On uncertainty in fusion neutronics

Frederick Samuel Thomas

University of York
Physics

Doctor of Philosophy

September 2018

Abstract

This work explores the nature of several sources of uncertainty in the neutronics of nuclear fusion devices. Background information, radiation-matter interactions and their descriptions as 'nuclear data' are discussed before an introduction to several useful methods in nuclear analyses. After a discussion of the consequences of certain uncertainties in nuclear fusion systems, there is analysis which quantifies aspects of uncertainty in tritium breeding. This analysis uses Total Monte Carlo sampling of nuclear parameters to determine the spread in Tritium Breeding Ratio (TBR) values for the proposed DEMO reactor, due to uncertainty in our knowledge of neutron interactions with the lead nucleus. The TBR is found to have a standard deviation (1 σ) of 1.2% of the mean value. The higher-order distribution moments are also determined, highlighting a low-value tail. Subsequently, modelling approximations are investigated, specifically the practice of spatial homogenisation of radiation transport geometries. The effects of spatial homogenisation in radiation shielding for the ITER experiment are quantified for on-load and shut-down dose rates. These are found to contribute a relatively modest (22% maximum deviation) change on the most likely value. More general conclusions about the applicability of spatial homogenisation as a technique are drawn. The thesis then explores how the energy domain is discretised in current analyses and whether an improved method, delivering greater accuracy in activation calculations, might be created. By first exploring self-shielding theory and distributions of nuclear resonances, a series of functions for assembling an optimum bin density are devised. A nuclear data processing workflow is developed to produce data on an arbitrary energy 'group structure' and employed to create nuclear data libraries in an optimised group structure. These optimised discretisations are tested against and outperform current standard group structures in an example of relevance to the JET experiment. Recommendations for future work on improving the optimisation algorithm are given. Lastly, conclusions comparing the investigated contributions to analysis uncertainty are drawn.

Table of contents

List of figures

List of tables

Acknowledgements

I want to thank the York Plasma Institute and more widely the Fusion Centre for Doctoral Training for giving me the opportunity to pursue this research. It has been a pleasure to study, learn and then gradually make my own contribution to the field. My supervisors John Pasley, initially Lee Morgan and latterly Michael Fleming have all been generous with their time, giving guidance and reassurance when I needed it.

Thanks to Jon Shimwell, Andrew Turner, Jon Naish, Bethany Colling, Tim Eade and all the other staff at the Culham Centre for Fusion Energy (CCFE). Their skills, advice, favours, amusing cynicism and questionable softball tactics have made my time at Culham all the more enjoyable. I give my thanks and best wishes to the current and past students of CCFE. Let us pray that if Charlie ever achieves a position of power he exercises it benevolently. I must credit the role that lunchtime football has had in facilitating my otherwise sedentary work. I will miss Darren's bombast, Alan's sarcasm and Jorge's finesse. Long may those games continue.

Thanks to Igor Lengar of the Jožef Stefan Institute for sharing his radiation transport models with me. I am very appreciative to Michael Loughlin for his advice and for arranging the research placement at ITER Organisation in Provence, it was a privilege. Jakhar Shrichand and Vladimir Barabash of ITER were also a great help to my work.

Outside of professional life, thank you to my girlfriend, Caitlin, for her silly dancing, serious counsel and shared jokes over the past four years. Also, to my parents, Max (who reminds me to work) and Gill (who reminds me there is much else, too). My friends in Oxford have been a persistent and growing joy since not long after I moved here and I credit our community with keeping me calm and content. You're all wonderful.

This work was funded by the Engineering and Physical Sciences Research Council (EPSRC) through grant EP/L01663X/1.

With four parameters I can fit an elephant,
and with five I can make him wiggle his trunk.
– John von Neumann

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, university. All sources are acknowledged as references. Research presented in this thesis has also been published in the following papers:

- Thomas, F. & Fleming, M. (2017). Quantifying TBR uncertainty due to lead nuclear data in HCLL blanket modelling by the Total Monte-Carlo method. *Fusion Engineering and Design*. 124:814-817. doi: 10.1016/j.fusengdes.2017.03.118
- Thomas, F. & Fleming, M. (in press). Optimised energy group structures for fusion activation calculations. *Fusion Engineering and Design*. doi: 10.1016/j.fusengdes.2018.05.038

Frederick Samuel Thomas
September 2018

Chapter 1

Introduction

1.1 Motivation

Hominids have always manipulated and altered their environment. Ancient people replanted wild flora and helped drive various megafauna such as woolly mammoth and steppe bison to extinction [88] [106]. Today, we are still changing our surroundings, but at a historically unprecedented rate. Paul Crutzen popularised the term anthropocene in the early 2000s, defining it as the current geological period of Earth's existence, one where humanity has a significant impact on the planet's various systems, in many ways out-competing natural processes [30]. Our rapacious desire for resources has felled forests, excavated vast holes in the Earth, driven countless species extinct and burnt fossil fuels on a tremendous scale, releasing greenhouse gases (GHG) including carbon dioxide into the atmosphere. This last process has driven atmospheric CO_2 concentrations from 280ppm to 410ppm in two centuries [95]. The rate of accumulation is currently 2ppm and accelerating [95]. This accumulation of GHGs is driving climate change, leading to a near-surface temperature increase of 0.8°C since the start of the 20th century [51]. The consequences of the increased energy retained in the Earth system are many and varied. Sea levels have risen by 19cm in the same period [24]. Arctic sea ice coverage has decreased by $\approx 10.4\% \pm 1.7\%$ per decade in the last 30 years [93]. The condition of the planet will undoubtedly continue to deteriorate in the coming decades. Continuing sea level rises will threaten tens of millions living at sea level in Bangladesh and Indonesia, amongst many other countries [40]. Incidence of tropical diseases will worsen in equatorial regions and spread to the warming temperate zones [126]. Extreme precipitation events have already become more intense by 3.3% and this trend is modelled to continue with a 5.2% increase for each degree of warming [140]. Many of these conditions are now 'baked-in' and will happen regardless of humanities efforts to avert them. This is because there is a lag between addition of GHG to the atmosphere and a new temperature equilibrium

being reached of at least a decade [110], but perhaps as much as a century for large GHG emissions [143]. Two degrees of atmospheric warming over pre-industrial temperature is now extremely likely, with four or more degrees a possibility [61]. We must now mitigate the risks posed by our changed and changing planet and act to prevent warming worse than two degrees, for the consequences of further increases are dire.

The main sources of GHG emission that humans are responsible for are combustion of fossil fuels for heat (industrial and domestic heating and cement production), transport (internal combustion engines and jet engines) and electricity production (typically with a boiler and steam turbine) and agriculture (deforestation and animals' methane production) [87]. We must rapidly reduce our GHG emission from current levels to near-zero or even negative rates [87]. In order to achieve this, heating, transport and electricity production must be decarbonised. Unfortunately, removing fossil fuels from heating and transport is most likely to be achieved with their electrification. This, coupled with a growing population and desire for increasing material living standards, means low-carbon electricity demand is going to dramatically increase in the near future. One study suggests the UK's annual electricity demand will grow from 340 to 423 TWh [13].

Low-carbon electricity production is currently produced from nuclear fission and renewables such as solar photo-voltaic, solar thermal, wind, hydro, tidal, geothermal and biofuels. The acronym WWS (Wind, Water and Solar) summarises the renewable methods with greatest potential for widespread adoption.

While nuclear fission is very safe by the metric of energy produced per death caused [89] it has spawned several major accidents, suffers from issues of public acceptability, radioactive waste storage and, more recently, construction cost. Additionally, the current trend for an open fuel cycle where nuclear fuel is not reprocessed and then reused is wasteful. This practice drastically limits the potential of nuclear fission at scale, with current uranium reserves running out by the end of this century [96]. WWS have tremendous potential and are currently gaining market share, with costs falling rapidly. For instance, some utilities are now bidding to supply electricity for as little as 3¢ kWh^{-1} [60]. For comparison, the Hinkley Point C nuclear fission plant will have a minimum supply price of 9.25p kWh^{-1} [91]. While some researchers have advocated for 100% WWS by 2050 [62] this seems unlikely with current energy storage technology as intermittancy is a significant problem for these technologies, specifically wind and solar [26]. Much work is required to fully exploit their potential: drastic improvements in energy storage methods, a commensurate deployment of such technology, intelligent demand optimisation and improved distribution networks. Even then, it is unclear if WWS and/or nuclear fission could provide enough electricity for our future society.

1.2 Nuclear fusion

A potential future alternative to both WWS and nuclear fission is nuclear fusion, the process of combining nuclei by which stars shine. Nuclear fusion power has several desirable attributes:

- Low CO_2 emissions
- High fuel energy density (reduced resource extraction burden)
- Abundant fuel (or feedstock for the manufacture of fuel)
- Dispatchable power generation
- No criticality (runaway) risk, unlike nuclear fission power
- High plant power density (GW scale facilities on the order of 1 km^2).

However, to date the goal of electricity generation from nuclear fusion has eluded scientists and engineers.

1.2.1 Discovery

Einstein's energy-mass equivalence, $E = mc^2$ led Arthur Eddington to hypothesise that the energy released by the sun was through the fusion of hydrogen nuclei into helium nuclei. However, contemporary predictions of the temperature required for the sun's power output did not match observations. Indeed, the temperatures observed seemed too low for nuclei to overcome the Coulomb energy barrier - the repulsive force experienced by two nuclei. The development of a theory of quantum tunnelling in the late 1920s by Gamow and others, predicted thermonuclear temperatures in accordance with observations of the sun and resolved the problem of stellar temperatures.

In the 1930s Mark Oliphant, Paul Harteck and Ernest Rutherford conducted particle beam experiments, bombarding various targets with deuterons. After several less interesting combinations, they observed "on bombarding heavy hydrogen with diplons[1] an enormous effect was produced" [97]. This included the emission of high energy neutrons. Soon thereafter, plans for fusion weapons and controlled fusion power were developed.

[1] Diplon was a contemporary name for what we now call a deuteron.

1.2.2 Possible reactions to harness

Trying to fuse nuclei is difficult, as scattering via the Coulomb force is far more likely than fusion for all potential reactants. This Coulomb barrier is proportional to the product of the charges of the reactants as shown in equation 1.1, where k_e is the Coulomb constant, Z_i the respective atomic numbers, e the charge on the electron and r, the interaction radius.

$$U_{coul} = k_e \frac{Z_1 Z_2 e^2}{r} \tag{1.1}$$

Given this, it is only plausible to employ light nuclei for energy production. Of the light nuclei, certain reactions are more likely than others, with a greater cross-section for a given collision energy. They are shown in table 1.1 below.

Fuel	Products	Q-value (MeV)	$\sigma(10\text{keV})(\text{barn})$	$\sigma(100\text{keV})(\text{barn})$
D + T	α + n	17.6	2.72×10^{-2}	3.43
D + D	T + p	4.04	2.81×10^{-4}	3.3×10^{-2}
D + D	^{3}He + n	3.27	2.78×10^{-4}	3.7×10^{-2}
D + ^{3}He	α + p	18.3	2.2×10^{-7}	0.1
p + ^{11}B	3α	8.68	4.6×10^{-17}	3×10^{-4}

Table 1.1 Comparison of potential fusion fuels. The Q-value is the combined kinetic energy of the reaction products. σ(E) is the cross-section of the reaction at a given energy, E, with centre-of-mass energies.

At an energy of 10 keV the D-T fusion reaction has a likelihood 100-fold of the next easiest reactions. It also has a large Q-value, with 17.6 MeV of binding energy released by the reconfiguration of the fuel nucleons. This energy is split between the products such that they have equal momentum, with the α-particle receiving $\frac{17.6}{5}\text{MeV} = 3.5\text{MeV}$ and the neutron the remaining 14.1 MeV. D-^{3}He and p-^{11}B fuels are aneutronic, with only charged products. These fuels would obviate several difficult problems associated with D-T fusion which arise from the energetic emitted neutrons. The principal benefit would be a reduced burden on materials science to develop radiation tolerant structural and plasma facing materials, qualified in time for usage in a reactor. There would also be a significant reduction in the quantity of radioactive waste generated, compared with a D-T fuelled reactor. Unfortunately, aneutronic fuels require extraordinary temperatures to fuse, far beyond anything achieved to date, or conceived as possible in a tokamak in the near future. D-T fuel is currently envisioned as the fuel mixture for nearly all government sponsored research programmes. Some private

efforts are attempting to use other fuels, especially those aforementioned aneutronic reactions [133].

1.2.3 Development

The 1950s saw the invention and development of so-called 'pinch' fusion concepts. As a current flows through a tube of conducting medium (such as a plasma) the magnetic field generated acts in tandem with the current to crush the material to a filament by a Lorentz, or $\mathbf{j} \times \mathbf{B}$ force. This is a kind of confinement, momentarily holding plasma ions in an increased density state. By bending a tube into a toroidal shape, matter might be circulated about a device, permanently confined. In reality, the gradient of the electric field across the torus acts to break confinement and lose material. The pinch concept was developed by Soviet scientists Sakharov and Tamm, into the tokamak. This was similar to previous designs of toroidal pinch machines, such as the British ZETA, but the relative strength of fields was different. Rather than having the dominant field be produced by the plasma current, in a tokamak the dominant field is generated by the toroidal field coils which wrap around the toroidal plasma chamber. Over the coming decades this would prove to be a more stable confinement system than other approaches.

Throughout the rest of the 20th century, the tokamak concept came to dominate magnetic confinement fusion research. Plasma physicists around the world continued to design devices and plasma scenarios which attempt to avoid plasma instabilities and maximise the plasma triple product, developed from Lawson's criterion [77]. The triple product, $nT\tau_E$, is the product of the plasma density, temperature and energy confinement time, respectively. Figure 1.1 shows a plot of this performance metric against time. One can see the progress in relation to the well known Moore's Law in computing performance per unit cost and the beam energy of various particle colliders. Viewed in this way, fusion power has achieved remarkable progress across the decades (logarithmic and temporal).

1.2.4 Future plans

There are many aspects of nuclear fusion power which require significant work before an electricity producing reactor could be constructed. In the field of plasma physics these include, improving energy and particle confinement, quelling various powerful emissions from the plasma (disruptions, ELMs, fast ions, runaway electrons) and increasing plasma temperatures. The ITER reactor under construction in the south of France will provide a test for plasma physics knowledge, through the equilibrium operation of a 'burning' plasma. However, the

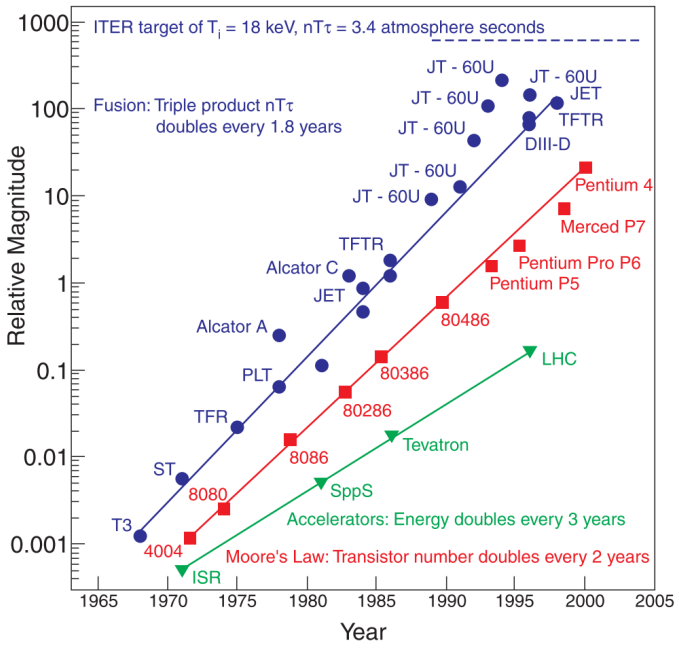

Fig. 1.1 Fusion triple product achieved plotted against time. Similar progress measures for computing and particle physics are also shown. Figure from [137].

plasma 'core' is only one aspect of a functioning fusion reactor. The surrounding systems of an electricity producing reactor necessarily include:

- Fuelling by gas or pellet injection
- Diagnostics to determine the state of the plasma and the plant
- Heating systems including micro and radio-waves and energetic particle beams
- Plasma heat and particle exhaust (divertor)
- Blanket for fuel breeding and capture of neutron energy
- Electricity production with heat exchangers, gas turbines and generators
- Tritium extraction, storage and recycling
- Cryoplant for the cooling of magnetic coils and other systems
- Remote handling for the maintenance of the other systems

While ITER will advance our understanding and experience of several of these ancillary systems, additional devices will be required before fusion electricity can be realised. These

might include one or several demonstration reactors. Progress is crucial in materials science, as the systems listed above must operate for years without being degraded by plasma etching, radiation damage and stresses induced by extreme temperature gradients. Consequently, it is commonly held that a device for the production of 14.1 MeV neutrons at a useful flux is also required for the development and testing of radiation tolerant materials.

1.3 Radiation-matter interactions

Nuclear fusion gives rise to various energetic particles. The neutrons are emitted with $\frac{4}{5}$ of the total energy liberated, to give 14.1 MeV of kinetic energy. Photons are constantly generated in a tokamak plasma. Low energy photons are created through the excitation and de-excitation of atomic electrons, while Bremsstrahlung radiation arises from the acceleration of charged particles. This mixed radiation field is then further complicated by numerous interaction process, as neutrons and photons from the plasma interact with the device. What follows is a short primer on the possible interaction processes for the principle particles, the neutron and photon.

1.3.1 Neutron

The neutron is a sub-atomic particle of mass $1.674927471(21) \times 10^{-27}$kg [92], composed of three quarks. It can exist as part of a nucleus or unbound, where it has a mean lifetime of 877.7 ± 0.7s [101]. While it is itself uncharged it can produce ionising radiation through interactions with other matter. These interactions are via the strong, weak, gravitational and electromagnetic[2] forces. When neutrons interact with nuclei, the outcome is a function of the nuclei rest mass and the kinetic energies of the reactants. For low energy neutrons, elastic scattering, diffraction, and compound nuclear reactions like neutron capture and sometimes fission are the dominant processes. At higher energies, inelastic scattering and direct reactions such as knockout or transfer reactions are more common. The neutron-matter interaction processes of interest for this work are outlined below.

1.3.1.1 Elastic scattering

Neutrons can scatter off nuclei. In the context of fusion neutronics an elastic scattering event is defined as a neutron-nucleus reaction where no kinetic energy is transferred into excitation of the nucleus. Both the kinetic energy and momentum of the reactants are

[2]While the neutron has zero electric charge, it does have a magnetic moment, and is therefore acted upon by electromagnetic fields.

conserved. Although scattering is a quantum-mechanical phenomenon, properly described by the interacting wave functions, in this case a 'billiard ball' treatment satisfactorily describes reactions.

The fast fusion ($>$ 1MeV) neutrons are significantly more energetic than the nuclei in condensed matter ($\approx$ 1eV), as such, the target nuclei may be thought of as at rest. In elastic scattering the energy lost by an incident neutron is gained by the target nucleus. The energy transferred is a function of the target nucleus mass, as shown below.

$$\frac{1}{2}m_n v_{n,i}^2 = \frac{1}{2}m_n v_{n,f}^2 + \frac{1}{2}m_a v_{a,f}^2 \tag{1.2}$$

Where m_n is the mass of the neutron, m_a the mass of the target nucleus, $v_{n,i}$ is the initial velocity of the neutron, $v_{n,f}$ the final velocity of the neutron and $v_{a,f}$ the final velocity of the nucleus. Conservation of momentum can be written as below.

$$m_n v_{n,i} = m_a v_{a,f} - m_n v_{n,f} \tag{1.3}$$

With equations 1.2 and 1.3 one can derive a quantity known as α, the fraction of the initial energy retained by a neutron in a head on collision, as a function of target nucleus mass, A. For the derivation see [53]. The expression for α is as follows.

$$\alpha(A) = \left(\frac{1-A}{1+A}\right)^2 \tag{1.4}$$

One can see that lighter nuclei, of closer mass to neutrons more effectively moderate the energy of neutrons. A collision with a single proton will halt the incident neutron, $\alpha(1) = 0$. Heavier nuclei such as ^{186}W give $\alpha(186) = 0.9787$ so scattered neutrons retain almost all of their energy.

1.3.1.2 Nuclear reactions

Nuclear reactions are distinct from elastic scattering. Rather than solely redistributing energy and momentum, reactions reconfigure nuclei and create new particles. There are two principle ways this can happen:

- Direct nuclear reactions - a single nucleon in the target particle interacts with the incident particle. The interaction time is very limited, around 10^{-21}s (i.e. the incident energy is high). Direct reaction products are anisotropically distributed, typically forward focused.

- Compound nucleus reactions - many nucleons interact together over a greater period of time, perhaps $10^{-18} - 10^{-16}$s. The incident particle and the target nucleus coalesce into a new, excited nucleus. The collection of nucleons within the nucleus reach thermal equilibrium after a series of collisions. At some point, the excited compound nucleus will decay, but the mode of decay will not depend on the method of compound nucleus formation. Instead, the decay mode, or 'exit channel' is dependent on the compound nucleus excitation energy and various nucleus-specific probabilities.

This dichotomy is not entirely accurate, as the incident particle always affects multiple nucleons, but for so-called direct nuclear reactions, this is minor. Reactions can also be exothermic or endothermic, liberating or requiring energy to occur, respectively. This is recorded by the 'Q-value' of a given reaction, say ^{208}Pb(n,2n)^{209}Pb, where Q = -7.37 MeV. This neutron multiplying reaction is endothermic and requires an energy input to occur. A plot of cross-section against energy will clearly show this 'threshold' behaviour, with zero probability of reaction until the minimum input energy is achieved. Exothermic reactions, by contrast can occur at any energy.

Inelastic scattering An inelastic neutron scatter involves the target nucleus absorbing the incident neutron, forming an excited, compound nucleus and then re-emitting the neutron. The nucleus remains excited and emits a gamma ray to dissipate excess energy and reach its ground state. Given this and radiative capture, neutron fields usually beget gamma fields.

Radiative capture Radiative capture is similar to inelastic scattering, with an incident neutron also forming an excited compound nucleus. However, here a neutron is not re-emitted, only a gamma ray. This reaction is more likely at lower energies, given the slower relative velocity of reactants and therefore longer period for interaction. The reaction acts as a neutron 'sink', removing neutrons from the system and is therefore an essential component of neutron shielding, along with moderating material and gamma-shielding.

Fission Inter-nucleon nuclear forces hold many nuclei together in a roughly spherical shape, however others are substantially distorted. These 'deformed nuclei' typically occur in the mass range $150 < A < 190$ and $A > 220$ [75]. Whilst binding energy goes like A, Coulomb repulsion amongst protons scales faster, more like Z^2. The fission process is governed by these competing nuclear and Coulomb forces. For large, unstable nuclei the balance between these forces is close and additional energy added to the system can act to raise the nucleus from its ground state out of its potential well, to a situation where separating into fragments becomes the most energetically favourable path. In reality, the situation is a

little more complex, with nuclear shell effects introducing a secondary potential well which must be crossed during nucleus deformation.

Nuclear fission is not normally associated with nuclear fusion reactors, however uranium is present in beryllium containing ores and even after refining, some remains as an impurity. One variant of the proposed European fusion reactor design, DEMO, is set to contain 560 t of beryllium. At 30 wppm uranium concentration (the ITER requirement) this results in 17 kg of uranium [66]. A small amount is fissile as ^{235}U, while the remaining ^{238}U is fissionable by high energy neutrons. The daughter nuclei from actinide fission and bred transuranics will undoubtedly complicate the process of reactor component recycling [19].

Multiplication An incident neutron can result in multiple neutrons being emitted, a so-called multiplication reaction, (n,2n), (n,3n), etc. the reaction is always endothermic. Multiplication reactions are important in the context of tritium breeding in nuclear fusion blankets, increasing the neutron flux and therefore the production of tritium. Higher yield multiplication reactions have a greater energy requirement and thus a higher energy threshold. Materials such as beryllium and lead have high multiplication cross-sections, with beryllium having the lowest energy threshold for (n,2n) by some margin at approximately 2.7 MeV.

1.3.2 Photon

The photon is an elementary particle, indivisible by nature. Aggregated, photons provide one way of conceptualising an electromagnetic field. Photons travel in a vacuum at 299,792,458 ms^{-1} and are created by various processes, across a wide spectrum of energies. The sources inside a fusion reactor include various plasma processes, neutron-matter interactions and radioactive decay amongst others.

1.3.2.1 Radioactive decay

The gamma ray, or gamma photon, is a photon emitted by a nucleus. They can have energies from several keV to many MeV, overlapping in energy with X-rays[3] The gamma ray is emitted from an excited nucleus, often preceded by α or β decay. As stated in section 1.3.1.2 gamma emission also occurs as a result of compound nucleus formation during radiative capture and inelastic scattering. Gamma photons are ionising radiation and thus biologically hazardous.

[3]X-rays are produced by interactions with and between electrons. These processes include Bremsstrahlung radiation, or X-ray fluorescence. Aside from an arbitrary energy and associated wavelength distinction, gamma photons and X-ray photons cannot be distinguished without knowledge of their source.

1.3.2.2 Bremsstrahlung

Literally 'braking radiation', Bremsstrahlung is the process whereby photons are generated as a charged particle is decelerated. As the braking particle is slowed, the kinetic energy lost is conserved to become a new photon. The energy spectrum of the emitted photons is continuous and will move to higher frequencies as the bombarding particle energy is increased. Bremsstrahlung is an important process in the plasma physics of nuclear fusion power. Unfortunately, Coulomb scattering is significantly more likely than fusion and these ion-ion and especially electron-ion scattering events generate significant Bremsstrahlung radiation, slowing the potential reactants and dissipating energy from the plasma. The expression for Bremsstrahlung power loss in a plasma is given as equation 1.5.

$$P_b \propto T_e^{\frac{1}{2}} n_e \sum_i n_i < Z_i >^2 \tag{1.5}$$

It can be seen that the power is proportional to the square root of the electron temperature and the square of the average ion charge state [121]. The latter relationship is one reason why so-called 'advanced' fusion fuels such as p-^{11}B are so much more difficult to fuse than purely hydrogenous fuels [111].

1.3.2.3 Photoelectric

The photoelectric effect is the phenomenon where incident photons are absorbed by atomic electrons and the electron is rapidly dislodged from the atoms electrostatic attraction. The energy threshold for the ejection of the most loosely bound electrons is known as the element's work function and is typically in the range 2–6 eV. It is in this low energy range where the photoelectric effect is most prevalent. An incident photon's energy can be described as $E = h\nu$ where h is Planck's constant and ν is the photon frequency. If this energy is smaller in magnitude than the work function, no electron emission is observed, regardless of the photon flux. Larger incident photon energies can eject more tightly bound electrons, where energies in the MeV range may be required for photoelectric emission in some elements. The photoelectric effect is a major mechanism for the absorption and therefore 'sinking' of photons.

1.3.2.4 Compton scattering

Compton scattering is a mid-energy range photon-matter interaction between photons and charged particles. Incident photons lose energy, transferring it to the recoiling charged particle, usually an electron. The reduced photon energy corresponds with an increase in

photon wavelength. The relationship between initial and final wavelength and scattering angle is shown as equation 1.6, where λ is the initial wavelength, λ' is the scattered wavelength, h is Planck's constant, m_e the mass of the electron and θ the scatter angle.

$$\lambda' - \lambda = \frac{h}{m_e c}(1 - cos\theta) \tag{1.6}$$

1.3.2.5 Pair production

Pair production is the general term for the creation of a particle and its antiparticle from an electrically neutral boson. In the context of photon interactions it is typically understood as the creation of an electron and positron from an energetic photon i.e., $\gamma \rightarrow e^- + e^+$. The photon is destroyed in the interaction. The particles are created from the incoming photon energy, which therefore must be at least equal to the rest mass of the particle pair, $2 \times 0.511\text{MeV} = 1.022\text{MeV}$. The process is the dominant high-energy photon-matter interaction.

1.4 Radiation transport methods

Finding the distribution of particle flux as a function of position, energy, time and angle is the essential step in computing nuclear responses such as reaction rates, volumetric heating, etc. This flux and its contributing and detracting terms can be represented by the Boltzmann transport equation, an integro-differential form of which is shown as equation 1.7.

$$[\boldsymbol{\Omega} \cdot \nabla + \sigma_t(\boldsymbol{r}, \boldsymbol{\Omega}, E)]\, \phi(\boldsymbol{r}, \boldsymbol{\Omega}, E) = \int_0^\infty \int_{4\pi} \kappa(\boldsymbol{r}, \boldsymbol{\Omega} \leftarrow \boldsymbol{\Omega}', E \leftarrow E') \phi(\boldsymbol{r}, \boldsymbol{\Omega}', E', t) d\boldsymbol{\Omega}' dE' + q(\boldsymbol{r}, \boldsymbol{\Omega}, E) \tag{1.7}$$

Where $\boldsymbol{r}$ is the spatial distribution of neutrons, $\boldsymbol{\Omega}$ the angular distributions of the neutrons and E the energy of the neutrons. ϕ is the angular neutron flux density. κ represents reactions where neutrons are scattered from E' and $\boldsymbol{\Omega}'$ in energy and angle to new values, E and $\boldsymbol{\Omega}$. σ_t is responsible for the loss of neutrons and q is a source term for neutrons [52]. Equation 1.7 is essentially a balance equation, conserving neutrons. It is difficult to solve for any real system because of the number of dimensions involved is large. There are essentially two methods for finding the flux distribution:

- Deterministic methods which numerically solve a version of the Boltzmann transport equation.

- Stochastic methods which run a large number of particle experiments to deduce the ensemble behaviours.

Deterministic methods are suited to solving simple geometries, perhaps 1D or 2D problems. They are often used for fission reactor problems, where a chain of different methods are used to solve for the flux distribution. These trade progressively larger physical scale for a simpler version of the physics, moving from the Boltzmann equation in the form above to a more simplified diffusion version.

Complicated 3D geometries are best suited to the stochastic method, where millions or billions of source particles are individually simulated. Their individual lifespans sample first from a source distribution, q, of the neutron phase space $(\boldsymbol{r}, \boldsymbol{\Omega}, E)$ and commence travelling through the problem. Random numbers are generated to sample from distributions which represent the physical behaviour of particles within materials. For example, the particle might scatter multiple times, drawing from distributions to determine the new scatter E' and $\boldsymbol{\Omega}'$, sampling distributions to determine the path length before another interaction. This method does not solve equation 1.7, but rather samples to estimate a behaviour. As a result, results of interest must be explicitly 'tallied'.

1.5 Material inventory methods

Tracking the nuclide inventory of a nuclear system is essential to determining its changing behaviour over time. This is especially true if the neutron flux is intense or particularly energetic (as is the case for fusion reactor concepts). Activation of previously stable nuclides into radioactive ones leads to a Shut-Down Dose Rate (SDDR) due to decay γ photons. Transmutation of elements from one to another can have substantial effects on the mechanical properties of materials. The transmutation of W into Ta, Re and Os in fusion reactor first walls is a particularly acute example. A pure W part will be transmuted such that it is >6% other elements by the end of its 5 year design life [48]. This can lead to embrittlement and other degradations. In order to simulate these effects and design more suitable reactor materials, one must be able to track the nuclide inventory as a function of time. This can be done with a series of Bateman equations, as demonstrated in equation 1.8.

$$\frac{dN_i}{dt} = -N_i(\lambda_i + \sigma_i\phi) + \Sigma_{j\neq i} N_j(\lambda_{ij} + \sigma_{ij}\phi) \tag{1.8}$$

Where ϕ is the incident particle flux. N_i, λ_i and σ_i are the number density, decay constant and cross-section for reactions with nuclide i. σ_{ij} is a cross-section for producing i from j, and λ_{ij} a decay constant for decay from j to i. N_j is the number density of nuclide j.

Equation 1.8 shows the balance of nuclei for a particular nuclide. Assembling a coupled system of these equations, with one for every nuclide, it is in theory possible to analytically solve for the nuclide inventories at any time. In practice, this is too difficult and numerical methods for the solution of the coupled differential equations are instead used.

