_{j+1,j}$;
        if $(\underline{\underline{\bar{H}}}_M)_{j+1,j} = 0$ **then**
            Set $M \coloneqq j$ and go to 12;
        Generate $\underline{v}_{j+1}^{\perp}$;
    End;
Compute $\underline{y}$ to minimise $\|\beta \underline{e}_1 - \underline{\underline{\bar{H}}}_M \cdot \underline{y}\|_{L_2}$ and $\hat{\underline{\Phi}} = \hat{\underline{\Phi}}^{(0)} + \underline{\underline{V}}_M \cdot \underline{y}$;
if $\|\vec{r}\|/\|\underline{b}\| > \tau$ **then**
    Set $\hat{\underline{\Phi}}^{(0)} \coloneqq \hat{\underline{\Phi}}$ and go to 1;
else
    Stop

Indeed, the storage requirements of the restarted GMRes(M) method can be fixed, *a priori*; however, for a non-positive $\underline{\underline{\tilde{A}}}$, the method can suffer from problems of *stagnation* and thus converge more slowly than the non-restarted algorithm. Nonetheless, since the problem is still represented as a matrix equation, standard preconditioning techniques can be employed; q.v. Section 7.3.2.

Moreover, Guthrie et al suggest that the *quality* of the Krylov subspace could be improved by performing L-number of transport sweeps within each GMRes iteration, for $L > 1$ [385]. At the cost of increased overheads, one would expect the basis vectors to contain *better* information about the higher-order components of the collided fluxes and multiply scattered sources; and thus, one could hope to reduce the dimension of the Krylov subspace, M.

7.3 Diffusion Acceleration

The SI procedure efficiently suppresses those error modes with strong spatial and angular variation; however, it can exhibit poor convergence properties for those scattering-dominated problems in the presence of optically thick regions. Thus, the SI procedure may require a large number of iterations to quash sufficiently any spurious errors introduced therewithin [55]. Therefore, one envisages an appropriate preconditioning of the linear system in the hope of obtaining one that is easier to solve. Specifically, one is interested in DSA schemes; indeed, these are analogous to the preconditioning of a matrix equation since their application has the intended consequence of reducing the spectral radius of the iteration matrix, $\rho(\underline{\underline{M}})$.

Diffusion acceleration schemes can effectively suppress those error modes with weak spatial and angular variation; and unsurprisingly so, since weak spatial and angular dependence are an inherent assumption of neutron diffusion theory; q.v. Section 2.6. Importantly, however, the effectiveness of a DSA scheme does not rely upon diffusion theory being a good approximation to transport theory; rather, it relies upon the appropriate *form* of the discretised NDE to accelerate effectively the neutron transport solutions [183, 185, 394].

Typically, the algorithm for DSA schemes sees the solution of a high-order NTE and a low-order NDE in alternating fashion; where the coupling between the two is achieved by either a correction to the diffusion coefficients or the removal cross-sections; or by a modification of the neutron diffusion sources. The former approach generally leads to the more exotic NDA schemes [395–397]. Whereas the latter gives rise to the standard linear scheme, which simply demands the application of a fixed operator representation of the diffusion acceleration equations at each iteration [244].

Morel demonstrated that DSA schemes, by conception, can effectively accelerate both the scalar neutron fluxes as well as the currents for problems where scattering is weakly anisotropic [398]. Indeed, this is sufficient for most reactor physics problems for neutral particles. However, in the interest of radiation shielding problems or modelling charged particles in the Fokker-Planck limit, the scattering tends to be more anisotropic and *forward-peaked*; in which case, it may be worth considering S_2 synthetic acceleration (S_2SA) or other transport synthetic acceleration (TSA) schemes [399, 400].

In the context of this work, one is content in the limit of linear DSA schemes that accelerate at most the zeroth-degree spherical harmonic flux moments; viz. the scalar neutron fluxes. Again, one emphasises that *acceleration* is understood as attaining the same discrete solution as per the SI procedure in just a fewer number of iterations.

7.3.1 Preconditioning the Source-Iteration Procedure

The general algorithm for a linear DSA scheme is a two-step procedure, the first of which consists of a single source iteration, or a full transport sweep of a *high-order* approximation of the NTE; e.g. the discrete SAAF-S_N equations or FONT-S_N equations. The result is an intermediate approximation of the angular neutron flux distribution; this is recited below, as per Equation 7.2b, ignoring the spatial discretisation and introducing a half-integer iteration index for $i \geq 0$:

$$L\hat{\varphi}^{(i+\frac{1}{2})} = S\Omega\hat{\varphi}^{(i)} + b. \tag{7.21}$$

One must now formulate an equation for the next full-integer iterate, such that $\hat{\varphi}^{(i+1)}$ is significantly more accurate than $\hat{\varphi}^{(i+\frac{1}{2})}$. Subtracting the above from the same equation as satisfied by the exact solution, φ, and rearranging yields:

$$\begin{aligned} L\big[\varphi - \hat{\varphi}^{(i+\frac{1}{2})}\big] &= S\Omega\big[\varphi - \hat{\varphi}^{(i+1)}\big], && (7.22a)\\ &= S\Omega\big[\varphi - \hat{\varphi}^{(i+\frac{1}{2})}\big] + S\Omega\big[\hat{\varphi}^{(i+\frac{1}{2})} - \hat{\varphi}^{(i+1)}\big]; && (7.22b)\\ \implies \big[L - S\Omega\big]\big[\varphi - \hat{\varphi}^{(i+\frac{1}{2})}\big] &= S\Omega\big[\hat{\varphi}^{(i+\frac{1}{2})} - \hat{\varphi}^{(i+1)}\big]. && (7.22c) \end{aligned}$$

The exact solution can now be expressed in terms of an additive correction, $\delta\hat{\varphi}$, on the previous iterate:

$$\varphi = \hat{\varphi}^{(i+\frac{1}{2})} + \underbrace{\big[L - S\Omega\big]^{-1} S\Omega\big[\Delta\hat{\varphi}^{(i+\frac{1}{2})}\big]}_{\delta\hat{\varphi}^{(i+\frac{1}{2})}}; \tag{7.23}$$

where $\Delta\hat{\varphi}^{(i+\frac{1}{2})} = \hat{\varphi}^{(i+\frac{1}{2})} - \hat{\varphi}^{(i)}$; and one expects $\delta\hat{\varphi}^{(i+\frac{1}{2})} \to 0$ as $i \to \infty$. At the cost of a direct solver, the above demands the inversion of the full coefficient matrix, Equation 7.1. However, it would be less expensive and thus, more efficient, to approximate this operator by some low-order approximation, such as that provided by diffusion theory; i.e. $[D]^{-1} \approx [L - S\Omega]^{-1}$, where the operator D denotes the left-hand side of Equation 4.1a, which is more readily invertible than $L - S\Omega$. Hence, the additive correction term, $\delta\hat{\varphi}$, from Equation 7.23 can be evaluated as the solution to a modified-source NDE:

$$\begin{aligned} D\delta\hat{\varphi}^{(i+1)} &= S\Omega\Delta\hat{\varphi}^{(i-\frac{1}{2})}; && (7.24a)\\ \implies \delta\hat{\varphi}^{(i+\frac{1}{2})} &= D^{-1}S\Omega\Delta\hat{\varphi}^{(i-\frac{1}{2})}. && (7.24b) \end{aligned}$$

Thus, re-expressing Equation 7.21 as per Equation 7.4c, the two-step procedure can be summarised for $i \geq 0$:

$$\begin{cases} \hat{\varphi}^{(i+\frac{1}{2})} = (I - L^{-1}A)\hat{\varphi}^{(i)} + L^{-1}b; & (7.25a)\\ \hat{\varphi}^{(i+1)} = \hat{\varphi}^{(i+\frac{1}{2})} + D^{-1}S\Omega(\hat{\varphi}^{(i+\frac{1}{2})} - \hat{\varphi}^{(i)}). & (7.25b) \end{cases}$$

In Section 7.2.1.1, the SI procedure was expressed as a Richardson iteration in terms of matrix operators; in the same vein, the DSA scheme can also be expressed as a (left-)preconditioned Richardson

iteration, although, in practice, the preconditioner need never be constructed explicitly. Eliminating the half-integer iterates by substitution of Equation 7.25a into Equation 7.25b yields:

$$\underline{\hat{\varphi}}^{(i+1)} = \left[(\underline{\underline{I}} - [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{A}}) \cdot \underline{\hat{\varphi}}^{(i)} + [\underline{\underline{L}}]^{-1} \cdot \underline{b}\right] \cdot \left[\overbrace{\underline{\underline{I}} + [\underline{\underline{D}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}}}^{\underline{\underline{P}}}\right] - [\underline{\underline{D}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}} \cdot \underline{\hat{\varphi}}^{(i)}, \tag{7.26a}$$

$$= \left[\underline{\underline{P}} - \underline{\underline{P}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{A}} + \underline{\underline{I}} - \underline{\underline{P}}\right] \cdot \underline{\hat{\varphi}}^{(i)} + \underline{\underline{P}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{b}; \tag{7.26b}$$

$$\implies \underline{\hat{\varphi}}^{(i+1)} = (\underline{\underline{I}} - \underline{\underline{P}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{A}}) \cdot \underline{\hat{\varphi}}^{(i)} + \underline{\underline{P}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{b}; \tag{7.26c}$$

where $\underline{\underline{P}} := \underline{\underline{I}} + [\underline{\underline{D}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}}$ serves as a *good* preconditioner, so long as $[\underline{\underline{D}}]^{-1} \approx [\underline{\underline{L}} - \underline{\underline{S}} \cdot \underline{\underline{\Omega}}]^{-1}$; i.e.:

$$\begin{aligned}
\underline{\underline{P}} &= \underline{\underline{I}} + [\underline{\underline{D}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}}, \\
&\approx \underline{\underline{I}} + [\underline{\underline{L}} - \underline{\underline{S}} \cdot \underline{\underline{\Omega}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\underline{\Omega}}, \\
&\approx [\underline{\underline{L}} - \underline{\underline{S}} \cdot \underline{\underline{\Omega}}]^{-1} \cdot \left([\underline{\underline{L}} - \underline{\underline{S}} \cdot \underline{\underline{\Omega}}] + \underline{\underline{S}} \cdot \underline{\underline{\Omega}}\right), \\
&\approx [\underline{\underline{A}}]^{-1} \cdot \underline{\underline{L}};
\end{aligned}$$

whence one infers that $\underline{\underline{P}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{A}} \approx \underline{\underline{I}}$; although, since Equation 7.24 is solved iteratively, the preconditioning could never be *exact*. Nevertheless, for an appropriate discretisation of both the high and the low-order approximations, the DSA-preconditioned system should sufficiently suppress the maximum number of slowly converging error modes at minimal cost and without inducing other more slowly converging or unstable modes; viz. $\rho(\underline{\underline{I}} - \underline{\underline{P}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{A}}) < \rho(\underline{\underline{I}} - [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{A}})$. Undeniably, each accelerated source iteration is more expensive; howbeit, the extra work per iteration is expected to be compensated by an a reduction in the total number of iterations required to attain convergence.

7.3.1.1 Consistency

One recalls that for a standard symmetric S_N quadrature set for $N > 2$, the angular discretisation has only a slight effect upon the convergence properties of a DSA scheme, since its effectiveness depends upon rotationally invariant discrete diffusion acceleration equations. Conversely, however, the spatial discretisation does introduce some more severe complications. The ensuing discussion pertains chiefly to the FONT-S_N equations because the SAAF-S_N equations are diffusion-like by nature and for a diffusion acceleration scheme thereof the elliptic systems of both the high-order and low-order approximations can enjoy the same spatial discretisation; qq.v. Chapter 4 and Chapter 5.

To enforce rigorously the balance at the end of each iteration, Equations 7.25, the discrete acceleration equations should be directly derived from the discrete neutron transport ones [376]. In doing so, one should expect to attain an unconditionally *effective* acceleration scheme for homogeneous problems, as predicted analytically by means of Fourier analysis [380]. Yet, inconsistencies can arise when projecting

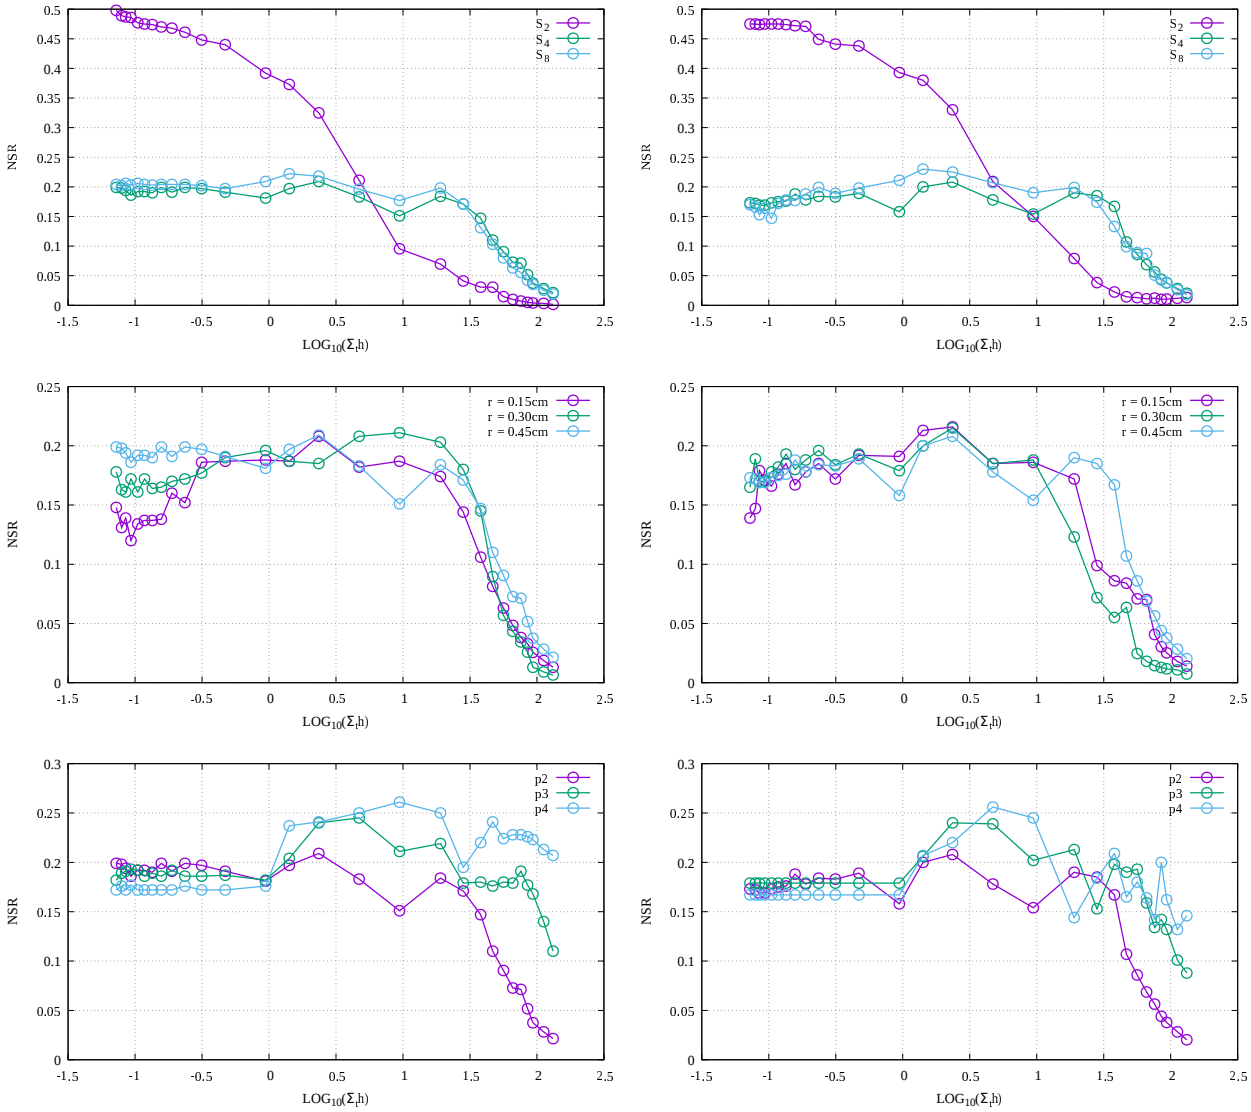


Figure 7.1: A parametric study of the NSR v optical thickness (MFP) for the CBG-IGA and MIP-DG-IGA DSA-SI schemes for the CBG-IGA (left) and the SIP-DG-IGA (right) SAAF-S_N equations, respectively, for the example of a 2D unit pin-cell mesh of an infinite homogeneous medium with $c = 0.9999$, formed of *nine* non-degenerate equal-volume bi-quadratic NURBS patches, unless otherwise stated; for the S_2, S_4 and S_8 LS quadrature sets (top); for the pin radii $r = 0.15, 0.30, 0.45$ cm for the S_4 equations (middle); and for a NURBS basis of $p = 2, 3, 4$ for the S_4 equations (bottom). (*Vide* the web-based version for reference to colour.)

the high-order equations onto a basis of spherical harmonics and then truncating the approximation as a linear anisotropic flux distribution. Consistently, these full P_1 equations can be difficult to derive and they cannot be collapsed into a single discretised NDE; moreover, they can become computationally prohibitive to solve in multiple dimensions [236]. With that in mind, the performance of a DSA scheme must be gauged equally by its *efficiency* as well as its effectivity; viz. perhaps, to the detriment of the latter, the solution to a given neutron transport problem should be computed to some specified accuracy at minimal computational effort.

Indeed, some degree of inconsistency should be largely compensated by any gains in simplicity and computation efficiency. Operating directly upon the analytical NDE and mimicking the spatial dis-

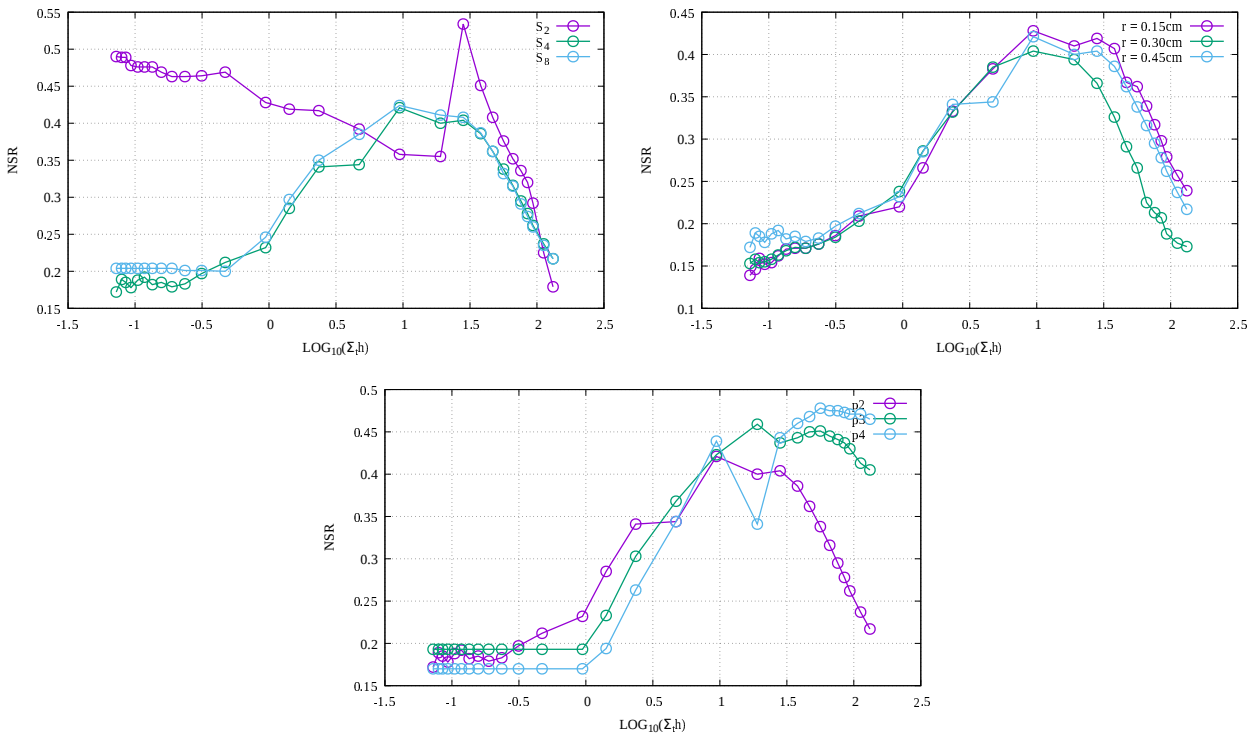


Figure 7.2: A parametric study of the NSR v optical thickness (MFP) for the MIP-DG-IGA DSA-SI scheme for the sweep-based DG-IGA FONT-S_N equations for the example of a 2D unit pin-cell mesh of an infinite homogeneous medium with $c = 0.9999$, formed of *nine* non-degenerate equal-volume bi-quadratic NURBS patches, unless otherwise stated; for the S_2, S_4 and S_8 LS quadrature sets (top left); for the pin radii $r = 0.15, 0.30, 0.45$ cm for the S_4 equations (top right); and for a NURBS basis of $p = 2, 3, 4$ for the S_4 equations (bottom). (*Vide* the web-based version for reference to colour.)

cretisation of the NTE and then employing Fick's law to eliminate all first-order flux moments, Adams and Martin modified Larsen's four-step method and developed an inconsistent, or rather partially-consistent, DSA scheme whose acceleration equations inadvertently took the form of an IIP-DG discretisation of the NDE for the constant penalisation of $\mu_f = ¼$ to eliminate the jump in the neutron current, expressed as partial currents, Equation 2.59, across all internal faces [185, 401]. This was found to be rapidly convergent in 1D slab and spherical geometries; and especially so for optically thick meshes, where the spectral radius is monotonically decreasing for an increasingly large mesh spacing. However, upon further investigation, this asymmetric modified four-step acceleration (M4S) scheme was also found to be less performant and even divergent for distorted meshes in multiple dimensions [380, 402].

Based upon an SIP-DG discretisation of the NDE, employing the same discontinuous finite elements as those employed for the discretisation of the transport approximation, Wang and Ragusa proposed the modified interior penalty (MIP) scheme for the DG discretisation of the diffusion acceleration equations; q.v. Section 4.4 [387]. Although *optimised* penalty parameters were derived rigorously for *any* element type in Section 4.4.1, they are often assigned heuristically in the literature [349]. Nevertheless, $\mu_f \propto h^{-1} \propto (\Sigma_t h)^{-1}$, where cell optical thickness is characterised by $\Sigma_t h$; viz. $\Sigma_t \to 0$

or $h \to 0$ in optically thin regions where MFP$:= 1/\Sigma_t$ becomes much larger than a cell's diameter, h; i.e. $\Sigma_t h \ll 1$; and $\Sigma_t \to \infty$ or $h \to \infty$ in optically thick regions where the MFP becomes much smaller than h; i.e. $\Sigma_t h \gg 1$; and an intermediate range covers that in between; i.e. $\Sigma_t h \approx 1$. Where the penalty parameters become vanishing and the SIP-DG DSA scheme becomes increasingly unstable for increasingly thick cells, Wang and Ragusa enforce a minimal penalisation of $\mu_f = ¼$, as *borrowed* from a consistent diffusion-conforming form (DCF) of the acceleration equations, which is not dissimilar to the M4S method that was found to be somewhat unstable for cells of intermediate thickness [387]. The stability and effectiveness of the MIP-DG DSA scheme has been demonstrated over many different grid types for the full range of cell thicknesses; and since the SIP-DG discretisation yields SPD linear systems, the acceleration equations can be solved efficiently by the CG method with multi-grid preconditioners [403–405].

For the MIP-DG-IGA discretisation of the diffusion acceleration equations for the *one-speed* FONT-S_N equations, the penalty parameters, μ_f^*, are evaluated as [387]:

$$\mu_f^* = \max\left\{\mu_f, ¼\right\}, \qquad \forall \partial V_f \in \partial V_{IP}. \tag{7.27}$$

As calculated by Equation 4.97, μ_f will sufficiently penalise any optically thin regions without being arbitrarily large. The conditioning of the diffusion acceleration equations varies linearly with the size of the penalty parameters; and thus, any optimisation of the latter will reduce the additional computational demands of each source iteration; q.v. Section 4.4.3.

In the same vein, for the MIP-DG-IGA discretisation of the diffusion acceleration equations for the *one-speed* SAAF-S_N equations, the penalty parameters, μ_f^*, are evaluated as:

$$\mu_f^* = \max\left\{\mu_f, \max_{V_e \in \bigcap_{e\prime} V_{e\prime} \ni \partial V_f} \left\{\frac{C^0_{e,f} F_e}{3\Sigma^g_{t,e}}\right\}\right\}, \qquad \forall \partial V_f \in \partial V_{IP}; \tag{7.28a}$$

$$\approx \max_{V_e \in \bigcap_{e\prime} V_{e\prime} \ni \partial V_f} \left\{D_e C^0_{e,f} F_e\right\}, \qquad \forall \partial V_f \in \partial V_{IP}. \tag{7.28b}$$

Again, μ_f are calculated as per Equation 4.97 with C^1; where C^0 and C^1 are the TIC introduced in Section 4.4.2. Although the above has been derived rather heuristically, numerical experiments do suggest that this assignment of penalty parameters does yield an effective DSA scheme; howbeit, for the sake of consistency, the diffusion acceleration equations are generally over-penalised over the faces of most regular NURBS patches, since $C^0 \gtrsim C^1$.