1.6 Nuclear data

Nuclear data (ND) is the general term for nuclear physics quantified and recorded for later retrieval. The behaviour of nuclei, in particular when under bombardment by nucleons and light nuclei can be probed by a variety of experiments. These utilise beams of particles and arrays of detectors to determine reaction likelihood (cross-section) and information on the reaction products (emitted energy & angle). Once experimental nuclear interaction findings are generated, they are typically submitted to and stored in the EXFOR database. This is a publically queryable dataset of experiments to date. While the EXFOR dataset covers many nuclides, reaction channels and energies, this parameter space is very large and expensive to explore with nuclear physics experiments. In some instances, for example with short-lived radioisotopes such as metastables, it is not possible to perform the appropriate experiment. In this case, any extant experimental data is augmented with information derived from theoretical nuclear physics models. These data, from experiment and theory are used to develop or 'evaluate' a bundle of nuclear data typically grouped by incident projectile and target nuclide. The evaluated files include likelihoods of interaction for various reactants (cross-sections) but also reaction product energy spectra, angular distributions, decay data, uncertainty information and more. There is a globally standard format known as Evaluated Nuclear Data File (ENDF) which specifies the scope, layout and precision of these files. The current version of the ENDF format is ENDF-6 [56]. As well as the reaction information listed above and detailed below, the files also contain information on contributing experiments and modifications which occurred during data evaluation. There are a series of different nuclear data libraries which constitute different attempts to capture and record the nuclear behaviour of a variety of nuclides. These may be general purpose (e.g. JENDL [123], TENDL [115] and the confusingly named ENDF/B [17]) or more specialist (e.g. EAF[98]). Many countries with nuclear programmes have developed their own nuclear data libraries over time.

ENDF type files are typically stored as ASCII text. They are arranged hierarchically as displayed in figure 1.2. The MAT or material is typically but not exclusively a nuclide, it can also be elemental or even compound in nature. A MAT code is a unique reference for a material, such as 125 for ^{1}H or 7443 for ^{186}W. Each material contains several 'files' referred

to by an MF number. These contain the data listed in the sections below. Each MF, or file, is split into MTs.

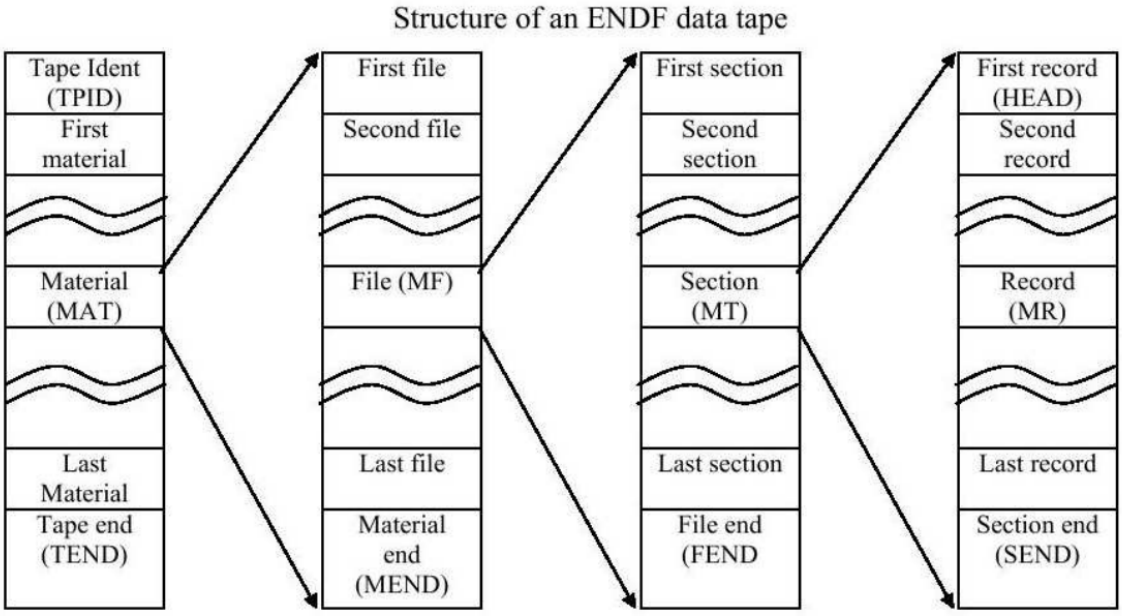

Fig. 1.2 The layout of an ENDF 'tape', containing materials (typically only one material per tape), which contain files, which contain sections, which contain records.

Nuclear data 'processing' is required to turn the contents of an ENDF formatted file into a form that radiation transport and activation-transmutation codes can use. This is accomplished with codes such as NJOY [85] or PREPRO [32]. These takes ENDF type files and process the wealth of information present, discarding some, performing various types of interpolation on others, etc. The resulting processed files, of use to nuclear analysts and other end users, are explained in chapter 4, section 4.3.1.

A brief introduction to the nuclear physics encoded in ENDF type data and its various uses are given below.

1.6.1 Resonance parameters

The cross-section as a function of interaction energy is often an irregular, unsmooth shape due to the presence of nuclear resonances. These resonances are compound nuclei. That is, compound nuclear reactions where a short lived particle, a compound nucleus is temporarily formed. These nuclear resonances occur because the nucleus is a quantum mechanical system, with certain, discrete energy levels permitted. Methods for predicting nuclear reactions based on forces within the nucleus are a work in progress (see [9] for an example of current state of

the art *ab initio* modelling). Inter-nucleon forces are many-body and very complex, especially for larger nuclei [94].

Insights on reaction behaviour can reliably be gained from 'black-box' models of the nucleus, where internal understanding is not assumed. One such approach is the R-Matrix interaction framework. In this method, unknown internal properties are parameters which can be deduced by comparison with experimental results. This method can be performed to varying degrees of complexity, with or without particle spin dependence, multiple reaction channels, etc. Other resonance formalisms are typically simplifications of the R-Matrix approach.

Rather than trying to record the cross-section for all interaction energies in a tabulated way, it is easier to reconstruct cross-sections from some 'background' (σ, E) tuples and a series of summed resonance curves parameterised by resonance theory. These curves have a functional description, dependent on energy with various parameters determining their shape. A simple Breit-Wigner distribution is shown as equation 1.9. This form is appropriate for non-overlapping resonances. The more sophisticated R-Matrix approach is required with overlapping resonances.

$$\sigma(E) = \frac{\pi\lambda\!\!\!^{-2}(2J+1)}{(2s_n+1)(2s_t+1)} \frac{\Gamma_i\Gamma_f}{[(E-E_r)^2+\Gamma_t^2/4]} dE \tag{1.9}$$

Where E is centre-of-mass energy of the system, Γ_i is the partial resonance width to decay to the initial state, Γ_f is the partial width to decay to the final state, Γ_t is the total width, $\lambda\!\!\!^{-}$ the reduced particle wavelength, E_r the rest mass energy of the resonance, J the resonance total angular momentum, s_n the neutron spin and s_t the target spin [80]. The Γ values, or widths, are inversely related to the stability of the resonant nucleus, those with a very narrow width will correspondingly exist for longer before decay than those with very large widths.

The shape of resonance peaks in the interaction likelihood and energy space is a function of temperature. This is best explained by considering a monoenergetic neutron flux incident on a block of material. Despite the flux beam being monoenergetic in the lab frame, the neutron-nucleus interactions have a distribution of energies. This is dependent on the temperature of this material (higher temperatures act to widen the spread of interaction energies). As such, resonances are Doppler broadened as material temperature increases. An advantage of reconstructing cross-sections from resonance parameters is that the temperature dependence of cross-sections can be handled in an elegant and flexible way.

While resonances at lower energies are detectable by experimental methods, at higher energies it becomes difficult to distinguish between neighbouring resonances. This is because the average spacing between nuclear levels, $\langle D \rangle$, decreases with energy while the average

resonance width, $\langle\Gamma\rangle$, increases. Where $\langle D\rangle \approx \langle\Gamma\rangle$ there exist bunches of partially overlapping resonant cross-section peaks. This is the transition from the Resolved Resonance Range (RRR) to the Unresolved Resonance Range (URR). An overview of the cross-section from thermal to MeV energies for a heavy nucleus is shown as figure 1.3. Resonances in the URR region are still important for self-shielding effects (see chapter 4). While no experimental data is available on the precise resonance locations in the URR, average level widths and spacings are given in the appropriate ND. This average approach can be improved on with so-called probability tables. Resonance parameters are stored within the MF=2 file of an ENDF formatted entry.

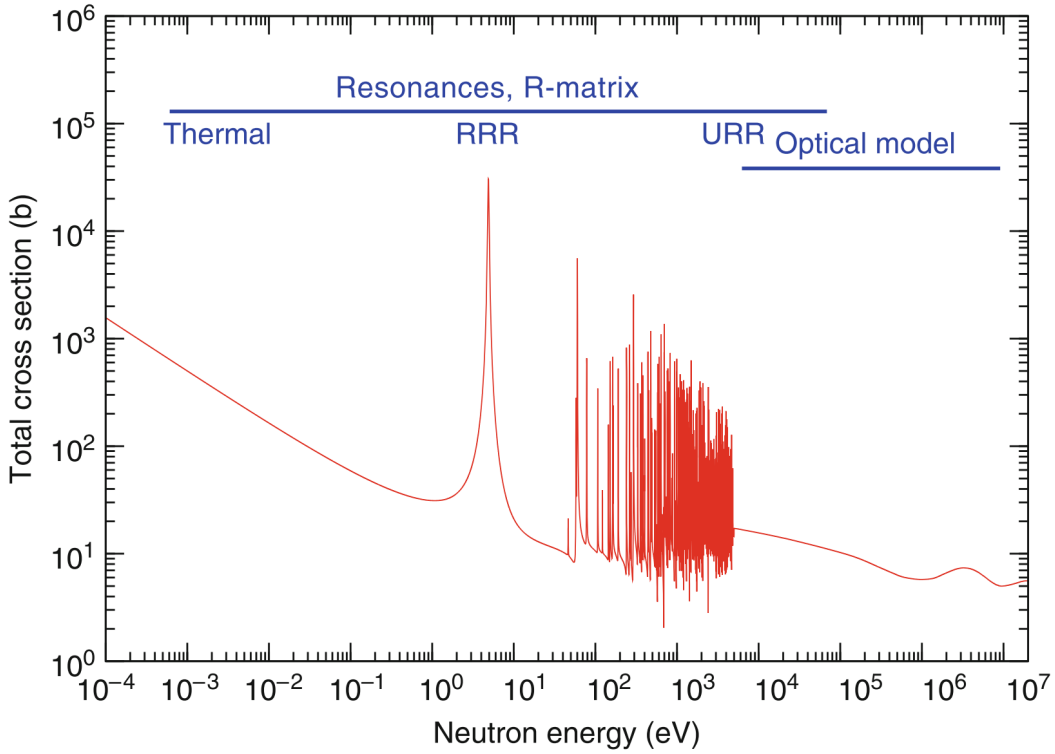

Fig. 1.3 The total cross-section as a function of energy for ^{197}Au. Note the smooth $1/v$ region where the probability of interaction decreases with energy (interactions are more likely to happen with greater time). Subsequently, resonances are the dominant phenomenon from eV–keV energies. As the level density increases past the energy resolution achievable in experiments, the RRR becomes the URR. Past this, a continuum region is obtained and methods other than resonance theory must be used to predict behaviour. Figure from [18].

1.6.2 Probability tables

As stated above, in the URR individual resonances cannot be discerned. However, by applying the appropriate nuclear reaction theory, it is possible to determine the probable values for various resonance parameters in the URR region. These include the probability distribution for the spacings (Wigner distribution) and distributions for the partial widths (χ^2 distributions for various degrees of freedom) [86]. An ENDF type evaluation may give these at various energies through the URR and different nuclear spins. This information can be used to generate Probability Tables (PTs). This concept was developed in the 1970s by Levitt, Nikolaev et al. Here the PDF as a function of energy for the total cross-section and conditional probabilities for individual reaction channels can be given. Several codes are capable of producing these PTs, including NJOY [85] and CALENDF [132]. These codes use a Monte Carlo approach for generating the PTs, first sampling a starting resonance, then sampling partial widths for that resonance, before sampling for the next resonance. This is continued to build a 'resonance ladder'. A series of these ladders are generated, with greater numbers converging to physical results [16].

1.6.3 Cross-sections

As explored above, it is beneficial to encode the cross-sections of resonant reactions in a parameterised way. Compared with a simple list of (σ, E) tuples it is more insightful and permits the reconstruction of cross-sections from resonance data for different material temperatures. However, past a certain point, resonances overlap because the aforementioned average resonance widths become much greater than the average level spacing, i.e. $\langle\Gamma\rangle \gg \langle D\rangle$ and so the combined, overlapping resonances form a continuum (see figure 1.3). Any scheme based on describing individual resonances is no longer practical.

For this high energy range then, where cross-sections change relatively slowly as a function of energy, simple tabulated data is used. This data can be obtained experimentally and with the aid of theoretical frameworks such as the 'optical' model and Hauser Feshbach theory [54]. This cross-section information is recorded in MF=3 with a single MT for each open reaction channel. For instance, MF=3, MT=102 is ^{i}X(n,γ)$^{i+1}$X for nuclide ^{i}X. Where cross-sections reconstructed from resonance parameters do not match experimentally measured totals, corrections are also recorded here.

1.6.4 Differential distributions

Both energy and angular distributions are important for the accurate simulation of radiation transport through matter. For simple interactions such as elastic scattering, the scattered angle and energy of a particular interaction can be calculated kinematically if the particle masses and incident energy are known. Repeating this process for a population of particles at low energies, one can deduce that elastic scattering is isotropic. However, nuclear reactions are not necessarily isotropic in the distribution of their products. Especially for direct nuclear reactions, where the energy is high, emitted particles may be forward biased. This is because at high energies, the wavelength of an incident particle is similar to the size of the scattering nucleus and so diffraction occurs. As with a light wave diffracting around a sphere, there is a peak directly behind the sphere. Forward biasing is of particular importance in fusion systems given the 14.1 MeV 'birth' energy of D-T fusion neutrons. The angular distributions of nuclear reaction products can be well represented by Legendre polynomials. For higher energies, higher order Legendre polynomials are required. The Rodrigues formula for generating Legendre polynomials is shown below as equation 1.10. An example angular distribution for elastic scattering off ^{186}W is shown as figure 1.4.

$$P_n(x) = \frac{1}{2^n n!}\frac{\mathrm{d}^n}{\mathrm{d}x^n}(x^2-1)^2 \tag{1.10}$$

Energy and angular distributions are found in MF=4,5,6 depending on the age and type of evaluation.

1.6.5 Decay

Unstable nuclei decay, that is, change by the emission of particles or by internal conversion of one nucleon to another. Radio-nuclides may be primordial, that is, manufactured in stars and supernovae or activated/transmuted in the present, into radio-nuclides from a stable parent nuclide. The process of attaining stability through decay is also described by ND. Half-lives, γ-photon, and other particle emission probabilities, and energies are recorded along with fission product yield data in MF=8.

1.6.6 Covariances

Without information on uncertainty, data may be worse than just wrong, it may be falsely reassuring. The practical utility of some ND is in part determined by its uncertainty, whether its quoted value is likely to be near the true value. For uncertainties to be quoted with engineering results, uncertainty must first be quantified within the inputs (ND). Then, ND

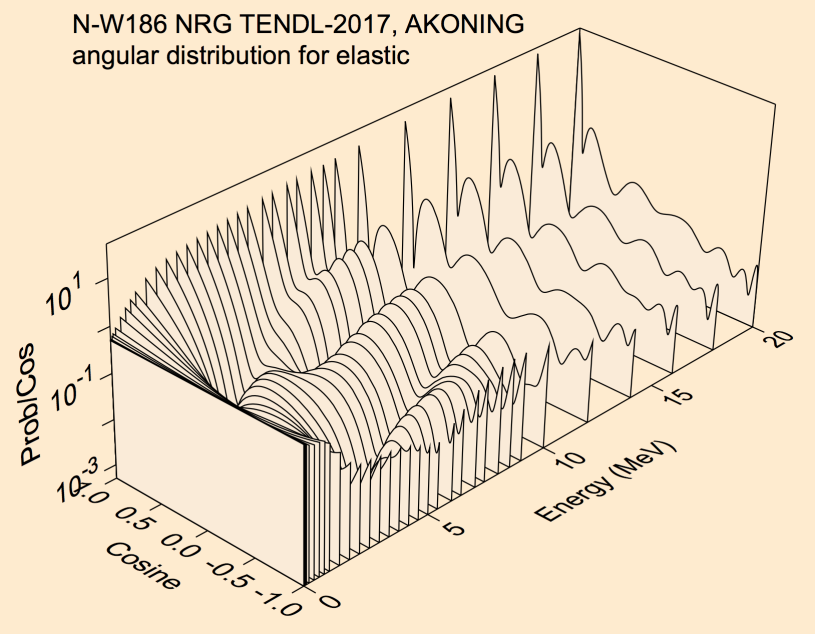

Fig. 1.4 The angular distribution of elastically scattered neutrons off ^{186}W. The probability is normalised by the cosine of the scattered angle. Note the isotropic nature at low energies, moving to a more anisotropic behaviour for greater energies. This plot was produced with NJOY [85] and released as part of the TENDL2017 library [115].

uncertainties can be propagated by various methods through to quantities of interest such as powers, neutron multiplication factors or material temperatures. The uncertainty data is a description of the experimental variance on certain data points, say, a cross-section for a given reaction channel and energy, but also with covariance information. These covariance data describe how correlated pairs of mean value data are. These correlations often arise from systematic measurement errors in nuclear physics experiments, where data for a whole energy range might be biased similarly in some way. The files containing variance and covariance information are colloquially known simply (and a little confusingly) as ‘covariances’ or ‘covariance files’. Until recently, uncertainty information was only released in nuclear data library documentation or accompanying references. The modern ENDF-6 standard groups variance and covariance information together with the mean value data [29].

1.7 Sources of uncertainty in fusion neutronics

The sources of uncertainty in fusion neutronics are many and varied. Currently envisioned fusion reactors such as ITER rank as some of the most complex machines ever. Their multitude of components, diversity of materials and range of temporal and spatial scales make

simulating their operation challenging. This complexity requires immense computational resources to model with a high fidelity. It is not currently possible to model any engineering system *ab initio*, solving the Schrödinger equation for each nucleon. Instead, approximations at each distance and time scale are made and fidelity traded off against simulation time. Examples of various modelling approximations commonly employed in fusion neutronics analyses follow below. These practices enable the solution of otherwise insoluble problems.

- Omission of material impurities. These are important for calculating expected activation of materials.
- Omission of small components. If mass is not conserved this could be deleterious for shielding analyses.
- Omission of fluid flow. Important for modelling γ dose from ^{16}N generated in water cooling systems.
- Spatial homogenisation of complex geometries. This very common practice combines adjacent volumes of differing materials to a mixture. When done correctly, mass will be conserved. However, the new homogenised materials alter the neutron spectrum and therefore the nuclear responses.
- Discretisation of the spatial domain. When particle fields and associated nuclear responses are computed using a stochastic, Monte Carlo approach, there is no functional description and instead the results are binned spatially. The bin resolution has an impact on the result.
- Discretisation of the particle energy domain. This can result in so-called self-shielding errors as discussed in a later chapter, the coarser the energy binning the more potentially erroneous the result.
- Assuming temporal invariance of plasma neutron source. This neglects the variability of plasma neutron emission from turbulence and other plasma physics effects.
- Assuming temporal invariance of material compositions. While some studies explicitly model material changes, others do not due to the additional complexity and computational burden involved.
- Assuming plasma neutron source axisymmetry. This is despite notable toroidal heterogeneity due to neutral beam on thermal plasma interactions.

- No multi-physics feedback with thermal processes. Example processes include Doppler broadening of materials' neutron cross-sections and thermalisation of neutron spectra.

Design uncertainty is also *de rigueur*, with large electricity producing plant several decades away. Neutronic analyses may take months to complete, during which time plant sub-assemblies designs' change, partially invalidating the analysis.

While the devices themselves are large and difficult to model, introducing uncertainties into calculations, other kinds of uncertainty are inherent in the practice of neutronics, regardless of the resources available. The knowledge of nuclear physics, colloquially referred to as 'nuclear data' (ND) is not precisely known for two main reasons. These are:

- Uncertainty in experimentally derived data. These uncertainties can be energy specific as with poor statistics for certain energy bins, or more general, such as 'normalisation uncertainties' like poor knowledge of sample mass or target geometry.
- Uncertainty in nuclear physics models. These are systematic errors where models depart from reality and introduce sources of uncertainty. Uncertainties on obervables stemming from inadequacies in nuclear models will be correlated.
- Uncertainty in nuclear physics model parameters. Where experimental data does not exist, theoretical models are used to estimate nuclear behaviour. These models require input parameters that can are either experimentally measured themselves–with some uncertainty–or are interpolated or extrapolated from other experimental information.

Our incomplete knowledge of nuclear physics coupled with limited computing power and complex, interrelated systems means many nuclear processes in controlled fusion cannot be simulated precisely. Instead they are approximated with the help of heuristics and simplifications as outlined above.

1.8 Implications of current uncertainties

The current engineering and physics uncertainties have a potentially large impact on realising fusion electricity.

1.8.1 Tritium breeding

To give an example, the amount of tritium consumed by a fusion power plant would be approximately 55.6kg FPY^{-1} GW^{-1} . Depending on the design, the uncertainty on the

Tritium Breeding Ratio (TBR) may be between 8% and 18% [38]. For a plant with $P_{fus} =$ 3GW this results in between $\pm$7.6 kg and $\pm$15kg uncertainty on tritium production when running at full power for a year. These are very large quantities of tritium, for comparison total world production in Heavy Water Reactors (HWRs) is currently 4 kg y^{-1} [74]. Under-producing tritium has the potential to rapidly curtail the operation of future fusion plant and to make it very difficult to provide a start-up inventory for subsequent devices. Overproducing is not ideal either, with regulatory agencies staunchly opposed to large stockpiles of tritium within the plant, given its danger as a highly permeable, radioactive gas with a function in nuclear weapons. Therefore reducing the uncertainty on TBR values and producing the exact amount required is highly desirable.

1.8.2 Cost of fusion electricity

Currently the Cost of Electricity (CoE) from future fusion plants is expected to be be approximated by equation 1.11.

$$CoE \propto \left(\frac{rF}{A}\right)^{0.6} \frac{1}{\eta_{th}^{0.5} P_e^{0.4} \beta_N^{0.4} N^{0.3}} \tag{1.11}$$

Where r is the discount rate, F the learning factor, A the plant availability, P_e the electrical power generated, β_N the normalised β and N the multiplier of the density compared to the density limit scaling. The cost of fuel is not an important factor, rather the capital cost as influenced by certain engineering parameters, the availability and the cost of money (interest) as represented by the discount rate. See [135] for further explanation of equation 1.11. The most expensive systems to construct in fusion power plants will be, in descending order, the magnets, the buildings / civil engineering work and the blanket assemblies [39]. A large degree of uncertainty in the nuclear responses relevant for these systems will result in increased costs through unnecessarily conservative shielding and the risk of reduced component lifetimes (thus lowering availability to facilitate early replacements).

It is worth drawing a parallel between the fission and future fusion experiences of regulation. The cost of nuclear fission reactors in countries such as the USA and France has increased over time [84]. Studies have associated between 14% and 21.5% of cost increases with an increase in regulatory activity from the domestic nuclear regulatory agency [20]. It seems unlikely that nuclear regulation will decrease in stringency in the future. Given this, uncertainty in the nuclear responses such as neutron and photon dose to workers, long-lived radioactive material production, tritium generation and nuclear heating of magnets will likely increase costs further by drawing the attention of regulators.

The above discussion has provided a series of motivations for quantifying and reducing the uncertainty currently encountered in the nuclear analyses of fusion systems.

1.9 Thesis outline

This introductory chapter has described the motivations for the development of controlled nuclear fusion energy. There was a discussion of the historical development of our nuclear physics understanding and the engineering development of magnetic confinement fusion. Following this, a brief review of radiation-matter interactions and their relevance. Some examples of uncertain phenomena in fusion technology were given, along with their implications for the realisation of a fusion power plant.

The next chapter is an exploration of Total Monte Carlo (TMC) applied to TBR simulations. Starting with a review of ND forms, this piece of work identifies potential contributions to TBR uncertainty and attempts to estimate the uncertainty due to lead nuclear data. The results highlight the capability of TMC in determining the distribution of uncertainty and how distribution skewness should be considered an important factor in risk analysis.

Chapter 3 details an investigation into the practice of spatial homogenisation for radiation transport. First, a review of relevant literature is presented. Then, homogenisation within the ITER reinforced concrete bio-shield is used to determine how significantly this approximation diverges from physical behaviour. Estimates of over and under-estimates of dose to workers are given for on-load and shut-down (activated) scenarios. Finally, advice on the appropriateness of this practice is given for various shielding scenarios and further work recommended.

Chapter 4 describes the development and testing of an algorithm to optimise the discretisation of the energy domain for activation and transmutation calculations. It begins with a discussion of relevant nuclear physics, developing an understanding of nuclear resonances and their distribution in the $\{E,Z,A\}$ space of interaction energy, nucleus proton number and mass. This information is used to develop a method for the targeting of energy resolution where it is most valuable. A series of tests are presented where optimised and logarithmically spaced Multi-Group (MG) results are compared with Point-Wise (PW) Reaction Rate (RR) values. Opportunities for improvement and further testing of the algorithm are presented.

The final chapter draws conclusions on the implications of the previously discussed sources of uncertainty. Their interrelated and compounded nature is discussed along with recommended future work in this area.

Chapter 2

Total Monte Carlo propagation of nuclear data uncertainties to nuclear fusion engineering parameters

2.1 Outline

This chapter describes the two principal methods for uncertainty propagation used in a nuclear engineering context, the sensitivity / perturbation theory approach and a sampling method known as Total Monte Carlo (TMC). The methods and their potential applications are described. Then, examples of uncertain nuclear fusion interaction data are given, before the TMC method is applied to propagate uncertainties from fundamental nuclear physics parameters for lead nuclei to the Tritium Breeding Ratio (TBR) in a proposed nuclear fusion reactor design, DEMO. The results of these simulations are then described, detailing the sampled TBR uncertainty and a fitted distribution. The strength of relationships between individual nuclear parameters and both intermediate data (cross-sections) and TBR results are inferred. Finally, comments on the relevance and consequences of the skewed TBR distribution are given in the conclusion.

2.2 Introduction

This chapter analyses an aspect of the uncertainty in producing hydrogen-3, or tritium, in a fusion device. The success of this process is crucial for the feasibility of nuclear fusion electricity, and the blanket systems which will contain the tritium producing materials and infrastructure are a significant projected fraction of any plant's capital cost. Reducing the

uncertainty on any tritium breeding estimates is of great value in developing fusion as a future power source.

2.2.1 Tritium breeding

A handful of different fusion fuels were discussed in the opening chapter. The cross-section and Q-value of the D-T reaction make a burning plasma the most easily realisable. However, sustainably liberating energy from the D-T fusion reaction requires a reliable supply of tritium fuel. Given tritium is not naturally occurring, with a $t_{\frac{1}{2}} = 12.3$y, it must be artificially produced.

It is worth considering whether this could be outsourced to a third party given the complexity of 'grow-your-own'. The HWR reactors of Canada, Romania and the Republic of Korea together produce about 4 kg y^{-1} and this is set to decrease in the near future. A GW class fusion reactor will require tens of kg y^{-1}. HWR reactors will in fact struggle to provide the relatively small amounts of tritium necessary for the start-up of a device. Some have suggested starting fusion reactors with a D-D fuel mix and creating tritium in the plasma through the D(D,p)T reaction, reincorporating the tritium as it is produced [141]. However, this is estimated to cost up to \$2 billion kg^{-1} T saved and is therefore not economically feasible [74].

Future plants must be tritium self-sufficient. Therefore, neutrons will initiate tritium producing reactions in ^{6,7}Li contained within a blanket. Depending on the blanket design, the tritium will either be purged or will naturally 'out-gas' before filtration and storage.

The ratio of tritium produced to tritium consumed is known as the Tritium Breeding Ratio (TBR). An equation for the TBR is given as equation 2.1 where T_{prod} is tritium produced in the blanket, T_{cons} tritium consumed in the plasma, N_s the number density of each nuclide, σ_r the r$^{\text{th}}$ tritium producing reaction cross-section for that nuclide and $\phi(E)$ the neutron flux as a function of energy.

$$\text{TBR} = \frac{T_{prod}}{T_{cons}} = \frac{\sum_{s=1}^{S} N_s}{T_{cons}} \int_0^{\infty} \phi(E) \sum_{r=1}^{R} \sigma_r(E) dE \tag{2.1}$$

Two reactions in lithium produce the majority of tritium in the blanket and are shown below. Both naturally occurring lithium isotopes have an anomalously low nuclear binding energy per nucleon. This is why ^{6}Li can exothermically fission despite being such a light nuclide. With the reaction being exothermic, its likelihood gains with decreasing interaction energy, meaning neutrons of all energies may potentially contribute to tritium production. Conversely, the tritium producing ^{7}Li reaction is endothermic, with a threshold interaction

energy of approximately 2.47 MeV. Natural lithium is $\approx 7.5\%$ ^{6}Li with the remainder ^{7}Li. Fusion breeding blankets will likely require ^{6}Li enrichment to achieve an acceptable TBR.

$$\begin{aligned} {}^{6}\mathrm{Li} + \mathrm{n} &\rightarrow \alpha + \mathrm{T} + 4.8\mathrm{MeV} \\ {}^{7}\mathrm{Li} + \mathrm{n} &\rightarrow \alpha + \mathrm{T} + \mathrm{n} - 2.5\mathrm{MeV} \end{aligned} \tag{2.2}$$

There are other reactions for creating tritons. For instance, the tritium decay product, ^{3}He has a large thermal cross-section for (n,p) and can be transmuted back to tritium this way. ^{10}B is also present in steels and other materials in small quantities and may undergo neutron capture, producing tritium as a result.

$$\begin{aligned} {}^{3}\mathrm{He} + \mathrm{n} &\rightarrow \mathrm{p} + \mathrm{T} \\ {}^{10}\mathrm{B} + \mathrm{n} &\rightarrow 2\alpha + \mathrm{T} \end{aligned} \tag{2.3}$$

There also a variety of multi-step pathways which produce tritium via the creation of other intermediate nuclides.

Within all proposed breeding systems there are a variety of tritium loss mechanisms: absorption in materials, leakage in the tritium extraction system and radioactive decay. To accommodate these losses and still retain a TBR in excess of unity, a margin, M is employed: $\mathrm{TBR} = 1 + M$. There is expected to be a legal constraint on the maximum allowable tritium inventory at a given facility on the order of kilograms. The window of adequate tritium supply is therefore relatively narrow and TBR should be precisely known and/or adjustable to keep within this window. Unfortunately there are many sources of uncertainty within TBR calculations. These can broadly be categorised as: poor/missing nuclear data, modelling simplifications and Monte-Carlo statistical uncertainty. Nuclear data often contributes the greatest uncertainty to TBR [38]. However, the effect of these uncertainties is rarely reported alongside calculated TBR values. Methods such as sensitivity-perturbation require covariance data, which is often incomplete or unavailable for these analyses.

A sensitivity analysis of Helium Cooled Lithium Lead (HCLL) type breeder blankets for the ITER Test Blanket Modules (TBM) has identified ^{6}Li, ^{56}Fe and the Pb cross-sections as the most important for TBR uncertainty [78]. This chapter quantifies the TBR uncertainty introduced by Pb nuclear data on the HCLL DEMO blanket design by employing the Total Monte Carlo (TMC) uncertainty propagation methodology.

2.2.2 Uncertainty propagation

Those in the nuclear engineering field who use nuclear data (ND) are concerned with how 'integral' quantities such as heating rates, neutron fluences, TBRs, etc. are impacted by ND uncertainties. To estimate this, one must 'propagate' the uncertainty from data to integral quantities. The main approaches are perturbation theory / sensitivity analysis and sampling methods. Perturbatory approaches were developed first and applied with great effect to fission reactor systems. Sampling methods only became available with the increases in computing power achieved in the 21st century. The basics of the methods and the merits and demerits of each are described below.

2.2.2.1 Perturbation and sensitivity

Perturbation theory is widely applied in many branches of the sciences, and generally consists of substituting an unsoluble equation for a related, soluble one plus some perturbatory series of terms. These increasing order terms have less and less impact on the result and the series is truncated at some point.

Work on applying this approach to reactor physics problems was undertaken by Wigner on the first nuclear pile [112]. A more sophisticated framework was developed by Gandini, the Generalised Perturbation Theory (GPT) [46] and subsequently into the Equivalent Generalised Perturbation Theory (EGPT) [47]. These efforts require obtaining the forward and adjoint solutions to the Boltzmann transport equation such that one can identify the magnitude of responses to any input perturbation. This means GPT and EGPT naturally lend themselves to deterministic methods for radiation transport, where full solutions to the Boltzmann equation are obtained, rather than a subsection of its phase space sampled (as in stochastic methods). This was overcome relatively recently for fission with codes like TSUNAMI-3D by Rearden, which can calculate uncertainties in k_{eff} with deterministic or Monte Carlo methods [108]. An alternative method, not specific to fission problems, is to couple a radiation transport code such as MCNP [50] with the SUSD or SUS-3D codes [65]. MCNP is used with perturbed cross-sections to determine how sensitive, S, a response parameter, R, is to changes, $\delta\sigma_g$, in a cross-section group, σ_g, as shown in equation 2.4.

$$S = \frac{(\delta R)/R}{(\delta\sigma_g)/\sigma_g} \tag{2.4}$$

SUS-3D can be used to combine processed covariance information with these sensitivity profiles to determine the uncertainty in various response parameters. Further detail on perturbation theory applied to radiation transport systems can be found in [119].

A disadvantage of perturbation theory methods is that they embody a linear relationship between input and output uncertainty (see equation 2.4). Also, covariance matrices are based on normally distributed uncertainties, whether or not this is in reality the case. The method is very much dependent on the quality and comprehensiveness of covariance matrices available. Very large and difficult to compile covariances matrices are required to record the potential correlations between all open reaction channels.

2.2.2.2 Sampling

A variety of sampling based techniques for propagating nuclear data uncertainty are available.

Koning and Rochman developed a new, integrated method of uncertainty propagation in the late 2000s [72]. This approach, known as Total Monte Carlo (TMC), relies on repeated sampling of varied nuclear data. Whichever problem is being investigated is simulated multiple times with different data, building a distribution of outcomes. It is a fundamentally different method to the sensitivity / perturbation approach.

The foundation is a reliable methodology for generating nuclear interaction data. The software package T6 [71] collects together various tools for calculating, formatting and validating nuclear data. Using T6 and a set of input parameters, reaction likelihoods and outcomes can be simulated for a variety of incident particles: neutron, photons, protons, deuterons, tritons helions and alpha-particles and for a range of energies from 10^{-5}eV to 200MeV. The interactions in the 1 keV – 200 MeV incident energy range are computed with the TALYS component of T6. This includes the optical model, level densities, direct reactions, compound reactions, pre-equilibrium reactions and fission reactions [68].

The theoretical models require parameters which may be directly experimentally measured, or are estimated through fitting to other, robust experimental data such as total cross-sections, σ_t. The RIPL-3 nuclear parameter database [21] contains reference values and confidence estimates for many of these parameters. These and other sources have been assimilated and parsed to provide default values for TALYS, with variances.

A set of nuclear parameters, $\vec{A} = [A_1, \ldots, A_n]$, are used as input to T6 to generate a complete ND library, with full cross-section, resonance, angular distribution, double-differential distribution, covariance, etc. information. These data are self-consistent. An example of this behaviour is where a total cross-section is well understood, it is to be expected that a reduction in one reaction channel cross-section, say elastic scattering, would correspond with an increase in another open reaction channel. This correlation can be produced with simple adjustment or perturbatory approaches but it requires sufficient covariance information, which is often lacking. Once the ND have been generated they are processed and formatted

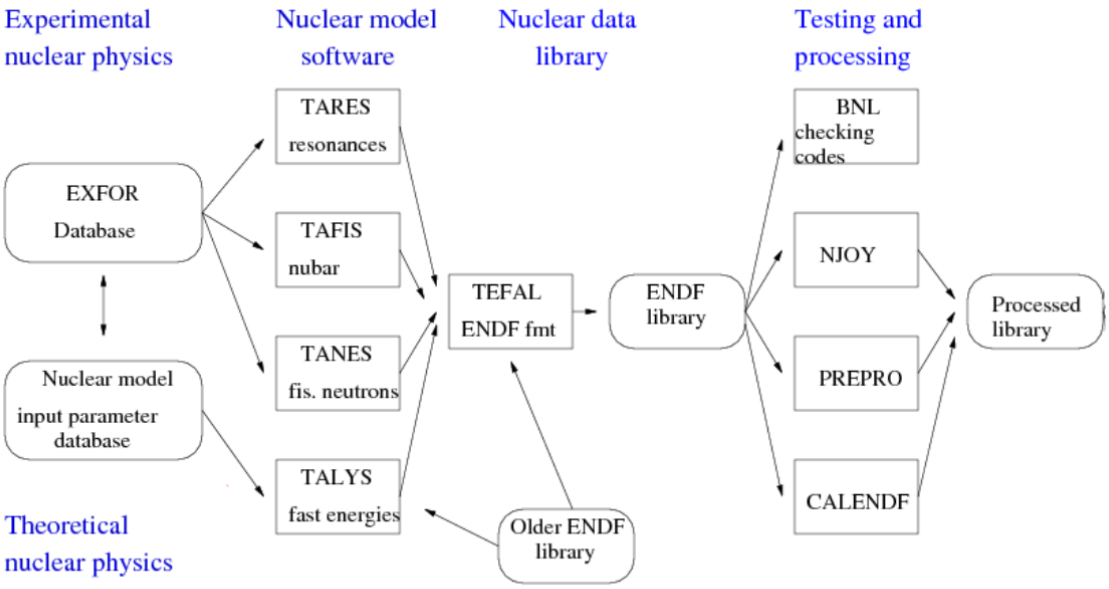

Fig. 2.1 Nuclear data generation flowchart for the T6 software package. Figure from [69].

to the ENDF standard with the TEFAL component of the T6 package, permitting use by all ENDF compliant codes [73]. The ND generation process is shown as figure 2.1.

Using the above nuclear physics parameter database and models, it is possible to repeat the nuclear data generation process to create a 'library of libraries', implicitly containing amongst their variance the uncertainty in the underlying nuclear physics parameters, $\vec{A}$. This data is now used as input to whichever nuclear system is to be simulated. Numerous simulations are launched, each picking a new set of input data, an evaluation from the library of libraries. Whichever reaction rates or spectral quantities can be tallied for inspection and subsequent analysis. Call the quantity of interest q. As the number of simulations, n, increases, any histogram of the values of $\vec{q} = [q_1, \ldots, q_n]$ will converge to a probability distribution function of the most likely value and its associated uncertainty, $\sigma_{observed}$.

In an ideal case, this observed uncertainty, $\sigma^2_{observed}$, would all be due to the variations in nuclear data, σ^2_A, as in equation 2.5.

$$\sigma^2_{observed} = \sigma^2_A \tag{2.5}$$

Indeed this is the case for deterministic methods. With Monte Carlo methods however, answers are inherently uncertain. Rather than solving an equation to find the exact flux and associated reaction rates at every point, the system behaviour is sampled by virtual particles

(see section 1.4 in chapter 1). Monte Carlo methods depend on the Law of Large Numbers (LLN), using a large number of simulated particles to converge on some mean behaviour. Far fewer particles populate the simulation than the real system. Short of an infinite number of 'source' particles, there is always an associated uncertainty.

Therefore, in this work the observed uncertainty, $\sigma_{observed,i}$ from a simulation i is composed of two components: uncertainty from the simulation method, the standard deviation $\sigma_{stat.,i}$ and uncertainty from the nuclear data, the standard deviation being $\sigma_{A,i}$. So long as these are independent, their variances sum to give the total variance, as in equation 2.6.

$$\sigma^2_{observed,i} = \sigma^2_{A,i} + \sigma^2_{stat.,i} \tag{2.6}$$

Assembling a histogram of the simulated values, $\vec{q}$, approximates the underlying PDF of the quantity, q of interest. We can modify equation 2.6 from the single simulation case to extract information on the nuclear data uncertainty contained within the PDF of all simulations. The average statistical standard deviation, $\overline{\sigma}^2_{stat.}$, shown in equation 2.7 is simply the mean value from all n simulations. Substituting this for $\sigma_{stat.,i}$ we have equation 2.8.

$$\overline{\sigma}^2_{stat.} = \frac{1}{n}\sum_{i=1}^{n} \sigma^2_{stat.,i} \tag{2.7}$$

$$\sigma^2_{observed} \approx \sigma^2_{A} + \overline{\sigma}^2_{stat.} \tag{2.8}$$

If the variance due to nuclear data, σ^2_A, is to be known, one must either enforce $\overline{\sigma}^2_{stat.} = 0$, or otherwise find the difference between the observed and 'statistical' simulation variance.

In terms of determining this statistical simulation uncertainty, Monte Carlo codes provide an estimate of $\sigma_{stat.,i}$. Comparing this uncertainty amongst differing values of q can be achieved by normalising by q. The Relative Standard Deviation (RSD)[1], R, is given below as equation 2.9 and its functional dependence on particle count as equation 2.10.

$$R = \frac{\sigma_{stat.}}{q} \tag{2.9}$$

$$R \propto \frac{1}{\sqrt{m}} \tag{2.10}$$

Where q is the reported quantity mean value and m is the simulation particle population count. Keeping R to an arbitrarily small value, say <0.005, allows us to ignore the $\overline{\sigma}^2_{stat.}$ term in equation 2.8 [118].

[1] Also known as the Fractional Standard Deviation (FSD)

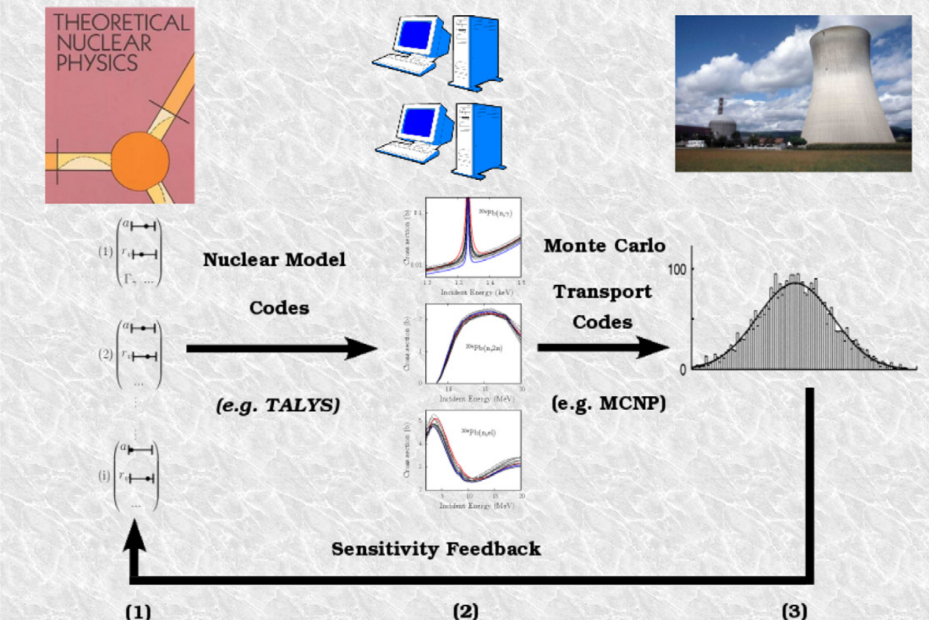

Fig. 2.2 Schematic overview of the TMC process. 1) Uncertainties in fundamental nuclear parameters estimated. 2) Many sets of ND generated with T6 software, sampling from fundamental parameters. 3) Nuclear system is simulated many times with generated ND, value of observable is added to PDF, eventually converging on an observable with a mean value and characteristic distribution. Figure from [72].

If only one nuclear parameter, A_i, is varied then the nuclear data uncertainty, σ_A is wholly attributable to that parameter. Similarly, if only one nuclide has its nuclear parameters and hence interaction data varied, σ_A represents uncertainty on q from that nuclide alone. When more parameters are varied for more nuclides, the uncertainty is an ensemble of their variation.

The general TMC process is diagrammatically outlined in figure 2.2. Compared to other methods, this approach has a variety of benefits. One advantage of this system is that one relates uncertainty in the earliest possible parameters, say the real volume potential or imaginary surface potential of the optical model to the integral quantities of most interest to end-users such as TBR or nuclear heating in a fusion context. This allows future targeting of nuclear physics research–where are resources best allocated to reduce system uncertainties? This is the idea represented by the sensitivity feedback loop in figure 2.2. Of course, this methodology does not have to only vary nuclear physics parameters, one can choose to vary parameters in the simulation model, or even methodology, to determine their contribution to uncertainty. Engineering parameters, such as component dimensions, operating temperatures, nuclide atomic densities, etc. can be varied to see how uncertainty on them effects quantities of interest.

Besides TMC, there do exist other sampling based methods. Typically these sample from the uncertainty represented in covariance matrices. For instance, SANDY [42] and NUSS [142]. SANDY is a program written in Python which can read ENDF-6 formatted nuclear data files, extract variance-covariance information as random variables and then use these distributions to create a selection of new nuclear data files which embody the underlying uncertainty. NUSS performs a similar function. These methods can of course use whichever covariances are available to generate their perturbed files. As with TMC they rely on repeating the radiation transport part of any calculation multiple times, and so are less appropriate for the most computationally intensive problems.

2.2.2.3 Comparison of methods

Sections 2.2.2.1 and 2.2.2.2 have discussed the methods of two broad uncertainty propagation schemes. It is worth noting the advantages and disadvantages of each, drawing some comparisons between them to elucidate where each method might be most appropriate.

The perturbation method, using covariance information and sensitivity profiles, has much experience associated with it. It has been used to estimate uncertainties for k_{eff} and other important parameters in fission reactor physics for several decades. By contrast, whilst the idea of repeatedly simulating a system with different inputs is not new, the integrated TMC methodology, with a sophisticated ND generation package, is novel. Published in 2008, it has seen substantially less use in industry than perturbatory approaches. The nuclear industry is often slow to adopt new technology as there can be a significant volume of work required to validate results, especially where their application is safety-critical. Studies in a variety of applications have sought to benchmark and demonstrate TMC, bringing it into wider use [72][117][127][73][5][6][116].

Assuming appropriate nuclear data is already available, the slowest component of either method is radiation transport by some Monte Carlo code. Combining covariance information with sensitivity profiles, or parsing results from multiple TMC simulations to assemble a PDF of some target quantity, q, is of minor computational burden. If the time taken for a single simulation is T time units, then TMC will take $n \times T$ where convergence on some q and its $\overline{\sigma}_{observed}$ typically takes $n \geq 500$ simulations to sufficiently sample input uncertainties [118]. This factor of several hundred slowdown is clearly a burden for many analyses, especially those requiring many billions of source particles such as full-core fission reactor simulations, shielding analyses or similar.

In the past 3 or 4 years, progress has been made in reducing the TMC time penalty. If a single calculation requires m particles for convergence, so-called 'Fast TMC' uses m/n particles in each of the n simulations. This dedicates fewer particle histories to each change

in nuclear data, exploring the A parameter-space more quickly, but with a larger $\sigma_{stat.}$ for each simulation. Rochman et al. indicate that ensuring the inequality $\frac{\sigma_{stat.}}{\sigma_{observed}} \leq 0.5$ holds for each simulation, equation 2.8 is still valid and the scheme still computes approximately the same uncertainties as n runs each with m particles. This is in contrast to traditional TMC where the rule of thumb is $\frac{\sigma_{stat.}}{\sigma_{observed}} \leq 0.05$.

Perturbatory methods require covariance matrices, but often only covariance information for certain phenomena are readily available. Along with cross-sections and resonance parameters there are angular distributions of scattered and emitted particles, double-differential distributions, $\overline{\nu}$ (average neutron yield from a fission event), etc. Evaluations vary in their completeness. As TMC calculates and varies a full complement of nuclear data with its set of nuclear models, it can claim to have a more comprehensive estimate of the true uncertainty. While perturbations to generate sensitivity profiles are for a single cross-section group value (for example), TMC varies parameters at an earlier step, as a fundamental parameter for the behaviour of the nucleus. Because of this, when a parameter is varied in the TMC approach, it in turns varies all correlated phenomena. Other areas where the TMC scheme embodies a more complex approach are the assumed linearity of the perturbation method, along with the inability of perturbation methods to accommodate non-Gaussian observable distributions. Despite these advantages, the TMC method has attracted certain criticisms. It can be argued that the all-important nuclear parameter distributions have not been rigorously estimated from experimental data [55]. Both Capote et al. and Rising note the arbitrariness of some parameter distributions. Recently, work has been undertaken to implement Bayes' theorem through the assignment of weights to generated ND [67].

The ease of application of the two methods is varied. While TMC is conceptually easy to understand and incorporates many advantageous features, there are limitations in its application currently, such the failure of TALYS to correctly predict the behaviour of light nuclides (the cut-off being A=12) [113]. For perturbation methods, the paucity of covariance information is a problem for some calculations (principally those outside the thermal fission realm).

2.2.3 Uncertainty in fusion relevant data

There are many nuclides of importance in nuclear fusion studies whose behaviour are not known to sufficient accuracy [45]. Nuclear technology development programmes for ITER [11] and DEMO [1] have both identified the need for more widespread uncertainty estimates and more accurate nuclear data evaluations to reduce uncertainties. Batistoni et al. cites conservative uncertainties for neutron fluence at the ITER pressure vessel ($\pm$15%) and superconducting Toroidal Field (TF) coil magnets ($\pm$30%) due to nuclear data uncertainties.

An improved uncertainty treatment has been advocated for tritium breeding since the early days of detailed reactor design [2]. Youssef and Abdou performed a sensitivity / perturbation analysis on uncertainties in TBR due to ND for early blanket designs finding that ND contributed an uncertainty on TBR of between 2–6% of the mean value depending on the specific design.

A variety of different tritium breeding schemes are proposed for reactors. These have either solid (ceramic pebbles) or liquid (molten metal) breeding compounds which may or may not be accompanied by an additional neutron multiplying material. The solid systems have some sort of additional coolant loop, typically employing He or H_2O. For the liquid metal systems, the breeder also functions as the primary coolant. All systems are likely to have a secondary coolant loop for subsequent electricity generation. Youssef and Abdou compared the TBR uncertainty due to ND for several of these systems. He found that Li_2O ceramic systems typically have a large uncertainty contribution from ^{16}O, followed by ^{56}Fe and the lithium isotopes (their relative importance depending on the degree of ^{6}Li enrichment). Uncertainty in a $LiAlO_2$ system on the other hand was dominated by uncertainty in ^{9}Be multiplicity cross-section data. As such, there were significant differences between similar breeding concepts, which is a frustration when trying to generalise and also underscores the importance of this kind of work. For a Li-Pb, liquid metal system, the TBR uncertainty was dominated by variance in the lead multiplicity cross-sections, the (n,2n) and to a lesser extent (n,3n) reactions.

As well as pure computation, work is available comparing simulation and experimentally determined values for important reaction reaction channels in ceramic (HCPB) tritium breeding systems [10]. Less work has been undertaken to experimentally corroborate simulated nuclear responses in liquid metal blankets. This is not for lack of desire, these lithium-lead blankets are currently in development in tandem with the solid breeders and account for half of the Test Blanket Modules (TBM) to be tested in ITER [25]. There are good reasons to expect lithium-lead blankets to be yet further developed, reasons including: the greater resource availability (beryllium for ceramic blankets is in limited supply [14] [125]), the facility for on-load lithium enrichment and thus TBR adjustment [59], higher natural breeding ability [27] and the comparative ease of extracting bred tritium from a liquid metal vs. porous pebbles. A more thorough discussion on the advantages and disadvantages of the different technologies can be found in [1]. Given the above, it is worth investigating the impact of ND uncertainties on LiPb type blankets.

As briefly mentioned previously, Leichtle et al. has undertaken a recent study where nuclear data was adjusted to obtain the sensitivity of TBR in a modern HCLL system (the ITER HCLL TBM) to various nuclear reaction cross-sections [78]. This work identified ^{6}Li,

^{56}Fe and the Pb cross-sections as the most important for TBR uncertainty from the a ND perspective. The analysis below looks at variation in lead nuclear data.

2.3 Method

This chapter is investigating the propagation of lead nuclear data uncertainties to a corresponding PDF for TBR uncertainty in the DEMO reactor. The following section describes nuclear data selection and sampling, along with radiation transport and data analysis methods for the study.

2.3.1 Nuclear data

As discussed in section 2.2.2.2, traditionally cross-section uncertainties are represented as single values for a given energy and reaction channel. There might be an estimate of uncertainty, which is typically given as a 1 σ, the standard deviation of a Gaussian distribution centred about the quoted mean value. As an example, the ^{208}Pb(n,2n)^{207}Pb neutron multiplication reaction at 14 MeV is shown in table 2.1 for a variety of experiments and nuclear data libraries. Experimental results from the EXFOR library are also shown. There are a wide spread of experimental results, for example Arakita and Shunk are relatively low at 1.04 and 1.31 b respectively and with quoted uncertainties of 5.77 and 8.0% respectively. These values are among the older data experimental data and have been relegated in importance by nuclear evaluators. Modern evaluated data releases tend to agree on a figure of $\sigma_{n,2n}(14\text{ MeV}) \approx 2.1$ barns. It would appear that the experiments of Simakov and Frehaut along with model calculations have been favoured of to produce the data contained within contemporary evaluated nuclear data libraries. Data shown in table 2.1 are also shown graphically in figure 2.3 along with the 300 TENDL files used for this work.

Source	Energy [MeV]	$\sigma_{n,2n}$ [b]	$\pm\%\Delta\sigma_{n,2n}$
Experimental data			
Ngoc	14.6	1.27	7.61
Simakov	14.1	2.38	5.88
Anders	14.6	1.36	4.98
Arakita	14.2	1.04	5.77
Frehaut	13.8	1.96	8.11
Frehaut	14.3	1.94	8.36
Frehaut	14.8	1.97	8.52
Salaita	14.8	1.31	8.85
Shunk	14.0	1.31	8.00
Prasad	14.8	0.99	12.12
Garg	14.7	1.09	6.78
Glagolev	14.7	1.70	17.65
Amemiya	14.8	1.22	8.20
Evaluated data			
BROND 3.1	14.0	2.09	-
CENDL 3.1	14.0	2.08	-
ENDF/B-VII.1	14.0	2.22	8.15
JEFF 3.2	14.0	2.18	7.0
JENDL 4.0	14.0	2.15	10.1
TENDL 2015	14.0	2.19	7.4
RUSFOND 2010	14.0	2.18	-

Table 2.1 Cross-section values for the n,2n reaction channel, $\sigma_{n,2n}$ on ^{208}Pb, around 14 MeV for various experiments and libraries. The experimental results were retrieved from the online EXFOR database [58]. The ENDF utility code Inter [36] was used to extract the values for each library. The uncertainties, $\Delta\sigma_{n,2n}$ are presented as percentages above and below the reported value.

This study used the TENDL-2015 library. It is a comprehensive, general-purpose nuclear data library which contains data on interactions between 7 projectiles and over 2,800 target nuclides. As described above, the library is produced by a suite of codes known as T6 and an adherence to a strict methodology of reproducability [115]. It uses the TALYS nuclear reaction code to model high-energy reactions, with TARES handling the lower energy, resonant region. The inputs are fundamental parameters, including data from the RIPL-3 database [21], each with their own probability distributions that reflect their uncertainties. The TENDL-2015 release contains 'average' or best-guess evaluations, but also a collection of 'random' files, those generated using the sampling method. These have been pre-generated

by the TENDL team. Comprehensive information concerning the input parameters and plots of the ND are available online [114].

It is instructive to inspect the nuclear data used as input for this TMC simulation. The sampled nuclear data files were downloaded in both ACE (293K) and ENDF format from the TENDL-2015 website [114]. Figure 2.3 shows the (n,2n) cross-sections as a function of energy for ^{208}Pb. Cross-sections from a given file at a certain energy are shown as grey circles. The mean value of the random files and a standard deviation is also shown, along with available EXFOR data. The shape of this endothermic multiplication reaction is visible, with an energy threshold of 7.36 MeV. The cross-section rises to approximately 2 b in the 14 MeV fusion neutron energy range before falling away at higher energies–here the n,3n reaction becomes more probable. No experimental data is available from the EXFOR database beyond 14.8 MeV. This lack of higher energy information means constraining model parameters for T6 data generation is more difficult and so the uncertainties here are larger.

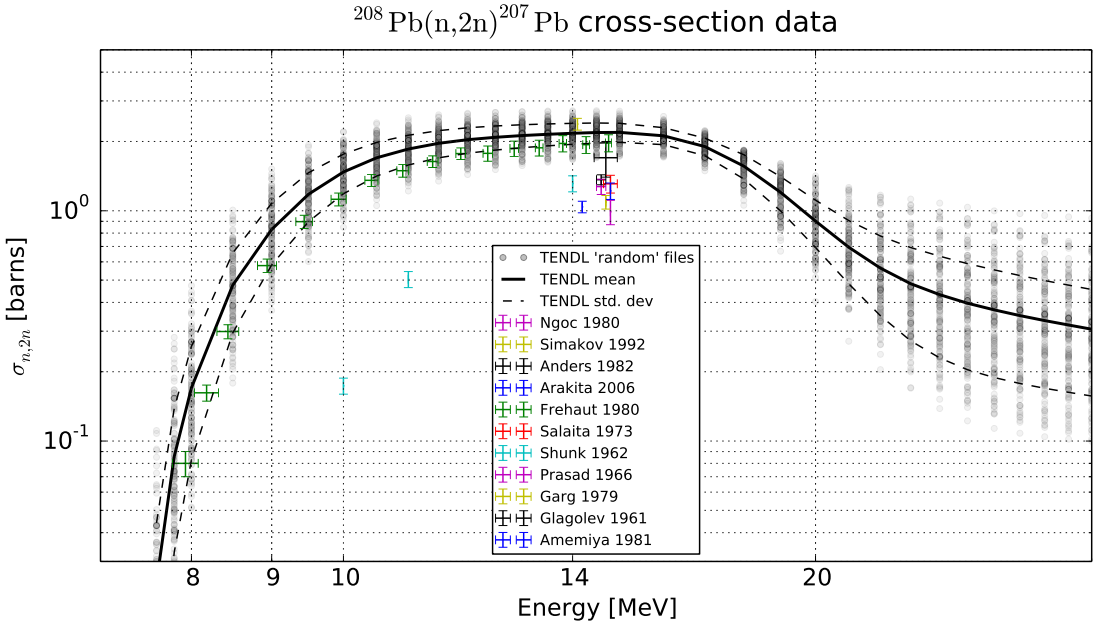

Fig. 2.3 Shown here are ^{208}Pb(n,2n)^{207}Pb cross-sections as a function of energy. The plot shows the TENDL 'random' data extracted from processed ACE files. Additionally, the few available experimental results in EXFOR are plotted for comparison. The relative paucity of experimental data is one reason for the relatively large uncertainty on this particular reaction channel.

In an unmoderated fusion neutron spectrum the dominant reaction channels for Pb are neutron multiplication (n,2n) and elastic scattering, at approximately 2 and 3 barns respectively. The ENDF utility code Inter was used to extract cross-section values at a given

energy from the ENDF files. The resulting distributions and correlations were plotted as figures 2.4, 2.5 & 2.6. Figures 2.4 and 2.5 were examined for their first 3 statistical moments.

Nuclide	Mean, μ [b]	RSD, $\frac{\sigma}{\mu}$	Skewness
Pb206	2.11	8.1%	-0.065
Pb207	2.12	10.7%	-0.868
Pb208	2.17	9.7%	-0.120

Table 2.2 TENDL2015 'random' Pb n,2n neutron multiplicity channel cross-section distribution statistics, also shown in figure 2.4.

Figure 2.4 and table 2.2 show the distributions of TENDL-2015 cross-sections for (n,2n) at 14 MeV. Lower values for $\sigma_{n,2n}$ will reduce the neutron multiplication of the blanket and, other quantities being equal, result in a reduced total neutron flux and therefore a lower TBR. The first two moments of each distribution in figure 2.4 were analysed for convergence testing. A normalised measure of the standard deviation, the relative standard deviation, $R(n) = \frac{\sigma(n)}{\mu(n)}$, where $\sigma(n) = \sqrt{\frac{\sum_{i=0}^{n}(x_i-\mu_i)^2}{n-1}}$ and $\mu(n) = \frac{\sum_{i=0}^{n} x_i}{n}$ was plotted as a function of sample size. $R(n)$ converges for $n \approx 150$ for both ^{206}Pb and ^{208}Pb at 8.1% and 9.7% respectively. For ^{207}Pb, $R(300) = 10.7\%$, but this figure is not completely converged, with some step changes in R due to extreme $\sigma_{n,2n}$ cross-sections from the low-value tail. Ideally more data would be available for this nuclide.

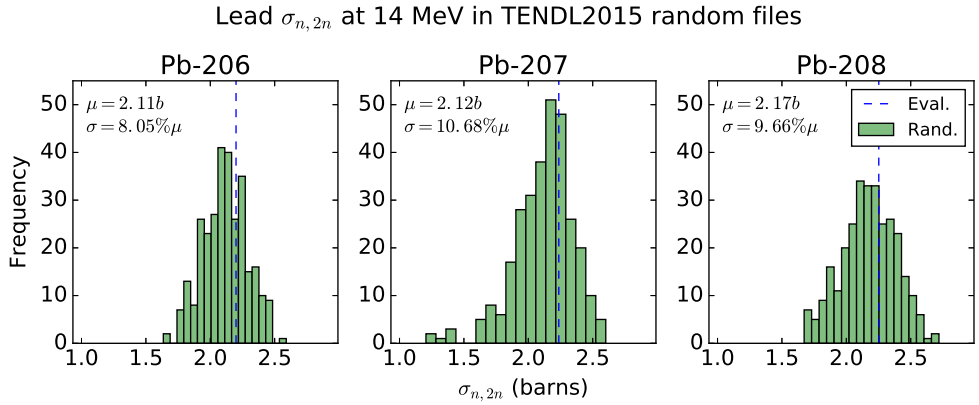

Fig. 2.4 Shown are histograms of $\sigma_{n,2n}$ at 14 MeV for the three major Pb isotopes. The histogram contains all 300 available files each Pb nuclide. It can be seen that while ^{206,208}Pb are approximately symmetrical with skewness values close to zero, ^{207}Pb has a skewness of -0.868 indicating a low-value tail. This may or may not be representative of the underlying distribution.

Nuclide	Mean, μ [b]	RSD, $\frac{\sigma}{\mu}$	Skewness
Pb206	2.84	7.5%	0.663
Pb207	2.80	8.0%	0.548
Pb208	2.81	8.1%	1.185

Table 2.3 TENDL2015 'random' Pb elastic scattering cross-section distribution statistics, also shown as figure 2.5.

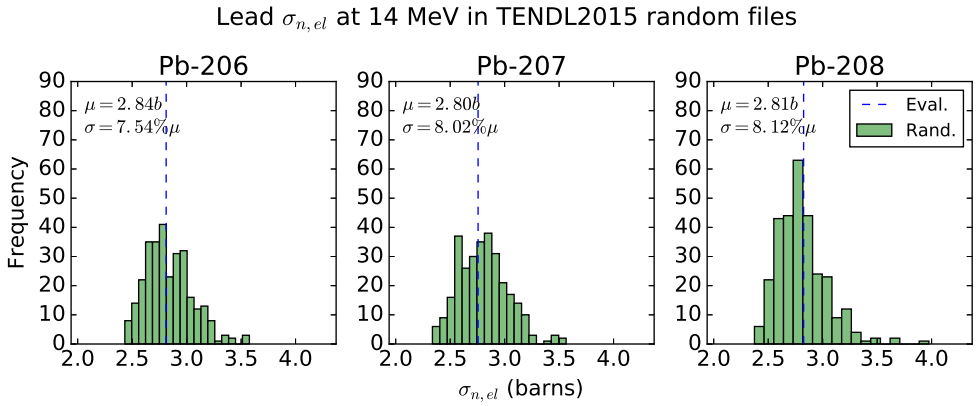

Fig. 2.5 The elastic scattering cross-sections for lead at 14MeV are all positively skewed, with ^{208}Pb having a skewness value of 1.185.

The elastic scattering data shown in figure 2.5 and table 2.3 are instead positively skewed, with high value tails for each isotope, with the most pronounced tail being ^{208}Pb. A greater elastic scattering cross-section will result in an increased likelihood of down-scattering neutrons to lower energies. The softer spectrum will increase triton production in ^{6}Li as the cross-section increases with decreasing energy, i.e. $\sigma_{n,t}(E) \propto \frac{1}{E}$. However, the elastic scattering cross-section is anti-correlated with the n,2n cross-section in the TENDL data (see figure 2.6). In other words, a high scatter cross-section is typically coupled with a low multiplicity cross-section. This inclusion of cross-channel correlation is one of the key advantages of TMC over more traditional uncertainty propagation (UP) methods.

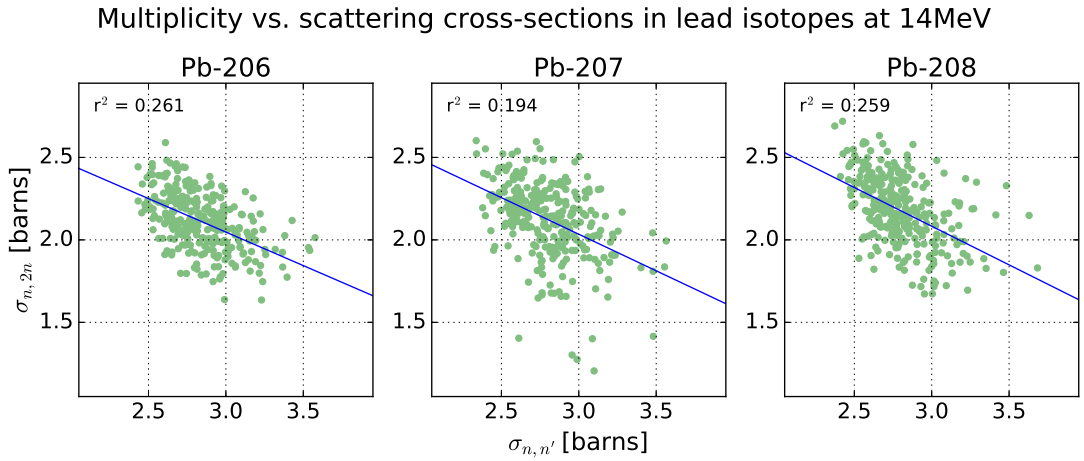

Fig. 2.6 Shown here are scatter plots of the relationship between the scattering and n,2n cross-sections in the TENDL data. The two channels are clearly anti-correlated. A least squares fit has been applied and the resulting best-fit linear relationship plotted.