To verify the effect of preconditioning the different forms of the discretised within-group NTE, a small parametric study was performed on the MIP-DG-IGA DSA schemes for both the 1G DG-IGA FONT-S_N equations and the 1G SIP-DG-IGA SAAF-S_N equations; v. Figure 7.2 and the right-hand side of Figure 7.1, respectively; the penalty parameters for the diffusion acceleration equations evaluated as

per Equation 7.27 and 7.28, respectively. For comparison, a similar study was also performed on a CBG-IGA DSA scheme for the 1G CBG-IGA SAAF-S_N equations; v. the left-hand side of Figure 7.1. Although the number of iterations required to converge the within-group equations will not necessarily vary linearly with the spectral radius, the conditioning of the linear system can be quantified by estimating the latter as the NSR, Equation 7.8; to obtain good estimates of the spectral radius, a sufficiently large number of iterations is required and thus, the self-scatter sources are converged to machine precision. With regard to the sweep-based approach for solving the FONT-S_N equations, any effect of SCC on convergence is mitigated by fully resolving every dependency in each digraph in a given transport sweep; q.v. Section 6.4.1.

For a 2D unit pin-cell geometry initially formed of *nine* equal-volume bi-quadratic NURBS patches, for an infinite homogeneous medium over a large range of cell thicknesses with $c = 0.9999$, the effects of preconditioning on the NSR are investigated by varying, *ceteris paribus*, the order, N, of the LS angular quadrature set; the radius of the central pin; and the polynomial degree of the NURBS basis. Optically thick cells, $\Sigma_t h > 1$ MFP, are modelled as a fixed number of elements, $h \approx 0.47$ cm, of an increasingly large macroscopic cross-section of total collision, $2\ \text{cm}^{-1} \leq \Sigma_t \leq 280\ \text{cm}^{-1}$; whereas optically thin cells are modelled as an increasingly large number of elements of reduced diameter and fixed cross-section data, $\Sigma_t \approx 2\ \text{cm}^{-1}$.

Overall, the findings are in good agreement with the Fourier analyses performed over Cartesian grids in the literature; that is, over a large range of optical thicknesses, the proposed DSA schemes effectively reduce the spectral radius of a homogeneous problem where $c \to 1$; and in the fine-mesh limit, NSR$\to 0.5c$ for $N = 2$ and NSR$\to 0.2247c$ for $N > 2$ [40, 403]. Indeed, one would expect the angular quadrature with the least rotational variance to deviate the most from the theory on analytical operators; moreover, the effect of angular discretisation is minimal for $N > 2$. Both discretisations of the DSA schemes for the SAAF-S_N equations present similar convergence properties, NSR≤ 0.25; becoming immediately convergent for $\Sigma_t h \geq 2$ for $N \geq 2$. Most noticeably, yet expected, the effectivity of the DSA scheme for the FONT-S_N equations is markedly degraded in the intermediate range of optical thicknesses; nevertheless, NSR < 0.45 for $N > 2$ [387]. In all cases, the effect of distorting the mesh by varying the radius of the central pin appears to be minimal; and whilst relatively insensitive in the fine-mesh limit, the systems become increasingly ill conditioned for increasingly large polynomial degrees of the NURBS basis for increasingly thick cells.

Clearly, partially-consistent diffusion acceleration schemes can still exhibit good convergence properties of scattering-dominated homogeneous problems, whilst being simple to derive and implement; and, in the current case, the linear systems of diffusion acceleration equations are SPD and thus, can be solved efficiently by the CG method with multi-grid preconditioners [403–405]. As it will be demonstrated

in Section 7.4.1, however, DSA-SI schemes are not inherently robust in their application to those problems that present a high degree of material heterogeneity and where $c \to 1$; in which case, it would certainly not be worth the effort of devising one that is fully consistent [380]. Nonetheless, partially-consistent DSA schemes can still serve as effective preconditioners for the more sophisticated GMRes method; q.v. Section 7.2.2.1.

7.3.2 Preconditioning the GMRes Method

As was first identified by Azmy, any DSA-SI scheme will become increasingly ineffective for heterogeneous problems in multiple dimensions for an increasingly large scattering ratio [381]. It was observed that no discretisation of the diffusion acceleration equations could effectively accelerate the solution of such problems in the analytical limit of $\Sigma_t h \to 0$; q.v. Section 7.4.1. Thus, for more realistic problems, a given DSA-SI scheme may serve neither efficiently nor reliably; and particularly so in those instances where it would be needed the most; i.e as $c \to 1$. However, such a virtually intractable problem may be resolved by substituting the SI procedure for the more powerful GMRes method. Generally, the GMRes method alone can accelerate the numerical solution of a discrete system of the NTE, yet, it may not always be as performant as a well-performing DSA-SI scheme; and whilst otherwise ineffective, Warsa et al demonstrated through numerical experiments that a partially-consistent DSA-GMRes scheme could always serve as an effective and efficient accelerator on unstructured tetrahedral meshes [380].

The order of the discrete linear systems of equations of the two-step DSA scheme are reduced by re-expressing them in terms of $\underline{\hat{\Phi}}$; i.e. by pre-multiplying Equations 7.25 by the matrix operator $\underline{\underline{\Omega}}$:

$$\begin{cases} \underline{\hat{\Phi}}^{(i+\frac{1}{2})} = \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\hat{\Phi}}^{(i)} + \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{b}; & (7.29a) \\ \underline{\hat{\Phi}}^{(i+1)} = \underline{\hat{\Phi}}^{(i+\frac{1}{2})} + \underline{\underline{D}}^{-1} \cdot \underline{\underline{S}} \cdot \left(\underline{\hat{\Phi}}^{(i+\frac{1}{2})} - \underline{\hat{\Phi}}^{(i)}\right); & (7.29b) \end{cases}$$

where $\underline{\underline{S}}$ no longer accounts for the mapping of scalar flux moments onto the angular neutron fluxes.

Eliminating the half-integer iterates by substitution of Equation 7.29a into Equation 7.29b yields:

$$\begin{aligned} \underline{\hat{\Phi}}^{(i+1)} &= \left[\underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\hat{\Phi}}^{(i)} + \underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{b}\right] \cdot \left[\overbrace{\underline{\underline{I}} + [\underline{\underline{D}}]^{-1} \cdot \underline{\underline{S}}}^{\tilde{\underline{\underline{P}}}}\right] - [\underline{\underline{D}}]^{-1} \cdot \underline{\underline{S}} \cdot \underline{\hat{\Phi}}^{(i)}, & (7.30a) \\ &= \left[\tilde{\underline{\underline{P}}} \cdot (\underline{\underline{I}} - \tilde{\underline{\underline{A}}}) + \underline{\underline{I}} - \tilde{\underline{\underline{P}}}\right] \cdot \underline{\hat{\Phi}}^{(i)} + \tilde{\underline{\underline{P}}} \cdot \tilde{\underline{b}}, & (7.30b) \\ &= (\underline{\underline{I}} - \tilde{\underline{\underline{P}}} \cdot \tilde{\underline{\underline{A}}}) \cdot \underline{\hat{\Phi}}^{(i)} + \tilde{\underline{\underline{P}}} \cdot \tilde{\underline{b}}; & (7.30c) \\ \implies \tilde{\underline{\underline{P}}} \cdot \tilde{\underline{\underline{A}}} \cdot \underline{\hat{\Phi}} &= \tilde{\underline{\underline{P}}} \cdot \tilde{\underline{b}}; & (7.30d) \end{aligned}$$

with $\tilde{\underline{\underline{A}}}$ and $\tilde{\underline{b}}$ defined as per Equations 7.15.

The same GMRes algorithm as that presented in Section 7.2.2.1 can be employed here; now, however, Arnoldi's method must also *wrap* around those parts of the numerical code that effect the diffusion acceleration; i.e. the matrix-matrix-vector multiplication $\underline{\underline{\tilde{P}}} \cdot \underline{\underline{\tilde{A}}} \cdot \underline{v}_j$, used to build the (left-)preconditioned Krylov subspace at the jth iteration, can be computed as the difference in the Arnoldi vector, $\underline{v}_j$, supplied across a transport sweep and a diffusion acceleration correction:

$$\underline{v}_j - \underline{\hat{\Phi}} = \underline{v}_j - (\underline{\underline{I}} - \underline{\underline{\tilde{P}}} \cdot \underline{\underline{\tilde{A}}}) \cdot \underline{v}_j, \tag{7.31a}$$

$$= \underline{\underline{\tilde{P}}} \cdot \underline{\underline{\tilde{A}}} \cdot \underline{v}_j. \tag{7.31b}$$

7.4 The Numerical Results

In the application to both the 1G SAAF-S_N equations and the 1G FONT-S_N equations, this sections presents some numerical results for the AMR-$h(p)$ strategies developed in Section 3.5 and Section 3.6, employed in conjunction with the acceleration techniques introduced within this current chapter. Unless otherwise stated, all operations are performed such as to optimise the L_2-error. For a couple of verification test cases and nuclear reactor physics benchmarks, the AMR strategies are compared against global uniform ones. Any data-plots that make reference to the DoF of the discretised system only account for that number of unknowns that were solved for, for each discrete-ordinate system of the high-order NTE; viz. $DoF = N$ for each K-number of ordinates solved over an energy-independent mesh of N-number of control-variables; n.b. this does not account for the DoF of any system of low-order diffusion acceleration equations.

Both CBG-IGA and MIP-DG-IGA diffusion acceleration schemes and the GMRes method have been implemented in the same new modern-Fortran code as that presented in the foregoing chapters; and thus, the same libraries, tolerances and parameters are employed here [73]. Equally, these acceleration techniques are employed to accelerate the adjoint equations too. When applicable, *one* diffusion correction is invoked for every transport sweep; and to avoid any excessive degradation of the efficiency of a given DSA scheme, HYPRE's BoomerAMG package is employed to accelerate the iterative solution of the diffusion acceleration equations and the convergence criteria are made an order-of-magnitude less strict. With regard to the partially-consistent DSA schemes for the SAAF-S_N equations, both discretisations of the high-order and low-order approximations over matching meshes will share and thus, can reuse, the same non-zero sparsity pattern for the CRS data structures held by PETSc [249]. Where poor quality FEM meshes have been shown to degrade the performance of preconditioners, the coarsest meshes are made to be as regular as possible and formed from non-degenerate NURBS patches [406]. When applicable, the restart parameter for the restarted GMRes(M) method will be specified for each problem presented; and the overall number of iterations will not account for the

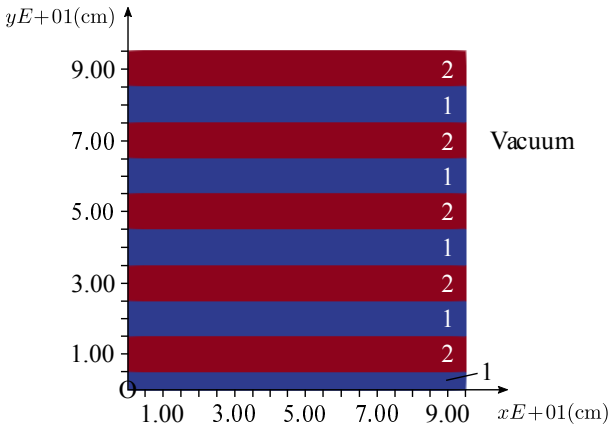


Figure 7.3: The geometry for the 2D 1G PHI test case [387]. The constituent materials are denoted by the numbers in parentheses; viz. thick material (1) and thin material (2). (*Vide* the web-based version for reference to colour.)

initial transport sweep that evaluates the uncollided scalar neutron flux, $\underline{\underline{\Omega}} \cdot [\underline{\underline{L}}]^{-1} \cdot \underline{b}$, that is used to determine the initial residual in the first line of Algorithm 1. Unless otherwise stated, only *one* transport sweep is performed per GMRes iteration.

To initiate the AMR procedure, the low-order NDE are solved to obtain an initial guess for the distribution of isotropic scatter sources; and then at each subsequent AMR iteration, the reference solution is *bootstrapped* from the previous iteration to serve as an initial guess over the current reference mesh. Again, all *jobs* are submitted to the general purpose Imperial College HPC cluster computing service (CX1).

7.4.1 A 2D PHI Test Case

Material	g	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to 1}$ (cm^{-1})	q_x^g (cm^{-3}.s^{-1})
(1) Thick	1	Σ	$c\Sigma_t^g$	1.000E+00
(2) Thin	1	$1/\Sigma$	$c\Sigma_t^g$	1.000E+00

Table 7.1: The nuclear data for a 2D 1G PHI test case [387].

	c, Scattering Ratios (-)							
Σ (cm^{-1})	9.000E−01		9.900E−01		9.990E−01		9.999E−01	
1.000E+01	0.89573	(0.37654)	0.98964	(0.50872)	0.99893	(0.52989)	0.99976	(0.53455)
2.000E+01	0.89641	(0.52188)	0.98937	(0.67447)	0.99893	(0.68789)	0.99975	(0.69073)
4.000E+01	0.89835	(0.58853)	0.98899	(0.78546)	0.99891	(0.81549)	0.99976	(0.82060)
8.000E+01	0.89909	(0.61993)	0.98932	(0.84615)	0.99885	(0.88499)	0.99972	(0.89099)
1.600E+02	0.89939	(0.63763)	0.98967	(0.88089)	0.99873	(0.92837)	0.99957	(0.93431)
3.200E+02	0.89973	(0.64929)	0.98983	(0.89869)	0.99883	(0.94946)	0.99965	(0.95719)
6.400E+02	0.89999	(0.65818)	0.98989	(0.90839)	0.99892	(0.96216)	0.99970	(0.97305)

Table 7.2: The NSR for a 2D PHI test case for the SI procedure for the CBG-IGA 1G SAAF-S_4 equations. Those results in parentheses are obtained for the CBG-IGA DSA-preconditioned system. A total of 9,025 bilinear NURBS patches were employed.

	c, Scattering Ratios (-)							
Σ (cm^{-1})	9.000E−01		9.900E−01		9.990E−01		9.999E−01	
1.000E+01	101	(10)	808	(17)	a	(21)	a	(24)
2.000E+01	104	(14)	794	(23)	a	(28)	a	(34)
4.000E+01	106	(18)	819	(33)	a	(43)	a	(51)
8.000E+01	106	(21)	862	(47)	a	(68)	a	(83)
1.600E+02	105	(24)	884	(63)	a	(105)	a	(131)
3.200E+02	104	(26)	891	(80)	a	(149)	a	(195)
6.400E+02	103	(28)	890	(95)	a	(200)	a	(292)

Table 7.3: The number of iterations to converge the self-scatter source for a 2D PHI test case for the SI procedure the CBG-IGA 1G SAAF-S_4 equations. Any number $> 5,000$ is denoted by a; and those results in parentheses are obtained for the CBG-IGA DSA-preconditioned system. A total of $9,025$ bilinear NURBS patches were employed.

	c, Scattering Ratios (-)							
Σ (cm^{-1})	9.000E−01		9.900E−01		9.990E−01		9.999E−01	
1.000E+01	24	(6)	79	(7)	314	(7)	1,051	(7)
2.000E+01	22	(7)	80	(8)	337	(9)	1,093	(8)
4.000E+01	20	(8)	76	(10)	329	(11)	1,137	(9)
8.000E+01	19	(9)	79	(11)	330	(12)	1,240	(12)
1.600E+02	20	(9)	74	(13)	331	(15)	1,361	(15)
3.200E+02	21	(9)	74	(14)	326	(15)	1,363	(17)
6.400E+02	20	(9)	64	(15)	319	(17)	1,240	(20)

Table 7.4: The number of iterations to converge the self-scatter source for a 2D PHI test case for the GMRes(20) method for the CBG-IGA 1G SAAF-S_4 equations. Those results in parentheses are obtained for the CBG-IGA DSA-preconditioned system. A total of $9,025$ bilinear NURBS patches were employed.

	c, Scattering Ratios (-)							
Σ (cm^{-1})	9.000E−01		9.900E−01		9.990E−01		9.999E−01	
1.000E+01	0.89372	(0.45122)	0.98975	(0.58657)	0.99896	(0.60486)	0.99975	(0.61454)
2.000E+01	0.89706	(0.54959)	0.98879	(0.71211)	0.99899	(0.72759)	0.99979	(0.72295)
4.000E+01	0.89834	(0.60606)	0.98887	(0.80777)	0.99900	(0.83927)	0.99985	(0.84229)
8.000E+01	0.89885	(0.63466)	0.98939	(0.85728)	0.99880	(0.89612)	0.99967	(0.90144)
1.600E+02	0.89915	(0.65263)	0.98963	(0.88680)	0.99875	(0.93321)	0.99959	(0.93927)
3.200E+02	0.89959	(0.66537)	0.98980	(0.90381)	0.99883	(0.95363)	0.99964	(0.96055)
6.400E+02	0.89998	(0.67478)	0.98987	(0.91394)	0.99890	(0.96399)	0.99969	(0.97408)

Table 7.5: The NSR for a 2D PHI test case for the SI procedure for the SIP-DG-IGA 1G SAAF-S_4 equations. Those results in parentheses are obtained for the MIP-DG-IGA DSA-preconditioned system. A total of $9,025$ bilinear NURBS patches were employed.

The periodic horizontal interface (PHI) test case is regularly employed in the literature to test the effectiveness of acceleration schemes over heterogeneous configurations [381, 387]. This problem presents *two* distinct regions, including horizontal strips of alternating thick diffusive, $\Sigma^1_{t,1} = \Sigma$, and thin transparent, $\Sigma^1_{t,2} = 1/\Sigma$, materials; as represented by Materials 1 and 2, respectively; v. Figure 7.3. The problem becomes increasingly heterogeneous for a sequence increasingly large values of $\Sigma = \{1.000\text{E}+01, 2.000\text{E}+01, 4.000\text{E}+01, 8.000\text{E}+01, 1.600\text{E}+02, 3.200\text{E}+02, 6.400\text{E}+02\}$;

	c, Scattering Ratios (-)							
Σ (cm^{-1})	9.000E−01		9.900E−01		9.990E−01		9.999E−01	
1.000E+01	101	(12)	791	(20)	a	(25)	a	(30)
2.000E+01	104	(15)	758	(27)	a	(36)	a	(42)
4.000E+01	105	(18)	826	(37)	a	(50)	a	(59)
8.000E+01	105	(21)	859	(50)	a	(75)	a	(91)
1.600E+02	104	(24)	878	(66)	a	(113)	a	(142)
3.200E+02	102	(27)	886	(83)	a	(161)	a	(210)
6.400E+02	104	(29)	885	(100)	a	(213)	a	(309)

Table 7.6: The number of iterations to converge the self-scatter source for a 2D PHI test case for the SI procedure for the SIP-DG-IGA 1G SAAF-S_4 equations. Any number $> 5,000$ is denoted by a; and those results in parentheses are obtained for the MIP-DG-IGA DSA-preconditioned system. A total of $9,025$ bilinear NURBS patches were employed.

	c, Scattering Ratios (-)							
Σ (cm^{-1})	9.000E−01		9.900E−01		9.990E−01		9.999E−01	
1.000E+01	23	(7)	85	(8)	378	(8)	1,182	(9)
2.000E+01	22	(6)	83	(10)	324	(11)	1,102	(11)
4.000E+01	20	(8)	82	(12)	316	(13)	1,329	(13)
8.000E+01	20	(9)	83	(11)	337	(15)	1,145	(16)
1.600E+02	18	(10)	67	(14)	337	(18)	1,366	(20)
3.200E+02	20	(10)	76	(14)	317	(19)	1,441	(21)
6.400E+02	19	(10)	76	(17)	297	(19)	1,284	(23)

Table 7.7: The number of iterations to converge the self-scatter source for a 2D PHI test case for the GMRes(20) method for the SIP-DG-IGA 1G SAAF-S_4 equations. Those results in parentheses are obtained for the MIP-DG-IGA DSA-preconditioned system. A total of $9,025$ bilinear NURBS patches were employed.

	c, Scattering Ratios (-)							
Σ (cm^{-1})	9.000E−01		9.900E−01		9.990E−01		9.999E−01	
1.000E+01	0.89732	(0.66456)	0.98909	(0.81647)	0.99894	(0.83279)	0.99973	(0.83607)
2.000E+01	0.89729	(0.71446)	0.98967	(0.90736)	0.99893	(0.92456)	0.99975	(0.92701)
4.000E+01	0.89878	(0.75344)	0.98945	(0.94712)	0.99891	(0.96442)	0.99974	(0.96665)
8.000E+01	0.89935	(0.78049)	0.98959	(0.96474)	0.99883	(0.98219)	0.99967	(0.98421)
1.600E+02	0.89935	(0.79374)	0.98982	(0.97188)	0.99897	(0.99037)	0.99966	(0.99225)
3.200E+02	0.89947	(0.79566)	0.98992	(0.97500)	0.99898	(0.99425)	0.99971	(0.99607)
6.400E+02	0.89992	(0.79140)	0.98994	(0.97727)	0.99898	(0.99613)	0.99974	(0.99793)

Table 7.8: The NSR for a 2D PHI test case for the SI procedure for the sweep-based DG-IGA 1G FONT-S_4 equations. Those results in parentheses are obtained for the MIP-DG-IGA DSA-preconditioned system. A total of $9,025$ bilinear NURBS patches were employed.

and the acceleration schemes are tested for a sequence of increasingly large scattering ratios, $c = \{9.000\text{E}{-}01, 9.900\text{E}{-}01, 9.990\text{E}{-}01, 9.999\text{E}{-}01\text{E}{-}01\}$; the corresponding macroscopic nuclear cross-section data are summarised in Table 7.1. One assumes that an isotropic source of unit strength is uniformly dispersed throughout the domain and that the surfaces of the physical boundary are bare; and as per the *one-quarter* symmetry in the problem, reflective BC are prescribed at $x = y = 0.00$ cm. The coarsest mesh is constructed from 9,025 bilinear NURBS patches. Without loss of generality,

	c, Scattering Ratios (-)							
Σ (cm^{-1})	9.000E−01		9.900E−01		9.990E−01		9.999E−01	
1.000E+01	143	(17)	1,252	(42)	a	(57)	a	(67)
2.000E+01	142	(20)	1,230	(81)	a	(124)	a	(151)
4.000E+01	140	(23)	1,243	(132)	a	(256)	a	(328)
8.000E+01	136	(24)	1,276	(194)	a	(489)	a	(671)
1.600E+02	130	(23)	1,274	(236)	a	(855)	a	(1,272)
3.200E+02	123	(21)	1,242	(267)	a	(1,259)	a	(2,205)
6.400E+02	119	(19)	1,192	(285)	a	(1,636)	a	(3,648)

Table 7.9: The number of iterations to converge the self-scatter source for a 2D PHI test case for the procedure for the sweep-based DG-IGA 1G FONT-S_4 equations. Any number $> 5,000$ is denoted by a; and those results in parentheses are obtained for the MIP-DG-IGA DSA-preconditioned system. A total of 9,025 bilinear NURBS patches were employed.

	c, Scattering Ratios (-)							
Σ (cm^{-1})	9.000E−01		9.900E−01		9.990E−01		9.999E−01	
1.000E+01	24	(7)	85	(13)	300	(14)	1,119	(14)
2.000E+01	21	(7)	75	(15)	316	(19)	968	(18)
4.000E+01	22	(9)	57	(14)	232	(40)	1,093	(64)
8.000E+01	17	(10)	40	(15)	127	(53)	568	(80)
1.600E+02	16	(11)	32	(16)	103	(61)	400	(85)
3.200E+02	15	(11)	28	(16)	95	(63)	226	(93)
6.400E+02	15	(11)	22	(17)	83	(63)	148	(107)

Table 7.10: The number of iterations to converge the self-scatter source for a 2D PHI test case for the GMRes(20) method for the sweep-based DG-IGA 1G FONT-S_4 equations. Those results in parentheses are obtained for the MIP-DG-IGA DSA-preconditioned system. A total of 9,025 bilinear NURBS patches were employed.

results are only presented for discretisations of the NTE-S_N with the LS angular quadrature set.

The NSR, Equation 7.8, for the different configurations of the PHI test case for each of the CBG-IGA 1G SAAF-S_4 equations, the SIP-DG-IGA 1G SAAF-S_4 equations and the sweep-based DG-IGA 1G FONT-S_4 equations are presented in Table 7.2, Table 7.5 and Table 7.8, respectively, along with the NSR for the corresponding DSA-preconditioned systems. Whilst each DSA-SI scheme does allow for significant acceleration, there is a clear degradation of their effectiveness for increasingly heterogeneous configurations as the thin material becomes void-like and for increasingly large scattering ratios; viz. NSR$\rightarrow c$, which is reflected by the increasingly large number of iterations demanded by the DSA-preconditioned systems to attain convergence; and this is most obviously so for the sweep-based DG-IGA FONT-S_4 equations; v. Table 7.3, Table 7.6 and Table 7.9, respectively.

Each discretisation of the different forms of the NTE also benefited from accelerated rates of convergence for the use of the more sophisticated GMRes(20) method, which is generally more performant than the DSA-SI schemes for each of the different configurations; and the DSA-GMRes(20) schemes even more so; v. Table 7.4, Table 7.7 and Table 7.10, respectively. Again, for a given heteroge-

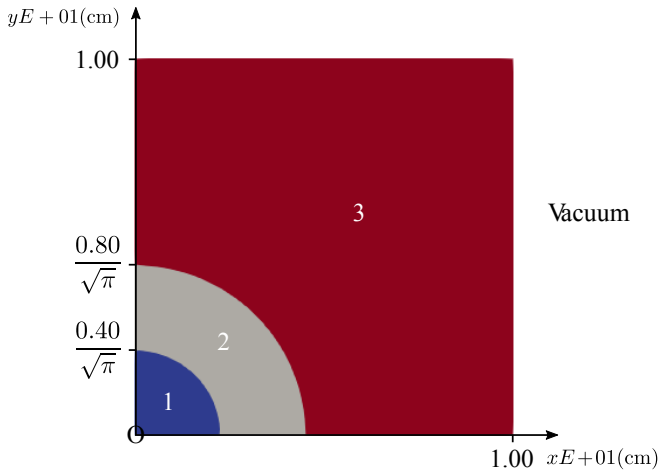


Figure 7.4: The geometry for Ragusa and Wang's 2D 1G radiation shielding problem, which has been adapted from the original Cartesian description to include for conic sections modelled by rational NURBS [70, 407]. The constituent materials are denoted by the numbers in parentheses; viz. source (1), shield (2) and absorber (3). (*Vide* the web-based version for reference to colour.)

neous configuration, Σ, a larger number of iterations is generally required to attain convergence for an increasingly large scattering ratio. For a given scattering ratio, the discretisations of the SAAF-S_4 equations appear to be relatively insensitive to the degree of heterogeneity in the problem. Conversely, however, for an increasingly heterogeneous problem for the FONT-S_4 equations, an increasingly fewer number of iterations is generally required of the unaccelerated GMRes method, whereas an increasingly larger number is required of the DSA-GMRes(20) scheme, although the latter was always more performant than the former. Although it has not been investigated here, one recalls that the convergence properties of the GMRes method are also dependent upon the restart parameter and potential issues of stagnation.

7.4.2 Ragusa and Wang's 2D Radiation Shielding Problem

Material	g	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to 1}$ (cm^{-1})	q_x^g (cm^{-3}.s^{-1})
(1) Source	1	1.000E+00	9.000E−01	1.000E+00
(2) Shield	1	1.500E+00	1.440E+00	0.000E+00
(3) Absorber	1	1.000E+00	3.000E−01	0.000E+00

Table 7.11: The nuclear data for Ragusa and Wang's 2D 1G radiation shielding problem [407].

Ragusa and Wang's 2D problem is representative of a heterogeneous radiation shielding problem with a localised source [408]. This problem presents *three* distinct regions, including a circular source region surrounded by an annular region of a highly scattering material, which is placed in a strong absorber of square shape with bare surfaces; as represented by Materials 1, 2 and 3, respectively; v. Figure 7.4. Whilst the geometry has been adapted from the original problem, the corresponding macroscopic nuclear cross-section data remain unchanged and are provided in Table 7.11. One assumes that an isotropic source of unit strength is uniformly dispersed across the fuelled region; and as per the *one-*

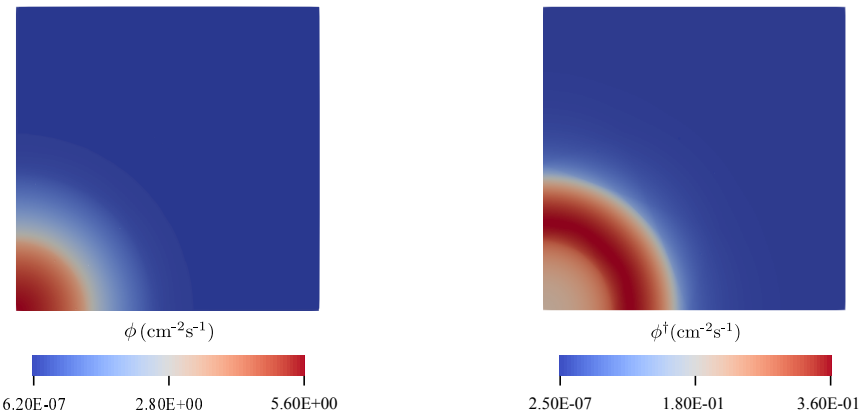
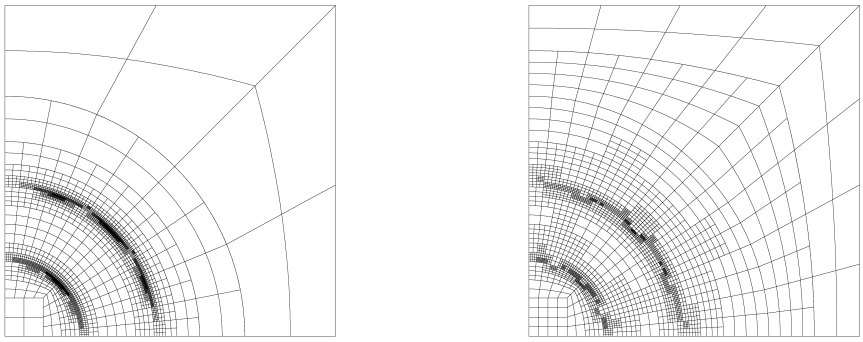


Figure 7.5: The spatial distribution of the mono-group component of the scalar neutron flux (left) and the scalar neutron importance (right) for the 1G FONT-S_4 equations for Ragusa and Wang's 2D radiation shielding problem. (*Vide* the web-based version for reference to colour.)



(a) $\mathcal{T}^{(9)}$ for SAAF-S_4: $57,408$ DoF **(b)** $\mathcal{T}^{(9)}$ for FONT-S_4: $39,552$ DoF

Figure 7.6: The meshes generated at the 9th AMR iteration for the DWR-AMR-h strategy for each of the SIP-DG-IGA 1G SAAF-S_4 equations (left) and the sweep-based DG-IGA 1G FONT-S_4 equations (right) for Ragusa and Wang's 2D radiation shielding problem; the QoI is the absorption in Material 2; v. Figure 7.4. (*Vide* the web-based version for reference to colour.)

quarter symmetry in the problem, reflective BC are prescribed at $x = y = 0.00$ cm. The coarsest mesh is constructed from *nine* bi-quadratic NURBS patches. A reference solution was obtained over the mesh generated after the 11th iteration of the DWR-AMR-hp strategy for the sweep-based DG-IGA spatial discretisation of the FONT-S_4 equations. The spatial distribution of the mono-group component of both the scalar neutron flux and the scalar neutron importance for the FONT-S_4 equations with the LS angular quadrature set are presented in Figure 7.5, for the QoI chosen as the absorption in Material 2; v. Figure 7.4.

For those acceleration techniques introduced in this chapter, the current problem and those to come investigate their performance in the application to those DWR-AMR-$h(p)$ strategies developed in Section 3.6.2 over the increasingly irregular meshes of complex rational geometries. To elucidate, the AMR procedure forms the outermost iteration loop, wrapped around those iterations of the DSA-SI scheme or the DSA-GMRes scheme; and from one AMR iteration to the next, the solution over the reference mesh from the previous one is *bootstrapped* onto that of the current AMR iteration to keep the number of DSA iterations relatively small [387].

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	36	361	0.915	(0.217)	163	(12)	(3.47)	21	(7)	(2.86)	-1.85
01	84	833	0.915	(0.225)	164	(12)	(2.67)	20	(8)	(2.57)	-2.23
02	252	2,425	0.915	(0.225)	165	(12)	(1.77)	19	(8)	(1.82)	-2.65
03	468	4,565	0.915	(0.222)	165	(12)	(1.77)	18	(8)	(2.08)	-2.88
04	708	6,919	0.915	(0.225)	166	(13)	(1.93)	18	(8)	(2.13)	-3.23
05	1,188	11,671	0.915	(0.226)	165	(13)	(1.89)	17	(8)	(2.03)	-3.30
06	1,692	16,703	0.915	(0.226)	165	(13)	(1.95)	15	(8)	(2.13)	-3.57
07	2,364	23,211	0.915	(0.228)	165	(13)	(2.11)	18	(8)	(2.07)	-3.60
08	2,508	24,655	0.915	(0.227)	165	(13)	(2.02)	18	(8)	(2.19)	-3.91
09	3,492	34,417	0.915	(0.227)	165	(13)	(2.22)	17	(8)	(2.48)	-4.24

Table 7.12: A summary of results for the DWR-AMR-h strategy for the SI procedure and the GMRes(15) method for the solution of the CBG-IGA 1G SAAF-S_4 equations for Ragusa and Wang's 2D radiation shielding problem; the QoI is the absorption in Material 2; v. Figure 7.4. Those results in parentheses are obtained for the CBG-IGA DSA-preconditioned system.

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	36	576	0.915	(0.217)	164	(12)	(2.30)	22	(7)	(3.01)	-2.03
01	84	1,344	0.915	(0.225)	165	(12)	(2.37)	20	(7)	(3.02)	-2.65
02	228	3,648	0.915	(0.225)	164	(12)	(1.97)	20	(7)	(2.45)	-3.12
03	252	4,032	0.915	(0.225)	166	(12)	(1.99)	18	(7)	(2.42)	-3.49
04	444	7,104	0.915	(0.226)	166	(12)	(2.17)	17	(7)	(2.46)	-3.79
05	852	13,632	0.915	(0.229)	166	(13)	(2.22)	17	(7)	(2.49)	-3.81
06	1,092	17,472	0.915	(0.228)	165	(12)	(2.39)	17	(7)	(2.78)	-4.08
07	1,668	26,688	0.915	(0.230)	165	(12)	(2.59)	17	(7)	(2.95)	-4.12
08	2,508	40,128	0.915	(0.232)	165	(12)	(2.72)	15	(7)	(3.09)	-4.31
09	3,588	57,408	0.915	(0.233)	165	(12)	(3.07)	15	(7)	(3.41)	-4.33

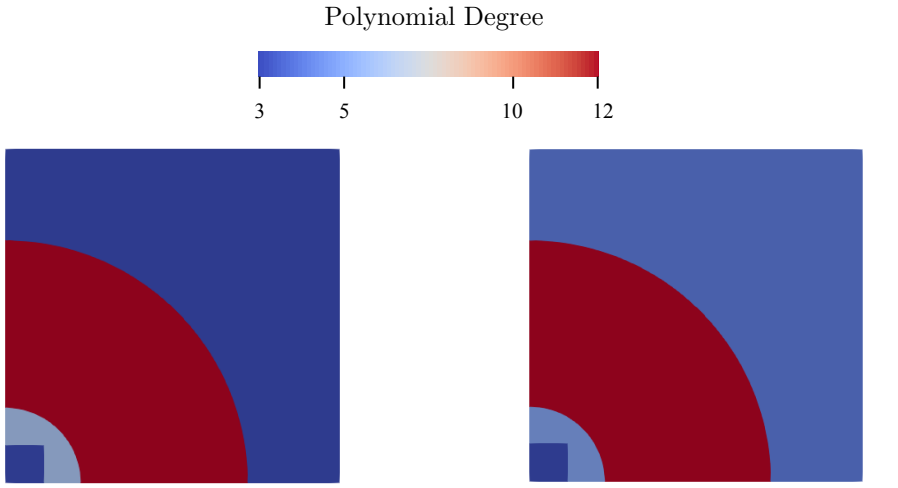
Table 7.13: A summary of results for the DWR-AMR-h strategy for the SI procedure and the GMRes(15) method for the solution of the SIP-DG-IGA 1G SAAF-S_4 equations for Ragusa and Wang's 2D radiation shielding problem; the QoI is the absorption in Material 2; v. Figure 7.4. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.

The meshes generated after the 9th iteration of the DWR-AMR-h strategy for the SIP-DG-IGA 1G SAAF-S_4 equations and the sweep-based DG-IGA 1G FONT-S_4 equations are compared in Figure 7.6; and similarly so, for the DWR-AMR-p strategy in Figure 7.7. As expected, for the chosen QoI, the refinement of the meshes has been focussed about the region of interest, Material 2. Interestingly, for the discretisation of the SAAF-S_4 equations, the DWR-AMR-h strategy has allocated almost twice as many DoF to the refinement of the mesh than that for the FONT-S_4 equations.

For each of the CBG-IGA 1G SAAF-S_4 equations, the SIP-DG-IGA 1G SAAF-S_4 and the sweep-based DG-IGA 1G FONT-S_4 equations, *two* tables are presented, one each for the corresponding DWR-AMR-h and the DWR-AMR-p strategies; v. Table 7.12 and Table 7.15; Table 7.13 and Table 7.16; and Table 7.14 and Table 7.17, respectively. Within each table, the number of elements and DoF

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	36	576	0.915	(0.283)	163	(14)	(6.64)	12	(8)	(9.21)	-3.54
01	84	1,344	0.915	(0.222)	165	(12)	(5.22)	13	(7)	(7.31)	-4.11
02	180	2,880	0.915	(0.227)	166	(12)	(4.60)	11	(7)	(5.67)	-4.31
03	276	4,416	0.915	(0.224)	165	(12)	(4.41)	14	(7)	(5.10)	-4.64
04	444	7,104	0.915	(0.225)	166	(12)	(4.20)	14	(7)	(4.92)	-4.90
05	684	10,944	0.915	(0.225)	165	(12)	(4.41)	15	(7)	(5.58)	-5.06
06	924	14,784	0.915	(0.225)	166	(12)	(5.07)	14	(7)	(4.98)	-6.16
07	1,392	22,272	0.915	(0.227)	165	(12)	(5.77)	17	(7)	(5.40)	-6.18
08	1,896	30,336	0.915	(0.227)	165	(12)	(6.03)	17	(7)	(5.93)	-6.38
09	2,472	39,552	0.915	(0.228)	165	(12)	(7.53)	20	(7)	(5.74)	-7.34

Table 7.14: A summary of results for the DWR-AMR-h strategy for the SI procedure and the GMRes(15) method for the solution of the sweep-based DG-IGA 1G FONT-S_4 equations for Ragusa and Wang's 2D radiation shielding problem; the QoI is the absorption in Material 2; v. Figure 7.4. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.



(a) $\mathcal{T}^{(9)}$ for SAAF-S_4: 3,288 DoF **(b)** $\mathcal{T}^{(9)}$ for FONT-S_4: 3,256 DoF

Figure 7.7: The meshes generated at the 9th AMR iteration for the DWR-AMR-p strategy for each of the SIP-DG-IGA 1G SAAF-S_4 equations (left) and the sweep-based DG-IGA 1G FONT-S_4 equations (right) for Ragusa and Wang's 2D radiation shielding problem; the QoI is the absorption in Material 2; v. Figure 7.4. (*Vide* the web-based version for reference to colour.)

belonging to the mesh at each AMR iteration are recorded, along with the convergence properties for each of the SI procedure and the GMRes method and their corresponding DSA-preconditioned systems. One observes that, naturally, the NSR is bounded from above and below by the largest and the smallest, respectively, of the scattering ratios for the materials present in the problem; viz. 0.30 <NSR< 0.96 [374]. Only the DWR-AMR strategies are investigated because they have been shown to reduce, consistently and reliably, the error in the QoI at each subsequent AMR iteration; and whilst the GMRes method does accelerate the convergence of the discrete problems, it does not always perform as well as the DSA-SI scheme. Nonetheless, the DSA-GMRes scheme is always be the most performant. Similar tables could be produced for the dual problem for each of the discretisations investigated, such that the corresponding discrete linear systems for the adjoint equations also benefit

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	36	361	0.914	(0.217)	163	(12)	(3.47)	21	(7)	(2.86)	-1.85
01	36	497	0.915	(0.207)	163	(9)	(3.70)	17	(8)	(1.42)	-2.02
02	36	649	0.915	(0.207)	163	(9)	(4.22)	15	(8)	(0.88)	-2.48
03	36	903	0.915	(0.206)	163	(9)	(4.58)	14	(8)	(2.47)	-2.82
04	36	1,119	0.915	(0.211)	164	(10)	(4.93)	13	(7)	(5.99)	-2.99
05	36	1,367	0.915	(0.205)	164	(10)	(2.44)	13	(7)	(5.28)	-3.19
06	36	1,725	0.914	(0.201)	164	(10)	(0.83)	13	(7)	(4.59)	-3.33
07	36	2,037	0.914	(0.214)	165	(11)	(0.37)	13	(7)	(4.22)	-3.41
08	36	2,381	0.914	(0.216)	165	(14)	(0.19)	11	(7)	(5.09)	-3.66
09	36	2,851	0.914	(0.218)	165	(14)	(0.10)	11	(7)	(3.43)	-3.93

Table 7.15: A summary of results for the DWR-AMR-p strategy for the SI procedure and the GMRes(15) method for the solution of the CBG-IGA 1G SAAF-S_4 equations for Ragusa and Wang's 2D radiation shielding problem; the QoI is the absorption in Material 2; v. Figure 7.4. Those results in parentheses are obtained for the CBG-IGA DSA-preconditioned system.

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	36	576	0.915	(0.217)	164	(12)	(2.30)	22	(7)	(3.01)	-2.03
01	36	720	0.915	(0.215)	166	(10)	(2.54)	20	(8)	(1.97)	-2.49
02	36	896	0.916	(0.218)	166	(10)	(2.76)	19	(8)	(1.57)	-2.82
03	36	1,176	0.916	(0.220)	165	(11)	(3.01)	20	(8)	(1.62)	-3.01
04	36	1,416	0.916	(0.221)	165	(11)	(2.43)	17	(8)	(2.95)	-3.19
05	36	1,688	0.915	(0.222)	165	(11)	(2.22)	18	(8)	(3.80)	-3.34
06	36	2,080	0.915	(0.222)	165	(11)	(2.11)	18	(8)	(4.11)	-3.41
07	36	2,416	0.914	(0.220)	165	(11)	(1.89)	17	(8)	(5.56)	-3.67
08	36	2,784	0.914	(0.223)	165	(12)	(1.23)	17	(8)	(6.95)	-3.94
09	36	3,288	0.914	(0.224)	165	(11)	(1.04)	15	(8)	(6.67)	-4.24

Table 7.16: A summary of results for the DWR-AMR-p strategy for the SI procedure and the GMRes(15) method for the solution of the SIP-DG-IGA 1G SAAF-S_4 equations for Ragusa and Wang's 2D radiation shielding problem; the QoI is the absorption in Material 2; v. Figure 7.4. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.

from similar accelerated convergence; however, for brevity, they are not.

Generally, for the DWR-AMR-h strategy, the relative CPU time taken to solve the diffusion acceleration equations is $< 5\%$ for the SAAF-S_4 equations and $< 10\%$ for the FONT-S_4 equations; and for the DWR-AMR-p strategy, $< 5\%$ for the SAAF-S_4 equations and $< 50\%$ for the FONT-S_4 equations. One recalls that the discrete linear systems for the SAAF-S_N equations can be rendered increasingly ill conditioned and a certain degree of instability is introduced for an increasingly large polynomial degree of the NURBS basis; and more so than the banded matrix of the corresponding diffusion acceleration equations, which thus, require increasingly less time to solve, relatively, for the DSA-SI scheme. In the case of the DWR-AMR-p strategy for the sweep-based DG-IGA FONT-S_N equations, there is relatively little additional computational effort demanded of the matrix-free approach with each

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	36	576	0.915	(0.283)	163	(14)	(6.64)	12	(8)	(9.21)	-3.54
01	36	720	0.914	(0.276)	163	(12)	(4.93)	13	(7)	(6.49)	-3.86
02	36	896	0.914	(0.238)	164	(11)	(6.89)	13	(7)	(9.89)	-4.53
03	36	1,104	0.913	(0.233)	164	(10)	(9.30)	12	(7)	(14.5)	-4.41
04	36	1,416	0.913	(0.231)	164	(10)	(19.4)	10	(7)	(25.3)	-4.78
05	36	1,688	0.914	(0.230)	165	(10)	(24.5)	10	(8)	(35.9)	-5.04
06	36	2,064	0.915	(0.229)	166	(11)	(44.1)	10	(8)	(44.6)	-5.42
07	36	2,488	0.915	(0.232)	166	(11)	(47.3)	10	(8)	(42.7)	-5.83
08	36	2,856	0.915	(0.232)	166	(11)	(48.1)	10	(8)	(42.9)	-6.26
09	36	3,256	0.915	(0.232)	166	(11)	(48.5)	11	(8)	(39.4)	-6.73

Table 7.17: A summary of results for the DWR-AMR-p strategy for the SI procedure and the GMRes(15) method for the solution of the sweep-based DG-IGA 1G FONT-S_4 equations for Ragusa and Wang's 2D radiation shielding problem; the QoI is the absorption in Material 2; v. Figure 7.4. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.

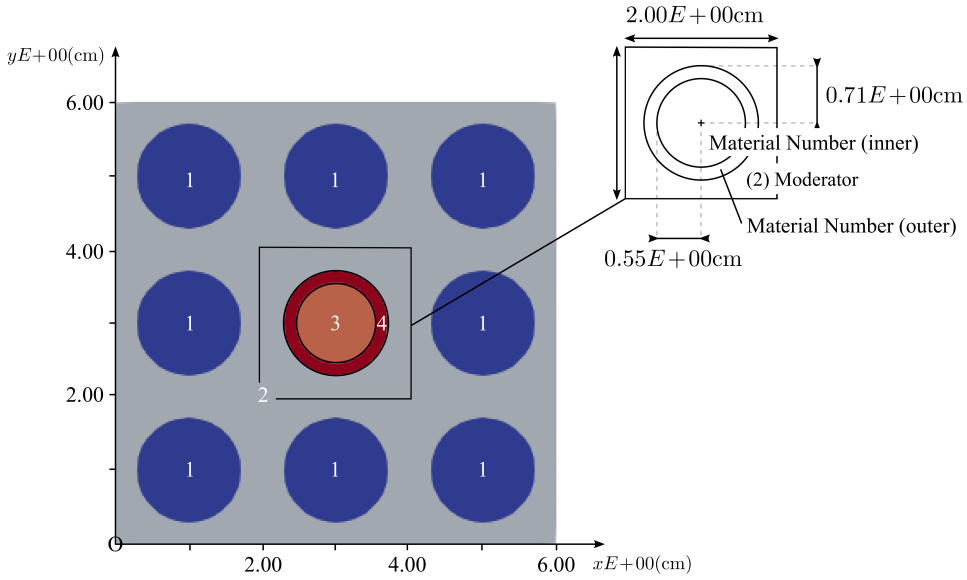


Figure 7.8: The geometry for Wood and William's 2D 1G BWR problem with a burnable poison. The constituent materials are denoted by the numbers in parentheses; viz. fuel (1), moderator (2), inner-poison (3) and outer-poison (4) [409]. (*Vide* the web-based version for reference to colour.)

subsequent AMR iteration, since the number of elements belonging to the mesh remains unchanged; and thus, the relative time spent solving the diffusion acceleration equations is relatively high by comparison; and increasingly so, for increasingly ill-conditioned banded matrices of the corresponding diffusion acceleration equations.

7.4.3 Wood and William's 2D BWR Problem

Wood and William's problem is a three-by-three simplified BWR supercell lattice with a central Gadolinium burnable poison pin [409]. This problem presents *four* physically distinct regions, includ-

Material	g	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to 1}$ (cm^{-1})	$\Sigma_{s0}^{g\to 1}$ (cm^{-1})	q_x^g (cm^{-3}.s^{-1})
(1) Fuel	1	5.145E−01	3.432E−01	0.000E+00	0.000E+00
(2) Moderator	1	7.557E−01	7.504E−01	2.500E−01	1.000E+00
(3) Poison (inner)	1	1.427E+01	3.400E−01	0.000E+00	0.000E+00
(4) Poison (outer)	1	2.163E+01	3.414E−01	0.000E+00	0.000E+00

Table 7.18: The nuclear data for Wood and William's 2D 1G BWR problem with a burnable poison [409].

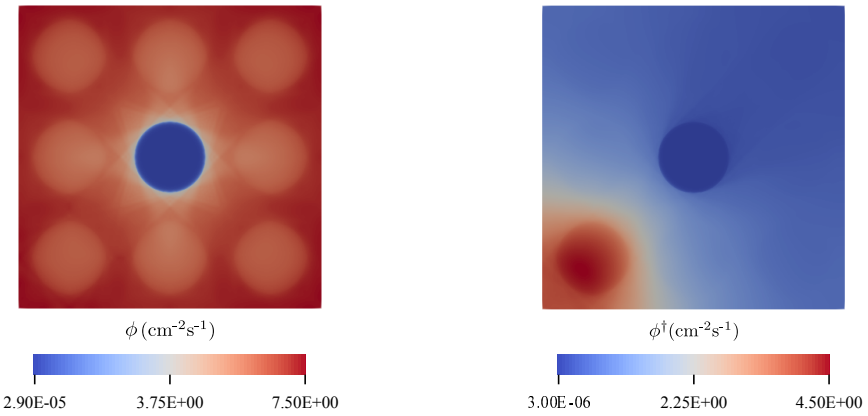
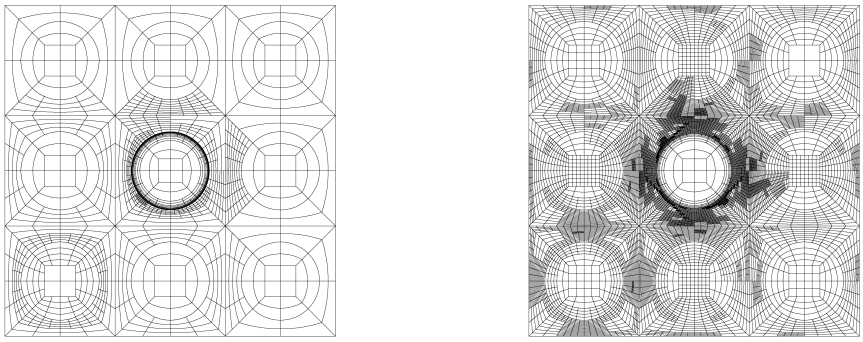


Figure 7.9: The spatial distribution of the mono-group component of the scalar neutron flux (left) and the scalar neutron importance (right) for the 1G FONT-S_6 equations for Wood and William's 2D BWR problem with a burnable poison [409]. (*Vide* the web-based version for reference to colour.)