2.3.2 Radiation transport

There exist several variants on the DEMO future fusion reactor concept. It is typically a several GW_{th} device, with a major radius, $R \approx 9$m. Different breeding blanket concepts are the topic of contemporary study and a radiation transport model is typically produced for each one. Of course the nuclear responses within the blanket are of particular interest, but the size and material composition of the blanket also determines the neutron spectra and flux behind them–throughout the rest of the machine.

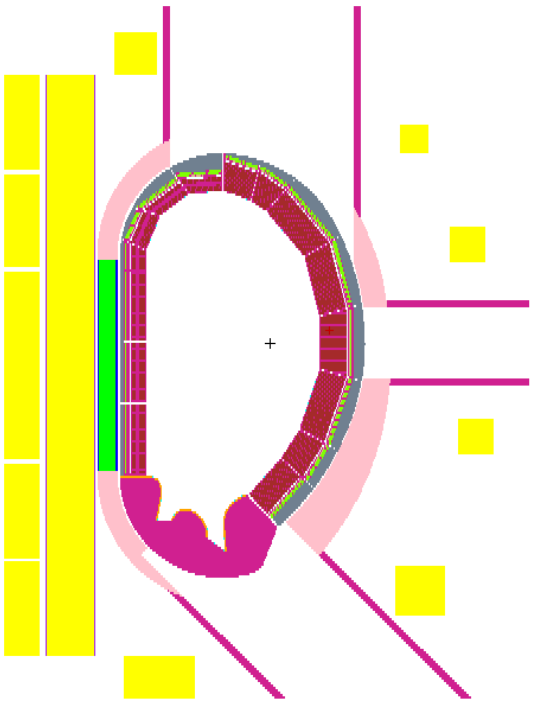

Fig. 2.7 A poloidal slice through the 2014 DEMO HCLL radiation transport model used for this study. Magnets are marked in yellow, the divertor is fuscia and the blanket in red. The blanket sectors are constructed using the MCNP 'universe' feature for repeating geometry and do not conform perfectly to the vacuum vessel shape. The three large voids above, below and to the right of the plasma chamber are access ports–extraction routes for blanket and divertor sectors in need of replacement. As noted previously, the distance from the left of this figure to the centre of the plasma chamber is approximately 9m.

An 11.25° toroidal sector of the 2014 DEMO HCLL MCNP model was modified to tally triton production per source particle (TBR). MCNP6.1 [50] was used for particle transport. The particle count for each simulation was set to 3×10^6 giving a relative standard deviation, RSD on TBR of ≈ 0.002 for each simulation. This is approximately 1/40th of the total (nuclear data & statistics) accumulated RSD and thus practically negligible. While other TMC analyses have attempted to limit to the statistical variation to only around half of the total variance and then subtract this from the observable variance, known as 'Fast TMC', the original TMC technique of minimising the statistical variance is used here.

Whilst the Pb nuclear data for sampling was from TENDL-2015, FENDL3.1b data was used for the neutron transport of other elements present in the reactor model. A series of simple Bash and Python scripts were used to sample different ^{206}Pb, ^{207}Pb & ^{208}Pb TENDL-2015 ACE files for the radiation transport runs, before creating and submitting jobs en masse

to the local cluster at the Culham Centre for Fusion Energy. Records of the input ND and tally data were kept, identified by a unique string.

The TMC process was run for a total of 1559 MCNP simulations, each 84 core-minutes across 32 cores. The total wall-clock time was 3 days, 15 hours.

2.4 Results & discussion

Python scripts were used to plot TBR convergence as a function of simulation count (see figure 2.8) and the distributions of simulated TBR (see figure 2.9).

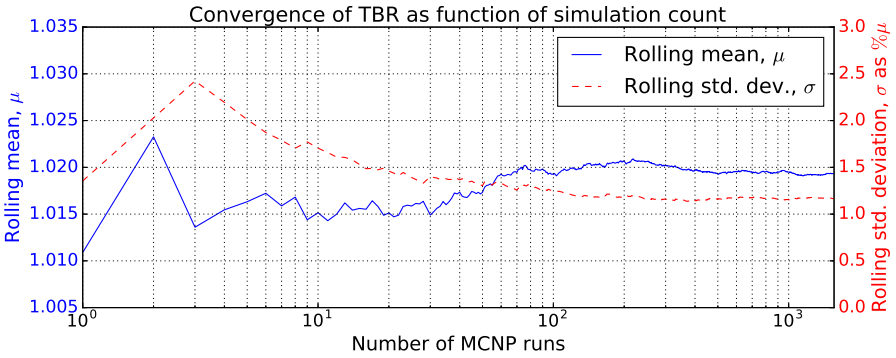

Fig. 2.8 The simple mean and standard deviation of the TBR results is presented as a function of the number of simulations.

The TMC simulation can be seen to converge in figure 2.8 at approximately 400 MCNP runs. $\mu(400) \approx 1.02$ and $\sigma(400) \approx 1.14$. The mean TBR value is 1.0193 whilst the median is 1.0200, with a one sigma standard deviation of 0.012 or 1.164% of the mean value.

The TBR distribution is shown as figure 2.9. Any results greater than 6σ from the mean TBR were deemed false and discarded before analysis. This filter rejected one simulation result, where the sampled TBR value was exceptionally far from the mean value: $0.826 \approx \mu - 15\sigma$. This simulation used an erroneous ^{207}Pb ACE file, created with unrealistic parameters.

The shape of the distribution appears somewhat Gaussian (see equation 2.11) although with an extended low-value tail. The Shapiro-Wilk test was employed to test if the TBR population have indeed been sampled from a normal distribution [122]. Formally, a null hypothesis, H_0, that the underlying distribution is normally distributed was proposed. This hypothesis was rejected as the computed p-value for the sample population was 0.003, less than the conventional threshold for acceptance (known as α) of 0.05[2].

[2] $\alpha = 0.05$ implies there is a 5% chance the null hypothesis is incorrectly rejected and is a common value for this statistical test.

The distribution is fitted well with the addition of a skewness parameter, γ to give the 'skew-normal distribution' as per equation 2.12.

$$f(x) = \frac{1}{\omega\sqrt{2\pi}} e^{-\frac{(x-\xi)^2}{2\omega^2}} \tag{2.11}$$

$$g(x) = f(x)\left[1 + \text{erf}\left(\frac{\gamma(x-\xi)}{\omega\sqrt{2}}\right)\right] \tag{2.12}$$

Equation 2.12 is the fit function plotted in figure 2.9. ξ is the location parameter, ω the shape and γ the skewness. If the data were symmetric, $\gamma = 0$. For these data, $\gamma = -1.432$, indicating the low-value TBR tail is longer than the high-value tail.

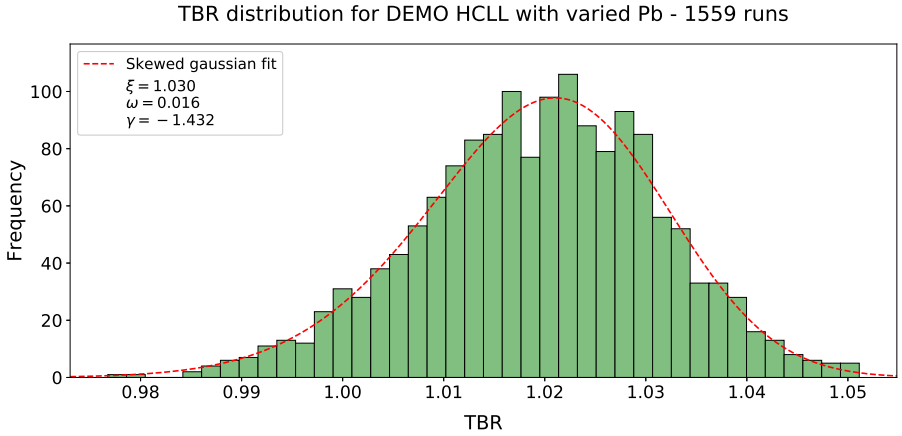

Fig. 2.9 Histogram of 1559 TBR values computed with the TMC methodology. The fit is a skew-normal distribution as described by equations 2.11 and 2.12. The standard deviation is 1.2% of the mean value and represents the variation from TENDL2015 lead nuclear data (specifically ^{206,207,208}Pb). Note that 5.8% of the distribution is less than unity.

Figure 2.10 shows the correlation between the n,2n cross-sections used in a given simulation and the TBR attained in that simulation. The nuclide-wise cross-sections were obtained with the Inter utility for a reaction energy of 14 MeV. As the cross-sections for each of the three lead isotopes were varied in every simulation, so the cross-sections plotted here are the elemental average values, that is, weighted by the natural abundance of each Pb isotope. A linear least squares fit has been applied to the data shown in figure 2.10 and the resulting relationship plotted in blue. As the cross-section for neutron multiplying reactions of the form ^{A}Pb(n,2n)$^{A-1}$Pb increases, so does the neutron flux available for the subsequent generation of tritium in the breeding blanket. As these reactions are endothermic, with an energy threshold relatively close to that of the D-T neutron creation energy, front loading a breeding blanket with neutron multiplying material is beneficial to TBR.

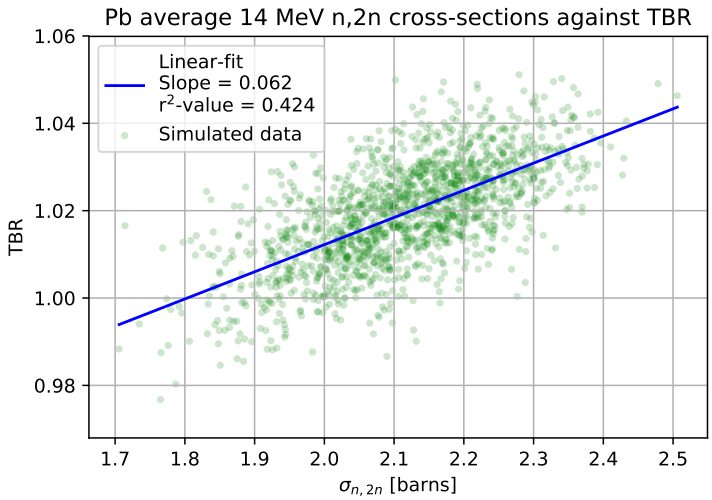

Fig. 2.10 Shown here are the ^{208}Pb(n,2n)^{207}Pb cross-sections at 14 MeV, from the files used as input ND for the TMC simulation. They are plotted against the resulting TBR value from each simulation. There is a reasonably strong positive correlation, with increasing n,2n cross-sections resulting in an increased TBR.

The relationship between n,2n cross-section and TBR shown in figure 2.10 is relatively strong, with a linear fit achieving an $r^2 = 0.42$, explaining 42% of the total variance in TBR. Elastic scattering (shown in figure 2.11) is more weakly correlated, with a very low $r^2 = 0.02$. Figures 2.10 and 2.11 plot relationships between the cross-section values and the simulated TBR result, but these cross-sections are only intermediate data in the TMC process. However it is possible to go one step further back and analyse for correlations between fundamental nuclear parameters and the observable of interest, TBR.

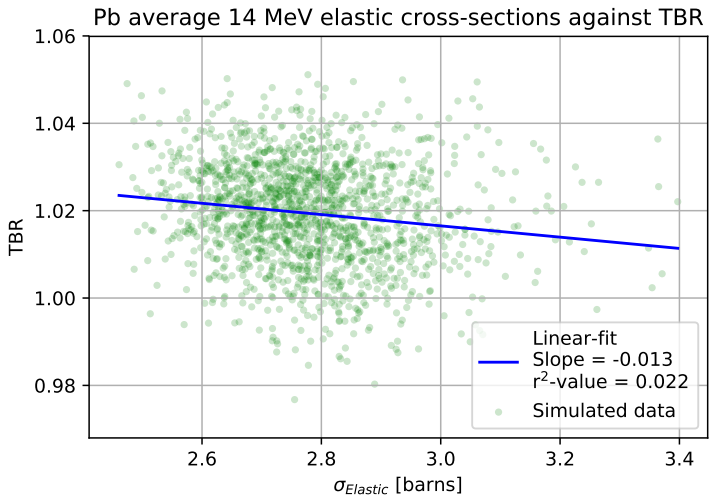

Fig. 2.11 Shown here are the ^{208}Pb(n,el)^{208}Pb cross-sections at 14 MeV, from the files used as input ND for the TMC simulation. They are plotted against the resulting TBR value from each simulation. There is a weak correlation, with a minor downward trend in TBR for increasing elastic cross-section.

The optical model is a nuclear model which aims to describe the interaction between an incident nucleon and a target nucleus. Rather than trying to solve particle behaviour *ab initio*, this approach uses a complex potential energy function or simply 'potential', $\mathscr{U}$ to represent nucleon-nucleus interactions. This potential energy function is a kind of description of the strong nuclear force, as a nucleon far from the nucleus, outside of the range of this force will experience no interaction. Closer, where the separation distance r becomes similar to the nuclear radius R and the particle may fall into the nucleus' potential well. Solving the Schrödinger equation with a given potential generates predictions for basic observables such as elastic scattering angular distributions and possibly the reaction and total cross-sections [57].

The optical model's name is due to an analogous case where a light wave is part refracted and part absorbed in a material of complex refractive index. The imaginary part of the complex refractive index is responsible for the absorption of light. In the nuclear case, the imaginary part of the nuclear potential represents all the non-elastic behaviour. An optical model potential (OMP) contains parameters and can be fit to experimental data. In this case the optical model is known as phenomenological. In part because of the wealth of experimental nuclear physics data now available, there has been a great advance in the capabilities and generality of the optical model since its first iteration by Fernbach et al. in 1949. Fits to experimental data can be global, determining parameter values which best fit a wide range of energies and nuclei, or local, where a fit is determined for a sole nucleus.

These local fits produce more accurate results but clearly have worse predictive power for other situations outside of their fit parameters.

The TENDL team used the T6 package to generate the TENDL2015 data employed for this work. The nuclear reaction program within this suite, TALYS, employs the optical model specified by Koning and Delaroche. The potential $\mathscr{U}$ has contributions from the terms in equation 2.13.

$$\mathscr{U}(r,E) = -\mathscr{V}_V(r,E) - i\mathscr{W}_V(r,E) - i\mathscr{W}_D(r,E) + \mathscr{V}_{SO}(r,E) + i\mathscr{W}_{SO}(r,E) \tag{2.13}$$

Where $\mathscr{V}_{V,SO}$ and $\mathscr{W}_{V,D,SO}$ are the real and imaginary components of the volume (V), surface (D), and spin-orbit (SO) potentials respectively. Each of these terms is separated into an energy dependent well-depth and an energy-independent radius dependent component[3]. The general form of the radius dependent part for all the potentials is a Woods-Saxon shape [136] as shown in equation 2.14.

$$f(r,R,a) = \left(1 + e^{\frac{r-R}{a}}\right)^{-1} \tag{2.14}$$

Where r is the separation, a a 'diffuseness' parameter and R the nuclear radius. r_V is equivalent to the nuclear radius parameter R in equation 2.14 for the $\mathscr{V}_V$ and $\mathscr{W}_V$ terms of the OMP shown as equation 2.13. The distribution of r_V values that were sampled in the creation of ^{207}Pb data for this study is shown as figure 2.12. It is uni-modal, non-normally distributed and positively skewed.

[3] See §4.1.1 in [68] for more detail

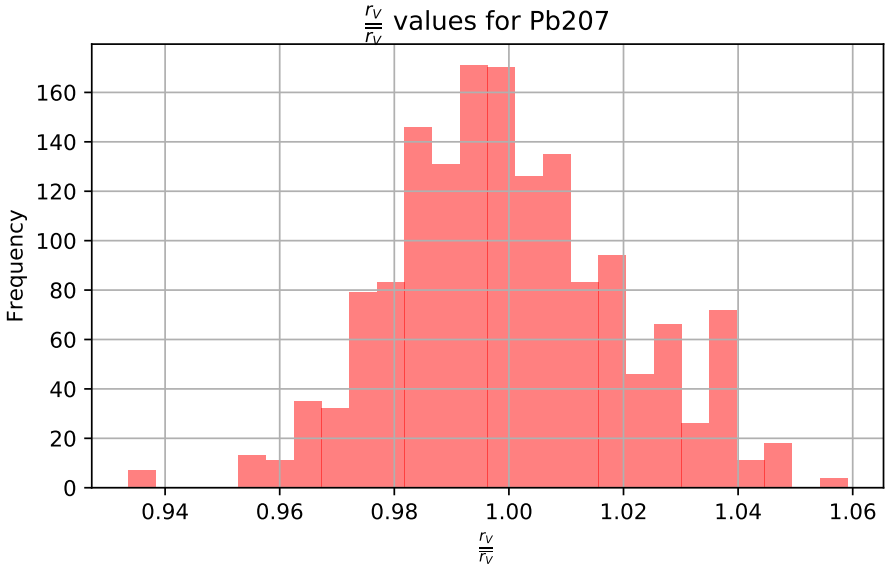

Fig. 2.12 This figure shows the distribution of nuclear radius, or r_V values sampled in the creation of ^{207}Pb cross-section data for this TMC work. r_V is a geometry parameter in optical model potentials. The distribution is slightly positively skewed. The data here have been normalised by their mean to centre the distribution on unity and allow for easier comparison between parameters of different scale.

Shown as figure 2.13 is the relationship between r_V as shown in figure 2.12, as applied to the generation of ^{207}Pb data, to the resulting TBR values. The data in figure 2.13 have a weak relationship. A linear least squares fit is shown, with a positive slope of $m = 0.06$. The r^2 value is very low at 0.01. This implies, as can be seen visually, that a great deal of variance in the data is left unexplained by the fit. However, this study has varied several nuclear parameters and nuclides simultaneously–were fewer inputs varied the relationship shown in figure 2.13 might be clearer.

Figure 2.14 shows TBR data plotted against the values of nuclear parameter d_1 used in the simulation input data. This parameter has the highest correlation coefficient, r^2 of any parameter for ^{207}Pb.

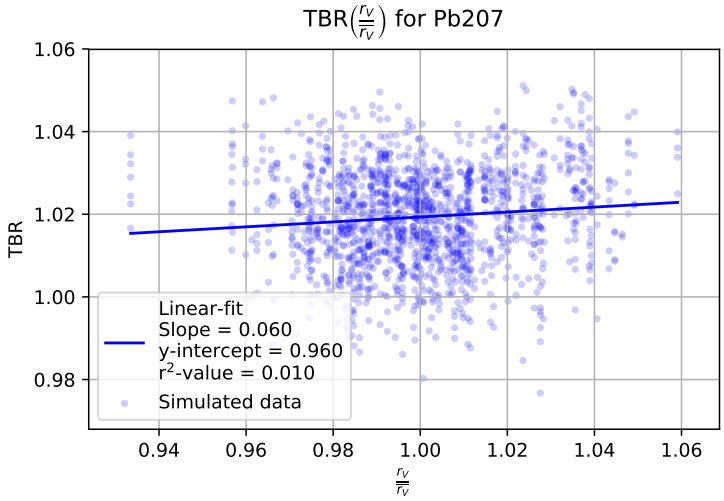

Fig. 2.13 Scatter plot of the r_V nuclear radius parameter against simulated TBR. The r_V values have been normalised to the mean of their distribution, centring the distribution on unity. The data has been fit with a linear relationship, with the fit data shown in the plot. A minor positive correlation is registered.

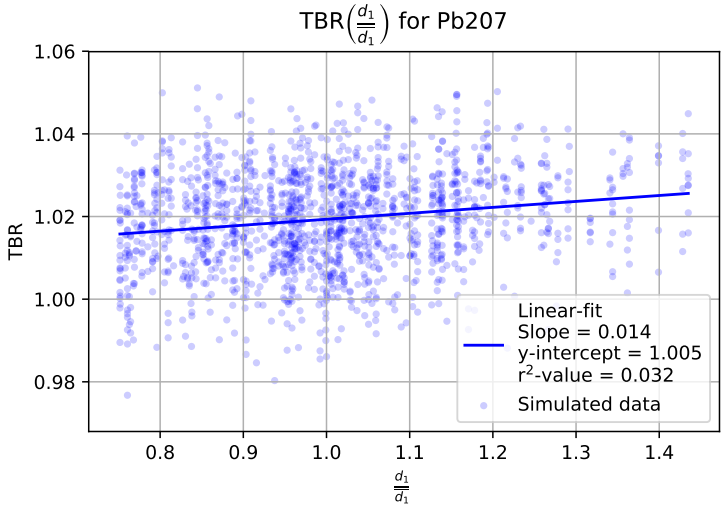

Fig. 2.14 Scatter plot of the d_1 parameter against simulated TBR. The d_1 values have been normalised to the mean of their distribution, centring the distribution on unity. The data has been fit with a linear relationship, with the fit data shown in the plot. A positive correlation is registered.

Performing a least-squares fit on all the varied nuclear parameters against the simulated TBR results, we can estimate which parameters TBR is most correlated with. For ^{207}Pb the parameter correlation bar chart is shown as figure 2.15, with the ordinate displaying the r^2 of

the linear regression. Prior to fitting, the parameters were normalised to their mean value. It can be seen that the sampled variation in the $d_{1,3}, w_2, r_V$ and a_v parameters for ^{207}Pb result in the most change in simulated TBRs. The parameter most strongly correlated with TBR for ^{206}Pb and ^{208}Pb is d_1, a multiplicative parameter in the energy dependent component of $\mathscr{W}_D$, the imaginary surface potential. Better constraining these parameters might act to reduce uncertainty in future lead based blanket analyses. This could be achieved with more experimental data or advances in model theory.

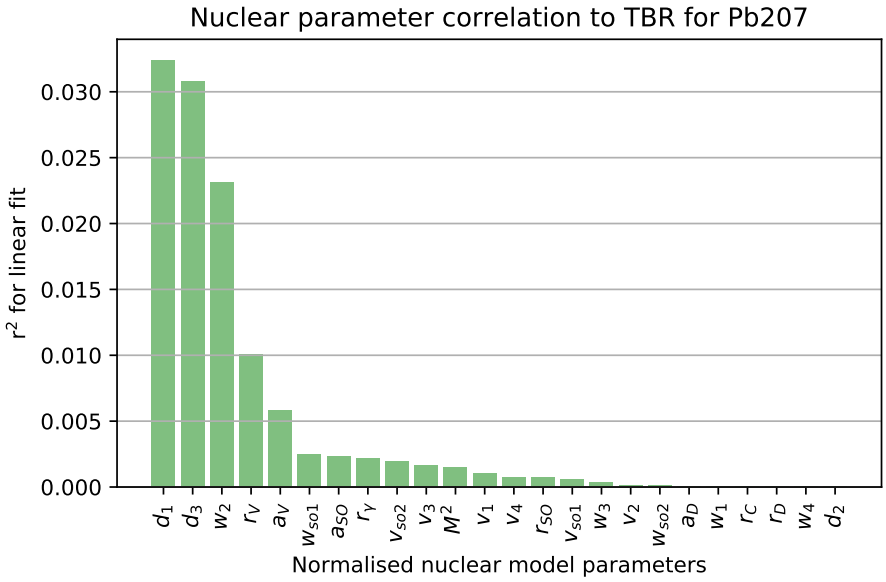

Fig. 2.15 This bar chart shows the relationship between various nuclear parameters and TBR simulation data. The abscissa shows a variety of model parameters, the vast majority from optical model potentials. The ordinate values for each parameter are measures of the correlation between sampled parameter data and the TBR. The correlation measure is the r-squared value of a linear fit to each dataset. The parameter data was normalised to its mean before fitting. The majority of parameters studied have little correlation with TBR, leaving a few optical model parameters such as r_V and d_1.

2.5 Conclusion

TBR uncertainty has been computed with the TMC technique, investigating the contribution of uncertain nuclear data from the three major lead isotopes. The standard deviation is 1.2% of the mean TBR. However, note that 5.8% of the distribution is less than unity. While the average value may appear to be feasible, it should be stressed that there is a non-negligible probability of a value below a required limit. The TBR values are relatively low as this

particular model has not been optimised for TBR and any practical design should have a TBR ≈ 1.1 [43], but in future design studies engineers should be aware of the probability of non-compliant operational parameters.

The TENDL-2015 nuclear data investigated in section 2.3.1, which is not necessarily Gaussian in shape, has yielded a TBR distribution with a small but finite negative skewness, an extended low-value tail. Decreasing the uncertainty in the aforementioned important optical model parameters, such as d_1 for Pb, would potentially act to reduce the uncertainty in future TBR analyses.

Uncertainty propagation in Monte-Carlo type radiation transport problems has often previously been computed using linear perturbation theory approaches. Unfortunately these are only applicable for small changes in the input data. They are also unable to reproduce probability distributions of the integral quantity of interest [112]. Whilst figure 2.9 shows a TBR distribution that is only slightly skewed, that is not to say other fusion quantities will not be. Koning & Rochman have demonstrated that fast and accelerator driven fission systems can have significantly skewed k_{eff} values, well described by an Extreme Value Fit (EVD) [72].

Future work on TBR in HCLL could include the effect of other nuclides and elements which TBR is sensitive to including iron and oxygen as well as a completely correlated uncertainty propagation method employing lithium if data for light nuclides becomes available.

More generally, when solving for integral quantities in nuclear fusion systems, thought should be given to fully-correlated uncertainty propagation and the form of the resulting probability distributions. In particular, whether a non-normal distribution with increased likelihood of extreme behaviour would have engineering design or safety implications for TBR, nuclear heating, fast flux, gas production or damage terms. Moreover, in all analyses for design applications the non-negligible probability of operational parameters in unacceptable regimes should be borne in mind.

Chapter 3

Quantifying received dose errors introduced by modelling approximations in reinforced concrete shielding

3.1 Outline

This chapter looks to determine the impact of a nuclear analysis modelling approximation known as spatial homogenisation–where complex geometries of differing materials are combined to form a homogeneous volume with a new material mixture. An introduction to radiation shielding is given. Then, undertaking a study for and of relevance to ITER, the effect of spatial homogenisation on the attenuation properties of reinforced concrete walls is investigated. A comparative method is described and subsequently the impact of spatial homogenisation on the shut-down dose rate to workers is also analysed. The effects of spatial homogenisation have at most a 22% underestimate for on-load doses in the cases modelled here. The shut-down dose rates on the inside face of the wall are slightly overestimated by spatial homogenisation. As a prelude to chapter 4, the importance of appropriately binning the energy domain is also explored.

3.2 Introduction

The ITER nuclear fusion experiment will begin DT operation in the 2030s. At 500MW fusion power the source rate of 14.1MeV neutrons will approximately 1.8×10^{20} s^{-1}. For comparison, JET's maximum source rate to date is 30 times less. But the real difference is

in the fluence. JET's lifetime neutron budget is now set at 2×10^{21} [81], or 10 seconds of operation for ITER-DT. The ITER life-time budget is approximately 3×10^{27} neutrons.

Designing a complex device such as a superconducting tokamak to operate in this radiation environment is a particularly challenging aspect of the ITER project. Care must be taken both to shield sensitive components like semiconductor devices from Single Event Effects (SEE) and to prevent excessive nuclear heating in large systems like the superconducting coils. The high flux and fluence of high energy neutrons means activation in areas like the Neutral Beam (NB) cells and in the Tokamak Cooling Water System (TCWS) will be significant. And of course, gamma from these activated components, and the prompt radiation of a thermonuclear plasma must be shielded against to protect the workforce operating and maintaining the plant.

Quantifying the absorbed dose to radiation sensitive components, or the equivalent dose to personnel requires crafting a computer model of the problem and typically making a series of assumptions and simplifications in the process. For reinforced concrete radiation shields, this often entails either homogenising the rebar with the concrete, or even neglecting the presence of rebar entirely. This process introduces systematic error into the reported solution of the problem. So, in addition to uncertainties from cross-section and secondary particle data as explored in chapter 2, modelling approximations play a role in thickening the fog of misunderstanding.

3.2.1 Radiation shielding

A fusion reactor harbours many different radiation fields. A variety of particles, across the energy spectrum, are present at one time or another. In order to preserve biological life and to extend the operable life of components, radiation shielding is employed to attenuate radiation fields, reducing their intensity by many orders of magnitude.

3.2.1.1 Radiation fields

The radiation fields of principal interest for shielding in fusion systems are the neutron and photon fields both during and after operation of the device.

- Prompt neutron–During a plasma shot, fusion neutrons will be emitted from the plasma at around 14.1 and 2.5 MeV for DT and DD reactions respectively. These neutrons propagate out of the plasma unheeded due to the very low plasma density. The uncharged particles travel in straight lines until they encounter a nucleus in the surrounding material. For a fusion reactor this will be the first wall, divertor or initial layers of the blanket. The neutrons then scatter elastic and inelastically off nuclei,

losing energy in the process by transferring it to the bombarded nuclei. Eventually the neutrons reduce in energy to the $1/v$ region of nuclear interaction, where lower energies (velocities, v) are favoured as the time for neutron-nuclei interaction is increased. At these low energies, reactions like radiative capture become very likely and neutron fluxes decrease through absorption.

- Prompt photon–While a fusion reactor is running, a large amount of Bremsstrahlung radiation will be emitted from the plasma at energies in the keV range and γ photons at MeV energies can be emitted by certain hydrogenous nuclear reactions. However, other radiation will be generated by the interaction of neutrons with surrounding matter. As the neutrons slow down their inelastic collisions excite nuclei which return to a ground state by emission of an energetic photon. This process generates an intense photon field.

- Shut-down photon–After cessation of plasma operations, there will still be intense radiation fields within and around a tokamak. The fusion neutrons will have activated and transmuted nuclei through reactions such as radiative capture. Many of the reconfigured nuclei will not be stable and so will decay through various emission mechanisms to a stable ground state. This often involves a chain of γ and other emissions. Despite the reactor being in 'shut-down' personnel and material will absorb a dose within the vicinty of the plant, hence the acronym Shut-Down Dose Rate (SDDR).

- Shut-down neutron–Despite the lack of fusion neutrons after the reactor has turned off, a neutron field still persists. Reactions such as photo-fission of actinide contaminants, ^{2}H(γ,n)^{1}H in cooling water [64] and ^{9}Be(γ,nα)^{4}He in the first-wall [33] will generate neutrons.

3.2.1.2 Shielding design

Shielding for a fusion system will have to attenuate the above fields to an acceptable level for the safety of humans and the wider environment, as well as any equipment which is sensitive to radiation. The basic principal for shielding photons is to use Compton scattering (see chapter 1, section 1.3.2.4) off electrons to reduce the photon energy until they are absorbed. Therefore the higher the electron density, the greater the γ shielding properties of a material. For this reason, high-Z materials such as heavy metals like steel, lead and tungsten are frequently used. Shielding for neutrons is similar in approach, moderating and then absorbing, although different materials are suitable. For moderation low-mass

nuclides are best, to maximise the energy transfer with each collision (see equation 1.4 in section 1.3.1.1, chapter 1). The best nuclide is ^{1}H, so any hydrogenous substance such as water provides an effective neutron moderator. Once moderated, absorption is best achieved by nuclides with a large radiative capture cross-section. There are many considerations in designing a long-lived, effective radiation shield, especially for high-flux operation. Any shield design will have to balance parameters such as cost, weight, volume, radiation damage, self-activation and nuclear heating. A much more complete discussion of radiation shielding design can be found in [105].

3.3 Method

The modelling simplification under investigation is that of spatial homogenisation and how it affects doses to workers, whether prompt or delayed. It is theorised that any errors incurred by this modelling approximation will be functions of geometry and material composition. For each set of variables interrogated, at least two simulations will be performed. One will have a high-fidelity model of the real-world problem geometry with reinforcing bar and stirrup rods faithfully reproduced. The other will smear the rebar across the concrete, homogenising the two materials into one with a density equal to the mass weighted sum of the concrete and steel constituents.

Simulations employing realistic ITER source particle distributions will tally transmitted or 'leakage' fluence at the rear of the shielding walls in both cases. This permits comparison between the modelling approximation and the higher-fidelity simulation.

The implications of spatial homogenisation prove to be different for on-load and shut-down (activated) dose rates. It is helpful to consider them separately. The following two sections detail the methods for determining these two dose rates.

3.3.1 Prompt neutron & gamma radiation

Determining the dose to personnel beyond a wall requires knowledge of the radiation fluxes there. The on-load gamma flux $\phi_\gamma(x,y,z,E,t)$ and neutron flux, $\phi_n(x,y,z,E)$, can be computed with Monte Carlo techniques, spawning neutrons from a prerecorded source distribution, $P(x,y,z,E,\Omega)$ at a particular point with a given energy and angle. These particles may propagate through the wall, scattering, being absorbed, exciting nuclei (which later relax through radiative emission) and perhaps leaking out of the back of the shield wall. Tallying this leakage flux of neutrons and photons is the main task of calculating a dose due to radiation.

3.3.1.1 Model geometry

The heterogeneous, realistic wall is constructed of a large concrete block, with two meshes of rebar, one buried below each face. The distance from wall exterior to rebar is known as the cover depth. Connecting the two meshes are a small number of narrow gauge 'stirrup' bars. The homogeneous model is the same external dimensions and total mass as the heterogeneous one, but without an internal steel structure and with homogenised materials.

A Python program has been written to take wall thickness, cover depth and other parameters as input variables to construct a corresponding MCNP model. This model utilises lattices to construct the repeated features of the reinforced concrete wall. The concrete block dimensions are $(x, y, z) = (10, y, 10)$ where y is the wall thickness. Embedded within the block are the major two meshes of steel reinforcement. These meshes are always constructed with 200mm spaced rebar extending in both the x and z directions [103]. The meshes are buried a small cover depth from the two y-perpendicular surfaces of the wall. Tying the two meshes together are some thinner 'stirrup' bars extending in the y direction, with $\approx 4m^{-2}$ of wall face [103]. A particular level of reinforcement may be referred to as 16HB200, i.e. 16mm diameter rebar, at 200mm spacing.

Simulations were run for the following wall thicknesses: 30, 50, 80, 120 & 210 cm. For each case, the fraction of total wall mass as steel was kept constant at 4.45%. This was achieved by increasing the rebar diameter, noting the relationship shown in equation 3.1.

$$\frac{d_r^2}{t_w} \propto f_s \tag{3.1}$$

That is, the rebar diameter, d_r squared divided by the wall thickness, t_w is proportional to the steel fraction, f_s. The value of 4.45% was chosen as this is a typical value computed from the description of reinforcement and dimensions in [103].

Material compositions, wall widths, rebar arrangements and steel fractions have been determined from ITER drawings and specifications for concrete walls in the tokamak building. The work has been conducted with a moderated ITER DT fusion neutron source spectrum. The results should give a good idea of the implications of the homogeneous modelling approximation at ITER but also more generally given the standard civil engineering design of the reinforced concrete walls.

3.3.1.2 Materials

Although much of the investigation will be by comparing values for homogeneous and heterogeneous responses, it is still important to supply the most accurate input information

Homogenising reinforced concrete

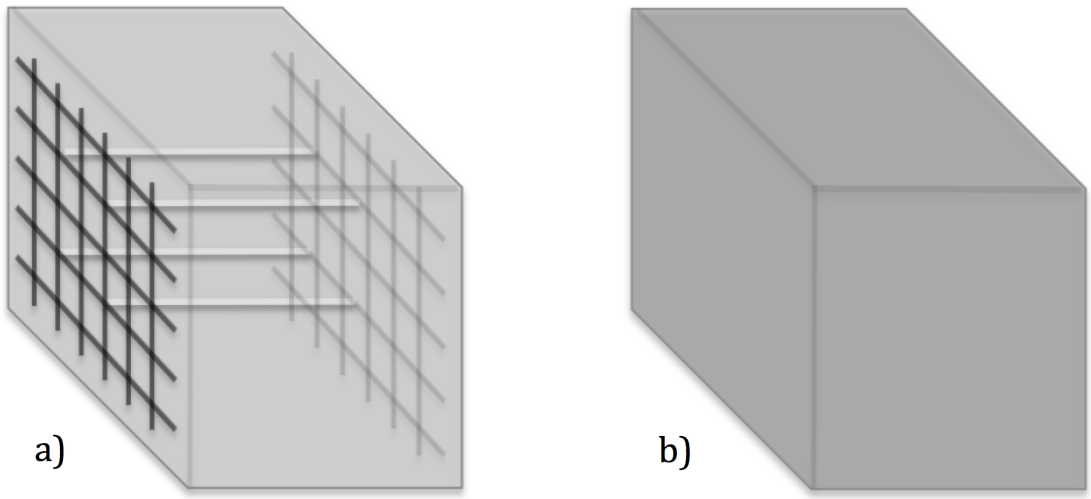

Fig. 3.1 The two modelling approaches considered in this investigation. a) Heterogeneous, where the steel reinforcing bar and stirrups are explicitly modelled. b) Homogeneous, where the mass of rebar is 'smeared' through the concrete. The new homogenised material has a greater density than plain concrete, conserving mass.

possible for the simulations. All the compositions that follow were mixed using the PyNE python package [134] employing natural abundance for isotopic distributions.

Concrete The density of concrete used is 2.2g cm^{-3} [63]. The composition is shown in table A.1 in appendix A.

Steel The density of steel used for rebar is 7.85g cm^{-3} [15]. The steel minor constituents Cu, Mn, Cr, Mo, V and Ni are not specified in [15] however the CEV (Carbon Equivalent Value) is given, a figure for quantitatively comparing the 'weld-ability' of steels. The CEV formula was used to determine likely % fractions of the aforementioned elements, giving a material composition as shown in table A.2 in appendix A.

3.3.1.3 Radiation sources

The effect of spatial homogenisation is considered for both neutron and photon radiation, hence likely distributions are required for both particle types.

No particle direction information was available, so it was assumed that all particles are born travelling perpendicular to the radiation shield. It was reasoned that this is a conservative assumption, liable to increase the discrepancy between heterogeneous and homogeneous

approaches as the heterogeneous model has features (the stirrup bars) which are aligned with this initial particle direction.

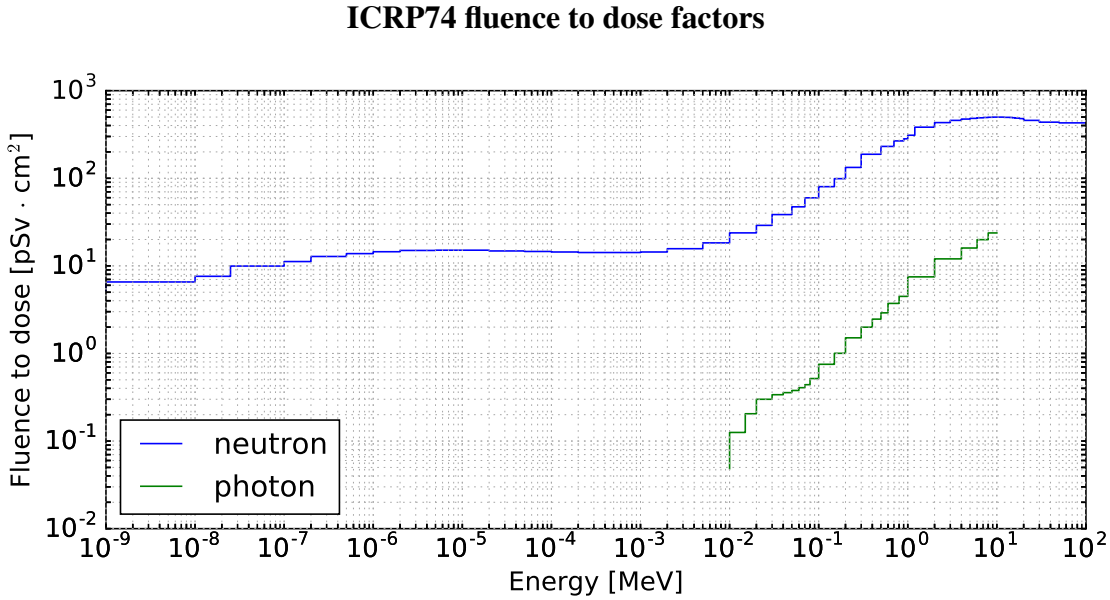

Fig. 3.2 The fluence to dose conversion factors for neutron and photon exposure.

The effective dose to humans is a function of the energy and type of radiation. It takes into account the varying susceptibilities of different tissues in the human body. For radiation protection, this is the typical figure quoted when specifying a dose rate. Where required, translation from fluence or flux to dose or dose-rate respectively is by functions as shown in figure 3.2.

Neutron The neutron source spectrum was obtained in previous work by Jakhar and is shown below as 3.3. The location tallied for this particle spectrum was behind a ITER neutral beam assembly, of neutrons incident upon the shield wall behind. In the ITER plant this area will receive an elevated flux over other areas of the bio-shield due to the penetrations from neutral beam injectors into the plasma chamber. The neutron spectrum was provided in the 175 VITAMIN-J group structure. Unfortunately this group structure provides very little information about the fine structure of the thermal neutron distribution below 0.1 eV. This is a deficiency in this study as the source spectrum is heavily thermalised. For the radiation transport simulations, no neutron source rates or wall loadings were specified. Instead, results were reported per source neutron. Many of the results later presented are ratios of the same quantity and hence unitless.

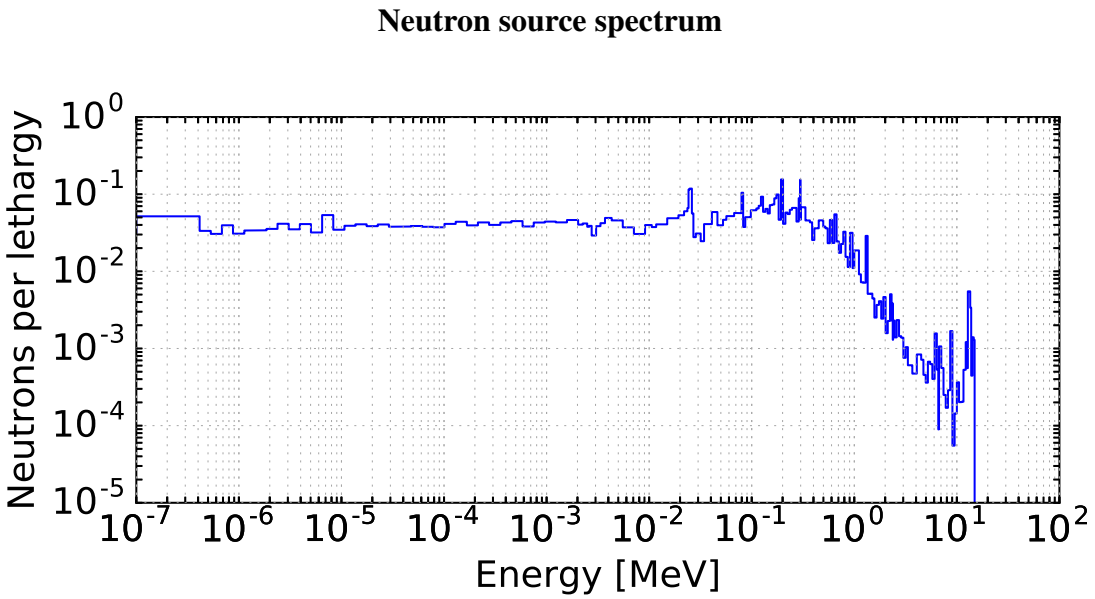

Fig. 3.3 Neutron spectra for incident radiation. This spectra was tallied at back of the ITER Neutral Beam assembly.

Photon The Tokamak Cooling Water System (TCWS) will move large quantities of water through an intense neutron flux during ITER's operation. This water leaves the tokamak through a network of penetrations and pipes in the surrounding facility. The water is activated mainly by the following reactions: ^{16}O(n,p)^{16}N and ^{17}O(n,p)^{17}N. These nitrogen isotopes rapidly decay as shown in table 3.1.

Activation product	$t_{\frac{1}{2}}$(s)	Decay mode
^{16}N	7.13	β
^{17}N	4.14	$\beta \rightarrow n$

Table 3.1 The parent, half-life and principal decay mode from activated nitrogen isotopes in the ITER water cooling system.

The nitrogen isotopes listed in table 3.1 decay through β and neutron emission to oxygen isotopes, however they may still be in an excited state and so will emit gamma radiation until they achieve their ground state. The gamma decay of ^{16}O is utilised as a source for several of the simulations presented in section 3.4.1.

3.3.1.4 Computation

The nuclear data employed was the continuous energy Joint Evaluated Fission and Fusion file 3.2 (JEFF3.2) for neutron transport and MCPLIB84 for photon transport. EAF2010 data was used at ITER Organisation's request for activation calculations. Radiation transport was conducted with MCNP6v1.0 [50].

The MCNP relative error, $R = \frac{\sigma_s}{\mu_s}$ which is the ratio of sample standard deviation to sample mean was kept to below 0.1 for all energy bins and mesh voxels where possible. However, the wall is a shield and heavily attenuating of source particles. For a 50cm thickness of concrete, the neutron attenuation is approximately 7 orders of magnitude. Therefore in order to converge the spectrum and solve the problem a significant number of histories must be run.

Converging results on a coarse energy grid or 'group structure' requires fewer source particles for a given R, but having fewer energy bins means the spectrum will be correspondingly poorly resolved, with large steps in flux across the energy domain. Initial simulations using VITAMIN-J 175 group structure did a good job of resolving the spectrum from 1 eV up to the fusion peak at 14MeV, however they completely fail to represent the fine structure of the thermal spectrum. This is demonstrated in figure 3.4. The two lowest bins are $[10^{-5}, 10^{-1})$ eV and $[10^{-1}, 4.1 \times 10^{-1})$ eV. Despite the lack of thermal energy resolution, VITAMIN-J extends to very low energies. This means that neutrons scored at thermal energies contribute equally across four orders of magnitude in the lowest bin. When a spectrum such as this is sampled later, for an activation calculation, say, fictitious neutrons at 10^{-5} and 10^{-4} eV are produced and will skew results for problems where reactions in the 1/E region are significant, potentially drastically overestimating reaction rates.

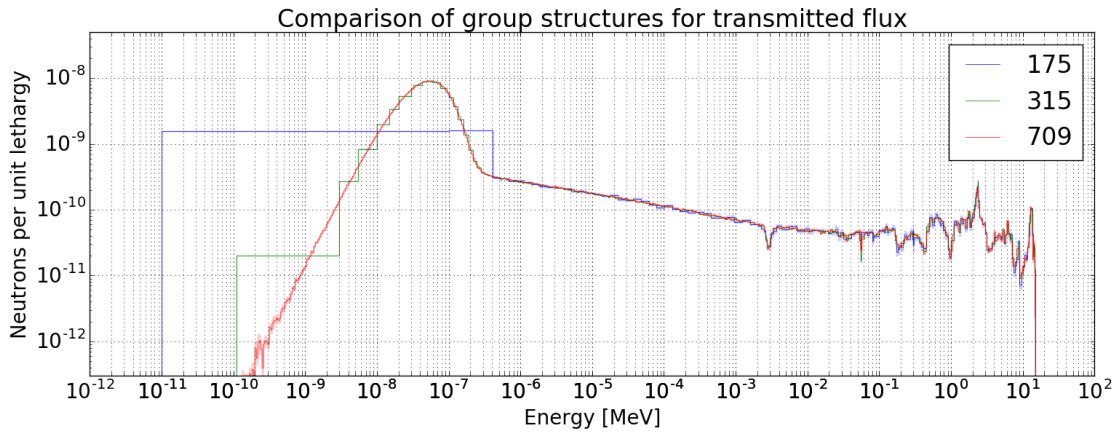

Fig. 3.4 This figure shows the same transmitted flux, as simulated by MCNP, binned in a variety of group structures. While the region down to approximately 1 eV is resolved broadly similarly, the VITAMIN-J 175 group structure has insufficient bins to resolve a Maxwellian distribution of thermal neutrons in the low-energy region.

The ideal group structure should contain enough bins in the low energy region to approximate a thermal distribution. However, as previously stated, all bins should converge to have a relative error of less than 0.1 as proof of convergence. For this particular problem, convergence is most difficult for the high energy region, as these particles are relatively rare. Hence, many fine bins in the high energy region will necessitate significantly longer run times. For the work presented here, the 315 neutron group structure was selected as a compromise between speed and fidelity.

3.3.2 Shut Down Dose Rate

The calculation of $\phi_\gamma(x,y,z,E,t)$ where t is some time after cessation of plasma operation comprises three main steps:

1. Neutron transport—as previously, compute the neutron flux during plasma operation, $\phi_n(x,y,z,E)$, recording the neutron flux binned by energy over a spatial mesh. Finer meshes will converge on true behaviour, with coarse meshes over or under-estimating fluxes depending on local geometry and flux gradients [1].

2. Activation—determine the appropriate irradiation scenario, then activate and transmute the materials present in the problem geometry. This involves assembling a system of differential equations, Bateman equations, to track the inventory of all the nuclides

[1] Using a mesh which conforms to the geometry, such as an unstructured mesh, has been shown to give significant improvements, especially for small features [37].

present. One can numerically solving this system for a series of time-steps using codes such as FISPACT-II [131]. The resulting nuclides and their abundances can be paired with decay data to generate a decay photon source on the original mesh.

3. Photon transport—using the distributed decay gamma source produced by step 2), conduct a radiation transport run to determine the photon flux, $\phi_\gamma(x,y,z,E,t)$, converting to effective dose as required.

3.3.2.1 Model geometry

The shut-down dose rate is calculated for the following scenario: it assumes a worker is inside the bio-shield at ITER during a shutdown, stood 30cm from the inner surface of a 150cm thick wall with a steel mass fraction of 4.5% the total wall mass. Reinforcing bar forms a mesh of squares 20cm in width and height at the front and back of the wall, 5cm from the surfaces.

3.3.2.2 Materials

For analysis of the Shut Down Dose Rate (SDDR), where small impurities can have large contributions to the dose rate, the steel composition was refined [8]. ITER limits for Co and Ni were imposed, 0.01%wt and 0.05%wt respectively. The resulting composition for steel is shown as table A.3 in appendix A.

3.3.2.3 Computation

As for before, radiation transport was conducted with MCNP6 v1.0. Neutron spectra were tallied on a mesh in the shield models. These spectra were used as input to the MCR2S activation linker code [34]. This program uses these spectra and a corresponding irradiation scenario to compute material changes within the meshed model. These new materials and decay information can be used to generate a photon source from the activated nuclides. The final steps in the calculation are to perform photon transport calculations from the activated wall to a target. The generation of photon sources and transport of emitted γ-rays must be computed for each decay time step of interest.

The irradiation scenario employed for the activation step was ITER's SA-2 [83], which approximates the ITER DT experimental programme total fluence and explicitly includes the final, end-of-life pulses for accurate estimation of short-lived nuclides.

As noted in section 3.3.1.4, the choice of energy group structure can be important for the accuracy of calculations. When a pre-sampled neutron flux is later sampled to calculate

a SDDR dose, the 175 group structure introduced a 20% increase in dose rate for steel at all time steps, and approximately the same increase for concrete until days after irradiation, when the discrepancy between 175 and 315 falls to zero.

The dominant contributors to dose rate in steel all originate through (n,γ) and as such are not threshold reactions, instead observing 1/E behaviour at low energies. Whilst the dominant contributor to dose rate in concrete is initially ^{24}Na formed though ^{23}Na(n,γ)^{24}Na, it decays with $t_{\frac{1}{2}} = 15$h before the dominant nuclide becomes ^{39}Ar, which is produced via ^{39}K(n,p)^{39}Ar. The (n,p) reaction is threshold, i.e. not effected by the low energy inaccuracy introduced by binning in 175. Therefore after a few days, the SDDR discrepancy between groups reduces to zero.

Further enquiry into the optimum energy group structures for activation calculations is presented as chapter 4.

3.4 Results & discussion

This section details the effect of spatial homogenisation on fluence and dose received by workers in several circumstances. First, on-load radiation: neutrons originating from the plasma and multiplication reactions within the reactor and photons emitted from neutron induced reactions such as inelastic scattering and activation of nitrogen in water. Subsequently, a case is investigated where neutrons are transported through the shielding wall, activate it and then the resulting gamma-rays transported, this time tallied on the inner side of the wall–estimating the dose received by workers performing a maintenance action between plasma shots.

3.4.1 Transmission of prompt radiation

As neutrons from the ITER plasma are incident upon the shielding wall, they advance through the material. Nuclides from steel and concrete materials scatter the neutrons, moderating their energy down. Figure 3.5 shows the neutrons have a mean free path of approximately 2cm in concrete, and most will scatter multiple times in the material–although the path length distribution is relatively broad, and some may travel 10cm or more without interaction. As the energy decreases, the likelihood of capture increases. Indeed, most neutrons are eventually absorbed, although a small fraction of the source particles propagate all the way through the shield and an even smaller proportion without any interactions at all–the DD and DT peaks visible in figure 3.6.

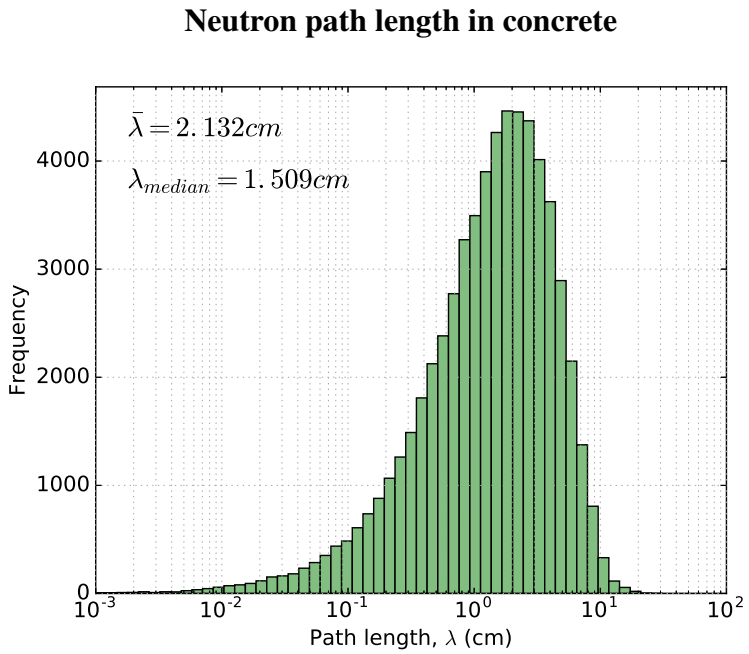

Fig. 3.5 The path lengths of neutrons in pure concrete are shown as a histogram. There is a wide distribution with a mean of approximately 2cm. This figure was compiled from MCNP ptrac data.

The transmitted, or 'leakage' neutron spectra for a thick (2.1m) concrete wall is shown below as figure 3.6. Note the DT and prominent DD reaction peaks at 14.1 and 2.5 MeV respectively. Other features are the flat slowing-down region with a few flux depressions and the significant thermal Maxwellian distribution. The flux has been severely reduced by its interaction with the wall, decreasing from the source by an order of magnitude each 20cm traversed. The most prominent difference between the heterogeneous and homogeneous models is in the thermal region.

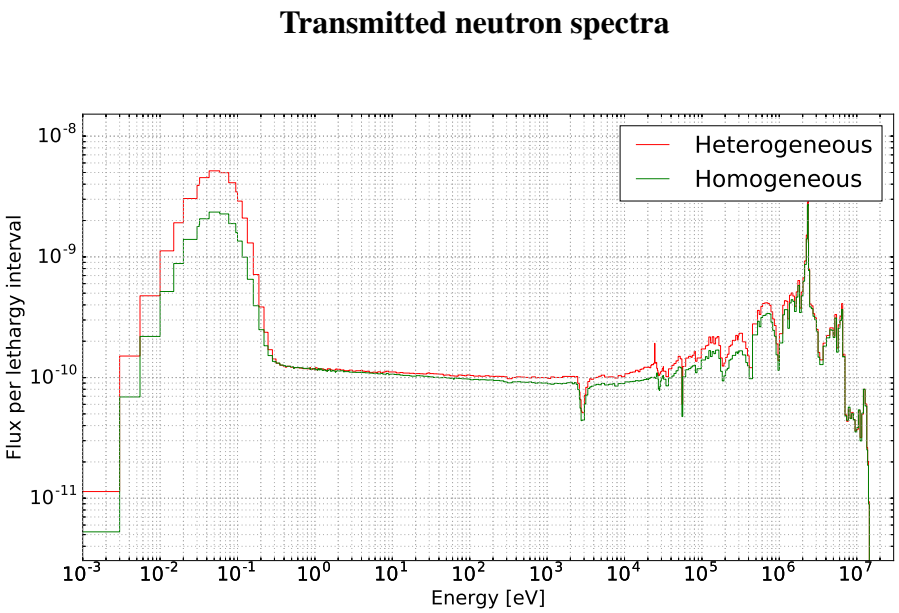

Fig. 3.6 The neutron spectra leaving the shield for the heterogeneous and homogeneously modelled cases. This example is for the thickest wall simulated, at 2.1m thickness. The rebar is 41mm in diameter at 200mm spacing, resulting in a steel mass fraction of 4.45% of the shield total. Note the substantially reduced thermal flux in the homogeneous simulation. The spectrum has been binned with the TRIPOLI 315 group structure.

The reduced transmitted homogeneous thermal flux in figure 3.6 is in part because steel is now available for neutron interactions throughout the depth of the wall, in the homogenised material, rather than solely being available for interactions near the surfaces. Steel has a greater macroscopic material capture cross-section than concrete, and so acts as a neutron sink for slow neutrons, preventing them from propagating through the shield. The material cross-sections for radiative capture are shown in figure 3.7.

Radiative capture probability

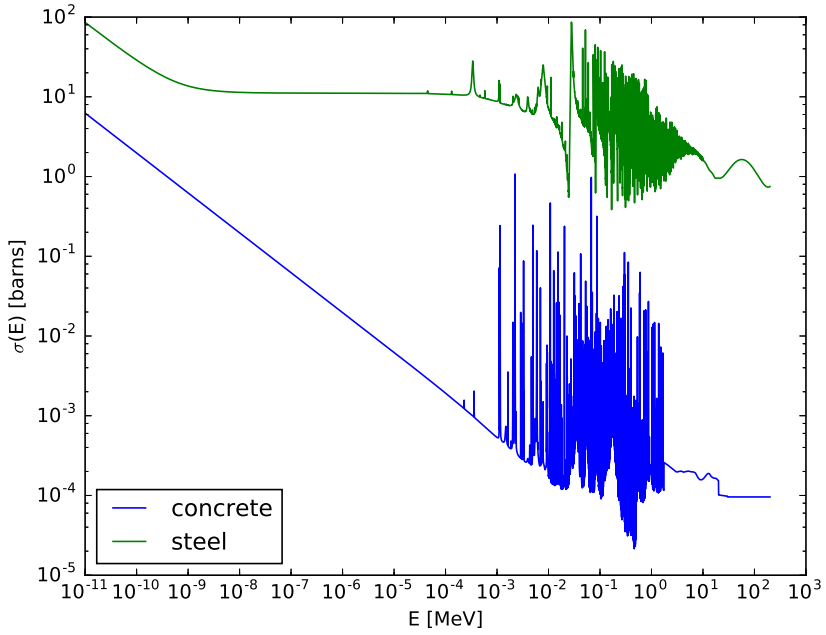

Fig. 3.7 The nuclear properties of steel and concrete are substantially different. Shown here is the cross-section for (n, γ) in both materials. This has been computed with MCNP6 [50], by performing an number density weighted sum of all the constituent nuclides for each material. At thermal energies steel has a material radiative capture cross-section two orders of magnitude greater than concrete.

Plotting the ratio of the leakage neutron spectra as figure 3.8, one can see the differences in flux more easily. Thermal flux is underestimated by more than a factor 2. But also, fast flux in the 1 keV–1 MeV range is underestimated by ~ 1.2. The fast flux discrepancy is correlated with wall thickness and is negligible for thin (<1m) walls.

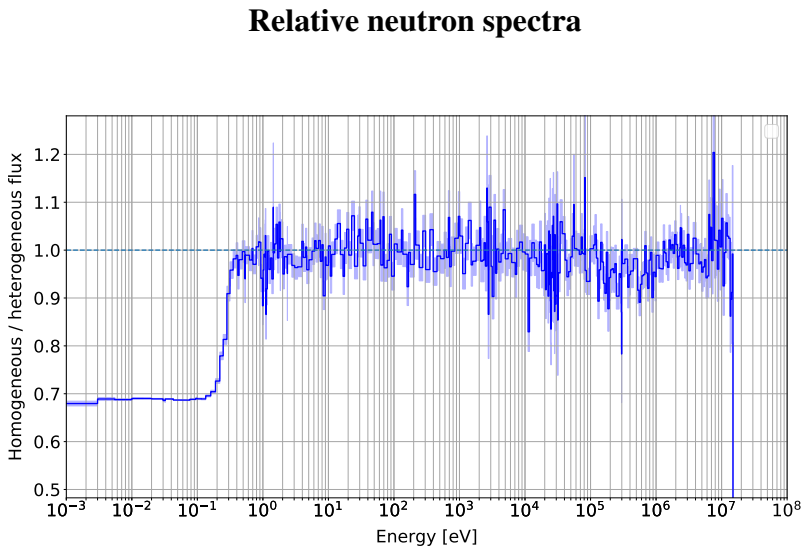

Fig. 3.8 The homogeneous flux decreases are readily visible in this figure, principally in the thermal region, but also in the 1 keV – 1 MeV range. The ratio is of the two spectra from figure 3.6. MCNP statistical errors have been combined in quadrature and are plotted as the light shading surrounding the mean value.

Figure 3.9 shows how the received dose discrepancy (as received by someone stood behind the wall) due to the homogeneous approximation varies with wall thickness. Thin walls have a relatively small discrepancy–to be expected as the heterogeneous model is at its most similar to the homogeneous in this case. As the walls become thicker and the rebar meshes become separated by a larger volume of concrete, the discrepancy increases. It plateaus at $\sim 1m$ wall thickness. Similar behaviour is observed for dose due to prompt photons. It can be seen that the effect is greatest for neutrons, with a maximum dose underestimate of 22%. The great majority of steel in the heterogeneous simulations is orientated perpendicular to the incident neutron flux, although a small number of 'stirrup' bars to run parallel with the flux. Pampin et al. studied water and steel radiation shields, where the orientation of water pipes was varied. Where the pipes were parallel to the incident neutron flux, the spatial homogenisation underestimate was as much as 375%, as the homogeneous model missed neutron streaming effects [100]. However, for the perpendicular configuration as is modelled here, Pampin et al. found the difference was 13%, not too dissimilar to the 22% underestimate reported in this work.

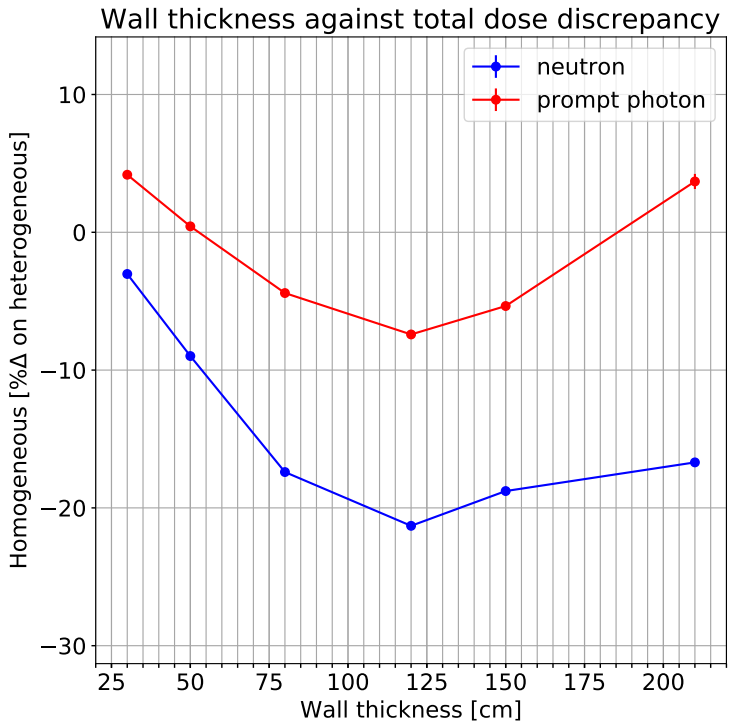

Fig. 3.9 This figure shows the relationship between the wall thickness and the difference between the homogeneous and heterogeneously modelled walls. The quantity plotted is the energy integrated dose for each case, for neutrons and prompt photons. Error bars due to radiation transport statistics are shown but may not be visible on this scale. The discrepancy between modelling approaches is a function of the wall thickness, greater for neutrons than photons.

The effect of the steel mass fraction has also been investigated. It was found that the discrepancy between homogeneous and heterogeneous approaches is effectively zero for transmitted photon flux. However, the homogeneous method underestimates neutron flux to a greater degree with an increased steel mass fraction. The greater availability of steel nuclides through the homogeneous mixture results in greater absorption inside the wall and thus a reduced leakage flux (and therefore dose due to neutrons).