(a) $\mathcal{T}^{(9)}$ for SAAF-S_6: 184,512 DoF **(b)** $\mathcal{T}^{(9)}$ for FONT-S_6: 167,296 DoF

Figure 7.10: The meshes generated at the 9th AMR iteration for the DWR-AMR-h strategy for each of the SIP-DG-IGA 1G SAAF-S_6 equations (left) and the sweep-based DG-IGA 1G FONT-S_6 equations (right) for Wood and William's 2D BWR problem with a burnable poison; the QoI is the absorption in Material 1 in the bottom left-hand corner; v. Figure 7.8. (*Vide* the web-based version for reference to colour.)

ing *eight* identical fuel pins of radius 0.71 cm arranged symmetrically about *one* burnable poison pin with an annular region of the same outer-radius and an inner-radius of 0.55 cm; each pin is moderated by water and set in a cell of pitch 2.00 cm; as represented by Materials 1, 4, 3 and 2; v. Figure 7.8. This problem has been adapted from the original one to account for isotropic scatter only; the corresponding one-group macroscopic cross-section data are provided in Table 7.18 [409]. To model an infinite lattice, reflective BC are specified at every external physical boundary. The coarsest mesh is constructed from 85 bi-quadratic NURBS patches. A reference solution was obtained over the mesh generated after the 11th iteration of the DWR-AMR-hp strategy for the sweep-based DG-IGA spatial discretisation of the FONT-S_6 equations. The spatial distribution of the mono-group component of

(i)	No. Elem	DoF	NSR		SI No.	SI Its	SI time (%)	GMRes(15) No.	GMRes(15) Its	GMRes(15) time (%)	Log(ϵ_{QoI})
00	340	3,097	0.882	(0.189)	81	(8)	(1.57)	10	(6)	(2.02)	-1.28
01	388	3,585	0.882	(0.188)	81	(9)	(2.19)	11	(7)	(2.36)	-1.50
02	484	4,529	0.882	(0.193)	81	(9)	(2.39)	11	(7)	(2.85)	-1.58
03	676	6,417	0.882	(0.189)	81	(9)	(3.05)	11	(7)	(3.47)	-1.60
04	1,060	10,193	0.882	(0.186)	81	(9)	(3.16)	11	(7)	(3.90)	-1.69
05	1,828	17,745	0.882	(0.185)	80	(9)	(3.91)	11	(7)	(4.38)	-1.78
06	3,412	33,313	0.882	(0.185)	80	(9)	(4.28)	11	(7)	(4.72)	-1.85
07	5,836	57,221	0.883	(0.186)	80	(9)	(4.41)	10	(6)	(4.19)	-2.10
08	8,452	82,889	0.883	(0.187)	80	(9)	(3.63)	10	(6)	(3.51)	-2.19
09	11,836	116,127	0.883	(0.188)	80	(9)	(2.96)	10	(6)	(2.83)	-2.51

Table 7.19: A summary of results for the DWR-AMR-h strategy for the SI procedure and the GMRes(15) method for the solution of the CBG-IGA 1G SAAF-S_6 equations for Wood and William's 2D BWR problem with a burnable poison; the QoI is the absorption in Material 1 in the bottom left-hand corner; v. Figure 7.8. Those results in parentheses are obtained for the CBG-IGA DSA-preconditioned system.

(i)	No. Elem	DoF	NSR		SI No.	SI Its	SI time (%)	GMRes(15) No.	GMRes(15) Its	GMRes(15) time (%)	Log(ϵ_{QoI})
00	340	5,440	0.885	(0.189)	83	(9)	(2.46)	11	(7)	(1.21)	-1.21
01	440	7,040	0.884	(0.191)	82	(9)	(2.38)	11	(7)	(1.37)	-1.33
02	548	8,768	0.884	(0.197)	82	(9)	(2.22)	11	(7)	(1.43)	-1.45
03	836	13,376	0.883	(0.201)	82	(9)	(2.39)	11	(7)	(1.45)	-1.56
04	1,212	19,392	0.883	(0.202)	82	(9)	(2.98)	11	(7)	(1.64)	-1.65
05	2,004	32,064	0.883	(0.201)	82	(9)	(3.21)	11	(7)	(1.38)	-1.73
06	3,176	50,816	0.883	(0.200)	81	(9)	(3.51)	11	(7)	(1.47)	-1.82
07	5,332	85,312	0.882	(0.200)	81	(9)	(3.43)	10	(7)	(1.75)	-2.04
08	7,292	116,672	0.882	(0.201)	81	(9)	(3.68)	10	(7)	(1.89)	-2.33
09	11,532	184,512	0.883	(0.202)	82	(9)	(3.76)	10	(7)	(2.31)	-2.62

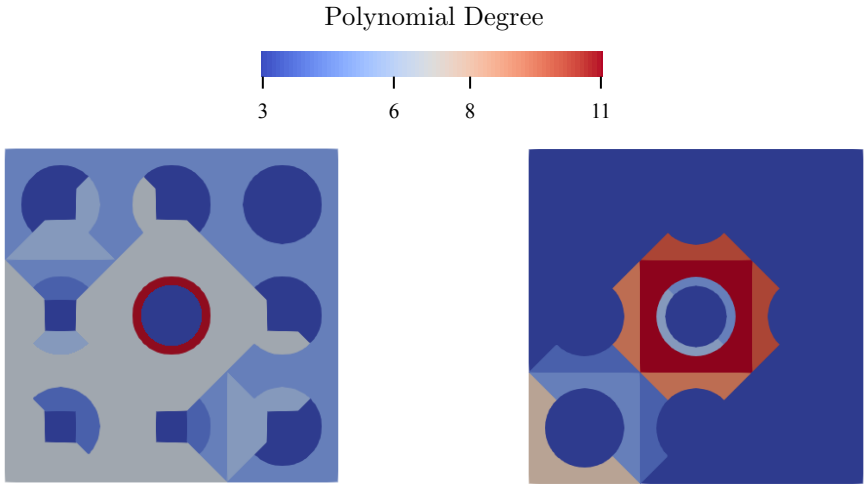
Table 7.20: A summary of results for the DWR-AMR-h strategy for the SI procedure and the GMRes(15) method for the solution of the SIP-DG-IGA 1G SAAF-S_6 equations for Wood and William's 2D BWR problem with a burnable poison; the QoI is the absorption in Material 1 in the bottom left-hand corner; v. Figure 7.8. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.

both the scalar neutron flux and the scalar neutron importance for the FONT-S_6 equations with a GC angular quadrature set are presented in Figure 7.9. The QoI is chosen as the absorption in Material 1 in the bottom left-hand corner; v. Figure 7.8.

The meshes generated after the 9th iteration of the DWR-AMR-h strategy for the SIP-DG-IGA 1G SAAF-S_6 equations and the sweep-based DG-IGA 1G FONT-S_6 equations are compared in Figure 7.10; and similarly so, for the DWR-AMR-p strategy in Figure 7.11. Along with some preferential refinement about the region of interest, the AMR strategies focus primarily on resolving the sharp gradients about the burnable poison pin. Interestingly, although perhaps not obvious, the DWR-AMR strategies for the SAAF-S_6 equations have, again, allocated more DoF to the refinement of their respective meshes mesh than that for the FONT-S_6 equations. In the case of the latter, one can

(i)	No. Elem	DoF	NSR		SI No. Its	SI time (%)	GMRes(15) No. Its	GMRes(15) time (%)	Log(ϵ_{QoI})
00	340	5,440	0.890	(0.208)	90 (9)	(2.18)	9 (7)	(1.86)	-2.45
01	436	6,976	0.890	(0.209)	89 (9)	(2.39)	9 (7)	(1.91)	-3.13
02	628	10,048	0.890	(0.209)	89 (9)	(2.80)	9 (7)	(1.87)	-3.25
03	892	14,272	0.890	(0.210)	89 (9)	(2.78)	9 (7)	(1.82)	-3.46
04	1,684	26,944	0.889	(0.210)	88 (9)	(2.57)	9 (7)	(1.79)	-3.61
05	2,188	35,008	0.889	(0.209)	88 (9)	(2.42)	9 (7)	(1.82)	-3.65
06	2,956	47,296	0.889	(0.208)	88 (9)	(2.60)	8 (7)	(1.79)	-3.99
07	5,872	93,952	0.889	(0.196)	87 (9)	(3.13)	8 (7)	(3.03)	-4.03
08	7,624	121,984	0.889	(0.199)	87 (9)	(4.08)	8 (7)	(2.63)	-4.32
09	10,456	167,296	0.889	(0.214)	87 (9)	(3.38)	7 (7)	(3.09)	-4.70

Table 7.21: A summary of results for the DWR-AMR-h strategy for the SI procedure and the GMRes(15) method for the solution of the sweep-based DG-IGA 1G FONT-S_6 equations for Wood and William's 2D BWR problem with a burnable poison; the QoI is the absorption in Material 1 in the bottom left-hand corner; v. Figure 7.8. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.



(a) $\mathcal{T}^{(9)}$ for SAAF-S_6: 12,896 DoF **(b)** $\mathcal{T}^{(9)}$ for FONT-S_6: 10,928 DoF

Figure 7.11: The meshes generated at the 9th AMR iteration for the DWR-AMR-p strategy for each of the SIP-DG-IGA 1G SAAF-S_6 equations (left) and the sweep-based DG-IGA 1G FONT-S_6 equations (right) for Wood and William's 2D BWR problem with a burnable poison; the QoI is the absorption in Material 1 in the bottom left-hand corner; v. Figure 7.8. (*Vide* the web-based version for reference to colour.)

observe that the ray effects are beginning to dominate the DWR-AMR-h strategy; v. Figure 7.10.

For each of the CBG-IGA 1G SAAF-S_6 equations, the SIP-DG-IGA 1G SAAF-S_6 and the sweep-based DG-IGA 1G FONT-S_6 equations, *two* tables summarise the convergence properties of each of the corresponding DWR-AMR-h and the DWR-AMR-p strategies; v. Table 7.19 and Table 7.22; Table 7.20 and Table 7.23; and Table 7.21 and Table 7.24, respectively. Qualitatively, the findings are similar to those made previously and thus, for brevity, the discussion is not repeated here; q.v. Section 7.4.2.

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	340	3,097	0.882	(0.189)	81	(8)	(1.57)	10	(6)	(2.02)	-1.28
01	340	3,249	0.882	(0.188)	81	(8)	(1.44)	9	(7)	(2.69)	-1.89
02	340	3,401	0.882	(0.185)	81	(8)	(2.09)	9	(7)	(1.74)	-2.16
03	340	3,585	0.883	(0.185)	81	(8)	(1.39)	9	(7)	(1.83)	-2.32
04	340	4,245	0.884	(0.185)	81	(7)	(1.04)	9	(7)	(1.39)	-2.41
05	340	5,009	0.885	(0.185)	82	(7)	(0.98)	9	(6)	(1.27)	-2.45
06	340	5,481	0.884	(0.185)	82	(7)	(0.91)	8	(6)	(1.13)	-2.49
07	340	6,479	0.885	(0.185)	82	(7)	(0.68)	8	(6)	(0.99)	-2.54
08	340	8,401	0.885	(0.185)	82	(7)	(0.59)	8	(6)	(0.85)	-2.63
09	340	11,003	0.885	(0.185)	83	(7)	(0.47)	8	(6)	(0.75)	-2.74

Table 7.22: A summary of results for the DWR-AMR-p strategy for the SI procedure and the GMRes(15) method for the solution of the CBG-IGA 1G SAAF-S_6 equations for Wood and William's 2D BWR problem with a burnable poison; the QoI is the absorption in Material 1 in the bottom left-hand corner; v. Figure 7.8. Those results in parentheses are obtained for the CBG-IGA DSA-preconditioned system.

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	340	5,440	0.885	(0.189)	83	(9)	(2.46)	11	(7)	(1.21)	-1.21
01	340	5,656	0.889	(0.190)	85	(8)	(2.88)	11	(8)	(1.55)	-1.80
02	340	5,872	0.887	(0.190)	85	(8)	(2.74)	11	(7)	(1.73)	-2.25
03	340	6,336	0.887	(0.191)	85	(8)	(3.44)	10	(7)	(2.17)	-2.30
04	340	6,808	0.887	(0.191)	84	(8)	(3.02)	10	(7)	(2.41)	-2.36
05	340	7,562	0.888	(0.191)	84	(8)	(2.36)	10	(7)	(1.99)	-2.41
06	340	8,414	0.888	(0.191)	84	(8)	(2.21)	10	(7)	(1.75)	-2.47
07	340	9,044	0.888	(0.191)	84	(8)	(2.88)	10	(6)	(1.18)	-2.54
08	340	10,912	0.888	(0.191)	84	(8)	(1.75)	10	(6)	(1.03)	-2.66
09	340	12,896	0.888	(0.191)	84	(8)	(1.21)	10	(6)	(0.89)	-2.79

Table 7.23: A summary of results for the DWR-AMR-p strategy for the SI procedure and the GMRes(15) method for the solution of the SIP-DG-IGA 1G SAAF-S_6 equations for Wood and William's 2D BWR problem with a burnable poison; the QoI is the absorption in Material 1 in the bottom left-hand corner; v. Figure 7.8. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.

7.4.4 The 2D IAEA CEA-2 Benchmark

The 2D IAEA CEA-2 benchmark problem is an infinite seven-by-seven lattice of pitch 1.25 cm that represents the heterogeneous configuration of a simplified PWR-like fuel assembly that contains fuel and poison pins of radius 0.45 cm and a central water gap; as represented by Materials 1, 2 and 3, respectively; v. Figure 7.12 [152]. The corresponding macroscopic cross-section data are provided in Table 7.25, where one assumes that the density of water is constant throughout the problem [152]. Modelling for only a single thermal energy group, one assumes a uniformly dispersed and isotropic slowing-down source across the moderator region. Furthermore, to model an infinite lattice, reflective BC are specified at every external physical boundary; and as per the *one-quarter* symmetry in the problem, reflective BC are prescribed at $x = y = 0.00$ cm. The coarsest mesh is constructed from 124

(i)	No. Elem	DoF	NSR		SI No. Its	SI time (%)	GMRes(15) No. Its	GMRes(15) time (%)	Log(ϵ_{QoI})
00	340	5,440	0.890	(0.208)	90 (9)	(2.18)	9 (7)	(1.86)	-2.45
01	340	5,728	0.890	(0.205)	89 (9)	(2.60)	9 (7)	(2.06)	-2.64
02	340	6,280	0.890	(0.202)	89 (9)	(2.88)	9 (7)	(2.00)	-3.16
03	340	6,752	0.890	(0.198)	89 (9)	(3.17)	10 (7)	(2.09)	-3.30
04	340	7,288	0.890	(0.200)	88 (10)	(3.53)	9 (7)	(2.34)	-3.49
05	340	7,800	0.890	(0.199)	88 (10)	(3.76)	8 (7)	(2.44)	-3.64
06	340	8,376	0.890	(0.200)	88 (10)	(4.12)	9 (7)	(2.70)	-3.96
07	340	9,208	0.888	(0.197)	87 (10)	(4.36)	9 (7)	(2.76)	-4.47
08	340	9,864	0.891	(0.197)	89 (10)	(4.94)	9 (7)	(2.81)	-4.70
09	340	10,928	0.890	(0.218)	88 (10)	(5.13)	9 (7)	(3.26)	-5.10

Table 7.24: A summary of results for the DWR-AMR-p strategy for the SI procedure and the GMRes(15) method for the solution of the sweep-based DG-IGA 1G FONT-S_6 equations for Wood and William's 2D BWR problem with a burnable poison; the QoI is the absorption in Material 1 in the bottom left-hand corner; v. Figure 7.8. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.

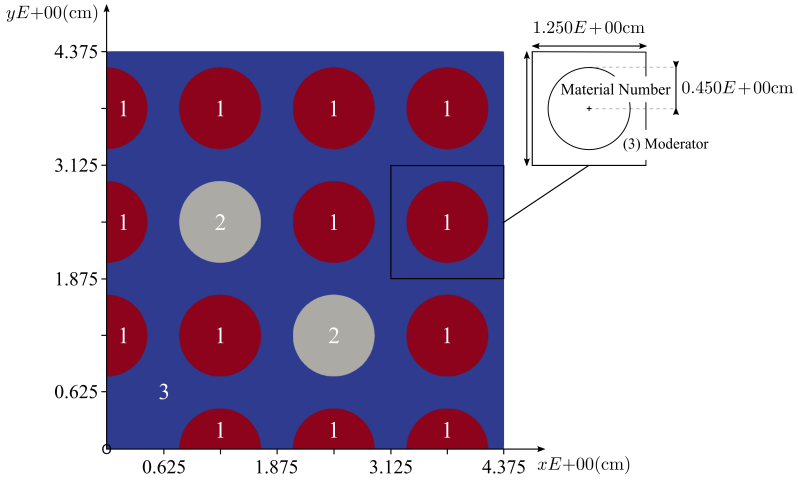


Figure 7.12: The geometry for the 2D 1G IAEA CEA-2 benchmark [152]. The constituent materials are denoted by the numbers parentheses; viz. fuel (1), poison (2) and moderator (3). (*Vide* the web-based version for reference to colour.)

Material	g	Σ_t^g (cm^{-1})	$\Sigma_{s0}^{g\to 1}$ (cm^{-1})	q_x^g (cm^{-3}.s^{-1})
(1) Fuel	1	6.250E−01	3.550E−01	0.000E+00
(2) Poison	1	1.400E+01	0.000E+00	0.000E+00
(3) Moderator	1	1.250E+00	1.242E+00	1.000E+00

Table 7.25: The nuclear data for the 2D 1G IAEA CEA-2 benchmark [152].

bi-quadratic NURBS patches. A reference solution was obtained over the mesh generated after the 11th iteration of the DWR-AMR-hp strategy for the sweep-based DG-IGA spatial discretisation of the FONT-S_8 equations. The spatial distribution of the mono-group component of both the scalar neutron flux and the scalar neutron importance for the FONT-S_8 equations with the LS angular quadrature set are presented in Figure 7.13. The QoI is chosen as the absorption in Material 1 in the top right-hand corner; v. Figure 7.12.

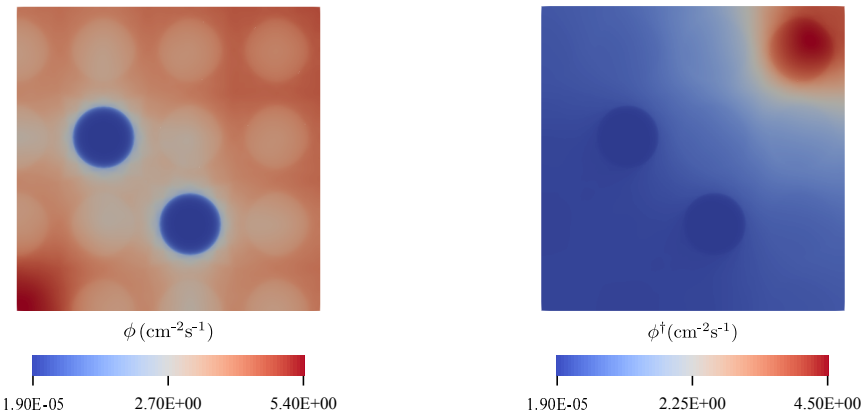
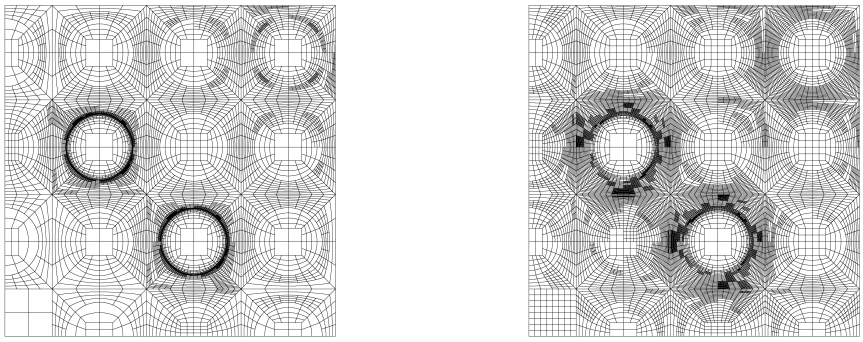


Figure 7.13: The spatial distribution of the mono-group component of the scalar neutron flux (left) and the scalar neutron importance (right) for the 1G FONT-S_8 equations for the 2D IAEA CEA-2 benchmark [152]. (*Vide* the web-based version for reference to colour.)



(a) $\mathcal{T}^{(9)}$ for SAAF-S_8: 125,523 DoF **(b)** $\mathcal{T}^{(9)}$ for FONT-S_8: 198,976 DoF

Figure 7.14: The meshes generated at the 9th AMR iteration for the DWR-AMR-h strategy for each of the CBG-IGA 1G SAAF-S_8 equations (left) and the sweep-based DG-IGA 1G FONT-S_8 equations (right) for the 2D IAEA CEA-2 benchmark; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 7.12. (*Vide* the web-based version for reference to colour.)

The meshes generated after the 9th iteration of the DWR-AMR-h strategy for the CBG-IGA 1G SAAF-S_8 equations and the sweep-based 1G FONT-S_8 equations are compared in Figure 7.14; and similarly for the DWR-AMR-p strategy in Figure 7.15. Again, along with some preferential refinement about the region of interest, the AMR strategies focus primarily on resolving the sharp gradients about the poison pins.

For each of the CBG-IGA 1G SAAF-S_8 equations, the SIP-DG-IGA 1G SAAF-S_8 and the sweep-based DG-IGA 1G FONT-S_8 equations, *two* tables summarise the convergence properties of each of the corresponding DWR-AMR-h and the DWR-AMR-p strategies; v. Table 7.26 and Table 7.29; Table 7.27 and Table 7.30; and Table 7.28 and Table 7.31, respectively. The findings are qualitatively similar to those of the previous problems and thus, for brevity, the discussion is not repeated here; q.v. Section 7.4.2.

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	496	4,585	0.866	(0.225)	110	(12)	(0.48)	15	(9)	(0.51)	-2.11
01	592	5,561	0.867	(0.221)	111	(12)	(0.54)	15	(9)	(0.56)	-2.37
02	784	7,449	0.867	(0.217)	111	(12)	(0.64)	15	(9)	(0.68)	-2.45
03	1,168	11,225	0.867	(0.211)	111	(12)	(0.69)	15	(9)	(0.74)	-2.46
04	2,104	20,457	0.867	(0.209)	111	(12)	(0.67)	15	(9)	(0.91)	-2.70
05	3,712	36,259	0.867	(0.208)	110	(12)	(0.90)	15	(9)	(0.90)	-3.12
06	4,096	40,003	0.867	(0.208)	110	(12)	(0.96)	15	(9)	(0.98)	-3.31
07	5,104	49,835	0.867	(0.206)	110	(12)	(1.13)	15	(9)	(1.19)	-3.39
08	7,504	73,307	0.867	(0.206)	110	(12)	(1.67)	15	(9)	(1.72)	-3.45
09	12,832	125,523	0.867	(0.206)	109	(12)	(2.39)	14	(9)	(2.46)	-3.49

Table 7.26: A summary of results for the DWR-AMR-h strategy for the SI procedure and the GMRes(15) method for the solution of the CBG-IGA 1G SAAF-S_8 equations for the 2D IAEA CEA-2 benchmark; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 7.12. Those results in parentheses are obtained for the CBG-IGA DSA-preconditioned system.

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	496	7,936	0.869	(0.225)	112	(12)	(1.06)	17	(9)	(1.38)	-1.99
01	592	9,472	0.869	(0.224)	111	(12)	(1.08)	15	(9)	(1.47)	-2.17
02	736	11,776	0.869	(0.223)	111	(12)	(1.13)	15	(9)	(1.63)	-2.21
03	1,048	16,768	0.868	(0.216)	111	(12)	(1.22)	15	(9)	(1.68)	-2.22
04	1,528	24,448	0.868	(0.215)	110	(12)	(1.42)	15	(9)	(1.91)	-2.34
05	2,368	37,888	0.867	(0.213)	110	(12)	(1.60)	15	(9)	(2.38)	-2.42
06	3,040	48,640	0.867	(0.211)	110	(12)	(1.64)	15	(9)	(3.01)	-2.62
07	5,104	81,664	0.867	(0.212)	110	(12)	(2.21)	15	(9)	(3.73)	-2.83
08	6,472	103,552	0.867	(0.211)	110	(12)	(2.94)	14	(9)	(3.87)	-3.07
09	10,240	163,840	0.867	(0.211)	110	(12)	(3.07)	14	(9)	(4.17)	-3.68

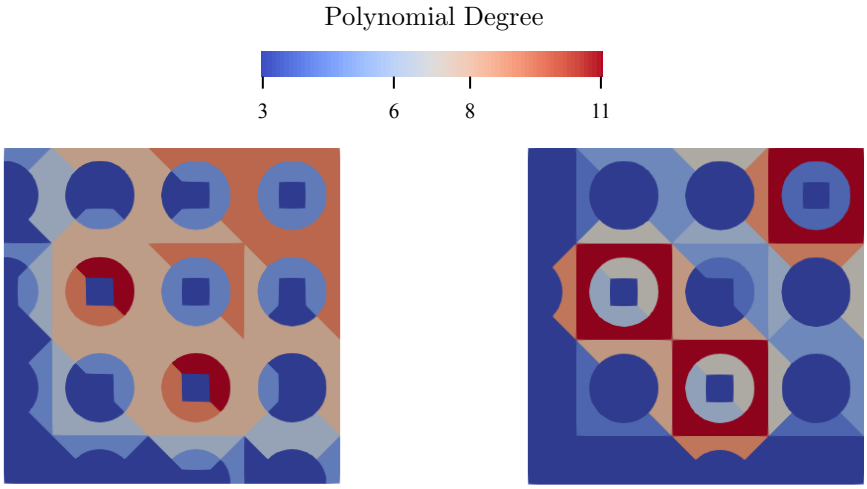
Table 7.27: A summary of results for the DWR-AMR-h strategy for the SI procedure and the GMRes(15) method for the solution of the SIP-DG-IGA 1G SAAF-S_8 equations for the 2D IAEA CEA-2 benchmark; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 7.12. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.