3.4.2 Shut Down Dose Rate

After operation of a tokamak, repairs and maintenance are often necessary. This section explores the shut-down dose due to be received from the wall itself. It assumes a worker is inside the bio-shield at ITER during a shutdown, stood 30cm from the inner surface of a 150cm thick wall with a steel mass fraction of 4.5% the total wall mass. Reinforcing bar

forms a mesh of squares 20cm in width and height at the front and back of the wall, 5cm from the surfaces. Dose rates due to neutron activation in the reinforced concrete shielding have been calculated for when the reinforced concrete is homogenised and when it is modelled faithfully.

The effect of the spatial homogenisation modelling approximation generally acts to overestimate the SDDR by a small amount. The behaviour is shown as figure 3.10. On the timescale of seconds and minutes, and from days out to decades, the overestimate is approximately 10%. On the timescale of hours, the discrepancy between homogeneous and heterogeneous approaches closes, briefly becoming reversed, with the homogeneous approximation underestimating the SDDR by 4%. The absolute dose due to photons from the reinforced wall will have fallen by an order of magnitude in the period elapsed between a minute and a week after the final DT shot and it is unlikely anyone will be entering the facility before then.

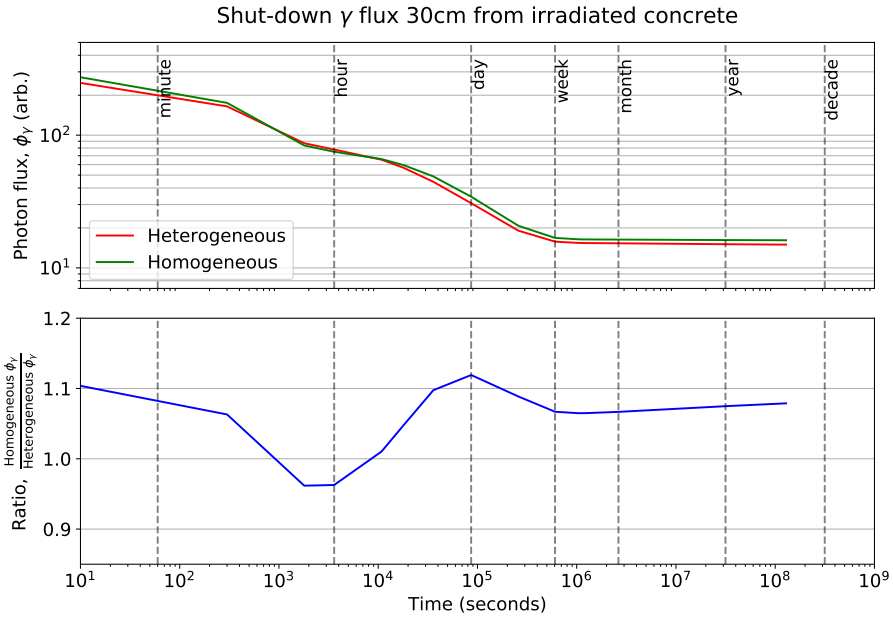

Fig. 3.10 This figure displays total ϕ_γ received at a distance of 30cm from the activated wall. The upper panel displays the absolute values in the heterogeneous and homogeneous modelling approaches, as a function of time since last irradiation. There are 16 time steps, for which the activation and subsequent photon transport has been carried out. The lower panel displays the ratio between the approaches, i.e. homogeneous values normalised by the heterogeneous values. One can see that the homogeneous approximation artificially increases the SDDR by $\approx 10\%$ at most time steps, bar those around an hour.

The complex behaviour displayed in figure 3.10 is a result of the many different nuclides which contribute to the decay γ field. These nuclides have a range of half-lives and inspecting their relative emissions as a function of time is instructive in understanding the shape of figure 3.10. The plot shown as figure 3.11 displays estimates for a contact dose with concrete and steel. As in the real case the reinforcing steel is buried within concrete, the numbers plotted here are not a substitute for transporting decay γ photons from their emission to their absorption (as was done for figures 3.10 & 3.12). Figure 3.11 reports specific contact dose rates, i.e. per unit mass. Note that 2.25% of the wall mass is steel close to the activated surface, while the mass of concrete between this reinforcing steel and the wall surface is $\frac{5cm}{150cm}(1-0.045) = 3.18\%$ of the total wall mass–so the amount of each material near the surface is very roughly equal and therefore specific contact dose rates are informative by themselves, without being weighted by mass.

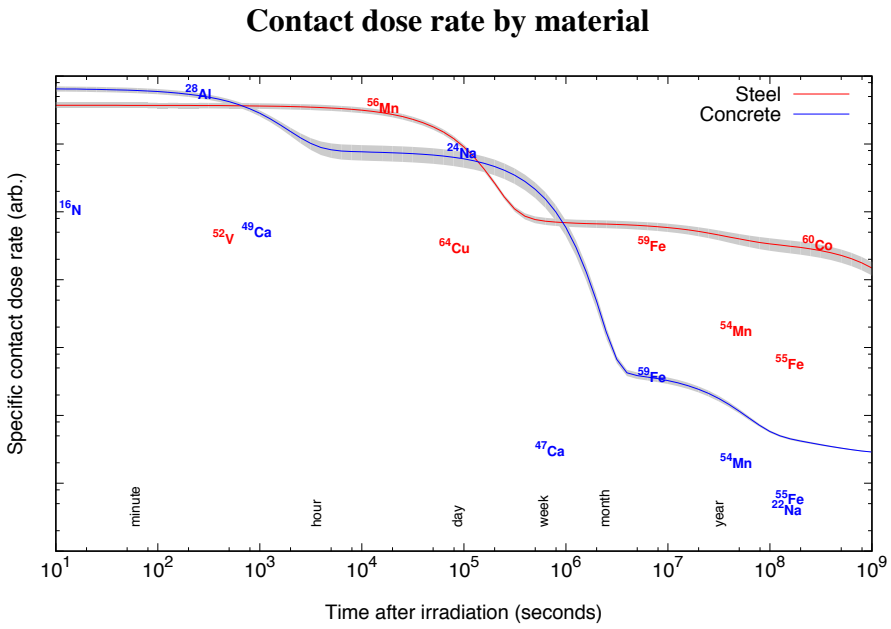

Fig. 3.11 This plot shows an estimate for the specific contact dose rate due to steel and concrete irradiated under the ITER SA-2 scenario. The total specific dose rate is given by the line plots, with an uncertainty band from EAF-2010 data included as the grey shading. Nuclides which contribute a significant fraction of the dose are shown with their abscissa value as their half-life, $t_{\frac{1}{2}}$ and their ordinate as their contribution to the specific dose rate. The specific dose rate behaviour is quite complex; the material with the highest specific activity changes three times in the simulation period due to various decays.

From figure 3.11 it is clear that the two materials have quite different nuclear properties, so it is not surprising that their spatial distribution influences the dose received. On a per-mass basis, concrete is more active on the timescale of minutes due to ^{28}Al, and as this decays the

homogeneous overestimate reduces. Subsequently, until a few hours have elapsed since last irradiation, gamma emission from ^{56}Mn makes steel the most active while the homogeneous approximation begins to overestimate SDDR again. After this, ^{24}Na in concrete is the dominant contributor to SDDR, making concrete the most active material on the timescale of days. After its decay, the homogeneous approximation falls slightly again as steel becomes the more active material out to very long timescales.

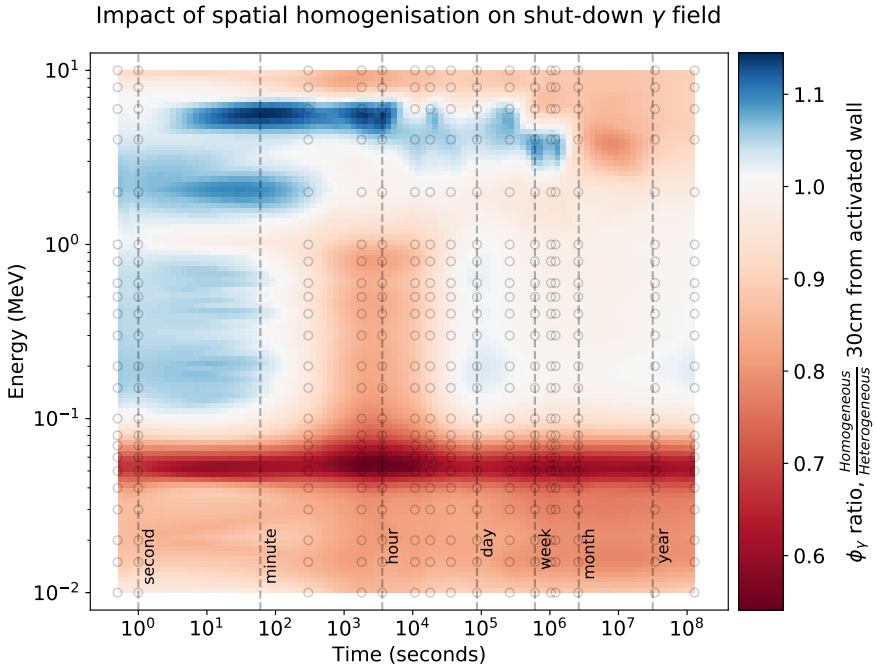

Fig. 3.12 This figure displays the ratio of homogeneous and heterogeneous ϕ_γ for a series of cooling times and all photon emission energies. Time since last irradiation is given in seconds on the abscissa, while the ordinate is photon emission energy in MeV. The grey circles indicate simulation data, i.e. flux ratios for a particular photon energy group at a particular cooling time-step. These data have been interpolated with the Clough Tocher method to generate the heat-map shown. A diverging colour-map helps identify where spatial homogenisation is overestimating dose (blue) and where it underestimates (red).

Homogeneous simulations tend to overestimate the number of high-energy gamma photons, as can be seen in figure 3.12. A homogeneous overestimate is shown as a blue band across much of the higher energy region. For the homogeneous case, some steel material is located everywhere in the material mixture, including the shallow 'cover depth' volume, outside of where the rebar is located in the heterogeneous models. Nuclides from steel then lie on the surface in the homogeneous model and will emit high-energy, gamma rays which can

leave the shielding unattenuated and are more likely to contribute to a worker's dose[2]. Work to corroborate this explanation could include varying the cover depth, to see if homogeneous models started to underestimate the received dose as the cover depth approached zero and the rebar was on the surface in the heterogeneous model.

It is worth noting that while not explicitly modelled here, it is likely that the SDDR on the exterior side of the shield will be artificially depressed by spatial homogenisation as the activating neutron flux is depressed by that approach. This is confirmed for similar circumstances in work by Sanz [120].

3.5 Conclusion

This chapter began with the basic principles of radiation shielding. A modelling approximation known as spatial homogenisation, in frequent use by nuclear analysts was introduced. Methods for investigating the effect of this approximation on dose rates at the ITER facility were described, before reporting on these results and discussion of their significance.

A series of simulations have been performed to assess the impact of the spatial homogenisation modelling approximation on various fluxes and associated doses in a radiation shielding context. The shielding was parametrically varied to explore how the modelling approximation impacts different geometries.

The specific circumstances investigated were firstly prompt neutron and photon emission from the the ITER plasma and its surroundings, transmitted through a variety of different thickness reinforced concrete walls. In this case, spatial homogenisation was found to underestimate received dose on the outside of the wall by up to 22%, with the discrepancy a function of wall thickness. For photons, there was less effect, with 10% the greatest divergence between approaches. The relatively small discrepancies are to be expected for reinforcing bar arrangements as used here.

The second case investigated how spatial homogenisation affected the dose to workers during a maintenance operation, the SDDR at the inside face of a wall activated by neutron irradiation. This particular configuration is relatively unaffected by spatial homogenisation, with the greatest deviation from the true case being an overestimate of 12% for the homogenised case. The discrepancy is a function of the cooling time, with a smaller discrepancy at timescales beyond a day, down to an overestimate of approximately 8%. As the homogenised results are dose overestimates, they are in this case conservative, and thus do not pose a risk to workers' health, rather providing an upper bound on the likely dose rate.

[2]High energy gamma-rays (between 1 and 2 MeV) constitute the majority of the dose received at all times, however beyond 2 MeV the flux and therefore contribution to dose falls off sharply.

Where the practice of spatial homogenisation likely results in a conservative estimate as is the case for SDDRs on the inside face, it can be tolerated, as it reduces model complexity and is safe. However, where neutron streaming paths are available, especially if features contain routes with low absorption or poor moderating properties, or if features are larger than the mean free path of the neutrons, then spatial homogenisation should be treated with caution and avoided if possible [100]. One imagines that with the increased demand for multi-physics analysis, in the not distant future details such as reinforcing bar and other small but potentially important features will be explicitly modelled in radiation transport so that accurate reaction rate information can be input to other simulations, otherwise one is introducing an error right at the start of any multi-physics calculation chain (or repeatedly within a loop).

Chapter 4

Optimising energy group structures for neutron activation calculations in fusion systems

4.1 Outline

This study utilises self-shielding factors as a means to optimise energy group structures for fusion activation calculations. Informed by an analysis of important fusion resonances and a survey of relevant incident particle spectra, we develop and test two new group structures designed to more accurately represent the physics of nuclear reactions. They are compared to group structures commonly used in fusion research & analysis. When used in a JET activation foil scenario, the optimised group structures outperform reference group structures, such as CCFE-709, while requiring fewer energy groups.

4.2 Introduction

Nuclear simulations for fusion devices are essential to determine material damage, activation-transmutation and to perform dose rate analyses. While continuous energy Monte Carlo modelling can be used to directly calculate reaction rates, due to the very large number (potentially tens of thousands) of possible reaction channels, it is impractical to compute all nuclear reactions of interest by this so-called point-wise approach. A separate inventory code, such as FISPACT-II [130], utilising a discretised incident particle spectrum is typically used to calculate all of these reaction rates, solve for the time-dependent nuclide inventory

and provide various responses and source terms. This multi-group method is computationally efficient, but introduces self-shielding errors.

4.2.1 Group structure optimisation

The likelihood of interaction between an incident particle and a nucleus is a a non-linear function of energy in the resonant energy region. The cross-section, $\sigma(\mathrm{E})$ and the associated reaction rate, RR(E), may vary by several orders of magnitude within a few electron-volts. The number of bins that the energy domain is discretised over has a strong influence on the accuracy of the result. While increasing the number of bins increases the accuracy it also increases the time required to converge an input spectrum and the memory required to record this information[1]. Figure 4.1 shows a reaction rate converging to the true result as a function of the number of cross-section groups. These results are from a simple model where a 14 MeV point neutron source is located within a series of concentric, spherical shells made of steel, water and tungsten. The tallied reaction rate is radiative capture in ^{186}W in the outermost, tungsten shell. The neutrons are moderated by the water and the energy spectrum will receive sharp flux depressions from resonant nuclides within the steel material. This flux is then incident upon the tungsten where resonant peaks from the tungsten isotopes will further impinge on the spectrum. Without adequately resolving these depressions, an accurate estimation of the desired reaction rate cannot be made. In this example it can be seen that getting from a 10% error to $\approx 1\%$ error requires a jump from 256 to more than 8192 bins. Any real problem will have many hundreds of nuclides present, with their nuclear data all discretised the same way, whether for a deterministic radiation transport calculation or an activation-transmutation calculation. Adopting such a fine discretisation for so many nuclides is not feasible given the memory requirements of such a scheme.

[1] A concern for simulations with millions of voxels, common in distributed activation-dose calculations.

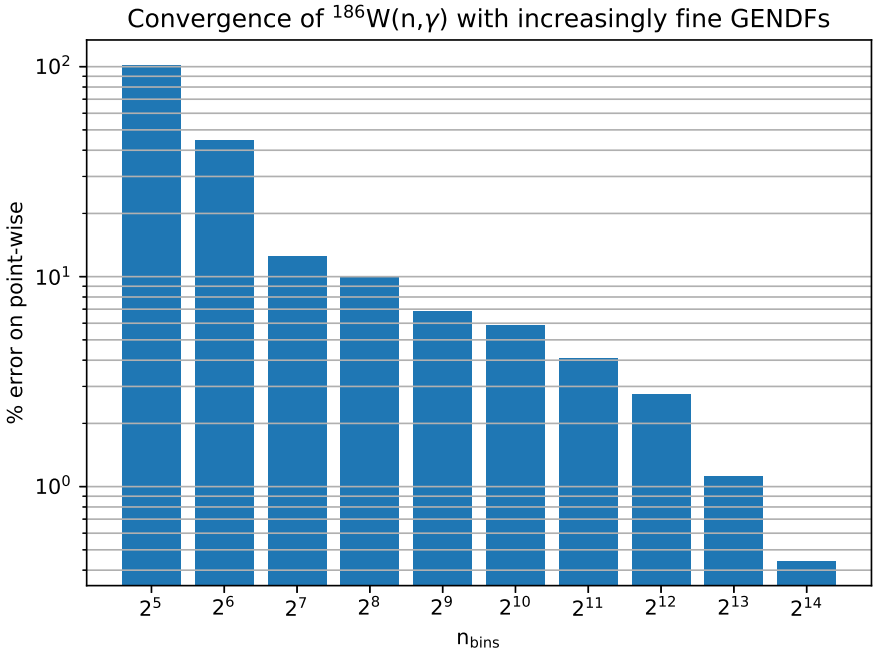

Fig. 4.1 This figure shows the results from a series of calculations with increasingly fine energy discretisations, or group structures. The bins for all groups are equally log-spaced. The calculation is to determine the (n,γ) reaction rate in a shell of ^{186}W due to some incident neutron flux. The dependent variable is the % change on a point-wise (true) value calculated with MCNP.

In addition to the number of bins, the bound locations have a strong relationship to the result accuracy. The integer bin count number and bound locations energy vector are together known as an energy group structure. In the example shown in figure 4.1 the bin boundaries have been uniformally logarithmically distributed. This is a very simple starting point (and many group structures largely adhere to it) and can be improved upon, even if only to adopt one or two regions of different bin density.

There have been several previous efforts to optimise neutron energy group structures. Certain applications lend themselves to highly targeted group structures. For example, fission reactor lattice physics calculations are typically interested in determining a few reaction rates to very high accuracy. The Studsvik team who develop the CASMO-5 lattice physics code employ a neutron group structure with very fine resolution around the principal ^{238}U and ^{240}Pu resonances [109]. Covering these areas with a fine energy grid results in a more accurate estimate of key reaction rates and hence calculation of the neutrons absorbed and consequently lost from the system. Particle Swarm Optimisation (PSO) has been used

multiple times to improve bin bound placement for multi-group libraries used for reactor physics applications [138] [4] [3] [44]. Morgan et al. [90] have explored hyper-fine multi-group (MG) data as an alternative to the interpolation of point-wise (PW) data as typically employed by particle transport codes such as MCNP6.1 [50]. Some attention has been directed towards the refinement of reaction rate calculation for specific elements and nuclides in fusion scenarios. For instance, work on the spatial heterogeneity of tungsten transmutation has been undertaken by Gilbert et al. [49] who used the CCFE 709 bin group structure and careful application of self-shielding factors to accurately determine reaction rates.

4.2.2 Resonance behaviour

It is clear that the placement of bin boundaries and the number of bins in an energy group structure has a significant effect on the potential accuracy of any calculation employing that group structure. The ideal location of these bin bounds is not equal spacing, but will be some complex result influenced by the kind of spectrum encountered and the nuclei present in the problem. This latter consideration can be explored by analysing the distributions of nuclear resonances for known nuclides within the {A,E,j,Γ} space, where A is the atomic mass, E_r the interaction energy, j the nuclear spin and Γ the resonance width. With knowledge of these distributions, especially in the mass and energy domains, perhaps energy bins might be targeted intelligently for a given problem.

To undertake this study, a set of scripts were written to parse resonance parameter information from nuclear data libraries. These can read ENDF-6 formatted entries, including several resonance formalisms for maximum nuclide coverage. The nuclear data library chosen to study was ENDF/B-VII.1 [23], the United States' reference nuclear data library. Figure 4.2 is a scatter plot of 41,112 nuclear resonances from nuclides contained within this library.

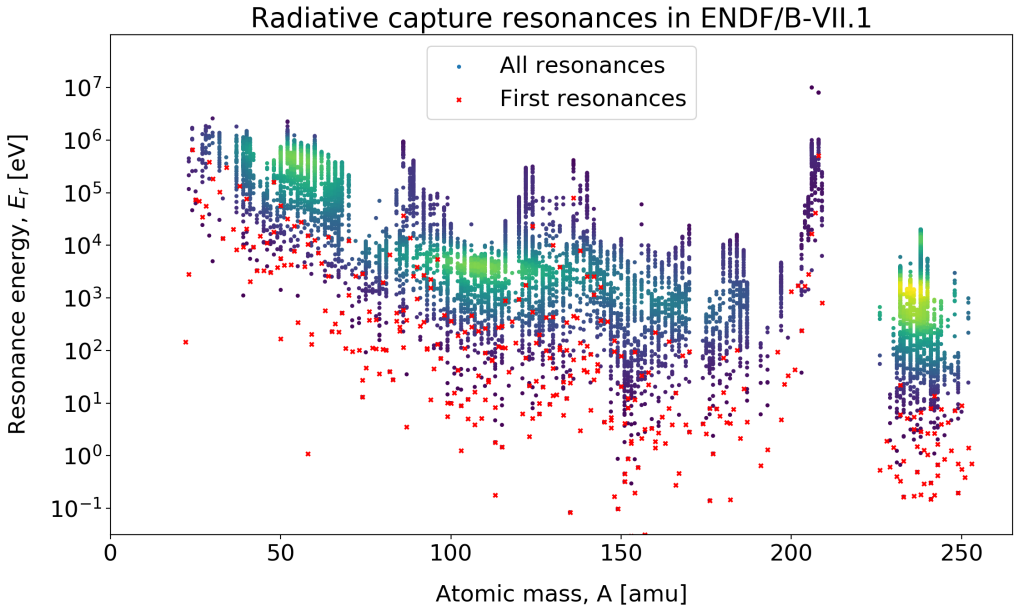

Fig. 4.2 This scatter plot shows the energy of resonances as a function of atomic mass, for the majority of resonances recorded in the ENDF/B-VII.1 nuclear data library. The data are coloured by density, with blue as lowest density, green medium and yellow highest. The density has been calculated in $\{\log_{10}(\mathrm{E}), \mathrm{A}\}$ space with a Kernel Density Estimator (KDE) approach. The first resonance for each nuclide is marked by a red cross for ease of identification.

Figure 4.2 shows the general trend for lower energy resonances with increasing atomic mass. This trend is most pronounced up to $\mathrm{A} \approx 150$. After this point, the relationship flattens off for very large mass nuclides ($\mathrm{A} > 220$) and is completely invalidated for nuclides with $\mathrm{A} \approx 210$, the region of Pb and Bi. Nuclides in this range are especially stable, close to the 'magic numbers' of stable nuclei predicted by the nuclear shell model [128]. Here, nucleons are especially tightly bound and a large energy input is required to reconfigure the nucleus. The highest density region is that of the actinides, where each nucleus has very many nuclear resonances, at comparatively low energies. Looking more generally, it can be seen that resonances are typically present for perhaps as much as 5 orders of magnitude of the energy spectrum for a given mass, a very large window. Taking all nuclides into account, resonances appear within the $10^{-1} < E_r[\mathrm{eV}] < 10^7$ range. If we assume that all resonances are equally important to consider when designing a group structure, then a fine grid might be required across this entire range, from practically thermal neutrons to unmoderated DT neutrons.

We can quantify the 'size' of a particular resonance by integrating a relevant descriptive function across some characteristic energy range, this so-called Resonance Integral (RI) is specified by equation 4.1.

$$RI = \int_{E_r-4\Gamma}^{E_r+4\Gamma} \frac{\pi\lambda\!\!\!^{-2}(2J+1)}{(2s_n+1)(2s_t+1)} \frac{\Gamma_i\Gamma_f}{[(E-E_r)^2+\Gamma_t^2/4]} dE \tag{4.1}$$

Where E is centre-of-mass energy of the system, Γ_i is the partial resonance width to decay to the initial state, Γ_f is the partial width to decay to the final state, Γ_t is the total width, $\lambda\!\!\!^{-}$ the reduced particle wavelength, E_r the rest mass energy of the resonance, J the total angular momentum of the resonance, s_n the neutron spin and s_t the target spin [80].

RIs give an indication of the resonance's contribution to reaction rate. Computing the RI for all the resonances parsed in ENDF/B-VII.1 we can identify the largest and thus the most important to correctly represent in any energy discretisation scheme. These largest resonances are plotted as a function of energy in figure 4.3. Given these data are a subset of those presented in figure 4.2, many of the same trends as visible, such as the inverse relationship between energy and atomic mass. These largest resonance data do tend to be at the lower energy of the complete set and indeed further analysis shows that 20% of the lowest energy neutron scattering resonances are also those with the largest RI for that nuclide. For radiative capture resonances, this figure is 41%[2]. As such, it is not the case that the lowest energy resonances for each nuclide have always or even more often the largest RI. Any energy group structure optimisation routine should not simply target these resonances.

[2]The difference is because neutron scattering partial widths are correlated with interaction energy, whereas partial gamma widths are roughly constant (see figure 4.4).

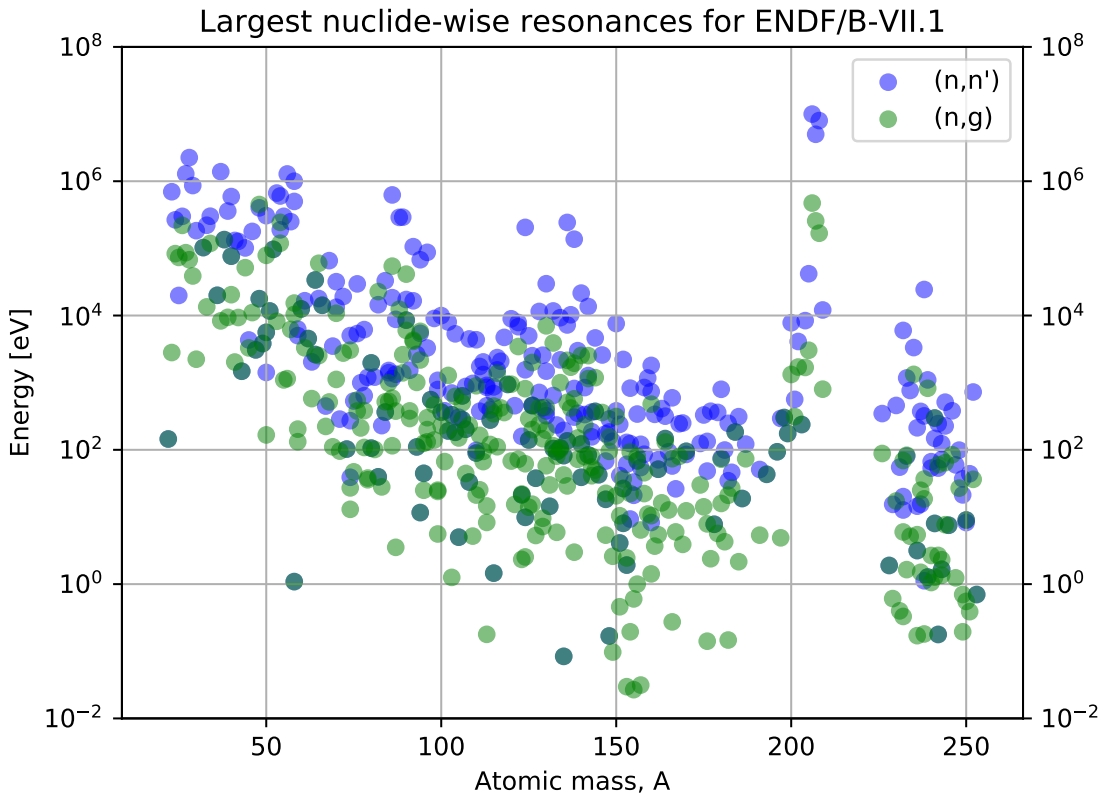

Fig. 4.3 The energy of the largest resonances for ENDF/V-BII.1 as a function of atomic mass. The 'largest' is determined by the RI as defined in equation 4.1. Two datasets are shown, RI calculated using widths for neutron emission, Γ_n in blue and for gamma emission, Γ_γ in green.

It is instructive to inspect the correlations between various parameters recorded in the assembled resonance database. As previously stated, this database contains information on some 42,000 nuclear resonances, read from the ENDF/B-VII.1 nuclear data library. The database takes the form of a row-column matrix with each row representing a resonance and each column a parameter of interest. Computing the covariance between each column-vector, it is possible to see how these parameters change with respect to each other. Normalising the covariance of any two parameters by the product of their standard deviations results in their correlation, a measure of the range -1 to +1.

The correlation matrix of selected resonance parameters is shown as figure 4.4. A is atomic mass, E_r resonance centre-of-mass energy, l neutron orbital angular momentum, J the total angular momentum, Γ_t the total width, Γ_f the fission partial width, Γ_n the neutron emission partial width, Γ_γ the γ-ray emission partial width, $\int R_n(E)$ the resonance integral for Γ_n and $\int R_\gamma(E)$ the resonance integral for Γ_γ. As previously shown in figure 4.2 resonance

energy and atomic mass are inversely correlated. Γ_n is correlated with energy, while Γ_γ is not. Given $\Gamma_t = \Gamma_n + \Gamma_\gamma + \Gamma_f$, but Γ_γ is largely constant and Γ_f is only important for a few resonances in particular nuclides, Γ_t displays largely the same relations as Γ_n.

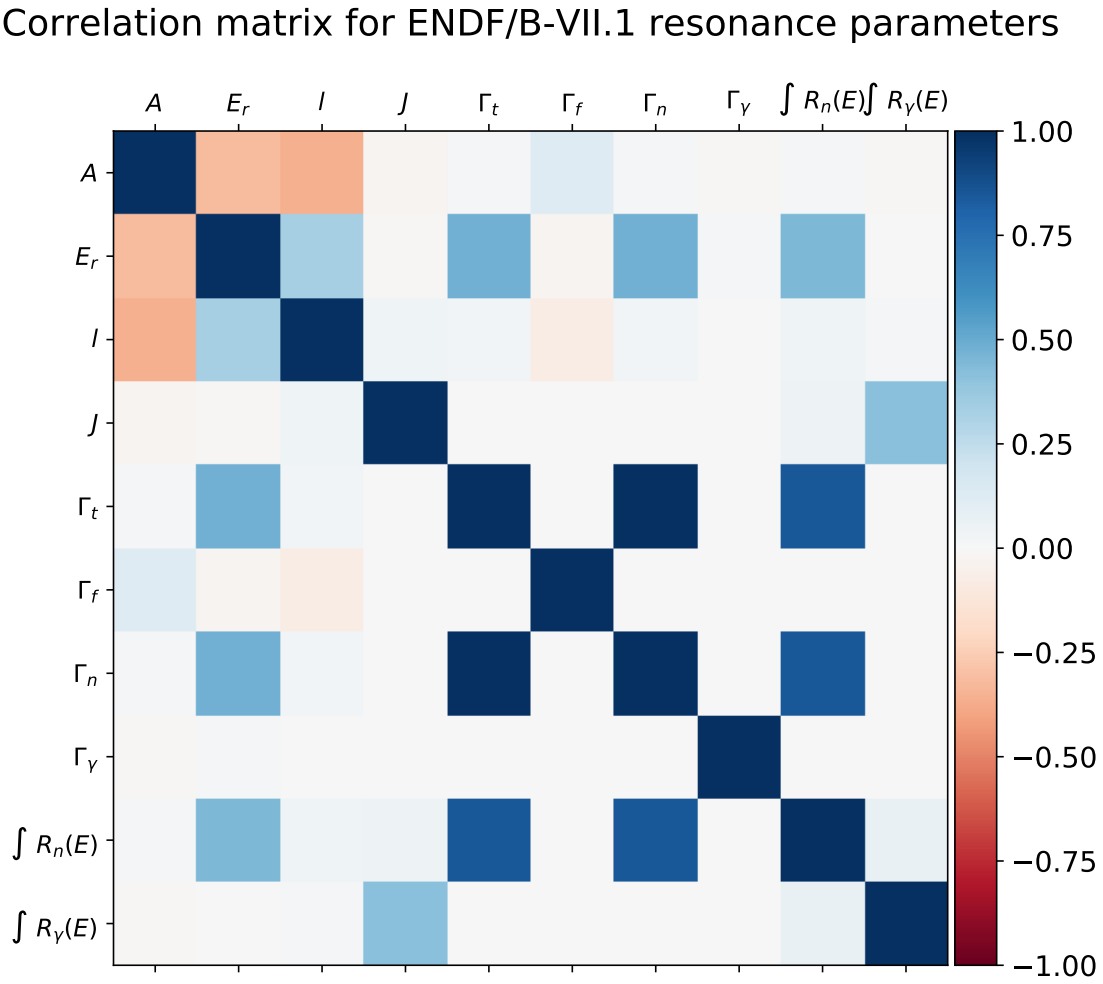

Fig. 4.4 This correlation matrix is assembled from the resonance parameter data contained within the ENDF/B-VII.1 nuclear data library. It is constructed from the covariance matrix, showing how pairs of variables vary together, but normalised such that a meaningful comparison between relationships can be undertaken. The normalisation is given here [28].