7.5 A Discussion

To those IGA spatial discretisations of the first and second-order forms of the NTE presented in the foregoing chapters, this chapter applied techniques to accelerate the convergence of the within-group self-scattering sources for isotropic scattering-dominated problems; namely, these were the GMRes method and diffusion *preconditioning* schemes. Rather than deriving them consistently from the discretised neutron transport operator, partially-consistent DSA schemes were established by directly discretising the analytical diffusion correction operator, yielding a CBG-IGA DSA scheme for the CBG-IGA SAAF-S_N equations; and separate MIP-DG-IGA DSA schemes for the sweep-based DG-IGA FONT-S_N equations and the SIP-DG-IGA SAAF-S_N equations. In each case, the same approximation space as spanned by NURBS was employed in the spatial discretisation of both the high-order transport

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	496	7,936	0.882	(0.224)	113	(11)	(1.29)	15	(7)	(2.21)	-3.49
01	592	9,472	0.882	(0.223)	113	(11)	(1.09)	15	(7)	(2.19)	-3.57
02	832	13,312	0.882	(0.225)	113	(11)	(1.05)	15	(7)	(2.71)	-3.82
03	1,336	21,376	0.882	(0.228)	112	(11)	(1.13)	14	(7)	(2.69)	-4.28
04	1,960	31,360	0.882	(0.228)	112	(11)	(1.32)	13	(7)	(3.05)	-4.41
05	2,896	46,336	0.881	(0.226)	112	(11)	(1.49)	13	(7)	(3.49)	-4.46
06	4,372	69,952	0.880	(0.224)	111	(11)	(3.05)	12	(7)	(4.81)	-4.47
07	6,136	98,176	0.880	(0.225)	111	(11)	(3.94)	11	(7)	(6.26)	-4.57
08	8,932	142,912	0.880	(0.223)	111	(11)	(5.51)	12	(7)	(6.99)	-4.61
09	12,436	198,976	0.880	(0.223)	111	(11)	(5.71)	11	(7)	(7.93)	-4.71

Table 7.28: A summary of results for the DWR-AMR-h strategy for the SI procedure and the GMRes(15) method for the solution of the sweep-based DG-IGA 1G FONT-S_8 equations for the 2D IAEA CEA-2 benchmark; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 7.12. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.



(a) $\mathcal{T}^{(9)}$ for SAAF-S_8: 15,328 DoF **(b)** $\mathcal{T}^{(9)}$ for FONT-S_8: 19,792 DoF

Figure 7.15: The meshes generated at the 9th AMR iteration for the DWR-AMR-p strategy for each of the CBG-IGA 1G SAAF-S_8 equations (left) and the sweep-based DG-IGA 1G FONT-S_8 equations (right) for the 2D IAEA CEA-2 benchmark; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 7.12. (*Vide* the web-based version for reference to colour.)

approximation and the low-order diffusion one, over matching meshes.

These partially-consistent DSA schemes did exhibit good convergence properties over scattering-dominated homogeneous problems and they were compatible with the aforementioned DWR-AMR-$h(p)$ algorithms over the meshes of a variety of complex rational geometries. Indeed, their introduction of any numerical inconsistencies was largely compensated by their simplicity and computational efficiency; viz. being SPD, the banded matrices of the discretised diffusion acceleration equations were solved by the PCG method. However, as would be inherently true of any fully-consistent or partially consistent DSA scheme, increasingly degraded performance was exhibited for increasingly scattering-dominated problems over increasingly heterogeneous configurations [380]. Nevertheless, the partially-

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	496	4,585	0.866	(0.225)	110	(12)	(0.48)	15	(9)	(0.51)	-2.11
01	496	4,889	0.866	(0.224)	111	(11)	(0.44)	13	(7)	(0.68)	-2.32
02	496	5,245	0.867	(0.222)	111	(11)	(0.64)	13	(7)	(0.44)	-2.44
03	496	5,804	0.867	(0.222)	112	(11)	(0.44)	13	(7)	(0.21)	-2.59
04	496	6,885	0.867	(0.222)	112	(11)	(0.32)	13	(7)	(0.10)	-2.72
05	496	8,355	0.867	(0.222)	112	(11)	(0.11)	12	(7)	(0.07)	-3.01
06	496	9,792	0.867	(0.221)	113	(11)	(0.07)	12	(7)	(0.07)	-3.19
07	496	11,850	0.867	(0.222)	112	(11)	(0.05)	11	(7)	(0.05)	-3.33
08	496	13,492	0.867	(0.223)	112	(12)	(0.05)	11	(7)	(0.03)	-3.45
09	496	15,328	0.867	(0.224)	112	(12)	(0.04)	11	(7)	(0.03)	-3.57

Table 7.29: A summary of results for the DWR-AMR-p strategy for the SI procedure and the GMRes(15) method for the solution of the CBG-IGA 1G SAAF-S_8 equations for the 2D IAEA CEA-2 benchmark; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 7.12. Those results in parentheses are obtained for the CBG-IGA DSA-preconditioned system.

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	496	7,936	0.869	(0.225)	112	(12)	(1.06)	17	(9)	(1.38)	-1.99
01	496	8,224	0.870	(0.226)	114	(13)	(0.88)	16	(8)	(1.79)	-2.21
02	496	8,720	0.871	(0.224)	114	(13)	(1.23)	16	(8)	(1.05)	-2.34
03	496	9,456	0.870	(0.226)	114	(13)	(0.87)	16	(8)	(0.48)	-2.51
04	496	10,992	0.870	(0.226)	113	(13)	(0.68)	16	(8)	(0.21)	-2.66
05	496	13,072	0.871	(0.226)	114	(13)	(0.19)	15	(8)	(0.19)	-2.97
06	496	15,014	0.872	(0.225)	114	(13)	(0.12)	15	(8)	(0.22)	-3.17
07	496	17,806	0.871	(0.225)	114	(12)	(0.09)	16	(8)	(0.16)	-3.34
08	496	19,868	0.871	(0.225)	114	(12)	(0.09)	15	(8)	(0.09)	-3.50
09	496	22,120	0.870	(0.225)	114	(12)	(0.05)	15	(8)	(0.07)	-3.62

Table 7.30: A summary of results for the DWR-AMR-p strategy for the SI procedure and the GMRes(15) method for the solution of the SIP-DG-IGA 1G SAAF-S_8 equations for the 2D IAEA CEA-2 benchmark; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 7.12. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.

consistent DSA schemes still served as effective preconditioners for the more sophisticated GMRes method. Alone, the GMRes method did well to accelerate consistently the convergence of the SI procedure; although, it did not always perform as well as the DSA-SI scheme. The DSA-GMRes scheme was always the more performant for a variety of test cases.

The time spent solving the linear systems of algebraic equations was reduced by tightly converging the transport solution on the coarsest initial reference mesh, at low-computational cost, and then *bootstrapping* it from one AMR iteration to the next; similarly, this helped to reduce the relative CPU time spent solving the diffusion acceleration equations [387]. However, one recalls that the banded matrices were rendered increasingly ill conditioned for an an increasingly large polynomial degree of the NURBS basis; and multi-grid preconditioners developed for the general bases of classical FEM

(i)	No. Elem	DoF	NSR		SI No. Its		SI time (%)	GMRes(15) No. Its		GMRes(15) time (%)	Log(ϵ_{QoI})
00	496	7,936	0.882	(0.224)	113	(11)	(1.29)	15	(7)	(2.21)	-3.49
01	496	8,224	0.882	(0.221)	114	(10)	(1.10)	13	(7)	(1.64)	-3.61
02	496	9,048	0.882	(0.222)	114	(10)	(0.95)	13	(7)	(0.57)	-3.77
03	496	9,824	0.882	(0.218)	114	(10)	(0.94)	13	(7)	(0.49)	-3.89
04	496	11,352	0.882	(0.214)	112	(9)	(0.83)	13	(7)	(0.45)	-4.07
05	496	12,600	0.882	(0.212)	112	(9)	(0.80)	12	(7)	(0.44)	-4.19
06	496	13,816	0.882	(0.211)	112	(10)	(1.24)	11	(7)	(0.42)	-4.30
07	496	15,176	0.882	(0.213)	112	(11)	(0.84)	11	(8)	(0.44)	-4.42
08	496	17,712	0.881	(0.211)	111	(11)	(0.86)	11	(8)	(0.47)	-4.56
09	496	19,792	0.880	(0.219)	110	(10)	(0.98)	11	(8)	(0.58)	-4.63

Table 7.31: A summary of results for the DWR-AMR-p strategy for the SI procedure and the GMRes(15) method for the solution of the sweep-based DG-IGA 1G FONT-S_8 equations for the 2D IAEA CEA-2 benchmark; the QoI is the absorption in Material 1 in the top right-hand corner; v. Figure 7.12. Those results in parentheses are obtained for the SIP-DG-IGA DSA-preconditioned system.

appeared to lose their effectiveness. The effect was more pronounced for the discretised SAAF-S_N equations; and thus, with each subsequent AMR iteration, increasingly more time and more computational effort was spent solving them, rather than the diffusion acceleration equations. Conversely, with respect to the DWR-AMR-p strategy for the sweep-based DG-IGA FONT-S_N equations, there was relatively little additional computational effort demanded of the matrix-free solution approach, since the number of elements belonging to the mesh would remain unchanged; and thus, with each subsequent AMR iteration, increasingly more time was spent solving an increasingly ill-conditioned linear system of the diffusion acceleration equations.

For demonstrative purposes, this chapter focussed expressly on the acceleration of the SI procedure for one-group, isotropic, fixed-source problems. Nonetheless, another acceleration technique has been implemented in the new modern-Fortran neutron transport code developed for the purpose of this current research. Analogous to a Krylov projection method for a subspace of reduced and limited size, $M = 3$, a simple Chebyshev k-eigenvalue extrapolation was employed to accelerate the convergence of the k_{eff} for each of the eigenvalue problems in the foregoing chapters [410]. For the eigenvalues of the N-dimensional discrete system, $|k_1| > |k_2| > \ldots |k_N|$, with $k_1 = k_{\text{eff}}$, the dominance ratio, $r = k_2/k_1$, determines the convergence properties of the PI procedure, for $r < 1$. Not dissimilar to the behaviour of the SI procedure as the scattering ratio approaches unity, the PI procedure can become arbitrarily slow to converge for optically-thick fission regions where $r \to 1$. For a good estimation of the dominance ratio, as approximated by Equation 7.8, for each of the problems investigated, the k_{eff} converged in at least twice as many unaccelerated power iterations as many accelerated ones.

Chapter 8

The Conclusion

This final chapter summarises the findings of the current thesis, upon which is based the proposal for some prospective research.

8.1 Some Concluding Remarks

This thesis developed IGA-based spatial discretisation methods and AMR-$h(p)$ algorithms for the first-order and second-order forms of the NTE, the multi-group FONT-S_N and the multi-group SAAF-S_N equations, respectively; as implemented within a new modern-Fortran neutron transport code.

Chapter 1 set the scope for the current research and provided the motivation for the development of high-fidelity computational neutron transport codes, the basics of which were introduced along with some useful quantities. Chapter 2 introduced the NTE in several of its most common partial integral-differential forms; and in the theme of deterministic methods, a few legacy techniques were presented for their discretisation with respect to each of the independent variables. Whilst it was not the intention for this thesis to compare and contrast amongst themselves the different first and second-order forms of the NTE, some points are worth calling to attention.

An inherent defect of the S_N method, the ray effects, are due to a lack of rotational invariance in the discretised streaming operator; q.v Section 2.5.2.1. Unsurprisingly, for low-order angular quadrature sets, both the multi-group SAAF-S_N and the multi-group FONT-S_N equations were shown to exhibit such aberrations and anomalies in the presence of scatter and fission sources. Nevertheless, the S_N method does lend itself to a relatively simple fixed-point solution strategy; q.v. Section 2.8. Whilst the overall SI procedure does not change under different spatial discretisations, its convergence properties are strongly affected by it; q.v. Section 2.8.1. Importantly, the manner in which the within-group

equations are solved for in a given direction of neutron flight is strongly dependent upon the the spatial discretisation, which itself, is determined by the nature of the governing equations; viz. this affects *how* the group-ordinate coefficient matrices of the corresponding discrete system are inverted.

Second-order forms of the steady-state NTE are elliptic PIDE; they have complex characteristics and govern equilibrium problems. An appropriate solution strategy must ensure that all internal requirements and auxiliary conditions are satisfied simultaneously since the proposed discretisation methods yield *irreducible* matrices. Yet, employing such *smoothing* techniques can numerically violate causality, a rigorous property of neutron transport; q.v. Section 2.3.1; and the use of efficient sparse-matrix data storage structures become increasingly important for increasingly refined problems, lest their computational memory requirements become prohibitive; q.v. Section 2.8.1. For the proposed AMR-*hp* strategies, it was observed that the resulting banded matrices became increasingly ill conditioned for increasingly large polynomials of the NURBS basis; thus, demanding increasingly more computational time and effort to invert the group-ordinate-wise coefficient matrices. Furthermore, without some special treatment for voids, second-order forms of the NTE applied to advection-dominated problems can become increasingly undefined in increasingly optically-thin regions; q.v. Section 5.5.2.

The FONT equation is a hyperbolic PIDE. Unfortunately, such equations are not amenable to a CBG spatial discretisation; q.v. Section 6.3. Whilst DG spatial discretisations accrue at least twice as many DoF as their CBG counterpart over a similar mesh, they also yield *reducible* matrices; and thus, the resulting discrete system can be solved for in a matrix-free fashion; q.v. Section 2.8.1. Moreover, since hyperbolic equations have real characteristics and govern propagation problems, it becomes natural to *follow the neutron* as per a marching scheme. For the proposed AMR-*hp* strategies, it was observed that little additional computational effort was demanded at each subsequent AMR iteration, for there was no dramatic increase in the number of elements to sweep through. Nevertheless, for meshes that contain skewed, curvilinear or other re-emergent elements, the digraph of a given direction of neutron flight can become acyclic; whereby, permutations of the corresponding group-ordinate coefficient matrices will not be strictly block lower-triangular; and therefore the angular neutron flux distribution could not be completely resolved in one transport sweep. Furthermore, the FONT equation is not self-adjoint, which means that only saddle-point principles can be derived; and not extremum-variational ones [99].

The actual spatial discretisation of choice was reserved for Chapter 3, where NURBS-based IGA was introduced. A global uniform *hp*-refinement strategy was presented as a point of reference. Furthermore, *two* different and cheap interpolation-based error measures were devised, with the intent of driving *two* different AMR-$h(p)$ strategies; each with the capacity for energy-dependent meshes. IGA provides a good framework within which to implement such AMR strategies because the parametrisa-

tion does not vary under refinement and thus, the underlying geometry is always exactly represented. The NRG-AMR strategy heuristically minimises any local contributions to the global discretisation error in the primal problem, as per some appropriate user-prescribed norm; whereas the DWR-AMR strategy employs duality arguments to minimise *important* local contributions to the error as measured in some user-defined QoI; however, this comes at the additional cost of solving the dual problem. Whilst separately, both AMR-h and AMR-p strategies may perform better than their respective and corresponding global uniform refinement strategies, a combined AMR-hp strategy addresses both the non-smooth and the smooth parts of the solution. For those elements that were *flagged* for further refinement, the optimal choice for local refinement was determined as that which maximised the rate of decrease in the local error contributions per increase in the number DoF.

The application of the automised NRG-AMR-$h(p)$ and the DWR-AMR-$h(p)$ strategies to separate CBG-IGA and SIP-DG-IGA spatial discretisations of the multi-group NDE and the multi-group SAAF-S_N equations was presented in Chapter 4 and Chapter 5, respectively. For the continuous discretisations of these elliptic systems, the potential for higher-order continuity within a NURBS patch was sacrificed in order to facilitate the wholly-local nature of the AMR algorithms; where continuity in solution field across irregular interfaces was enforced via constraint equations based upon master-slave relationships and the conservative interpolation between *surface* meshes. A similar Galerkin projection was presented in Chapter 3 for the conservative interpolation between *volume* meshes when evaluating the group-to-group source terms over energy-dependent meshes; and similarly so for the calculation of the interpolation-based error measures. Yet again, a similar Galerkin projection was presented in Chapter 6 for the conservative interpolation between *surface* meshes for the projection of the upwinded fluxes in the sweep-based DG-IGA discretisation of the FONT-S_N equations. For the SIP scheme for the discontinuous discretisations, *optimised* penalty parameters, based upon rigorous coercivity analyses, were derived for general element types to ensure sufficient, yet not excessive, penalisation of the discrete solution field across the spatial mesh(es). Allowing for the imposition of a certain degree of regularity to avoid any multiply constraints, the AMR-$h(p)$ algorithms for both of the CBG-IGA and the SIP-DG-IGA discretisations did generate qualitatively similar meshes for similar DoF. Inevitably, enforcing continuity of a discrete field over irregular meshes does introduce some errors into the discretisation; whilst never prohibitive, these errors do appear to be more significant for the constraint-based approach over non-Cartesian grids. Moreover, one recalls that DG discretisations typically do offer more flexibility in terms of allowing for local mesh refinement.

The application of the automised NRG-AMR-$h(p)$ and the DWR-AMR-$h(p)$ strategies to a sweep-based DG-IGA spatial discretisation of the multi-group FONT-S_N equations was presented in Chapter 6. Due to the nature of hyperbolic systems, CBG discretisations of the FONT-S_N equations are unstable for heterogeneous configurations; rather, they are more amenable to upwinding schemes that

follow the propagation of information in specific directions. Any SCC arising in the digraphs of the spatial mesh were resolved by systematically deleting the weakest outgoing edges until an acyclic directed path was obtained; and any unresolved dependencies were simply lagged within the inner-iteration procedure. Although it is not possible to *invert* exactly the discrete linear system by simply visiting each element belonging to the mesh only once per modified transport sweep in a given direction of neutron flight, it was found to be the most simple and computationally efficient option [231, 247].

In each case, consistency of the different discretisations of the different forms of the NTE considered in this thesis, as well as the different mechanisms of the AMR algorithms, were verified by means of the MMS. For more-complicated heterogenous configurations, several nuclear reactor physics and shielding problems were also presented. For a given localised QoI, thus, with the exception of global quantities such as the k_{eff}, the NRG-AMR-h strategy tended to be less performant than the global uniform h-refinement one, since the refinement of the mesh is focussed on resolving any sharp gradients in the solution, regardless of where they may be. In that respect, the DWR-AMR strategies consistently and reliably reduced the error in the QoI at each subsequent AMR iteration; whereas the NRG-AMR strategies would do so only sporadically. Howbeit, the AMR-hp strategies were shown to yield the best accuracy per DoF; with the DWR-AMR-hp being the most performant. That being said, the individual group-ordinate-wise banded matrices for the SAAF-S_N equations do become increasingly more ill conditioned for increasingly large polynomial degrees of the NURBS basis; and with each AMR iteration, increasingly more computational effort is required to perform the physical transport sweeps, since the algebraic multi-grid preconditioner developed for the general bases of classical FEM appear to become increasingly less effective. Conversely, for no dramatic increase in the number of elements belonging to the mesh, little additional computational effort was demanded of similar AMR-$h(p)$ strategies for the matrix-free approach of the sweep-based DG-IGA FONT-S_N equations.

Chapter 7 saw the application of acceleration techniques to those IGA spatial discretisations of the FONT-S_N equations and the SAAF-S_N equations; namely, the GMRes method and a diffusion *preconditioning* (DSA) scheme were introduced to accelerate the convergence of the within-group self-scattering sources for isotropic scattering-dominated problems. One recalls that any acceleration was understood as that with respect to the SI procedure; viz. converging to the same solution in a fewer number of iterations. Rather than deriving them consistently from the discretised neutron transport operator, partially-consistent DSA schemes were established by directly discretising the analytical diffusion correction operator, yielding a CBG-IGA DSA scheme for the CBG-IGA SAAF-S_N equations; and separate MIP-DG-IGA DSA schemes for the sweep-based DG-IGA FONT-S_N equations and the SIP-DG-IGA SAAF-S_N equations [387]. In each case, the same approximation space as spanned by NURBS was employed in the spatial discretisation of both the high-order transport approximation and the low-order diffusion one, over matching meshes.

The partially-consistent DSA schemes did exhibit good convergence properties over scattering-dominated homogeneous problems and they were compatible with the aforementioned DWR-AMR-$h(p)$ algorithms over the meshes of a variety of complex rational geometries. Indeed, their introduction of any numerical inconsistencies was largely compensated by their simplicity and computational efficiency; viz. being SPD, the banded matrices of the discretised diffusion acceleration equations were solved by the PCG method. However, as would be inherently true of any fully-consistent or partially consistent DSA scheme, increasingly degraded performance was exhibited for increasingly scattering-dominated problems over increasingly heterogeneous configurations [380]. Nonetheless, the partially-consistent DSA schemes still served as effective preconditioners for the more sophisticated GMRes method. Whilst, alone, the GMRes method did well to accelerate consistently the convergence of the SI procedure, it did not always perform as well as the DSA-SI scheme; although the DSA-GMRes scheme was always the more performant for a variety of test cases.

Generally, the relative time spent solving the diffusion acceleration equations was kept low for the DWR-AMR-h strategy; this was thanks, in part, to the *bootstrapping* of the reference solution from one AMR iteration to the next. However, the banded matrices were rendered increasingly ill conditioned for an an increasingly large polynomial degree of the NURBS basis and multi-grid preconditioners developed for the general bases of classical FEM appeared to lose their effectiveness. The effect was more pronounced for the discretised SAAF-S_N equations; and thus, with each subsequent AMR iteration, increasingly more time and computational effort was spent solving them than the diffusion acceleration equations. Conversely, in the case of the DWR-AMR-p strategy for the sweep-based DG-IGA FONT-S_N equations, there was relatively little additional computational effort demanded of the matrix-free solution approach, since the number of elements belonging to the mesh would remain unchanged; and thus, with each subsequent AMR iteration, increasingly more time was spent solving an increasingly ill-conditioned linear system of diffusion acceleration equations.

8.2 Prospective Research

Finally, one proposes some areas of prospective research based upon the work developed and presented in the current thesis.

8.2.1 Improved Geometry Creation

The first and most obvious extension of the current research is that in the application to full 3D problems. Whilst the underlying discretisation methods and the refinement algorithms would remain

unchanged, the task of resolving any dependencies within the CBG-IGA discretisations over irregular meshes would become evermore complicated; viz. the elementary stencil can become large when establishing constraint equations between those patches that share an irregular face, as well as those that share an irregular edge.

For the relatively simple geometries of each of the test cases presented, the coarsest meshes were constructed with bespoke software. As per the ultimate goal of IGA, these and more general geometries should instead be created with some CAD software and their descriptions passed directly to the analysis program; e.g. in neutral IGES or STEP file format. In order to simplify the initial connectivity information, the coarsest meshes were made to be conforming and as regular as possible. Of course, this constraint could be removed to facilitate the processing of *dirty* CAD descriptions.

8.2.2 Alternative Bases for Isogeometric Methods

Whilst the industry standard for CAD software is to use NURBS surfaces in the boundary representation of 3D objects, they are not exclusive to IGA methods. Indeed, other more-exotic bases should be investigated; e.g. LR B-splines, T-splines, U-splines and Trigonometric NURBS [261–263,285,411,412]. Derived from generalised B-splines, the latter would allow for rough surfaces to be modelled as a series of trigonometric functions [412]. This could not be represented by a basis of NURBS; yet, trigonometric NURBS could represent exactly any NURBS description of a conic section or quadric surface.

8.2.3 Additional Spatial Refinement Strategies

Based upon local H^1 and L_2-norms, computationally *cheap* error measures were devised to drive the proposed AMR-$h(p)$ algorithms. These required only the isotropic components of the numerical solution to be stored. Rather than refining any given element in each parametric direction, an anisotropic AMR-$h(p)$ algorithm could be driven by conveying any anisotropies in the error measures through their direction derivatives [257]. However, no formal error estimate was derived for the error remaining in the spatial discretisations of the different forms of the NTE; and such an approach may prove to be effective for anisotropic shielding problems. In which case, one should make comparisons with more rigorous, angular-dependent, residual-based error estimates [255].

At the most, DSA schemes have the potential to accelerate only the zeroth and the first-order components of the $\mathrm{P_L}$-expansion of the scattering source; viz. the isotropic and linearly anisotropic components. To accelerate the higher-order components of the spherical harmonic flux moments for forward-peaked scattering problems, one could envisage a multi-level angular multi-grid approach,

where an additive correction term for arbitrary-order spherical harmonic flux moments is built from recursively solving the $S_{N/2}$ equations, with the potential of *filtering out* any high-frequency portions with the lowest-level diffusion acceleration equations [400, 413].

One could also incorporate the capacity for mesh coarsening into the AMR-$h(p)$ strategies.

8.2.4 Adaptive Angular Refinement Strategies

The proposed AMR-$h(p)$ strategies were devised to minimise iteratively those errors due to the spatial discretisation. However, these can be strongly coupled to the angular discretisation; and especially so for dominating ray effects.

To allow for improved fidelity in the use of energy-dependent meshes, one could solve for the multi-group components of the scalar neutron flux distribution over each spatial mesh with different N-order angular quadrature sets. The underlying algorithms and mesh-to-mesh interpolation routines would remain unchanged, since the group-to-group sources are only constructed from the spherical harmonic flux moments. Angular-dependent error measures could then be devised to determine those energy groups that require better resolution of their angular discretisation.

A more sophisticated approach to ray-effect mitigation would be to incorporate spherical NURBS-based IGA methods in the discretisation of the angular variable, which would allow for the development of full space-angle AMR-$h(p)$ strategies [164, 414].

8.2.5 Non-Linear Diffusion Acceleration

Naturally, the acceleration techniques proposed in Chapter 7 could be extended to multi-group formulations and eigenvalue problems; however, it becomes increasingly more important to accelerate the group-to-group sources in the outer (power) iterations for an increasingly fine group structure.

Where linear DSA schemes for fixed-source problems were implemented as additive corrections from source-corrected NDE, non-linear diffusion acceleration (NDA) schemes should be devised to accelerate the k_{eff} for eigenvalue problems [80, 183, 395].

8.2.6 Parallelisation

A rather basic distributed-parallel computing architecture was devised for the those methods developed in this thesis and implemented in a new modern-Fortran neutron transport code. A more sophisticated

approach would include for spatial domain decomposition that would allow for massive scalability on large modern HPC systems [172, 415]. The local adaptive h-refinement algorithms naturally give rise to a concentration of elements about specific regions of interest; this could easily lead to load balancing issues unless the domain were decomposed differently between AMR iterations [416].

The Bibliography

[1] C. Böhringer. *The Kyoto Protocol: A Review and Perspectives*. Report No. 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] J. Ambrose. *National Grid warns of short supply of electricity over next few days*. URL https://www.theguardian.com/business/2020/oct/14/national-grid-warns-of-short-supply-of-electricity-over-next-few-days, 2020. The Guardian, Online; last accessed on 16/10/2020.

[5] *The UK's Nuclear Future*. Report No. BIS/13/67, HM Government, 2013.

[6] *Projected Costs of Generating Electricity:* 2005 *Update*. Technical report, OECD/IEA NEA, 2005.

[7] A. Einstein. *Ist die Trägheit eines Körpers von seinem Energiehalt abhängig? Annalen der Physik*, 323(13):639–641, 1905.

[8] J. Chadwick. *Existence of a Neutron. Proc. of the Royal Society*, 136(830):692–708, 1932.

[9] O. Hahn and F. Strassman. *Über den Nachweis und das Verhalten der bei der Bestrahlung des Urans mittels Neutronen entstehenden Erdalkalimetalle. Naturwissenschaften*, 27:11–15, 1939.

[10] L. Meitner and O. R. Frisch. *Disintegration of Uranium by Neutrons: A New Type of Nuclear Reaction. Nature*, 143:239–240, 1939.

[11] H. Poincaré. *La théorie de Lorentz et le principe de réaction. Archives Néerlandaises des Sciences Exactes et Naturelles*, 5:252–278, 1900.

[12] A. Einstein. *Zur Elektrodynamik bewegter Körper. Annalen der Physik*, 17:891–, 1905.

[13] *The Economics of Nuclear Power*. Information Papers, World Nuclear Association, 2008.

[14] *International Energy Outlook* 2019 *with projections to* 2050. Technical report, U.S. Energy Information Administration, 2019.

[15] M. Marshall. *Hinkley Point reactors herald new nuclear age for UK. New Scientist*, 220:6–7, 2013.

[16] *Generic design assessment of Hitachi-GE Nuclear Energy limited's UK Advanced Boiling Water Reactor: Decision Document*. Technical report, Environment Agency and Natural Resources Wales, 2017.

[17] J. Pickard, D. Sheppard, and K. Inagaki. *Hitatchi preparing to pull out of nuclear project in blow to UK climate ambitions*. URL https://www.ft.com/content/1a079f1e-6cfe-4d27-ab8f-df3b5b93e198, 2020. The Financial Times, Online; last accessed on 16/09/2020.

[18] J. McCurry. *Hitachi to pull plug on north Wales nuclear power station*. URL https://www.theguardian.com/business/2020/sep/15/hitachi-to-pull-plug-on-north-wales-nuclear-power-station, 2020. The Guardian, Online; last accessed on 16/09/2020.

[19] *Advanced Modular Reactor (AMR) Feasibility and Development Project: Abstracts from the Applicant Organisation's Proposals*. Technical report, Department for Business, Energy & Industrial Strategy, 2018.

[20] G. Locatelli, M. Mancini, and N. Todeschini. *Generation IV nuclear reactors: Current status and future projects. Energy Policy*, 61:1503–1520, 2013.

[21] *UK SMR - A National Endeavour*. Technical report, Rolls-Royce, 2017.

[22] *Royal Navy nuclear Reactor Test Facility Review.* Technical report, Ministry of Defence, 2014.

[23] 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.

[24] A. Santamarina, D. Bernard, P. Blaise, P. Leconte, C. Vaglio, and J-F. Vidal. *Recent advances in accurate neutronic calculation of GEN-3 LWR reactors. Annals of Nuclear Energy*, 87:68–75, 2016.

[25] D. Gaston, C. Newman, G. Hansen, and D. Lebrun-Grandie. *MOOSE: A parallel computational framework for coupled systems of nonlinear equations. Nuclear Engineering and Design*, 239:1768–1778, 2009.

[26] C. J. Permann, D. R. Gaston, D. Andrš, R. W. Carlsen, F. Kong, A. D. Lindsay, J. M. Miller, J. W. Peterson, A. E. Slaughter, R. H. Stogner, and R. C. Martineau. *MOOSE: Enabling massively parallel multiphysics simulation. SoftwareX*, 11:100430, 2020.

[27] G. C. Pomraning. *The Equations of Radiation Hydrodynamics.* Pergamon Press, 1973.

[28] R. A. Knief. *Nuclear Criticality Safety: Theory and Practice.* American Nuclear Society, first edition, 1985.

[29] Y. V. Nazarov and Y. M. Blanter. *Quantum Transport: Introduction to Nanoscience.* Cambridge University Press, 2009.

[30] K. J. Moody, P. M. Grant, and I. D. Hutcheon. *Nuclear Forensic Analysis.* CRC press, 2014.

[31] R. J. Vetter and M. S. Stoeva. *Radiation Protection in Medical Imaging and Radiation Oncology.* Medical Physics and Biomedical Engineering. CRC press, first edition, 2015.

[32] L. E. Boltzmann. *Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen. Sitzungsberichte der Mathematisch-Naturwissenschaftlichen Classe der Kaiserlichen Akademie der Wissenschaften*, 66:275–370, 1872.

[33] L. S. Ornstein and G. E. Uhlenbeck. *Some Kinetic Problems Regarding the Motion of Neutrons Through Parafine. Physica*, 4(6):478–486, 1937.

[34] T. Rockwell. *Reactor Shielding Design Manual.* Van Nostrand Reinhold Inc., 1956.

[35] V. P. Mashkovich. *Error Estimation in Reactor Shielding Calculations.* AIP Translation Series. American Institute of Physics, 1988.

[36] M. W. Thompson. *Defects and Radiation Damage in Metals.* Cambridge Monographs on Physics. Cambridge University Press, 1974.

[37] A. D. Amaev, I. V. Gorynin, V. A. Nikolaev, and N. N. Alekseenko. *Radiation Damage of Nuclear Power Plant Pressure Vessel Steels.* American Nuclear Society, first edition, 1997.

[38] J. A. Trangenstein. *Numerical Solution of Hyperbolic Partial Differential Equations.* Cambridge University Press, 2009.

[39] E. W. Larsen. *Asymptotic Methods for Continuous and Discretized Transport Problems.* University of Michigan, 1991.

[40] M. L. Adams and E. W. Larsen. *Fast Iterative Methods for Discrete Ordinates Particle Transport Calculations. Progress in Nuclear Energy*, 40:3–159, 2002.

[41] D. G. Cacuci. *Two dimensional corner singularities in neutron diffusion - Part 1: Analysis. Nuclear Science and Engineering*, 128:1–16, 1998.

[42] D. G. Cacuci. *Two dimensional corner singularities in neutron diffusion - Part 2: Application to the SNR-300 Benchmark. Nuclear Science and Engineering*, 128:17–26, 1998.

[43] M. M. R. Williams. *The Wiener-Hopf Technique: An alternative to the singular eigenfunction method. Advances in Nuclear Science and Technology*, 7:283–327, 1973.

[44] N. J. McCormick and I. Kuščer. *Singular eigenfunction expansions in neutron transport theory. Advances in Nuclear Science and Technology*, 7:181–282, 1973.

[45] J. Lewins. *IMPORTANCE - The Adjoint Function. The Physical Basis of Variational and Pertubation Theory in Transport and Diffusion Problems.* Pergamon Press, first edition, 1965.

[46] N. Wirth. *Algorithms and Data Structures.* Prentice-Hall, second edition, 1986.

[47] J. J. Duderstadt and L. J. Hamilton. *Nuclear Reactor Analysis.* John Wiley & Sons, Inc., 1976.

[48] J. C. Maxwell. *Illustrations of the dynamical theory of gases. Part 1. On the motions and collisions of perfectly elastic spheres. The London, Edinburgh and Dublin Philosophical Magazine and Journal of Science*, 19:19–32, 1860.

[49] J. C. Maxwell. *Illustrations of the dynamical theory of gases. Part 2. On the process of diffusion of two or more kinds of moving particules among one another. The London, Edinburgh and Dublin Philosophical Magazine and Journal of Science*, 20:21–37, 1860.

[50] J. W. Gibbs. *On the Fundamental Formula of Statistical Mechanics, with Applications to Astronomy and Thermodynamics*. 33:57–58, 1884.

[51] L. Pál. *On the Theory of Stochastic Processes in Nuclear Reactors. Il Nuovo Cimento*, 7(1):25–42, 1958.

[52] G. I. Bell. *On the Stochastic Theory of Neutron Transport. Nuclear Science and Engineering*, 21(3):390–401, 1965.

[53] J. R. Lamarsh. *Introduction to Nuclear Reactor Theory*. American Nuclear Society, Inc., 2002.

[54] C. Demazière. *Modelling of Nuclear Reactor Multi-Physics: From Local Balance Equations to Macroscopic Models in Neutronics and Thermal-Hydraulics*. Academic Press, 2020.

[55] B. G. Carlson and K. D. Lathrop. *Transport Theory: The Method of Discrete Ordinates*. Report No. LA-3251, Los Alamos Scientific Laboratory of the University of California, 1965.

[56] *The JEFF-3.1.1 Nuclear Data Library*. JEFF Report No. 22, OECD/NEA Data Bank, 2009.

[57] M. B. Chadwick, M. Herman, P. Obložinksý, 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. Buber, 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.

[58] R. E. MacFarlane and A. C. Kahler. *Methods for Processing ENDF/B-VII with NJOY. Nuclear Data Sheets*, 111:2739–2890, 2010.

[59] B. E. Watt. *Energy Spectrum of neutrons from Thermal Fission of U235. Physical Review*, 87(6):1037–1041, 1952.

[60] R. E. MacFarlane, D. W. Muir, R. M. Boicourt, and A. C. Kahler. *The NJOY Nuclear Data Processing System, Version 2012*. Report No. LA-UR-12-27079, Los Alamos National Laboratory, 2012.

[61] E. E. Lewis. *Fundamentals of Nuclear Reactor Physics*.

[62] 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.

[63] M. P. Knight, P. Bryce, and S. K. Hall. *WIMS/PANTHER Analysis of $UO2/MOX$ Cores Using Embedded Supercells. Nuclear Technology*, 183:398–408, 2012.

[64] A. Santamarina. *From Integral Experiments to Nuclear Data Improvement*. In *Proc. of the International Conference on Nuclear Data for Science and Technology*, 2007.

[65] A. Sood, R. A. Forster, and D. K. Parsons. *Analytical Benchmark Test Set for Criticality Code Verification. Progress in Nuclear Energy*, 42(1):55–106, 2003.

[66] A. Santamarina, D. Bernard, P. Blaise, P. Leconte, J-M. Palau, B. Roque, C. Vaglio, and J-F. Vidal. *Validation of the New Code Package APOLLO2.8 for Accurate PWR Neutronics Calculations*. In *Proc. of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering (M&C2013)*, 2013.

[67] K. Salari and P. Knupp. *Code Verification by the Method of Manufactured Solutions*. Report No. SAND2000-1444, Sandia National Laboratories, 2000.

[68] J. S. Warsa. *A continuous finite element-based, discontinuous finite element method for S_N transport. Nuclear Science and Engineering*, 160:385–400, 2008.

[69] Y. Wang and J. Ragusa. *On the convergence of DGFEM applied to the discrete ordinates transport equation for structured and unstructured triangular meshes*. *Nuclear Science and Engineering*, 163:56–72, 2009.

[70] A. R. Owens. *Discontinuous Isogeometric Analysis Methods for the First-Order Form of the Neutron Transport Equation with Discrete Oridnate Angular Discretisation*. PhD Thesis, Imperial College London, 2017.

[71] J. A. Welch. *Isogeometric Analysis for Second-Order Forms of the Neutron Transport Equation with Applications to Reactor Physics*. PhD Thesis, Imperial College London, 2018.

[72] C. Latimer. *Self-Adaptive Isogeometric Discretisations of the Second-Order Forms of the Neutron Transport Equation with Dual-Weighted Residual or Goal-Based Error Measures*. PhD Thesis, Imperial College London, 2020.

[73] M. Metcalf, J. Reid, and M. Cohen. *Modern Fortran Explained*. Numerical Mathematics and Scientific Computation. Oxford University Press, 2011.

[74] A. Hébert. *Applied Reactor Physics*. Presses Internationales Polytechnique, second edition, 2016.

[75] E. E. Lewis and W. F. Miller, Jr. *Computational Methods of Neutron Transport*. John Wiley & Sons, Inc., 1984.

[76] P. Reuss. *Précis de neutronique*. EDP Sciences, 2003.

[77] J. J. Duderstadt and W. R. Martin. *Transport Theory*. John Wiley & Sons, Inc., 1979.

[78] H. Bateman. *Solution of a system of differential equations occuring in the theory of radioactive transformations*. In *Proc. of the Cambridge Philosophical Society*, volume 15, pages 423–427, 1910.

[79] D. G. Cacuci, editor. *Handbook of Nuclear Engineering: Nuclear Engineering Fundamentals*, volume 1. Springer, 2010.

[80] R. D. O'Dell and R. E. Alcouffe. *Transport Calculations for Nuclear Analyses: Theory and Guidelines for Effective Use of Transport Codes*. Report No. LA-10983-MS, Los Alamos National Laboratory, 1987.

[81] G. L. Bell and S. Glasstone. *Nuclear Reactor Theory*. Van Nostrand Reinhold Company, 1970.

[82] I. Lux and L. Koblinger. *Monte Carlo particle transport methods: neutrons and photons calculations*. 1991.

[83] D. B. Fromowitz. *Batch and History Sampling for Fixed-Source Monte Carlo Problems*. *Nuclear Science and Engineering*, 187(2):142–153, 2017.

[84] S. D. Richards, C. M. J. Baker, A. J. Bird, P. Cowan, N. Davies, G. P. Dobson, T. C. Fry, A. Kyrieleis, and P. N. Smith. *MONK and MCBEND: Current status and recent developments*. *Annals of Nuclear Energy*, 82:63–73, 2015.

[85] T. Shuttleworth, M. Grimstone, and S. Chucas. *Application of Acceleration Techniques in MCBEND*. *Journal of Nuclear Science and Technology*, 37:406–411, 2000.

[86] E. Brun, F. Damian, C. M. Diop, E. Dumonteil, F. X. Hugot, C. Jouanne, Y. K. Lee, F. Malvagi, A. Mazzolo, O. Petit, J. C. Trama, T. Visonneau, and A. Zoia. *TRIPOLI-4®, CEA, EDF and AREVA reference Monte Carlo code*. *Annals of Nuclear Energy*, 82:151–160, 2015.

[87] W. M. Stacey. *Nuclear Reactor Physics*. Wiley-VCH, third edition, 2018.

[88] *WIMS - A General Purpose Code for Reactor Core Physics*. Technical report, The ANSWERS Software Service WIMS, AEA Technology plc, 1995.

[89] A. Jonsson. *THESEUS: A One Group Collision Probability Code with General Boundary Conditions*. Report No. AEEW-R 253, IAEA, 1963.

[90] H. H. W. Pitcher. *A program for calculating collision probabilities in R-Z geometry*. Report No. AEEW-M 469, IAEA, 1964.

[91] D. W. Anderson. *The WIMS-E Module W-PIJ: A Two-Dimensional Collision Probability Code with General Boundary Conditions*. Report No. AEEW-R 862, IAEA, 1973.

[92] M. J. Roth. *The estimation of collision probabilities in complicated geometries*. Report No. AEEW-M 857, IAEA, 1969.

[93] J. R. Askew, F. J. Fayers, and P. B. Kemshell. *General Description of the Lattice code WIMS*. *Journal of the British Nuclear Energy Society*, 4(1):564–585, 1966.

[94] J. R. Askew. *A Characteristics Formulation of the Neutron Transport Equation in Complicated Geometries*. Report No. AEEW-M 1108, UK Atomic Energy Establishment, 1972.

[95] W. Boyd, S. Shaner, L. Li, B. Forget, and K. Smith. *The OpenMOC method of characteristics neutral particle transport code*. *Annals of Nuclear Energy*, 68:43–52, 2014.

[96] T. Newton, G. Hosking, L. Hutton, D. Powney, B. Turland, and T. Shuttleworth. *Developments within WIMS10*. In *Proc. of the International Conference on the Physics of Reactors, Nuclear Power: A Sustainable Resource*, 2008.

[97] J. G. Hosking, B. A. Lindley, and P. J. Smith. *CACTUSOT: A 3D Method of Characteristics Solver in WIMS*. In *Proc. of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2017)*, 2017.

[98] D. Sciannandrone, S. Santandrea, and R. Sanchez. *Optimized tracking strategies for step MOC calculations in extruded 3D axial geometries*. *Annals of Nuclear Energy*, 87:49–60, 2016.

[99] 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.

[100] W. C. Fan, C. R. Drumm, and J. L. Powell. *Discrete Ordinates Approximations to the First- and Second-Order Radiation Transport Equations*. Report No. SAND2002-1880, Sandia National Laboratories, 2002.

[101] V. S. Vladimirov. *Mathematical Problems in the One-Velocity Theory of Particle Transport*. Report No. AECL-1661, Atomic Energy of Canada Ltd. Translated from Transactions of the V. A. Steklov Mathematical Institute (1961), 1963.

[102] S. Kaplan and J. A. Davis. *Canonical and Involutory Transformations of the Variational Problems of Transport Theory*. *Nuclear Science and Engineering*, 28(2):166–176, 1967.

[103] E. E. Lewis. *Finite Element Approximation to the Even-Parity Transport Equation*. In J. Lewins and M. Becker, editors, *Advances in Nuclear Science and Technology*, volume 5 of *Advances in Nuclear Science and Technology*. Springer, 1981.

[104] T. Noh, W. F. Miller, Jr, and J. E. Morel. *The Even-Parity and Simplified Even-Parity Transport Equations in Two-Dimensional x-y Geometry*. *Nuclear Science and Engineering*, 123:38–56, 1995.

[105] R. T. Ackroyd. *Finite Element Methods for Particle Transport: Applications to Reactor and Radiation Physics*. John Wiley & Sons, Inc., 1997.

[106] B. Sigg. *The Extended P_L Approximation of the Multidimensional Neutron Transport Equation*. *Nuclear Science and Engineering*, 57(4):277–291, 1975.

[107] 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.

[108] C. Drumm, W. Fan, and W. Bielen. *Least-squares finite-element algorithms in the SCEPTRE radiation transport code*. In *Proc. of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C2011)*, 2011.

[109] 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(1-7):374–401, 2014.

[110] J. R. Peterson, H. Hammer, J. E. Morel, J. Ragusa, and Y. Wang. *Conservative Nonlinear Diffusion Acceleration Applied to the Unweighted Least-Squares Transport Equation in MOOSE*. In *Proc. of the Joint International Conference on Mathematics and Computation, Supercomputing in Nuclear Applications and the Monte Carlo Method*, volume 1. American Nuclear Society, 2015.

[111] 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(4):388–403, 2019.

[112] T. A. Manteuffel and K. J. Ressel. *Least-squares finite-element solution of the neutron transport equation in diffusive regions*. *SIAM Journal on Numerical Analysis*, 35(2):806–835, 2002.

[113] C. Latimer, M. D. Eaton, J. Kópházi, and R. G. McClarren. *A geometry conforming, isogeometric, weighted least squares (WLS) method for the neutron transport equation with a discrete ordinate (S_N) angular discretisation*. *Progress in Nuclear Energy*, 121:103238, 2020.

[114] W. Zheng. *Least-Squares and Other Residual Based Techniques for Radiation Transport Calculations*. PhD Thesis, Texas A&M University, 2016.

[115] M. A. Abbott. *An Introduction to the Method of Characteristics*. Thames and Hudson, 1966.

[116] J. D. Hoffman. *Numerical Methods for Engineers and Scientists*. Marcel Dekker, Inc., second edition, 2001.

[117] R. T. Ackroyd. *Least-Squares Derivation of Extremum and Weighted-Residual Meethods 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.

[118] J. E. Morel and J. M. McGhee. *A Self-Adjoint Angular Flux Equation. Nuclear Science and Engineering*, 132:312–325, 1999.

[119] C. Latimer, J. Kópházi, M. D. Eaton, and R. G. McClarren. *A geometry conforming isogeometric method for the self-adjoint angular flux (SAAF) form of the neutron transport equation with a discrete ordinate (S_N) angular discretisation. Annals of Nuclear Energy*, 136:107049, 2020.

[120] 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(2):201–225, 2014.

[121] T. Gozani. *Experimental neutron Kinetic Studies in a ^{238}U Sphere. Nuclear Science and Engineering*, 36:143–158, 1969.

[122] H. Hurwitz, Jr. and M. S. Nelkin. *The Thermal Neutron Spectrum in a Diffusing Medium. Nuclear Science and Engineering*, 3(1):1–10, 1958.

[123] K. M. Case and P. F. Zweifel. *Linear transport theory*. Addison-Wesley, 1967.

[124] H. Hurwitz, Jr. and R. Ehrlich. *Multigroup Methods for Neutron Diffusion Problems*. Report No. HECD-3595, Knolls Atomic Power Laboratory, 1953.

[125] 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.

[126] B. R. Bergel'son and B. Z. Torlin. *Effective isotropic scattering kernel and the multigroup transport equation. Soviet Atomic Energy*, 27:1112–1114, 1969.

[127] D. E. Cullen. *Why are the S_N and the P_N Methods Equivalent?* Report No. UCRL-ID-145518, Lawrence Livermore National Laboratory, 2001.

[128] G. C. Pomraning. *An Asymptotically Correct Approximation to the Multidimensional Transport Equation. Nuclear Science and Engineering*, 22:328–338, 1965.

[129] *On the Equivalence of the Spherical Harmonics Method and the Discrete Ordinates Method Using Gauss Quadrature for the Boltzmann Equation*. Report No. WAPD-TM-118, Bettis Atomic Power Laboratory, 1958.

[130] E. M. Gelbard. *Spherical Harmonics Methods: P_L and Double P_L Approximations*. In H. Greenspan, C. N. Kelber, and D. Okrent, editors, *Computing Methods in Reactor Physics*. Gordon and Breach, 1968.

[131] E. M. Gelbard. *Applications of spherical harmonics method to reactor problems*. 1960.

[132] G. C. Pomraning. *Asymptotic and Variational Derivation of the Simplified P_N Equations. Annals of Nuclear Energy*, 20(9):623–637, 1993.

[133] 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.

[134] J. Yvon. *La diffusion macroscopique des neutrons une methode d'approximation. Journal of Nuclear Energy (1954)*, 4(3):305–318, 1957.

[135] S. H. Ghazaie, A. Zolfaghari, and M. Abbasi. *An extended half-range spherical harmonics method for first-order neutron transport equation based on variational treatment. Progress in Nuclear Energy*, 100:389–405, 2017.

[136] J. P. Jenal, P. J. Erickson, W. A. Rhoades, D. B. Simpson, and M. L. Williams. *The Generation of a Computer Library for Discrete Ordinates Quadrature Sets*. Report No. ORNL/TM-6023, Oak Ridge National Laboratory, 1977.

[137] W. F. Walters. *Use of the Chebyshev-Legendre Quadrature Set in Discrete Ordinates Codes*. Report No. LA-UR-87-3621, Los Alamos National Laboratory, 1987.