4.2.3 Self-shielding

The phenomenon of self-shielding is that whereby a particle flux is depleted in certain energy regions through absorption by the containing medium. This then 'shields' the rest of the medium from those particle energies and decreases reaction rates. If this reduction in flux is not fed back then reaction rates will be overestimated. This effect is a particular problem

where cross-sections and spectra are discretised into group-wise formats, as in deterministic radiation transport and activation-transmutation calculations. If energy groups are too large, flux depressions cannot be resolved.

It is convenient to define a number which is the ratio of the unshielded and shielded cross-sections. This 'self-shielding factor', or SSF, is of the range [0,1] where 1 indicates there is no self-shielding occurring, so any reaction rates are the simple multiplication of particle flux and unmodified cross-section. For values closer to 0, the cross-section is revised down, reducing the reaction rate. Further detail and context on computing self-shielding factors is available in [35] and [12].

Figure 4.5 shows cross-sections from several tungsten isotopes, ^{182}W, ^{183}W, ^{184}W and ^{186}W which together account for 99.88% of the atomic mass of natural W. When computing the self-shielding factors for ^{183}W in a material of natural W, it can be seen that areas without a significant background cross-section such as $140 < E[eV] < 160$, the self-shielding factors make a greater change to the unshielded cross-section. However, some resonances from different isotopes overlap in energy. Where there is a large resonance from an isotope other than nuclide in question, it elevates the background cross-section, σ_0, and absorbs particles in this energy region. An example of this behaviour is the ^{183}W resonance at 220 eV, which is overshadowed by resonances in ^{182,186}W and so does not significantly self-shield. The self-shielding factor for this energy bin is correspondingly close to unity.

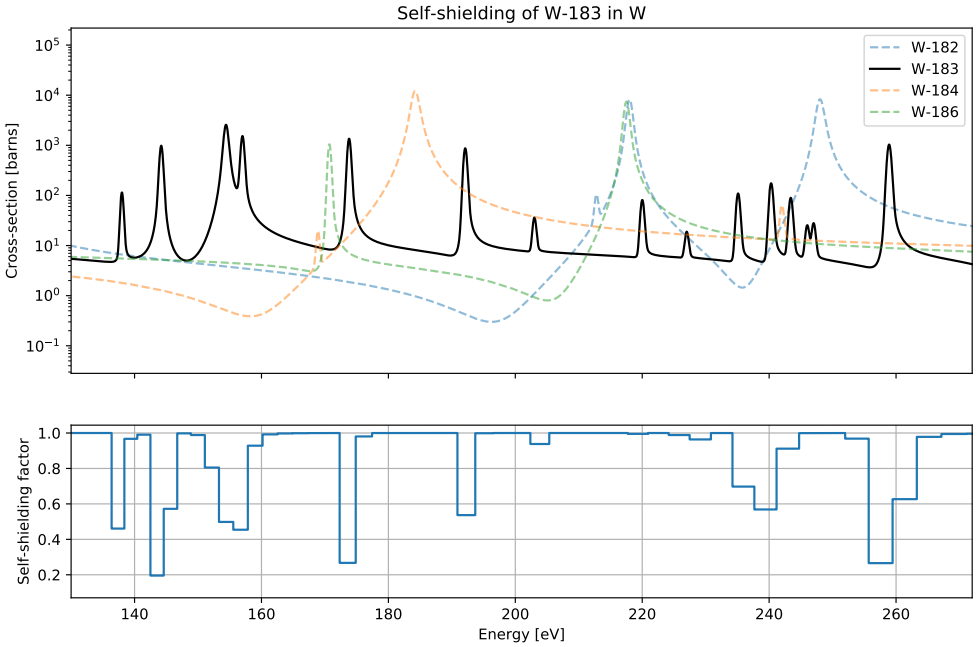

Fig. 4.5 This figure shows cross-sections and computed self-shielding factors for ^{183}W in elemental tungsten. The upper panel shows the total cross-section, $\sigma_t(E)$, for the major constituent isotopes of tungsten. The interaction data is processed from ENDF-B/VII.1 and was obtained from Shimwell's online cross-section plotting service [124]. The lower panel shows total self-shielding factors computed for group-wise ^{183}W data using the FISPACT-II code [131]. The group structure employed is a relatively fine 2,048 group with equal $\log_{10}$ spacing. Note how ^{183}W resonances are only shielded when the have a greater cross-section than other nuclides in the mixture (natural tungsten).

Calculating self-shielding factors by the above method, or other more sophisticated approaches, allows one to then take simple unshielded cross-sections, multiply then by the SSF and arrive at group cross-sections that are more appropriate for the nuclide given its environment. A convenient measure of the magnitude of this shielding is given by equation 4.2.

$$\overline{\text{SSF}} = \frac{\sum_{g=0}^{G} \text{RR}(E_g)\text{SSF}(E_g)}{\sum_{g=0}^{G} \text{RR}(E_g)} \tag{4.2}$$

An example application of group-wise self-shielding factors is shown in figure 4.6. Here self-shielding factors have been applied to a naïve, unshielded group-wise (n,γ) reaction-rate, reducing the total energy-integrated reaction rate to 0.364 of its unshielded value. Note the log-scale of the ordinate axis–the groups in the immediate vicinity of the three resonances

contribute much of the total reaction rate and significant shielding on one of these groups can severely change the result.

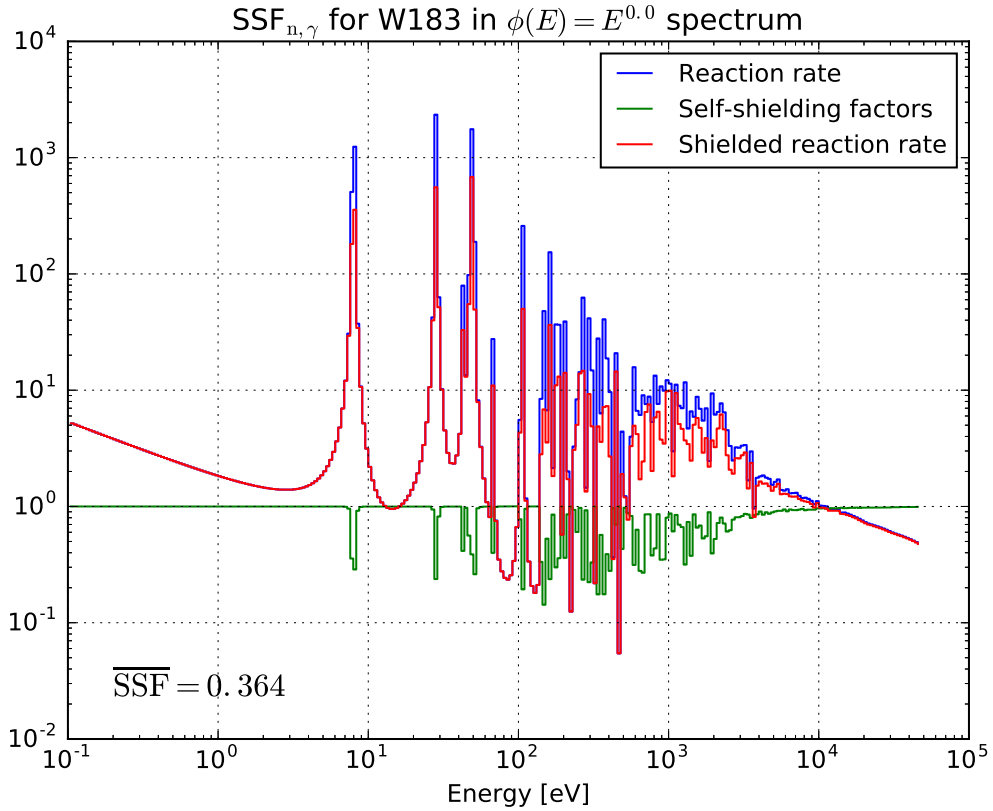

Fig. 4.6 Shown here are the unshielded and shielded ^{183}W(n,γ)^{184}W reaction rates and the associated macro-partial self-shielding factors for this reaction channel as a function of energy. The neutron spectrum in this example is a constant function of energy and so the reaction rates are a simple re-normalisation of the cross-section. The material is elemental tungsten.

The calculation of reaction rates should incorporate the effects of self-shielding. As indicated above, a resonant material of particular interest to fusion plant designers is tungsten. Figure 4.6 highlights the difficulties of determining reaction rates to a high accuracy. ^{183}W has a large resonance at 8 eV and the energy-integrated reaction rate within this and neighbouring resonances typically dominate the total reaction rate. With the coarse group structure shown in the figure, an unshielded approach to computing the reaction rate yields values several times the 'true' point-wise value. Pampin showed the importance of self-shielding for tungsten in fusion applications, finding that neglecting shielding effects can overestimate radiative capture in ^{186}W by a factor of 6 [99]. Gilbert et al.'s work built on this, sampling

flux at multiple depths in first-wall tungsten to predict accurate rates of transmutation to rhenium and osmium. Even with this more sophisticated modelling approach, total unshielded reaction rates were sometimes twice their true value [49]. The self-shielding of resonant materials other than tungsten has yet to receive much attention in nuclear analyses of fusion power systems. Examples from the fission realm also abound, using a relatively coarse (but commonplace) 175 group structure [104], individual groups for ^{232}Th cross-sections can be over-estimated by a factor of 50 [18].

Various methods exist for computing the effects of self-shielding, including the one laid out above. These methods have been validated in a variety of scenarios, especially with regards to fission calculations. However, the deployment of these methods to geometrically complex heterogeneous systems such as a fusion reactors is not a straightforward prospect. Undoubtedly self-shielding effects must be included in future studies, but reliance upon the computation of accurate self-shielding factors should be reduced where possible. Assuming $\mathrm{SSF}(\mathrm{E_g})$ is a self-shielding factor, where g is some group containing a large, self-shielded resonance–if $\mathrm{SSF}(\mathrm{E_g}) \approx 0.1$, a mere 0.01 error in this one factor results in a 10% error for the group cross-section and may dramatically change the total reaction rate.

To reduce the requirement for the accurate computation of self-shielding factors, one can adopt a higher-fidelity description of interaction and particle spectrum data. However, as seen earlier, many tens of thousands of groups are required for the accurate computation of reaction rates when used in an untargeted way. The remainder of this chapter uses concepts from self-shielding to *target* bin density resolution for optimised energy group structures.

4.3 Method

The general approach adopted here for optimising energy group structures is to use functions derived from group-wise cross-section and self-shielding data to construct a distribution that represents how poorly represented each part of the energy domain is for a given set of reactions. After modifications, this distribution is used as a bin density distribution, allowing a given number of bins to be apportioned by need throughout the energy domain.

This approach requires a robust, repeatable methodology for the generation of nuclear data on an arbitrary energy grid. Section 4.3.1 outlines this. Then, section 4.3.2 describes how a bin density distribution can be derived from self-shielding data. Finally, section 4.3.3 details the computational methods utilised to test the optimised group structures.

4.3.1 Nuclear data processing

As mentioned in chapter 2, a key to producing reliable nuclear data is having a processing architecture that is robust and repeatable [72]. For this group structure optimisation work, many kinds of nuclear data are required. The production of this data was coordinated by a series of Python modules and scripts written during the project. These launched other codes, parsing and checking outputs before inputting these to other codes. Below is a short description of the attributes and production process for each kind of relevant nuclear data.

- Group structure–The group structure is defined by an ebins file, as per FISPACT-II. These stipulate the position of energy bounds, at 7 digit precision with 2 digits for the order of magnitude[3]. The file is in descending order energy, with units of electron-volts. The ebins file is given a unique ID that is also used for all grouped data which adopt its group structure.

- Evaluated Nuclear Data Files (ENDF)–The ENDF files used for this work were those of the ENDF/B-VII.1 library [23].

- Point-wise ENDF (PENDF)–Point-wise, or PENDF files are created using NJOY [86]. Resonances were constructed with the reconr module. For all work presented here the resonances were Doppler broadened to room temperature (294K) with the broadr module.

- Group-wise ENDF (GENDF)–Group-wise data, or GENDF files were produced from the broadened PENDFs using the groupie module of the PREPRO processing code [32]. This operation takes a group structure specified by an ebins description.

- Probability tables–To enable self-shielding calculations within the unresolved resonance range, probability tables were employed (see section 1.6.2 in chapter 1). These were computed with the CALENDF code [129]. This is the most computationally expensive step of the ND generation and was parallelised to utilise multiple cores with the Python multiprocessing module [107].

- Self-shielding factors–The computation of self-shielding factors was accomplished with FISPACT II, version 3.2 [131]. Inputs include the group structure, ENDF file, group-wise interaction data, probability tables, nuclides to shield, containing material, material temperature and incident particle spectrum. Each calculation would compute the self-shielding factors for all available macro-partial cross-sections (neutron scattering, radiative capture, fission) and the total.

[3] See [31] for effects of ND numerical precision on calculations.

The target nuclide would be defined within its elemental material, i.e. ^{56}Fe would be shielded for a $^{nat.}$Fe mixture. Clearly, nuclides are present within materials of more than just their containing element, whenever they are compounded or alloyed. However, it is almost always true that nuclides are present alongside fellow isotopes. Given that the aim of this study is to use self-shielding factors to target bin resolution rather than accurately calculating a specific reaction rate or rates, this 'elemental composition' assumption was deemed acceptable, providing likely background cross-sections for the shielding calculation.

FISPACT can take an arbitrary particle spectrum as input. For the simulations described below two broad cases were used:

1. Realistic spectra, sampled from a stochastic radiation transport code (MCNP6). These were sampled in the relevant group structure and allow comparison between group-wise and point-wise calculations.
2. 'Parametric spectra' where a flux for each group structure bin was calculated from a relationship of the sort: $\log_{10}\phi(E) = m\log_{10}(E) + c$, a linear relationship in log-log space. If $c = 0$, this is simply $\phi(E) = E^m$. By varying m, the gradient of the slowing-down region can be modified and greater weight given to cross-sections (and resonances) within the low or high energy regions. This m parameter is an input to the optimisation process outlined below, in section 4.3.2.

- A Compact ENDF (ACE)–To enable testing of optimised group structures, point-wise calculations using MCNP6 are used as a reference. The underlying nuclear data used in this simulation should be the same as that used to create the discretised data. ACE files for MCNP6 were produced from the ENDF/B-VII.1 data at 294K, using the acer module of NJOY [86].

4.3.2 Group structure optimisation

Accurate nuclear simulations require cross-section data which preserve the true reaction rates even when discretised. An energy-dependent function of the reaction rate accuracy could be used to indicate where a higher-fidelity representation of the resonance would be useful. If this action could be performed for a variety of cases (materials, common spectra), this function could be used to target group structure bin density, reducing future errors. The process could perhaps be performed iteratively to converge on an ideal group structure. The group structure could be very specific (for a certain reaction channel and spectrum) or relatively general (for the whole table of nuclides, in a typical slowing down spectrum).

4.3.2.1 Towards an ideal bin density function

To assemble a distribution for bin density we define several piece-wise functions. These identify where in the energy domain self-shielding factor modifications are important. This, in turn, indicates where group structure resolution should be concentrated to converge on point-wise precision of representation.

Equation 4.3 describes the cumulative differences between reaction rates for an energy group as defined in section 4.2.3.

$$S(E_g) = \sum_{i=0}^{g} (\mathrm{RR}(E_i) - \mathrm{RR}(E_i)\mathrm{SSF}(E_i)) \text{ for } g = 0, 1, ..., G \tag{4.3}$$

$S(E_g)$ is the sum of differences between shielded and unshielded reaction rates from the start of the resolved resonance region to the g-th group. Repeating this procedure for all groups up to the final, G, assembles the cumulative difference between unshielded and shielded reaction rates. $S(E_G)$ is the last value of the cumulative sum and hence the total sum.

The cumulative, effective self-shielding distribution, $C(E_g)$ is shown by equation 4.4 and is normalised to start at 1 for $g = 0$ and reduce through the shielded energy region to $\overline{\mathrm{SSF}}$. This normalisation achieved through the division of $S(E_g)$ by its maximum value $S(E_G)$ and multiplication by the factor of $1 - \overline{\mathrm{SSF}}$. The normalisation means that effective self-shielding distributions from multiple nuclides can be compared in a way which is agnostic to the absolute values of individual cross-sections.

$$C(E_g) = 1 - \left(\left(1 - \overline{\mathrm{SSF}}\right) \frac{S(E_g)}{S(E_G)} \right) \tag{4.4}$$

This function, $C(E)$, gives a dimensionless measure of the importance of self-shielding modifications to interaction data. It allows comparison between nuclides of different cross-section values, as it is scaled between unity and the nuclide's $\overline{\mathrm{SSF}}$. Discontinuities in $C(E)$ show regions with large effective self-shielding. An example of a cumulative SSF distribution is shown in magenta in figure 4.7.

$$D(E_g) = C(E_g) - C(E_{g-1}) \tag{4.5}$$

The differences between discrete elements of the cumulative distribution provide an effective self-shielding distribution, $D(E_g)$. Equation 4.5 shows this, identifying the relative contributions from each bin to the overall self-shielding effect. It is shown as the effective self-shielding distribution in red in figure 4.7. In this example, the nuclide ^{186}W, the effective self-shielding distribution is almost entirely located within two bins, around the first resonance.

Despite the significant gap between shielded and unshielded reaction rates between 100 eV and 10 keV, this higher energy region does not contribute to a dip in $C(E_g)$ or an associated rise in $D(E_g)$ as the reaction rate values here are at least an order of magnitude less than those of the first resonance.

$D(E_g)$ can be thought of recording where the input group structure poorly represented the reaction rate. Consequently, $D(E_g)$ indicates where the group structure requires enhanced resolution.

This process of computing $D(E_g)$ can be repeated for a set of nuclides, $n = [0, 1, ..., N]$ and the distributions summed to give their combined effective self-shielding distribution, i.e. $\sum_0^{n=N} D_n(E)$. However, outside the resonance ranges, for example for E $< 10^{-2}$ or E $> 10^5$ eV, this distribution is identically zero. To enforce a minimum bin density, a constant, b, is added to all bins of $D(E)$ in the algorithm employed.

$$b = \frac{\int_{g=0}^{g=G} (\Delta E(E_g) \cdot \sum_0^{n=N} D_n(E_g)) dE}{\frac{n_{bins}}{bpd_{min}} - d} \tag{4.6}$$

The minimum bins per decade, bpd_{min} must be specified. Where d is the number of decades described by the group structure, n_{bins} is the total number of bins required of the new group structure and $\Delta E(E)$ is an array of the input group bin widths. The multiplication of $\Delta E(E)$ and $\sum_0^{n=N} D_n(E)$ is a scalar product, integrated over the energy domain.

$$\rho(E) = b + \sum_0^{n=N} D_n(E) \tag{4.7}$$

$\rho(E)$ is an array of proposed local bin densities, on the input group structure. This is then re-binned into equal 'areas' in ρ, E space, each with the value k as shown in equation 4.8. Here the right hand side shows the total energy-integrated value of the bin density distribution $\int \rho(E) dE$, divided by the number of desired bins to give k. Working from the starting energy, with k and $\rho(E)$ it is now possible to determine the new, optimised group bounds and thus the group structure.

$$k = \frac{1}{n_{\text{bins}}} \int_{g=0}^{g=G} \rho(E_g) dE \tag{4.8}$$

An example bin density distribution and the associated new group structure bin bounds are shown in figure 4.8. One can see a greater density of new bin boundaries (green) where the distribution (red) is peaked. This figure can be thought of as a nuclide-wise sum of effective self-shielding distributions, an example of which is the bottom panel of 4.7.

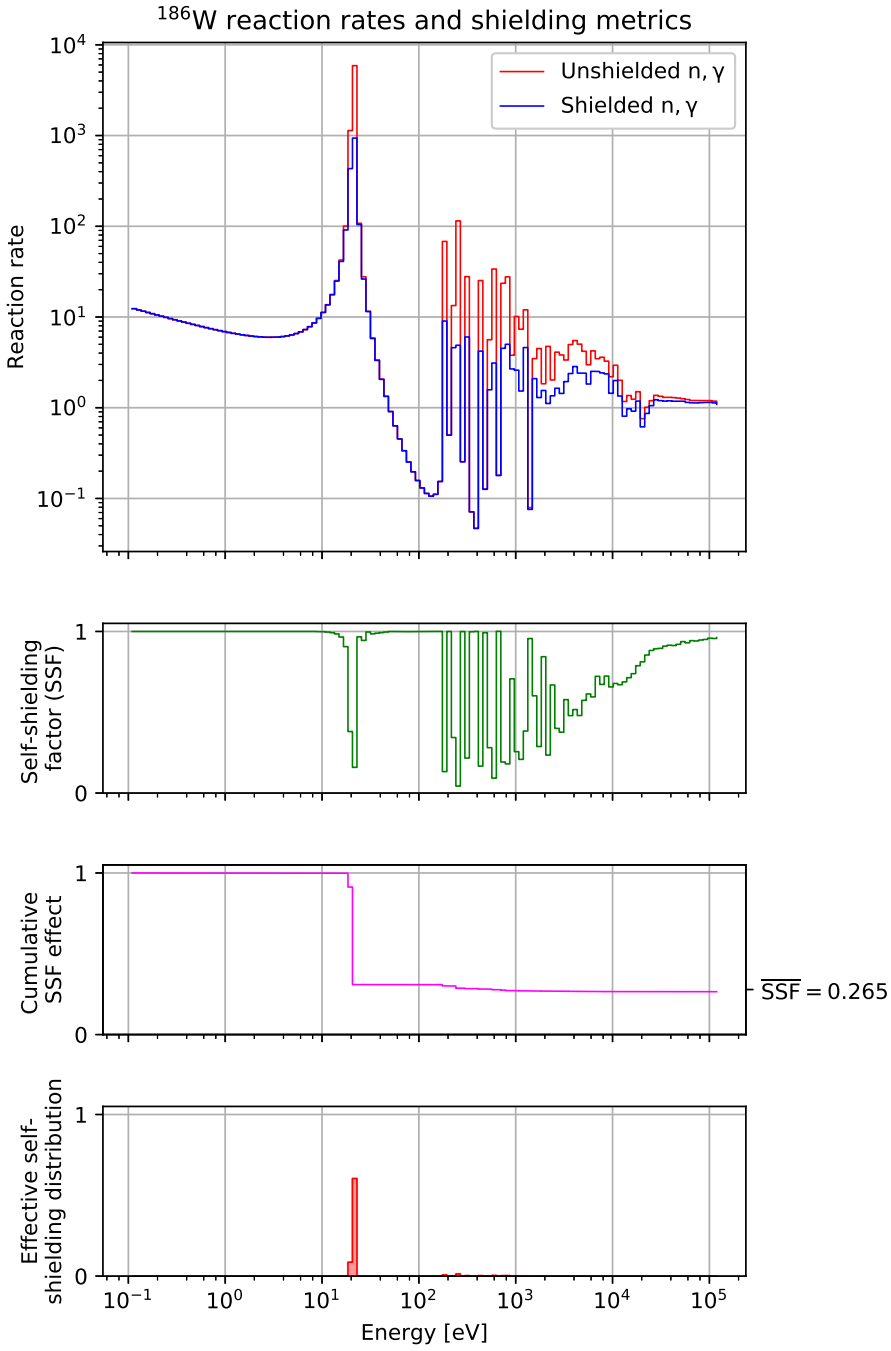

Fig. 4.7 Shown in the upper plot is a ^{186}W radiative capture reaction rate, both unshielded and shielded. The incident particle spectrum in this case is a power function of energy, i.e. $\phi(E) = E^m$ where m is some constant, in this case 0. The second plot shows the calculated self-shielding factors as a function of energy for this reaction, given an elemental W material composition. The third plot is the cumulative self-shielding distribution, $C(E_g)$ from equation 4.4. The fourth and final plot is of, $D(E_g)$, from equation 4.5. All of these quantities are for ^{186}W only.

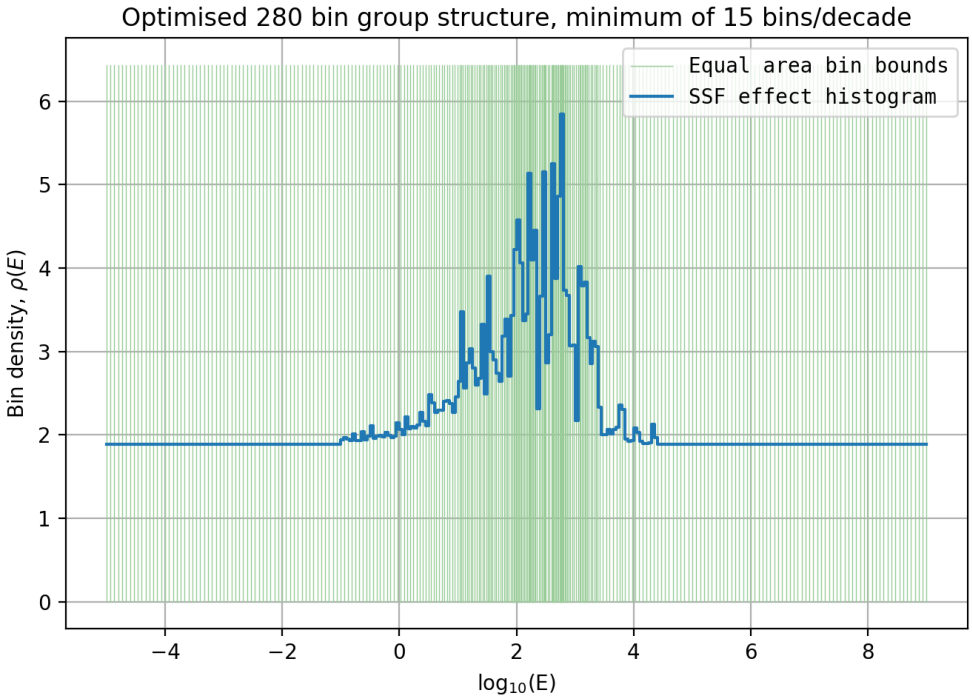

Fig. 4.8 This plot shows the bin density, $\rho(E)$ distribution as per equation 4.7. This is the summation of effective self-shielding distributions for all targeted nuclides, plus some constant to give a minimum bins per decade outside of the resonant region.

The above process provides a method to determining where a group structure provides a poor representation of the spectra information and nuclides' nuclear physics. It can be used to apportion a given number of bins to construct a new group structure, as shown in equation 4.8. Figure 4.8 shows a bin density function. The 280 log-spaced bin boundaries for this distribution constitute the original group structure. An optimised 280 bin group structure is determined by re-binning the $\rho(E)$ distribution and is shown by the green lines. There are a minimum of 10 bins per decade.

However, one need not stop at this point. With these distributions, other optimisation methods are conceivable. The process can be repeated, in an attempt to converge on an optimum group structure. Or, one might start with a very fine representation, with many tens of thousands of bins and attempt to coarsen based on the effective self-shielding distribution. These constitute future avenues of research for this group structure optimisation method.

4.3.3 Computation

The calculation of effective cross-sections were carried out using three different techniques:

1. Multi-group FISPACT-II unshielded collapse
2. Multi-group FISPACT-II shielded collapse
3. Point-wise MCNP6 Monte Carlo Estimator (MCE)

Methods 1 & 2 discretise the incident spectrum, over a $\phi(E)$ and $\sigma(E)$, taking the inner product of two vectors of constants,

$$\vec{\phi} \cdot \vec{\sigma} = \sum_i \phi(E_i, E_{i+1}) \sigma(E_i, E_{i+1}),$$

as the spectrum-averaged, effective cross section. However, as mentioned in section 4.2.3, for each group interval, $(E_i, E_{i+1}]$ the true functions $\phi(E),\ \sigma(E)$ are not constant. Regions with resonances in the cross section cause local decreases in the neutron flux, resulting in significant 'self-shielding' errors with the multi-group method. To address this, self-shielding factors, SSF(E_i), may be used to account for these resonance effects. An incident spectrum, material inventory and nuclear data are input to FISPACT-II. The spectrum and interaction nuclear data are then collapsed to generate one-group effective cross-sections. With the addition of probability table data generated with CALENDF-2010 [129], SSF(E_i) are calculated for both the resolved and unresolved resonance regions.

Method 3 employs nuclear cross-section data with tens of thousands of points per reaction channel (ACE files), interpolating between them to approximate a continuous energy treatment. The effective radiative capture cross-section can be determined from the point-wise calculated (n,γ) reaction rate as equation 4.9.

$$\sigma_{\text{eff}} = \frac{\int RR_{n,\gamma}(E)dE}{\int \phi(E)dE} \tag{4.9}$$

The effective cross-sections calculated for this study were a series of radiative capture reactions in metals. These reactions often have a large fraction of their total reaction rate within the neutron slowing down region between thermal and fast energies. Hence, an accurate resonance treatment is required for accurate results.

To test optimised group structures, both simple examples and faithful real-world radiation transport models have been used. These models allow the sampling of neutron spectra with MCNP6 for the computation of point-wise effective cross-sections, but also for the computation of effective cross-sections using multi-group methods with FISPACT-II. The following section briefly discusses the radiation transport geometries and models.

4.3.3.1 Convergence study

To explore the benefits of group structure optimisation by the methods outlined above, and to inspect the effects of bin count on accuracy, a convergence study was carried out. This utilised a simple geometry of concentric spheres as shown in figure 4.9. The concentric shells of material are stainless steel, water and tungsten. The steel and water act to locally depress and moderate the spectrum. There is a point neutron source of 14 MeV situated at the centre of the nested spheres.

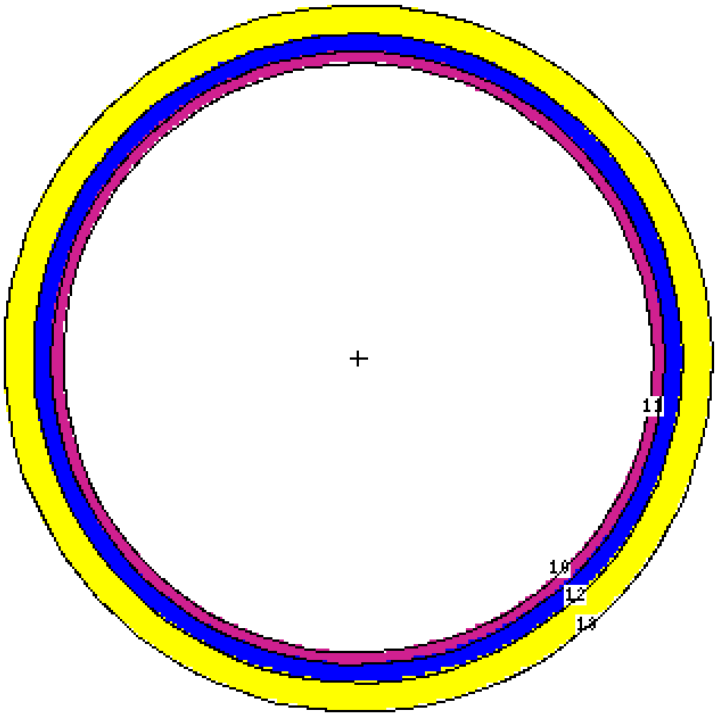

Fig. 4.9 A 2D slice of a simple MCNP test geometry, a series of concentric shells of material. These shells are steel, water and tungsten from inside to outside. There is a 14 MeV point source located in the centre.

Tallies sample the neutron spectrum in a variety of energy group structures within the outermost, tungsten shell. These are for six bin counts, $n_{\text{bins}} = [64, 128, 256, 512, 1024, 2048]$, each with a logarithmically spaced and an optimised version, giving 12 group structures in total. The group structures were optimised for all naturally occurring tungsten isotopes, with a flat spectrum. In addition to the multi-group tallies, there is also a single point-wise Monte Carlo estimator in the MCNP model, sampling the reaction rate of ^{186}W(n,γ)^{187}W.

4.3.3.2 JET activation foils

This study again tests various group structures against a point-wise reference. Two of the test group structures are optimised for a particular set of nuclides, in an attempt to accurately calculate reaction rates for these nuclides. However the radiation transport geometry for this example is more complex than used in section 4.3.3.1.

The radiation transport geometry utilised in this study is the JET tokamak located near Abingdon, Oxfordshire. Octant 8 of the tokamak has previously housed the Long Term

Irradiation Station (LTIS) for exposing activation foils to the JET neutron field. This facility is used for experimental work where foils are irradiated over the course of an entire campaign of plasma shots. After irradiation, techniques such as gamma spectroscopy can be used to determine the presence of nuclides generated by neutron activation. The relative and absolute strength of gamma emission from these nuclides can be used to unfold the neutron spectrum which was incident upon the foils.