[138] B. G. Carlson. *Solution of the Transport Equation by S_N Approximations*. Technical report, Los Alamos Scientific Laboratory of the University of California, 1953.

[139] K. D. Lathrop and B. G. Carlson. *Discrete Ordinates Angular Quadrature of the Neutron Transport Equation*. Report No. LA-3186, Los Alamos Scientific Laboratory of the University of California, 1964.

[140] C. Lee. *The Discrete S_N Approximation to Transport Theory*. Report No. LA-2595, Los Alamos National Laboratory, 1961.

[141] G. Sjoden and A. Haghighat. *PENTRAN Code System: Users Guide to Version 9.4X.5 Series*. Technical report, HSW Technologies, LLC., 2008.

[142] R. E. Alcouffe, R. S. Baker, J. A. Dahl, E. J. Davis, T. G. Saller, S. A. Turner, R. C. Ward, and R. J. Zerr. *PARTISN: A Time/Dependent, Parallel Neutral Particle Transport Code System*. Report No. LA-UR-17-29704, Los Alamos National Laboratory, 2008.

[143] 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 *Proc. of the Joint International Conference on Mathematics and Computation, Supercomputing in Nuclear Applications and the Monte Carlo Method*. American Nuclear Society, 2015.

[144] C. G. Wick. *Über ebene diffusionprobleme. Zeitschrift für Physik*, 121:702–718, 1943.

[145] S. Chandrasekhar. *On the radiative equilibrium of stellar atmosphere II. The Astrophysical Journal*, 100:76–86, 1944.

[146] W. H. Reed. *Spherical Harmonic Solutions of the Neutron Transport Equation from Discrete Ordinate Equations. Nuclear Science and Engineering*, 49(1):10–19, 1972.

[147] J. E. Morel. *A Collocation Method for the Solution of the Neutron Transport Equation with Both Symmetric and Asymmetric Scattering*. Report No. SAND80-1446, Sandia National Laboratories, 1981.

[148] E. W. Larsen, J. E. Morel, and W. F. Miller, Jr. *Asymptotic Solutions of Numerical Transport Problems in Optically Thick, Diffusive Regimes. Journal of Computational Physics*, 69:283–324, 1987.

[149] J. J. Jarrell and M. L. Adams. *Discrete-Ordinates Quadrature Sets Based on Linear Discontinuous Finite Elements*. In *Proc. of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2011)*. American Nuclear Society (Latin American Section), 2011.

[150] A. H. Stroud and D. Secrest. *Gaussian Quadrature Formulas*. Prentice-Hall, 1966.

[151] K. D. Lathrop. *Ray Effects in Discrete Ordinates Equations. Nuclear Science and Engineering*, 32(3):357–369, 1968.

[152] *Transport Theory and Advanced Reactor Calculations*. Report No. IAEA-TECDOC-254, International Atomic Energy Agency, 1981.

[153] C. Yang, T. Cheng, L. Deng, Z. Li, and L. Wen. *Development of 3-D parallel first-collision source method for discrete ordinates code JSNT-S. Annals of Nuclear Energy*, 135:106942:1–10, 2020.

[154] K. D. Lathrop. *Remedies for Ray Effects. Nuclear Science and Engineering*, 45(3):255–268, 1971.

[155] *Transport Synthesis*. Report No. WAPD-T-1849, Bettis Atomic Power Laboratory, 1965.

[156] L. L. Briggs, W. F. Miller, Jr, and E. E. Lewis. *Ray-effect mitigation in discrete ordinate-like angular finite element approximations in neutron transport. Nuclear Science and Engineering*, 57(3):2015–217, 1975.

[157] W. F. Miller, Jr, 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.

[158] W. R. Martin C. E. Yehnert, L. Lorence, and J. J. Duderstadt. *Phase-space finite element methods applied to the first-order form of the transport equation. Annals of Nuclear Energy*, 8:633–646, 1981.

[159] M. Mordant. *Phase-space finite elements encoded in Zephyr for X-Y and R-Z transport calculations. Progress in Nuclear Energy*, 18:27–37, 1986.

[160] G. G. M. Coppa, G. Lapenta, and P. Ravetto. *Angular finite element techniques in neutron transport. Annals of Nuclear Energy*, 17:363–378, 1990.

[161] G. Widmer, R. Hiptmair, and Ch. Schwab. *Sparse adaptive finite elements for radiative transfer. Journal of Computational Physics*, 227:6071–6105, 2008.

[162] H. Park and C. R. E. de Oliveira. *Adaptive angular resolution for the finite element-spherical harmonics method. Transactions of the American Nuclear Society*, 93:517–519, 2005.

[163] H. Park and C. R. E. de Oliveira. *Coupled Space-Angle Adaptivity for Radiation Transport Calculations. Nuclear Science and Engineering*, 162(2):216–234, 2009.

[164] J. Kópházi and D. Lathouwers. *A space-angle DGFEM approach for the Boltzmann radiation transport equation with local angular refinement. Journal of Computational Physics*, 297:637–668, 2015.

[165] J. E. Morel, T. A. Wareing, R. B. Lowrie, and D. K. Parsons. *Analysis of Ray-Effect Mitigation Techniques. Nuclear Science and Engineering*, 144:1–22, 2003.

[166] R. N. Blomquist and E. E. Lewis. *A Rigorous Treatment of Transverse Buckling Effects in Two-Dimensional Neutron Transport Computations. Nuclear Science and Engineering*, 73(2):125–139, 1980.

[167] K. D. Lathrop. *Elimination of Ray Effects by Converting Discrete Ordinates Equations to Spherical-Harmonic-Like Equations.* In *Proc. of the Conference on New Developments in Reactor Mathematics and Applications CONF-710302.* U.S. Atomic Energy Commission, 1971.

[168] W. F. Miller, Jr and W. H. Reed. *Ray-Effect Mitigation Methods for Two Dimensional Neutron Transport Theory. Nuclear Science and Engineering*, 62:10, 1977.

[169] K. D. Lathrop and F. W. Brinkley. *TWOTRAN-II: An Interfaced Exportable Version of the TWOTRAN Code for Two-Dimensional Transport.* Report No. LA-4848-MS, Los Alamos Scientific Laboratory, 1973.

[170] R. S. Baker and K. R. Koch. *An S_N Algorithm for the Massively Parallel CM-200 Computer. Nuclear Science and Engineering*, 128(3):312–320, 1998.

[171] R. S. Baker. *An S_N Algorithm for Modern Architectures. Nuclear Science and Engineering*, 185:107–116, 2017.

[172] J. I. C. Vermaak, J. C. Ragusa, M. L. Adams, and J. E. Morel. *Massively parallel transport sweeps on meshes with cyclic dependencies. Journal of Computational Physics*, 425:109892, 2021.

[173] E. W. Larsen. *The Asymptotic Diffusion Limit of Discretized Transport Problems. Nuclear Science and Engineering*, 112(4):336–346, 1992.

[174] E. W. Larsen and J. B. Keller. *Asymptotic solution of neutron transport problems for small mean free paths. Journal of Mathematical Physics*, 15:75–81, 1974.

[175] A. M. Weinberg and E. P. Wigner. *The Physical Theory of Neutron Chain Reactors.* University of Chicago Press, third edition, 1958.

[176] S. Choi, K. Smith, H. C. Lee, and D. Lee. *Impact of inflow transport approximation on light water reactor physics. Journal of Computational Physics*, 299:352–273, 2015.

[177] E. M. Gelbard. *Streaming in Lattices. Advances in Nuclear Science and Technology*, 15:223–400, 1983.

[178] R. T. Ackroyd. *On the Use of Diffusion-Theory for Transport Calculations.* In *Proc. of the British Nuclear Energy Society in conjunction with the Institute of Physical Society, International Conference*, 1967.

[179] J. C. Lee. *Nuclear Reactor: Physics and Engineering.* John Wiley & Sons, Inc., 2019.

[180] J. Wood. *Computational Methods in Reactor Shielding.* Pergamon Press, 1982.

[181] *Reactor Shielding for Nuclear Engineers.* Report No. TID-25951, U.S. Atomic Energy Commission, 1973.

[182] A. N. Prinja and E. W. Larsen. *General Principles of Neutron Transport.* In D. G. Cacuci, editor, *Handbook of Nuclear Engineering.* Springer, 2010.

[183] R. E. Alcouffe. *Diffusion Synthetic Acceleration Methods for the Diamond-Differenced Discrete-Ordinates Equations. Nuclear Science and Engineering*, 64:344–355, 1977.

[184] J. S. Warsa. *Analytical (S_N) solutions in heterogeneous slabs using symbolic algebra computer programs. Annals of Nuclear Energy*, 29:851–874, 2002.

[185] E. W. Larsen. *Unconditionally Stable Diffusion Synthetic Acceleration Methods for the Slab Geometry Discrete-Ordinates Equations. Part 1: Theory. Nuclear Science and Engineering*, 82:47, 1982.

[186] S. Schunert and Y. Y. Azmy. *A Two-Dimensional Method of Manufactured Solutions Benchmark Suite Based on Variations of Larsen's Benchmark with Escalating Order of Smoothness of the Exact Solution.* In *Proc. of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering.* American Nuclear Society, 2011.

[187] S. M. Ayyoubzadeh, M. J. Safari, and N. Vosoughi. *Solution of diffusion equation in deformable spheroids. Annals of Nuclear Energy*, 38:982–988, 2011.

[188] N. G. Sjöstrand and A. Rentze. *The "geometric buckling" of spheroidal nuclear reactors. Arkiv for Fysik*, 13(19):229–234, 1957.

[189] M. M. R. Williams. *Neutron Diffusion in Spheroidal, Bispherical and Toroidal Systems. Nuclear Science and Engineering*, 94:251–263, 1986.

[190] E. B. Becker, G. F. Carey, and J. T. Oden. *Finite Elements: An Introduction.* Prentice-Hall Inc., 1981.

[191] P. B. Bochev and M. D. Gunzburger. *Least-Squares Finite Element Methods*, volume 166 of *Applied Mathematical Sciences.* Springer, 2009.

[192] R. S. Jeffers, J. Kópházi, M.D. Eaton, F. Févotte, F. Hülsemann, and J. Ragusa. *Goal-based h-adaptivity of the 1-D diamond difference discrete ordinate method. Journal of Computational Physics*, 335:179–200, 2017.

[193] R. S. Jeffers, J. Kópházi, M.D. Eaton, F. Févotte, F. Hülsemann, and J. Ragusa. *Goal-Based Error Estimation, Functional Correction, h, p and hp Adaptivity of the 1-D Diamond Difference Discrete Ordinate Method. Journal of Computational and Theoretical Transport*, 46(6-7):427–458, 2017.

[194] B. Q. Li. *Discontinuous Finite Elements in Fluid Dynamics and Heat Transfer.* Computational Fluid and Solid Mechanics. Springer, 2006.

[195] W. H. Reed. *Triangular Mesh Difference Schemes for the Transport Equation.* Report No. LA-4769, Los Alamos National Laboratory, 1971.

[196] J. N. Reddy. *Applied Functional Analysis and Variational Methods in Engineering.* McGraw-Hill Book Company, 1986.

[197] O. C. Zienkiewicz and K. Morgan. *Finite Elements & Approximations.* Dover Publications, Inc., 2006.

[198] S. Kaplan. *Variational Methods in Nuclear Engineering. Advances Nuclear Science Technology*, 3(135):185–221, 1969.

[199] W. Ritz. *Theorie der Transversalschwingungen einer quadratische Platte mit freien Rändern. Annalen der Physik*, 28:737–786, 1909.

[200] J. Singer. *On the Equivalence of the Galerkin and Rayleigh-Ritz Methods. Journal of the Royal Aeronautical Society*, 66:592–592, 1962.

[201] B. A. Finlayson and L. E. Scriven. *The Method of Weighted Residuals - A Review. Applied Mechanics Reviews*, 19(9):735–748, 1966.

[202] S. Faedo. *Un nuovo metodo per l'analisi esistenziale e quantitativa dei problemi di propagazioni. Annali della Scuola Normale Superiore di Pisa*, 1(3):1–41, 1949.

[203] S. H. Crandall. *Engineering Analysis (A Survey of Numerical Procedures).* McGraw-Hill, 1956.

[204] R. A. Frazer, W. P. Jones, and S. W. Skan. *Approximations to functions and to the solutions of differential equations.* Report No. 1799, Aeronautical Research Committee of Great Britain, 1937.

[205] B. G. Galerkin. *Sterzhni i plastiny. Ryady v nekotorykh voprosakh uprugogo ravnovesiya sterzhnei i plastin. Vestnik inzhenerov i tekhnikov*, 19:897–908, 1915.

[206] S. G. Mikhlin. *Variational Methods in Mathematical Physics.* Pergamon & Macmillan, 1964.

[207] 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(3):371–390, 1977.

[208] F. Bassi and S. Rebay. *A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier-Stokes equations. Journal of Computational Physics*, 131:267–279, 1997.

[209] C. B. Laney. *Computational Gasdynamics.* Cambridge University Press, 1998.

[210] J. Jin. *The Finite Element Method in Electromagnetics.* John Wiley & Sons, Inc., second edition, 2002.

[211] 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.

[212] L. A. Semenza, E. E. Lewis, and E. C. Rossow. *Application of the finite element method to the multigroup neutron diffusion equation. Nuclear Science and Engineering*, 47(3):302–310, 1972.

[213] K. F. Hansen and C. M. Kang. *Finite element methods in reactor physics analysis. Advances in Nuclear Science and Technology*, 8:173–253, 1975.

[214] R. A. Lillie and J. C. Robinson. *A Linear Triangle Finite Element Formulation for Multigroup Neutron Transport Analysis with Anisotropic Scattering.* Report No. ORNL/TM-5281, Oak Ridge National Laboratory, 1976.

[215] R. T. Ackroyd. *The Why and How of Finite Elements. Annals of Nuclear Energy*, 8:539–566, 1981.

[216] 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.

[217] R. T. Ackroyd. *Foundations of Finite Element Applications to Neutron Transport. Progress in Nuclear Energy*, 29(1):43–56, 1985.

[218] T. Stolarski, Y. Nakasone, and S. Yoshimoto. *Engineering Analysis with ANSYS Software.* Butterworth-Heinemann, second edition, 2018.

[219] J. G. Issa. *Finite Element Method Applied to Neutron Transport Problems.* PhD Thesis, Imperial College of Science and Technology, University of London, 1982.

[220] P. J. Frey and P-L George. *Mesh Generation: Application to Finite Elements.* ISTE Ltd, 2007.

[221] 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.

[222] J. T. Oden and J. N. Reddy. *An Introduction to the Mathematical Theory of Finite Elements.* Dover Publications, 2011.

[223] E. Cuthill and J. McKee. *Reducing the bandwidth of sparse symmetric matrices.* In *Proc. of the 1969 24th National Conference (ACM '69).* Association for Computing Machinery, 1969.

[224] T. J. R. Hughes. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis.* Dover Publications, 2000.

[225] I. Ergatoudis, B. M. Irons, and O. C. Ziekiewicz. *Curved Isoparametric "Quadrilateral" Elements for Finite Element Analysis. International Journal of Solids and Structures*, 1968:31–42, 1968.

[226] O. C. Zienkiewicz. *The Finite Element Method in Engineering Sciences.* McGraw-Hill, 1971.

[227] G. Beer. *Advanced Numerical Simulation Methods: From CAD Data Directly to Simulation Results.* Taylor & Francis Group, 2015.

[228] 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.

[229] J. A. Cottrell, T. J. R. Hughes, and Y. Bazilevs. *Isogeometric Analysis: Toward Integration of CAD and FEA.* John Wiley & Sons, Inc., 2009.

[230] J-Y. Moller and J-J. Lautard. *MINARET, A deterministic neutron transport solver for nuclear core calculations.* In *Proc. of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2011).* Latin American Section (LAS) / American Nuclear Society (ANS), 2011.

[231] 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.

[232] C. Runge. *Über empirische Funktionen und die Interpolation zwischen äquidistanten Ordinaten. Zeitschrift für Mathematik und Physik*, 46:224–243, 1901.

[233] 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.

[234] W. H. Reed. *New Difference Schemes for the Neutron Transport Equation. Nuclear Science and Engineering*, 46(2):309–314, 1971.

[235] W. H. Reed, T. R. Hill, F. W. Brinkley, and K. D. Lathrop. *TRIPLET: A Two-Dimensional, Multigroup, Triangular Mesh, Planar Geometry, Explicit Transport Code.* Report No. LA-5428-MS, Los Alamos Scientific Laboratory, 1973.

[236] T. A. Wareing, J. M. McGhee, J. E. Morel, and S. D. Pautz. *Discontinuous Finite Element S_N Methods on Three-Dimensional Unstructured Grids. Nuclear Science and Engineering*, 138(3):256–268, 2001.

[237] D. A. Di Pietro and A. Ern. *Mathematical Aspects of Discontinuous Galerkin Methods*, volume 69 of *Mathémathiques et Applications*. Springer, 2012.

[238] W. H. Reed and T. R. Hill. *Triangular mesh methods for the neutron transport equation.* Report No. LA-UR-73-479, Los Alamos National Laboratory, 1973.

[239] W. H. Reed, T. R. Hill, F. W. Brinkley, and K. D. Lathrop. *TRIDENT: A two dimensional multigroup, triangular mesh, explicit neutron transport code.* Report No. LA-6735-MS, Los Alamos National Laboratory, 1977.

[240] J. E. Morel and J. S. Warsa. *An S_N spatial discretization scheme for tetrahedral meshes. Nuclear Science and Engineering*, 151:157–166, 2005.

[241] J. M. McGhee and T. A. Wareing. *Attila version 2: User's manual.* Report No. LA-UR-00-5778, Los Alamos National Laboratory, 2000.

[242] B. Rivière. *Discontinuous Galerkin methods for solving elliptic and parabolic equations: Theory and Implementation.* Frontiers in Applied Mathematics. SIAM, 2008.

[243] T. J. R. Hughes, G. Engel, L. Mazzei, and M. G. Larson. *The Continuous Galerkin Method is Locally Conservative. Journal of Computational Physics*, 163:467–488, 2000.

[244] W. H. Reed. *The effectiveness of acceleration techniques for iterative methods in transport theory. Nuclear Science and Engineering*, 45(3):245–254, 1971.

[245] M. A. Freitag. *Inner-Outer Iterative Methods for Eigenvalue Problems - Convergence and Preconditioning.* PhD Thesis, University of Bath, 2007.

[246] Y. Wang. *Adaptive Mesh Refinement Solution Techniques for the Multigroup S_N Transport Equation using a Higher-Order Discontinuous Finite Element Method.* PhD Thesis, Texas A&M University, 2009.

[247] 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.

[248] A. George and J. W-H. Liu. *Computer Solution of Large Sparse Positive Definite Systems.* Prentice-Hall Series in Computational Mathematics. Prentice-Hall, Inc., 1981.

[249] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. Gropp, D. Karpeyev, D. Kaushik, M. Knepley, D. May, L. C. McInnes, R. Mills, T. Munson, K. Rupp, P. Sanan, B. Smith, S. Zampini, H. Zhang, and H. Zhang. *PETSc Users Manual: Revision 3.9.* Report No. ANL-95/11 Rev. 3.9, Argonne National Laboratory, 2018.

[250] 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.

[251] H. A. Van Der Vorst and K. Dekker. *Conjugate gradient type methods and preconditioning. Journal of Computational and Applied Mathematics*, 16:73–87, 1988.

[252] Y. Azmy and E. Sartori. *Nuclear Computational Science: A Century in Review.* Springer, 2010.

[253] 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.

[254] 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*, 115:115–147, 2019.

[255] D. Lathouwers. *Goal-oriented spatial adaptivity for the S_N equations on unstructured triangular meshes. Annals of Nuclear Energy*, 38:1373–1381, 2011.

[256] D. Lathouwers. *Spatially adaptive eigenvalue estimation for the S_N equations on unstructured triangular meshes. Annals of Nuclear Energy*, 38:1867–1876, 2011.

[257] L. Demkowicz, W. Rachowicz, and Ph. Devloo. *A Fully Automatic hp-Adaptivity. Journal of Scientific Computing*, 17(1-4):117–142, 2002.

[258] P. Šolin, K. Segeth, and I. Dolezel. *Higher-Order Finite Element Methods.* Studies in Advanced Mathematics. Chapman & Hall/CRC, 2004.

[259] C. Bajaj, J. Chen, and R. Holts. *Energy formulations of A-splines. Computer-Aided Geometric Design*, 16:39–59, 1999.

[260] F. Cirak, M. Ortiz, and P. Schröder. *Subdivision surfaces: a new paradigm for thin shell analysis. International Journal for Numerical Methods in Engineering*, 47:2039–2072, 2000.

[261] D. F. Rogers. *An Introduction to NURBS with Historical Perspective.* Academic Press, 2001.

[262] R. K. C. Cheng. *Inside Rhinoceros 5.* Cengage Learning, 2014.

[263] E. Moss. *Autodesk AutoCAD 2019 Fundamentals.* SDC Publications, 2018.

[264] *Application Protocol for Product Life Cycle Support (PLCS).* Report No. AP 239, National Institute of Standards and Technology (NIST), 2015.

[265] R. N. Nagel, W. W. Braithwaite, and P. R. Kennicott. *Initial Graphics Exchange Specification IGES Version 1.0.* Report No. NBSIR 80-1978, National Bureau of Standards, 1980.

[266] L. Piegl and W. Tiller. *The NURBS Book: Monographs in Visual Communication.* Springer, 1997.

[267] T. J. R. Hughes, G. Sangalli, and M. Tani. *Isogeometric Analysis: Mathematical and Implementational Aspects, with Applications.* In *Splines and PDEs: From Approximation Theory to Numerical Linear Algebra*, volume 2219 of *Lecture Notes in Mathematics.* Springer, 2018.

[268] 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.

[269] I. Akkerman, Y. Bazilevs, V. M. Calo, T. J. R. Hughes, and S. Hulshoff. *The role of continuity in residual-based variational multi-scale modelling of turbulence. Computational Mechanics*, 41:371–378, 2008.

[270] J. Cottrell, T. Hughes, and A. Reali. *Studies of refinement and continuity in isogeometric structural analysis. Computer Methods in Applied Mechanics and Engineering*, 196:4160–4183, 2007.

[271] J. V. Venås. *Isogeometric analysis of acoustic scattering.* Master's Thesis, NTNU-Trondheim, 2015.

[272] A. Buffa, G. Sangalli, and R. Vasquez. *Isogeometric methods for computational electromagnetics: B-spline and T-spline discretisations. Journal of Computational Physics*, 257:1291–1320, 2014.

[273] Y. Bazilevs, V. M. Calo, Y. Zhang, and T. J. R. Hughes. *Isogeometric fluid-structure interaction analysis with applications to arterial blood flow. Computational Mechanics*, 38:310–322, 2006.

[274] Y. Bazilevs, M. C. Hsu, and M. A. Scott. *Isogeometric fluid-structure interaction analysis with emphasis on non-matching discretizations, and with application to wind turbines. Computer Methods in Applied Mechanics and Engineering*, 249:28–41, 2012.

[275] Z. An, T. Yu, T. Q. Bui, C. Wang, and N. A. Trinh. *Implementation of isogeometric boundary element method for 2-D steady heat transfer analysis. Advances in Engineering Software*, 116:36–49, 2018.

[276] J. A. 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, 2017.

[277] S. G. Wilson, J. Kópházi, A. R. Owens, and M. D. Eaton. *Interior penalty schemes for discontinuous isogeometric methods with an application to nuclear reactor physics.* In *Proc. of the International Conference of Nuclear Engineering (ICONE26)*, volume 3: Nuclear Fuel and Material, Reactor Physics and Transport Theory. ASME, 2018.

[278] A. R. Owens, J. Kópházi, 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.

[279] M. G. Cox. *The numerical evaluations of B-splines. IMA Journal of Applied Mathematics*, 10:134–149, 1972.

[280] C. de Boor. *On calculating with B-splines. Journal of Approximation Theory*, 6:50–62, 1972.

[281] G. H. Golub and J. H. Welsch. *Calculation of Gauss Quadrature Rules. Mathematics of Computation*, 23(106):221–230, 1969.

[282] T. J. R. Hughes, A. Reali, and G. Sangali. *Efficient quadrature for NURBS-based isogeometric analysis. Computer Methods in Applied Mechanics and Engineering*, 199:301–313, 2010.

[283] J. E. Cobb. *Tiling the Sphere with Rational Bézier Patches.* Report No. TR UUCS-99-009, University of Utah, 1988.

[284] M. A. Scott, X. Li, 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.

[285] T. W. Sederberg, G. T. Finnigan, X. Li, H. Lin, and H. Ipson. *Watertight Trimmed NURBS. Association for Computing Machinery Transactions on Graphics*, 27(3):1–8, 2008.

[286] H. Prautzsch, W. Boehm, and M. Paluszny. *Bézier and B-Spline Techniques.* Mathematics + Visualization. Springer, 2000.

[287] H. Al Akhras, T. Elguedj, A. Gravouil, and M. Rochette. *Towards an automatic isogeometric analysis suitable trivariate models generation - Application to geometric parametric analysis. Computer Methods in Applied Mechanics and Engineering*, 316:623–645, 2017.

[288] D. N. Arnold. *An interior penalty finite element method with discontinuous elements. SIAM Journal on Numerical Analysis*, 19:742–760, 1982.

[289] S. K. Hall, M. D. Eaton, and M. P. Knight. *The implementation of a simplified spherical harmonics semi-analytic nodal method in PANTHER. Annals of Nuclear Energy*, 57:280–293, 2013.

[290] Y. Wang and J. Ragusa. *Application of hp Adaptivity to the Multigroup Diffusion Equations. Nuclear Science and Engineering*, 161(1):22–48, 2009.

[291] M. A. Goffin, C. M. J. Baker, A. G. Buchan, C. C. Pain, M. D. Eaton, and P. N. Smith. *Minimising the error in eigenvalue calculations involving the Boltzmann transport equation using goal-based adaptivity on unstructured meshes. Journal of Computational Physics*, 242:726–752, 2013.

[292] P. E. Farrell and J. R. Maddison. *Conservative interpolation between volume meshes by local Galerkin projection. Computer Methods in Applied Mechanics and Engineering*, 200:89–100, 2011.

[293] J. Ragusa. *A simple Hessian-based 3D mesh adaptation technique with applications to the multigroup diffusion equations. Annals of Nuclear Energy*, 35:2006–2018, 2008.

[294] Y. Wang, W. Bangerth, and J. Ragusa. *Three-dimensional h-adaptivity for the multigroup neutron diffusion equations. Progress in Nuclear Energy*, 51:543–555, 2009.

[295] X. Jiao and M. T. Heath. *Common-refinement-based data transfer between non-matching meshes in multiphysics simulations. International Journal for Numerical Methods in Engineering*, 61:2402–2427, 2004.

[296] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide.* Technical report, Society for Industrial and Applied Mathematics, 1999.

[297] M. Ainsworth and J. T. Oden. *A posteriori error estimation in finite element analysis. Computer Methods in Applied Mechanics and Engineering*, 142:1–88, 1997.

[298] I. Babuška and B. Q. Guo. *The h, p, and h − p version of the finite element method; basis theory and applications. Advances in Engineering Software*, 15:159–174, 1992.

[299] L. Demkowicz. *Computing with hp-Adaptive Finite Elements: One and Two Dimensional Elliptic and Maxwell Problems*, volume 1 of *Chapman & Hall/CRC Applied Mathematics and Nonlinear Science Series.* Chapman & Hall/CRC, 2007.

[300] L. Demkowicz, J. Kurtz, D. Pardo, M. Paszynski, W. Rachowicz, and A. Zdunek. *Computing with hp-Adaptive Finite Elements: Frontiers Three Dimensional Elliptic and Maxwell Problems with Applications*, volume 2 of *Chapman & Hall/CRC Applied Mathematics and Nonlinear Science Series.* Chapman & Hall/CRC, 2008.

[301] B. Guo and I. Babuška. *The h − p version of the finite element method, Part 1: The basic approximation results. Computational Mechanics*, 1(1):21–41, 1986.

[302] B. Guo and I. Babuška. *The h − p version of the finite element method, Part 2: General Results and Applications. Computational Mechanics*, 1(3):203–220, 1986.

[303] M. Ainsworth and J. T. Oden. *A Posteriori Error Estimation in Finite Element Analysis.* Pure and Applied Mathematics. John Wiley & Sons, Inc., 2000.

[304] L. Demkowicz, P. Monk, L. Vardapetyan, and W. Rachowicz. *De Rhamm diagram for hp finite element spaces. Mathematics and Computers with Applications*, 39(7-8):29–38, 2000.

[305] Y. Wang and J. Ragusa. *Adaptive hp-Mesh Refinement Applied to 1-D, One-Group Diffusion Problems. Trans. of the American Nuclear Society*, 93:585, 2005.

[306] Y. Wang and J. C. Ragusa. *Goal-Oriented hp-Mesh Adaptation for 1-D Multigroup Diffusion Equations.* In *Proc. of the Joint International Toplical Meeting on Mathematics & Computation and Supercomputing in Nuclear Applications*. American Nuclear Society, 2007.

[307] W. Rachowciz, J. T. Oden, and L. Demkowicz. *Toward a universal h-p adaptive finite element strategy, Part 3. Design of h-p meshes. Computer Methods in Applied Mechanics and Engineering*, 77:181–212, 1989.

[308] J. P. de S. R. Gago, D. W. Kelly, O. C. Zienkiewicz, and I. Babuška. *A posteriori error analysis and adaptive processes in the finite element method: part II - Adaptive mesh refinement. International Journal for Numerical Methods in Engineering*, 19:1621–1656, 1983.

[309] W. Bangerth and R. Rannacher. *Adaptive Finite Element Methods for Differential Equations.* Lectures in Mathematics ETH Zurich. Springer, 2003.

[310] P. Šolin and L. Demkowicz. *Goal-Oriented hp-adaptivity for elliptic problems. Computer Methods in Applied Mechanics and Engineering*, 193:449–468, 2004.

[311] P. Hennig, M. Kästner, P. Morgenstern, and D. Peterseim. *Adaptive mesh refinement strategies in isogeometric analysis: A computational comparison. Computer Methods in Applied Mechanics and Engineering*, 316:424–448, 2017.

[312] K. P. S. Gahaluat, J. K. Kraus, and S. K. Tomar. *Multigrid methods for isogeometric discretisation. Computer Methods in Applied Mechanics and Engineering*, 253:413–425, 2013.

[313] G. Sangalli and M. Tani. *Isogeometric preconditioners based on fast solvers for the Sylvester equation. SIAM Journal of Scientific Computing*, 38(6):A3644–A3671, 2016.

[314] M. Ceze and K. J. Fidkowski. *Anisotropic hp-adaptation framework for functional prediction. AIAA Journal*, 51(2):492–509, 2013.

[315] C. Johnson and J. Pitkaranta. *An analysis of the discontinuous Galerkin method for a scalar hyperbolic equation. Mathematics of Computation*, 46(173):1–26, 1986.

[316] T. E. Peterson. *A Note on the Convergence of the Discontinuous Galerkin Method for a Scalar Hyperbolic Equation. SIAM Journal on Numerical Analysis*, 28(1):133–140, 1991.

[317] M. Renardy and R. C. Rogers. *An Introduction to Partial Differential Equations*, volume 13 of *Texts in Applied Mathematics*. Springer, second edition, 2004.

[318] Y. Wang. *hp-Mesh Adaptation for 1-D Multigroup Neutron Diffusion Problems.* Master's Thesis, Texas A&M University, 2006.

[319] L. Demkowicz. *Projection based interpolation.* Report No. ICES 04-03, University of Texas at Austin, 2004.

[320] R. Hartmann. *Adjoint Consistency Analysis of Discontinuous Galerkin Discretizations. SIAM Journal on Numerical Analysis*, 45(6):2671–2696, 2007.

[321] H. Park. *Coupled Space-Angle Adaptivity and Goal-Oriented Error Control for Radiation Transport Calculations.* PhD Thesis, Georgia Institute of Technology, 2006.

[322] L. Coox, F. Greco, O. Atak, D. Vandepitte, and W. Desmet. *A Robust Patch Coupling Method for NURBS-Based Isogeometric Analysis of Non-Conforming Multipatch Surfaces. Computer Methods in Applied Mechanics and Engineering*, 316:235–260, 2017.

[323] Z. Lei, F. Gillot, and L. Jezequel. *A C^0/G^1 multiple patches connection method in isogeometric analysis. Applied Mathematical Modelling*, 39:4405–4420, 2015.

[324] N. Collier, L. Dalcin, D. Pardo, and V. M. Calo. *The cost of continuity: performance of iterative solvers on isogeometric finite elements. SIMA Journal on Scientific Computing*, 35:A767–A784, 2013.

[325] J. A. Nitsche. *Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilrämen, die keinen Randbedingungen unterworfen sind*, volume 36 of *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*. Springer, 1971.

[326] M. F. Wheeler. *An elliptic collocation-finite element method with interior penalties. SIAM Journal on Numerical Analysis*, 15:152–161, 1978.

[327] P. Percell and M. F. Wheeler. *A Local Residual Finite Element Procedure for Elliptic Equations. SIAM Journal on Numerical Analysis*, 15(4):705–713, 1978.

[328] F. Brunero. *Discontinuous Galerkin methods for isogeometric analysis.* Master's Thesis, Università degli Studi di Milano, 2012.

[329] D. N. Arnold, F. Brezzi, B. Cockburn, and L. D. Marini. *Unified Analysis of Discontinuous Galerkin Methods for Elliptic Problems. SIAM Journal on Numerical Analysis*, 39(5):1749–1779, 2002.

[330] F. Brezzi, B. Cockburn, L.D. Marini, and E. Süli. *Stabilization mechanisms in discontinuous Galerkin finite element methods. Computer Methods in Applied Mechanics and Engineering*, 195:3293–3310, 2006.

[331] P. Castillo. *Performance of discontinuous Galerkin methods for elliptic PDEs. SIAM Journal on Scientific Computing*, 24:401–407, 2002.

[332] G. B. Alvarez, A. F. D. Loula, E. G. D. do Carmo, and F. A. Rochinha. *A discontinuous finite element formulation for Helmholtz equation. Computer Methods in Applied Mechanics and Engineering*, 195:4018–4035, 2006.

[333] J. Douglas, Jr and T. Dupont. *Interior Penalty Procedures for Elliptic and Parabolic Galerkin Methods*, volume 58 of *Lecture Notes in Physics.* Springer-Verlag, 1976.

[334] 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.

[335] G. I. Marchuk. *Numerical Methods for Nuclear Reactor Calculations.* Consultants Bureau, Inc. and Chapman & Hall, Ltd., 1959.

[336] B. A. Finlayson. *The Method of Weighted Residuals and Variational Principles with Application in Fluid Mechanics, Heat and Mass Transfer*, volume 87 of *Mathematics in Science and Engineering.* Academic Press, 1972.

[337] Y. Bazilevs, C. Michler, V. M. Calo, and T. J. R. Hughes. *Weak Dirichlet Boundary Conditions for Wall-Bounded Turbulent Flows. Computer Methods in Applied Mechanics and Engineering*, 196:4853–4862, 2007.

[338] A. Ern and J-L. Guermond. *Theory and Practice of Finite Elements*, volume 159 of *Applied Mathematical Sciences.* Springer, 2004.

[339] J. Donea and A. Huerta. *Finite Element Methods for Flow Problems.* John Wiley & Sons, Inc., 2003.

[340] Y. Bazilevs, V. M. Calo, J. A. Cottrell, J. A. 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.

[341] P. Kagan, A. Fischer, and P. Z. Bar-Yoseph. *Mechanically based models: Adaptive refinement for B-spline finite element. International Journal of Numerical Methods in Engineering*, 57:1145–1175, 2003.

[342] M. J. Grote, A. Schneebeli, and D. Schötzau. *Interior penalty discontinuous Galerkin method for Maxwell's equations: Energy norm error estimates. Journal of Computational and Applied Mathematics*, 204:375–386, 2007.

[343] M. Dryja. *On discontinuous Galerkin methods for elliptic problems with discontinuous coefficients. Computational Methods in Applied Mathematics*, 3(1):76–85, 2003.

[344] J. T. Oden, I. Babuška, and C. E. Baumann. *A discontinuous hp finite element method for diffusion problems. Journal of Computational Physics*, 146:491–519, 1998.

[345] J. S. Hesthaven and T. Warburton. *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis and Applications*, volume 54 of *Texts in Applied Mathematics.* Springer, 2010.

[346] Y. Epshteyn and B. Rivière. *Estimation of penalty parameters for symmetric penalty Galerkin methods. Journal of Computational and Applied Mathematics*, 206:843–872, 2007.

[347] E. Burman, A. Quarteroni, and B. Stamm. *Interior Penalty Continuous and Discontinuous Finite Element Approximations of Hyperbolic Equations. Journal of Scientific Computing*, 43:293–312, 2010.

[348] K. Shahbazi. *An explicit expression for the penalty parameter of the interior penalty method. Journal of Computational Physics*, 205:401–407, 2005.

[349] K. Hillewaert. *Development of the Discontinuous Galerkin Method for High-Resolution, Large-Scale CFD and Acoustics in Industrial Geometries.* PhD Thesis, Université Catholique de Louvain, 2013.

[350] M. Drosson and K. Hillewaert. *On the stability of the symmetric interior penalty method for the Spalart-Allmaras turbulence model. Journal of Computational and Applied Mathematics*, 246:122–135, 2013.

[351] T. Warburton and J. S. Hesthaven. *On the constants in hp-finite element trace inverse inequalities. Computer Methods in Applied Mechanics and Engineering*, 192:2765–2773, 2003.

[352] Y. Saad. *Numerical Methods for Large Eigenvalue Problems.* Classics in Applied Mathematics. SIAM, second edition, 2011.

[353] V. E. Henson and U. M. Yang. *A Parallel Algebraic Multigrid Solver and Preconditioner. Applied Numerical Mathematics*, 41:155–177, 2002.

[354] J. M. Melenk, A. Parsania, and S. Sauter. *General DG-Methods for Highly Indefinite Helmholtz Problems. Journal of Scientific Computing*, 57:536–581, 2013.

[355] *Argonne Code Center: Benchmark Problem Book.* Report No. ANL-7416, Argonne National Laboratory, 1977.

[356] G. Theler. *Unstructured Grids and the Multigroup Neutron Diffusion Equation. Science and Technology of Nuclear Installations*, 2013:1–26, 2013.

[357] A. Hébert. *Application of the Hermite Method for Finite Element Reactor Calculations. Nuclear Science and Engineering*, 91(1):34–58, 1985.

[358] A. Hébert. *A User Guide for Trivac Version-4.* Report No. IGE-293, Ecole Polytechnique de Montréal - Institut de Génie Nucléaire, 2019.

[359] G. H. Golub and C. F. Van Loan. *Matrix Computations.* John Hopkins University Press, 1989.

[360] S. G. Hong and N. Z. Cho. *CRX: A Code for Rectangular and Hexagonal Lattices Based on the Method of Characteristics. Annals of Nuclear Energy*, 25(8):547–565, 1998.

[361] G. Marleau, A. Hébert, and R. Roy. *A User Guide for Dragon Version 5.* Report No. IGE-335, Institut de Génie Nucléaire, 2019.

[362] X. Yang and N. Satvat. *MOCUM: A two-dimensional method of characteristics code based on constructive solid geometry and unstructured meshing for general geometries. Annals of Nuclear Energy*, 46:20–28, 2012.

[363] F. Brezzi, L. D. Marini, and E. Süli. *Discontinuous Galerkin Methods for First-Order Hyperbolic Problems. Mathematical Models and Methods in Applied Sciences*, 14(12):1893–1903, 2004.

[364] J. Ragusa and Y. Wang. *A High-Order Discontinuous Galerkin Method for the S_N Transport Equations on 2D Unstructured Triangular Meshes. Annals of Nuclear Energy*, 36:931–939, 2009.

[365] T. S. Haut, P. G. Maginot, V. Z. Tomov, B. S. Southworth, T. A. Brunner, and T. S. Bailey. *An Efficient Sweep-Based Solver for the S_N Equations on High-Order Meshes. Nuclear Science and Engineering*, 193(7):746–759, 2019.

[366] H. G. Stone and M. L. Adams. *A piecewise linear finite element basis with applications to particle transport.* In *Proc. of the Nuclear Mathematical and Computational Sciences: A Century in Review, A Century Anew.* American Nuclear Society, 2003.

[367] A. Gibbons. *Algorithmic Graph Theory.* Cambridge University Press, 1985.

[368] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *Data Structures and Algorithms.* Addison-Wesley, 1987.

[369] G. R. Richter. *An Optimal-Order Error-Estimate for the Discontinuous Galerkin Method. Mathematics of Computation*, 50(181):75–88, 1988.

[370] P. Lesaint and P. A. Raviart. *On a Finite Element Method for Solving the Neutron Transport Equation.* In *Proc. of the Symposium on Mathematical Aspects of Finite Elements in Partial Differential Equations.* Mathematics Research Center, University of Wisconsin, 1974.

[371] B. W. Patton and J. P. Holloway. *Application of preconditioned GMRES to the numerical solution of the neutron transport equation. Annals of Nuclear Energy*, 29:109–136, 2002.

[372] Y. Y. Azmy. *Comparison of three approximations to the linear-linear nodal transport method in weighted-diamond difference form. Nuclear Science and Engineering*, 100:190–200, 1988.

[373] H. Khalil. *A Nodal Diffusion Technique for Synthetic Acceleration of Nodal S_N Calculations. Nuclear Science and Engineering*, 90:263–280, 1985.

[374] M. R. Zika and E. W. Larsen. *Fourier Mode Analysis of Source Iteration in Spatially Periodic Media. Transactions of the American Nuclear Society*, 79:135–, 1998.

[375] H. J. Kopp. *Synthetic Method Solution of the Transport Equation. Nuclear Science and Engineering*, 17:65–74, 1963.

[376] R. E. Alcouffe. *A Stable Diffusion Synthetic Acceleration Method for Neutron Transport Iteration. Trans. of the American Nuclear Society*, 23:203–, 1976.

[377] E. M. Gelbard and L. A. Hageman. *The synthetic method as applied to the S_N equations. Nuclear Science and Engineering*, 37:288–298, 1969.

[378] J. S. Warsa, T. A. Wareing, and J. E. Morel. *Fully Consistent Diffusion Synthetic Acceleration of Linear Discontinuous Transport Discretizations on Unstructured Tetrahedral Meshes. Nuclear Science and Engineering*, 141(3):236–251, 2002.

[379] J. S. Warsa, T. A. Wareing, and J. E. Morel. *Solution of the discontinuous P_1 equations in two-dimensional cartesian geometry with two-level preconditioning. Journal of Scientific Computing*, 24:2093–2124, 2003.

[380] J. S. Warsa, T. A. Wareing, and J. E. Morel. *Krylov Iterative Methods and the Degraded Effectiveness of Diffusion Synthetic Acceleration for Multidimensional S_N Calculations in Problems with Material Discontinuities. Nuclear Science and Engineering*, 147(3):218–248, 2004.

[381] Y. Y. Azmy. *Unconditionally Stable and Robust Adjacent-Cell Diffusive Preconditioning of Weighted-Difference Particle Transport Methods is Impossible. Journal of Computational Physics*, 182:213–233, 2002.

[382] K. L. Derstine and E. M. Gelbard. *Use of the Preconditioned Conjugate Gradient Method to Accelerate S_N Iterations. Transactions of the American Nuclear Society*, 50:260–261, 1985.

[383] M. R. Hestenese and E. Stiefel. *Methods of conjugate gradients for solving linear systems. Journal of Research of the National Bureau of Standards*, 49:409–436, 1952.

[384] Y. Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, second edition, 2003.

[385] B. Guthrie, J. P. Holloway, and B. W. Patton. *GMRES as a multi-step transport sweep accelerator. Transport Theory and Statistical Physics*, 28(1):83–102, 1999.

[386] S. F. Ashby, P. N. Brown, M. R. Dorr, and A. C. Hindmarsh. *A Linear Algebraic Analysis of Diffusion Synthetic Acceleration for the Bolzmann Transport Equation. SIAM Journal on Numerical Analysis*, 32(1):128, 1995.

[387] Y. Wang and J. C. Ragusa. *Diffusion Synthetic Acceleration for High-Order Discontinuous Finite Element S_N Transport Schemes and Application to Locally Refined Unstructured Meshes. Nuclear Science and Engineering*, 166(2):145–166, 2010.

[388] A. Barnett, J. E. Morel, and D. R. Harris. *A Multigrid Acceleration Method for the One-Dimensional S_N Equations with Anisotropic Scattering. Nuclear Science and Engineering*, 102:1–21, 1989.

[389] S. L. Campbell, I. C. F. Ipsen, C. T. Kelley, and C. D. Meyer. *GMRES and the Minimal Polynomial. BIT Numerical Mathematics*, 36:664–675, 1996.

[390] Y. Saad and M. H. Schultz. *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. Journal on Scientific and Statistical Computing*, 7(3):856–869, 1986.

[391] B. W. Patton and J. P. Holloway. *Application of Krylov subspace iterative methods to the slab geometry transport equation.* In *Proc. of the Radiation Protection & Shielding.* American Nuclear Society, 1996.

[392] B. W. Patton and J. P. Holloway. *Some Remarks on GMRES for Transport Theory.* In *Proc. of the Nuclear Mathematical and Computational Sciences: A Century in Review, A Century Anew.* American Nuclear Society, 2003.

[393] W. E. Arnoldi. *The principle of minimized iterations in the solution of the matrix eigenvalue problem. Quarterly of Applied Mathematics*, 9:17–29, 1951.

[394] E. W. Larsen. *Diffusion-synthetic acceleration methods for discrete-ordinates problems. Transport Theory and Statistical Physics*, 13:107–126, 1984.

[395] Y. Wang. *Nonlinear Diffusion Acceleration for the Multigroup Transport Equation Discretized with SN and Continuous FEM with Rattlesnake.* In *Proc. of the Joint International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering.* American Nuclear Society, 2013.

[396] S. Schunert, H. Hammer, J. Lou, Y. Wang, J. Ortensi, F. Gleicher, B. Baker, M. DeHart, and R. Martineau. *Using Directional Diffusion Coefficients for Nonlinear Diffusion Acceleration of the First Order S_N Equations in Near-Void Regions.* Report No. INL/CON-16-38952, Idaho National Laboratory, 2016.

[397] S. Schunert, Y. Wang, F. Gleicher, J. Ortensi, B. Baker, V. Laboure, C. Wang, M. DeHart, and R. Martineau. *A Flexible Nonlinear Diffusion Acceleration Method for the S_N Transport Equations Discretized with Discontinuous Finite Elements.* Report No. INL/JOU-16-37745, Idaho National Laboratory, 2017.

[398] J. E. Morel. *A Synthetic Acceleration Method for Discrete-Ordinates Calculations With Highly Anisotropic Scattering. Nuclear Science and Engineering*, 82:34–46, 1982.

[399] L. J. Lorence, E. W. Larsen, and J. E. Morel. *An S_2 Synthetic Acceleration Scheme for the One-Dimensional S_N Equations. Trans. of the American Nuclear Society*, 52:417–, 1986.

[400] J. E. Morel and T. Manteuffel. *An Angular Multigrid Acceleration Technique for S_N Equations with Highly Forward-Peaked Scattering. Nuclear Science and Engineering*, 107:330–342, 1991.

[401] M. L. Adams and W. R. Martin. *Diffusion synthetic acceleration of discontinuous finite element transport iterations. Nuclear Science and Engineering*, 111:145–167, 1992.

[402] T. Wareing, E. W. Larsen, and M. L. Adams. *Diffusion Accelerated Discontinuous Finite Element Schemes for the S_N Equations in Slab and X,Y Geometries.* In *Proc. of the ANS Topical Meeting, Advances in Mathematics, Computations and Reactor Physics*, 1991.

[403] B. Turcksin and J. C. Ragusa. *Discontinuous diffusion synthetic acceleration for S_N transport on 2D arbitrary polygonal meshes. Journal of Computational Physics*, 274:356–369, 2014.

[404] B. O'Malley, J. Kópházi, R. P. Smedley-Stevenson, and M. D. Eaton. *Hybrid Multi-level solvers for discontinuous Galerkin finite element discrete ordinate diffusion synthetic acceleration of radiation transport algorithms. Annals of Nuclear Energy*, 102:134–147, 2017.

[405] B. O'Malley, J. Kópházi, R. P. Smedley-Stevenson, and M. D. Eaton. *P-multigrid expansion of hybrid multilevel solvers for discontinuous Galerkin finite element discrete ordinate (DG-FEM-S_N) diffusion synthetic acceleration (DSA) of radiation transport algorithms. Progress in Nuclear Energy*, 98:177–186, 2017.

[406] J. O'Malley. *Parallel, Hierarchical Solution Algorithms for Diffusion Synthetic Acceleration of the Neutron Transport Equation.* PhD Thesis, Imperial College London, 2018.

[407] J. C. Ragusa and Y. Wang. *A two-mesh adaptive mesh refinement technique for S_N neutral-particle transport using a higher-order DGFEM. Journal of Computational and Applied Mathematics*, 233:3178–3188, 2010.

[408] Y. Wang and J. Ragusa. *Standard and goal-oriented adaptive mesh refinement applied to radiation transport on 2D unstructured triangular meshes. Journal of Computational Physics*, 230:763–788, 2011.

[409] 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.

[410] L. A. Hageman. *The Chebyshev Polynomial Method of Iteration.* Report No. WAPD-TM-537, Bettis Atomic Power Laboratory, 1967.

[411] D. C. Thomas, L. Engvall, S. K. Schmidt, K. Tew, and M. A. Scott. *U-Splines: Splines Over Unstructured Meshes.* 2018.

[412] C. Manni, A. Reali, and H. Speleers. *Isogeometric collocation methods with generalized B-splines. Computers & Mathematics with Applications*, 70(7):1659–1675, 2015.

[413] S. D. Pautz. *Discrete Ordinates Transport Methods for Problems with Highly Forward-Peaked Scattering.* Report No. LA-13444-T, Texas A&M University and Los Alamos National Laboratory, 1998.

[414] J. Hoschek and G. Seemann. *Spherical splines. Modélisation mathématique et analyse numérique*, 26(1):1–22, 1992.

[415] K. R. Koch, R. S. Baker, and R. E. Alcouffe. *Solution of the first-order forms of the 3D discrete ordinates equation on a massively parallel processor.* Report No. LA-UR–91-4157, Los Alamos National Laboratory, 1992.

[416] G. Gorman. *Parallel Anisotropic Unstructured Mesh Optimisation and its Applications.* PhD Thesis, Imperial College of Science, Technology and Medecine, 2006.