In this work, the LTIS MCNP model [79] was integrated into a reference JET model. The LTIS was then populated with foils of the materials noted in table 4.1. The general layout of the JET device and the location of the LTIS holder is shown as figure 4.10. The neutron source used within the JET model was a custom written CCFE parametric source routine. This source takes plasma parameters and computes likely neutron emission distributions in space, energy and angle. This study used default parameters for a JET DT plasma shot.

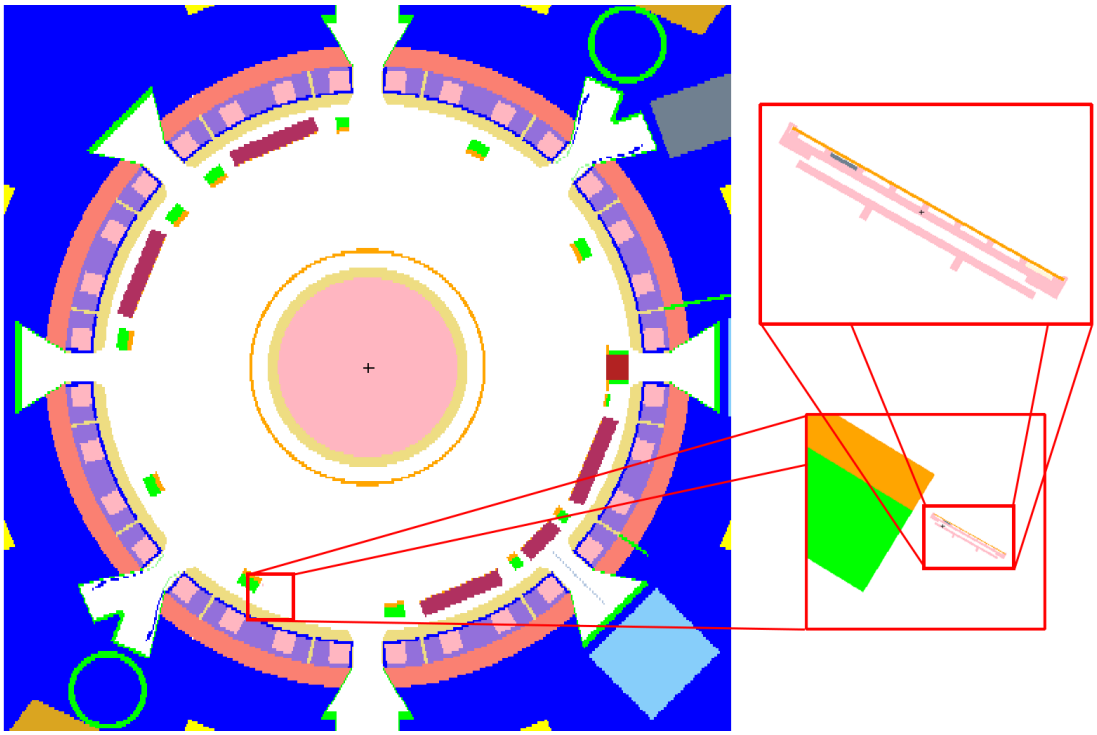

Fig. 4.10 The figure shows a plan view of the JET tokamak in Oxfordshire, UK and the location of the LTIS foil holder. The main panel is approximately 12m across, displaying the major radius, R = 3m of JET. The ports for each octant are visible as openings in the purple vessel. The LTIS foil holder is located in the lower left, in octant 8. It is shown at greater magnification in the inset panels. This figure is assembled from various MCNP plotter views.

The reaction channels chosen for study are listed in table 4.1. The nuclides listed are found in the designs of fusion power systems; molybdenum an important steel alloying element and tungsten is a current first choice for a plasma facing high-heat flux material. Each element to be tested was represented by a foil of this material held within the LTIS model.

Reaction	$E_{\text{first res.}}$ (eV)	E_{URR} (eV)
$^{95}\text{Mo}(n,\gamma)^{96}\text{Mo}$	44.7	5.87×10^4
$^{182}\text{W}(n,\gamma)^{183}\text{W}$	4.16	9.91×10^4
$^{186}\text{W}(n,\gamma)^{187}\text{W}$	18.8	1.21×10^5

Table 4.1 Shown above are the reactions simulated in this study. $E_{\text{first res.}}$ indicates the energy of the first resonant peak in the interaction cross-section. E_{URR} defines the end of the resolved resonance range (RRR) and the start of the unresolved resonance range (URR) where experimental energy resolution is insufficient to resolve individual resonances.

A variety of (n,γ) effective cross-sections are calculated for nuclides in the foil materials. Each reaction is calculated for 5 different group structures, some of these group structures are optimised for the materials present in the simulation, while others are standard group structures in common usage today. The group structures sampled in the LTIS foils are listed in table 4.2. Two groups, 280 and 650, generated by the process outlined above were optimised for all stable nuclides of the following common engineering elements: Fe, W, Mo, Nd, Sn, Zr, Cu, Co and Ta. This means performing the nuclide-wise sum of equation 4.5 for all naturally occurring isotopes of these elements before re-binning this distribution into equal areas, and using the boundary of these areas as the new group structure bin bounds.

n_{bins}	E_{min} (eV)	E_{max} (eV)	bpd_{res}	Description
175	10^{-5}	1.96×10^7	5–22	VITAMIN-J
280	10^{-5}	1.00×10^8	11–73	Optimised coarse
315	10^{-5}	1.94×10^7	24–50	TRIPOLI
650	10^{-5}	1.00×10^8	37–101	Optimised fine
709	10^{-5}	1.00×10^9	50	CCFE

Table 4.2 Comparison of group structures tested, noting their bin counts, n_{bins}, the energy range over which they are defined and the bins per decade, bpd_{res} they employ in the resonant region.

The spectrum used as input to the optimisation process is pictured in figure 4.11. Here the flux is shown in neutrons per source neutron per lethargy interval, ΔU. Lethargy, U, can be defined as $ln(\frac{E_0}{E})$ where E_0 and E are the starting and current energy, respectively. The lethargy interval ΔU is $U_2 - U_1 = ln(\frac{E_1}{E_2})$. Dividing each flux bin by its lethargy width allows easier comparison of flux values across a wide energy range.

The spectrum was sampled from a short initial run of the problem, determining roughly what shape the slowing-down region has in this problem. The exponent of the fitted relationship $\phi(E) = E^{0.192}$ is used to characterise the spectrum as input to FISPACT-II for calculating self-shielding factors prior to group structure optimisation. It effectively acts as a weighting function, determining whether lower or higher energy resonances will contribute most to the reaction rate (and so where reaction rate errors would be most problematic).

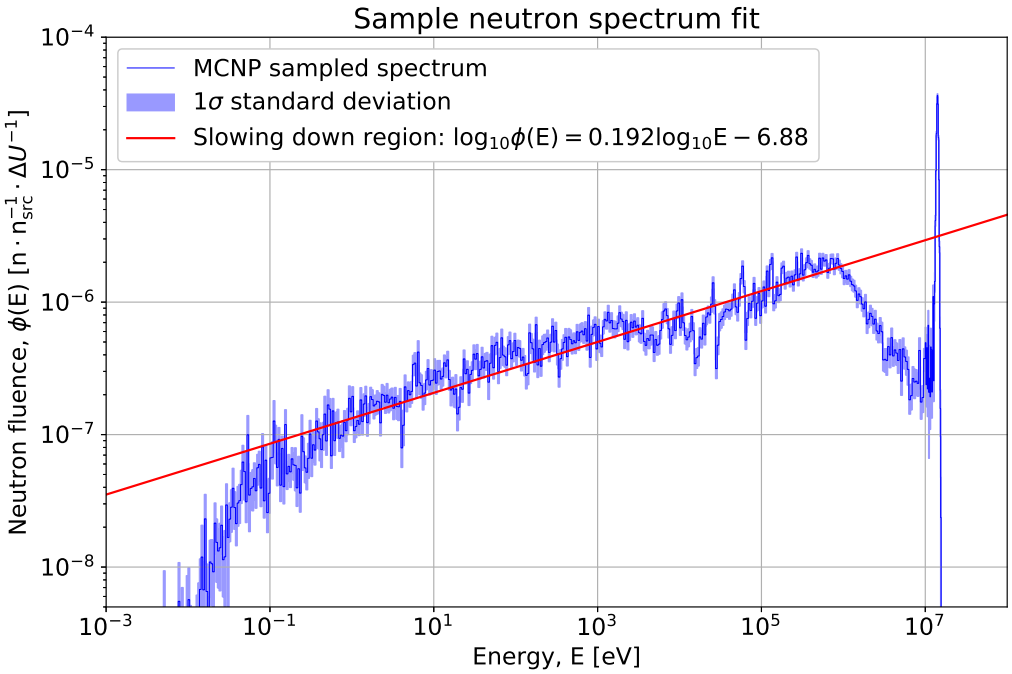

Fig. 4.11 The approximate neutron spectrum in the steel of the LTIS foil holder. This spectrum was obtained a short initial MCNP run. Statistical uncertainty from MCNP sampling is shown as light blue shading around the dark blue mean value. The slowing-down region between 1 eV and 100 keV has been fitted with the following relationship: $\phi(E) = E^{0.192} + 10^{-6.88}$. Fluence is given per source neutron, n_{src}, per lethargy interval, ΔU.

4.4 Results & Discussion

This section details the results from both test cases, a simple set of concentric shells and a more sophisticated radiation transport model of the JET tokamak. Both cases utilise utilise traditional log-spaced group structures and optimised ones. They also both probe the effects of self-shielding.

4.4.1 Convergence

For the first study 12 group structures were tested against a point-wise reference. The effective one group cross-sections of radiative capture in ^{186}W were compared with a Monte Carlo estimator using fine and interpolated data. The comparative results are shown in figure 4.12.

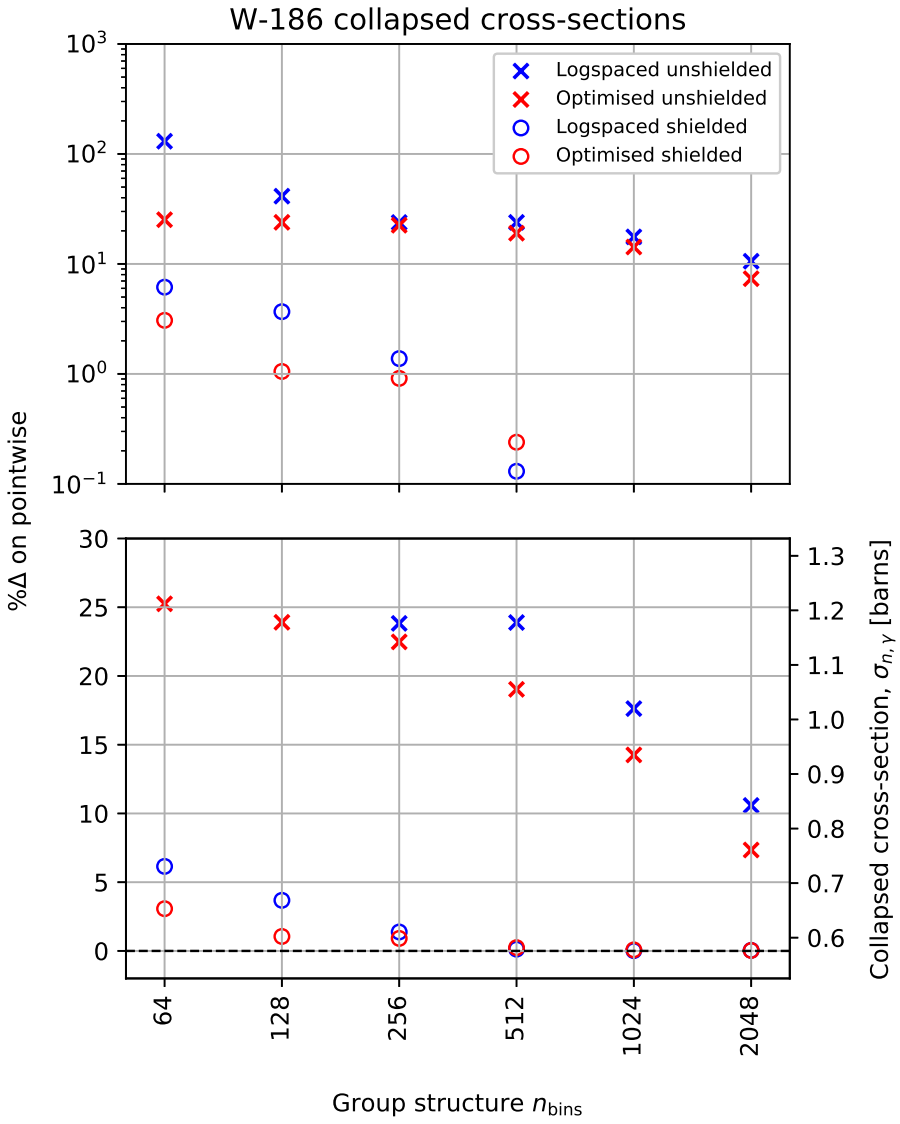

Fig. 4.12 This figure shows the results from the convergence study in tungsten. The two panels both show the same information, that is the percentage error between the collapsed multi-group cross-section and the point-wise result. The upper panel ordinate axis is log-scaled to show the breadth of results. For $\%\Delta < 0.05$, some of the shielded data was returned as marginally negative and so could not be displayed on the log-scaled plot. The lower panel adopts a linear-scaling. This panel also has a secondary axis to display the one-group cross-section values in barns and a dashed line to indicate the point-wise reference value. The bin count for the group structures is the abscissa axis for both panels. Blue markers indicate an energy group with a simple log-spaced arrangement while red markers indicate the group has been optimised for representing tungsten data. Crosses are unshielded results, while circles are shielded.

Figure 4.12 shows that increasing the bin count of a group structure will provide a more accurate description of spectral and cross-section information, as we have already seen at the start of this chapter in figure 4.1. We can also see in figure 4.12 that the multi-group computed results which were closest to the point-wise estimate were those which employed self-shielding factors in the calculation. For shielded results, by 512 groups the error to point-wise was $< 1\%$ whether the groups were optimised or log-spaced. This is a relatively small error, as the uncertainty on the ^{186}W(n,γ)^{187}W reaction is $\geq 5\%$ at all energies in the TENDL2017 ND library. The next most significant factor was the optimisation process. Here we see an improvement in performance in both the unshielded and shielded cases, however the difference is most stark for unshielded results, where with 64 groups, the error was reduced from 130% to 25%. For a finer group, 2048, the error was reduced from 10.6% to 7.34% for unshielded results.

4.4.2 JET activation foils

A set of neutron spectra from the LTIS assembly are as shown in figure 4.13. These particular spectra are sampled from within the tungsten foil. Resonances both within the tungsten and the surrounding materials have resulted in flux depressions. An expanded view of the ^{186}W 18 eV resonance is shown as figure 4.14. The degree to which the resonances are resolved by the different group structures is a function of the bin resolution around the resonance. The 709 group, with a standard 50 bins per decade reaches down into the 18 eV resonance, recording a fluence of 3.5×10^{-9} whilst the optimised group structures, 280 and 650 record 2.32×10^{-9} and 1.26×10^{-9}. The standard group structures 175 and 315 which are much coarser in this region record a minimum fluence more than an order of magnitude greater than the 650 group structure.

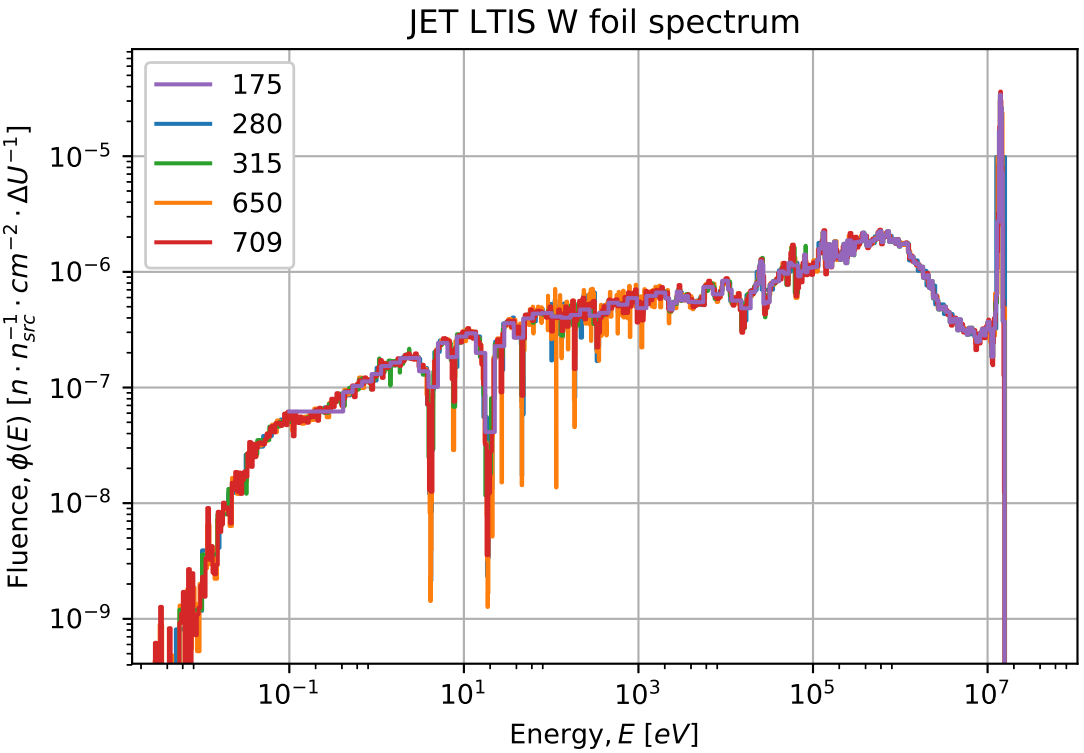

Fig. 4.13 Typical DT neutron spectrum within the LTIS activation foil holder. The fluence of each bin has been divided by the lethargy width for that bin, ΔU. The 14.1 MeV neutron emission peak from the DT plasma is clearly visible, along with a typical slowing-down spectrum, filled with flux depressions. The spectrum is generally 'hard', with little thermalisation as a consequence the foil's proximity to the plasma. The same neutrons have been binned according to various group structures, from 175 to 709 groups as listed in table 4.2.

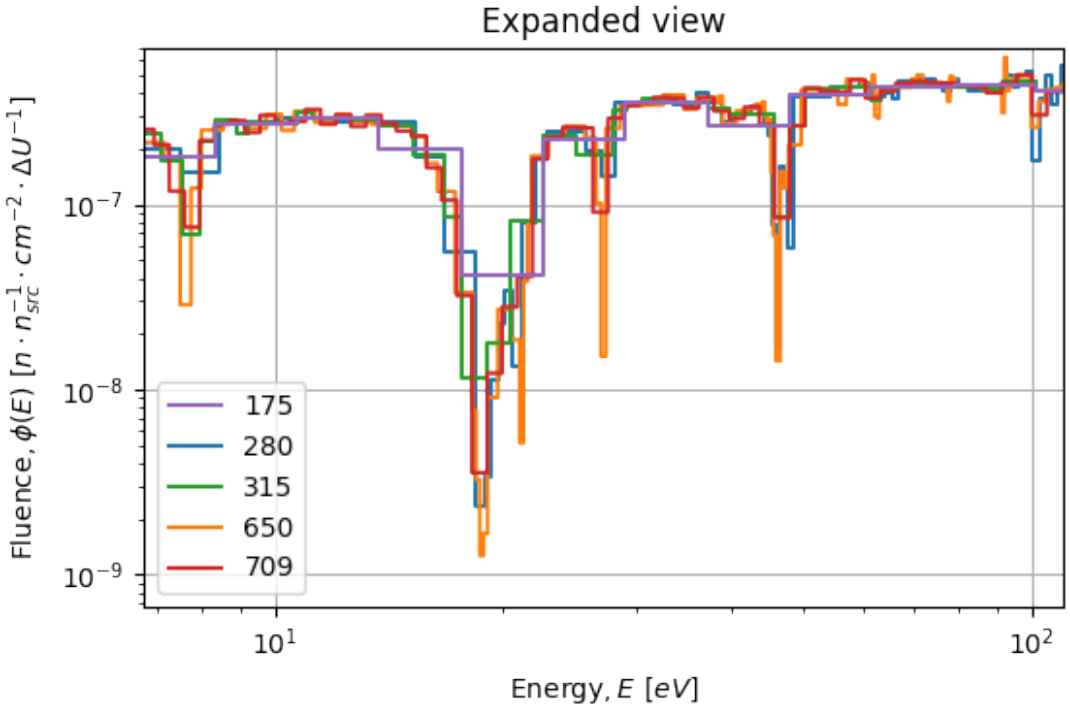

Fig. 4.14 Magnified section of spectrum showing how progressively finer group structures resolve flux depressions in a neutron spectrum. The flux depression corresponding to the ^{186}W 18 eV resonance is clearly visible, along with a variety of other resonances from ^{186}W and other W isotopes.

Figure 4.15 shows a comparison of collapsed cross-sections for ^{182}W. The abscissa lists each group structure tested. A pair of data points exist for each group structure, one with the use of self-shielding factors (shielded) and one without (unshielded). The abscissa order is a sorted such that the unshielded values decrease from left to right. The left ordinate indicates the absolute cross-section value in barns, while the right ordinate gives the cross-section as a percentage change over the point-wise (reference) value, i.e. +100% means the multi-group calculation has over-estimated the cross-section and therefore reaction rate by a factor of 2.

As would be expected, in a simple, unshielded case having more groups and therefore a larger number of groups per energy decade tends to give multi-group results closer to the reference point-wise result. The 709 group persistent performs better than the 175 group, for instance. Flux depressions are better resolved with narrower groups and therefore the reaction rate resulting from resonances is not overestimated to the same degree. However, the 650 optimised group, and even the 280 optimised group perform better than the 709 group in the unshielded regime. Their narrow groups are targeted where most useful, around the resonances which were poorly described by a equal logarithmically spaced group structure.

The optimisation process has, to some degree, worked. This optimisation was for 44 nuclides of 9 elements and yet has still conferred an advantage.

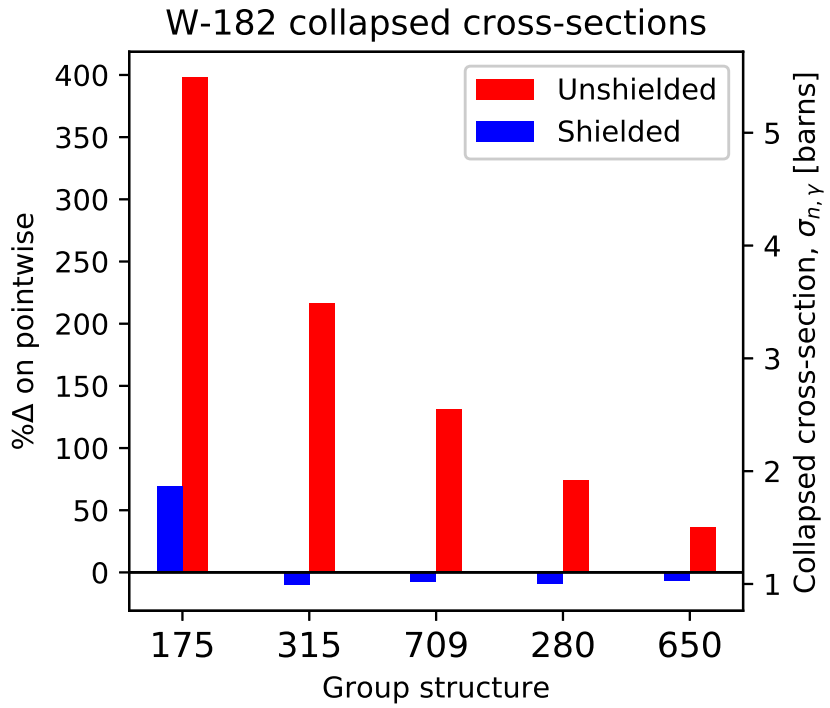

Fig. 4.15 Comparison of different calculation methods for radiative capture in ^{182}W under JET neutron irradiation.

With regards to individual groups, the 175 group used unshielded performs very poorly, overestimating the reaction rate by at least a factor of 3 in all cases. This group should be used with caution. As explored in chapter 3, it rarely suitable for heavily thermalised spectra because of its very poor resolution in the thermal energy region. However, this is not the cause of its inadequacy here as the spectrum is still very hard, having penetrated only the thin LTIS to interact with the test foils. The 650 optimised group is always within 50% of the point-wise result and is the best performing group. As might be expected the shielded results show far less variation between groups. 280, 315, 650 & 709 demonstrate largely similar behaviour when used in conjunction with self-shielding factors. The shielded results reinforce the utility of self-shielding factors. When possible, their usage can dramatically improve performance for multi-group calculations. 175 does again over-estimate, between 0% and 120% depending on the nuclide.

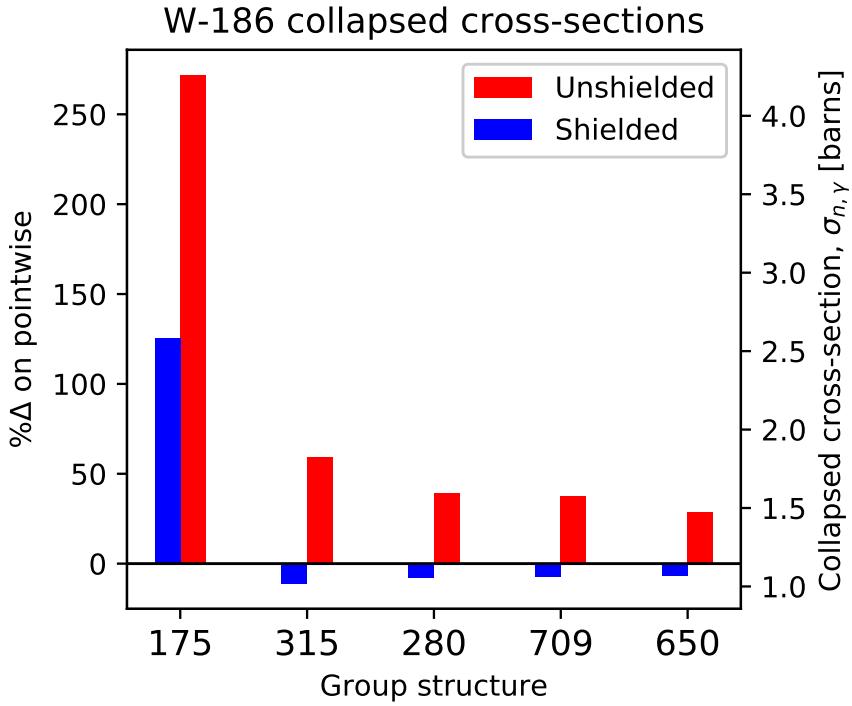

Fig. 4.16 Comparison of different calculation methods for radiative capture in ^{186}W under JET neutron irradiation.

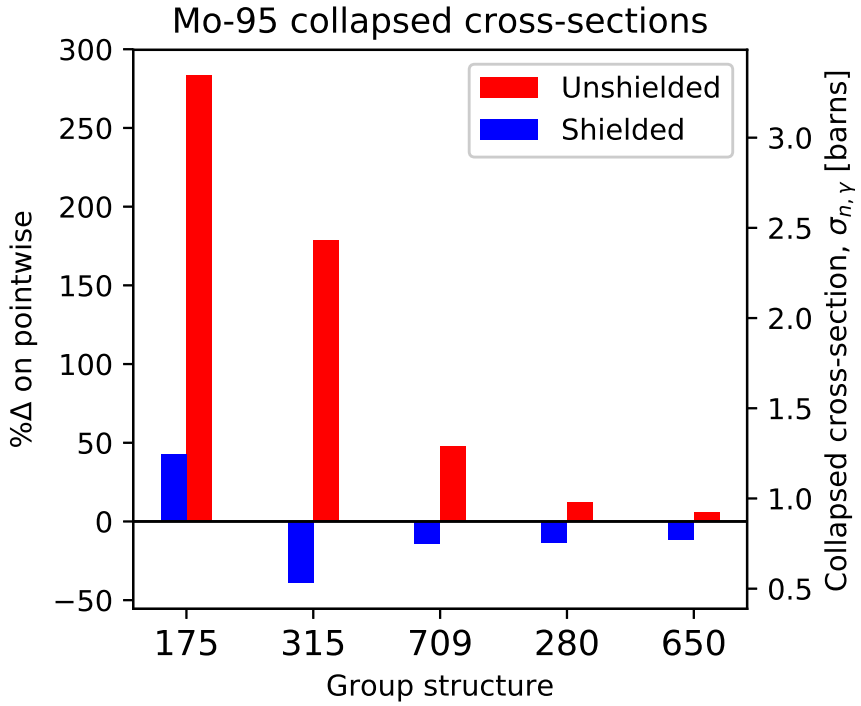

Fig. 4.17 Comparison of different calculation methods for radiative capture in ^{95}Mo under JET neutron irradiation.

The ^{95}Mo (figure 4.17) results indicate over-shielding, that is self-shielding factors over compensating for a poor description of the spectrum. This is particularly pronounced for the 315 group in ^{95}Mo, where the shielded result is 40% less than the point-wise result. This is possibly due to the application of one set of self-shielding factors over the entirety of the foil depth. Although only 500 microns, the compound geometric effects are, it seems, essential and are the subject of further work.

4.5 Conclusions

This study has investigated how an incident spectrum and nuclear data library discretisation in energy (i.e. the group structure) effects the accuracy of reaction rate calculations. This has relevance for deterministic radiation transport and activation-transmutation calculations.

A background in nuclear resonant behaviour and self-shielding was given. Nuclear data processing infrastructure was created to produce nuclear data on an arbitrary energy grid. It was then theorised that it might be possible to optimise the discretisation of the energy domain for nuclear analyses by using self-shielding theory to identify which energy regions

were poorly described. A novel method for assembling a bin density distribution, $\rho(E)$ was described in section 4.3.2. This distribution in energy was re-binned into equal-area segments, with the boundaries of these segments becoming the new optimised bin bounds. This allocates more bins into areas that were previously poorly described.

Testing was undertaken to examine the performance of these optimised group structures. Of the testing presented here, this was initially using groups optimised for only a narrow range of nuclides (for instance, only natural tungsten). Testing of optimised and log-spaced groups in a simple case showed the benefits of both correcting for self-shielding and optimisation. For a small number of groups, 64, the benefits were greatest, with a reduction in error from 130% to 25% by using an optimised group.

Subsequently, a pair of optimised group structures (280 & 650) were developed using the same method. These group structures were optimised for a characteristic slowing-down spectrum and for the naturally occurring isotopes of the following elements: Fe, W, Mo, Nd, Sn, Zr, Cu, Co and Ta. These group structures were compared with three 'industry standard' group structures (175, 315 & 709) in a real-world scenario: the Long Term Irradiation Station (LTIS) in the JET tokamak. Results from this testing showed the accuracy of multi-group reaction rate calculations is severely hampered by using the legacy VITAMIN-J 175 group structure. In the context of this activation foil problem, coarse group structures can give reasonable results if used in conjunction with modifying self-shielding factors. If used, self-shielding factors must be applied carefully, with potential geometric shielding effects considered. If used inappropriately, over-shielding may result, which may not be conservative depending on the context.

For the shielded results presented as figures 4.15 to 4.17, the optimisation process confers little advantage over a logarithmically spaced group structure. However, unshielded multi-group calculations using an optimised group structure with 280 bins more closely approximated the point-wise result than 709 bin unshielded multi-group calculations in the majority of cases examined. As such, the optimisation process outlined above is beneficial primarily in unshielded cases. This is typically how many analyses are conducted, with only individual reaction channels, if any, given a full shielded treatment, and the majority of reaction rates computed with an unshielded, naïve approach.

Additional work not presented here has explored the potential benefits of iteratively applying the optimisation process but has not yet returned any conclusive results. Suggested further work could include improving the optimisation algorithm and exploring the performance of optimised group structures in a wider array of circumstances, for a greater set of nuclides. It would be enlightening to see if optimising for a very large nuclide set (say, all

stable nuclides) negates the benefits of optimisation witnessed for a relatively constrained nuclide set.

Currently, nuclear analysts use nuclear data binned to a very general purpose grid, or even one devoid of 'design'. Nuclear data is only easily available in these long-lived group structures. Given a robust nuclear data processing infrastructure, it would not necessarily be difficult to generate nuclear data on a bespoke, optimised energy grid prior to some piece of nuclear analysis work. Processing a whole nuclear data library (all stable nuclides), including probability table data can be done in a day on a laptop, or within an hour on a cluster. With the recent 'open-sourcing' of the NJOY processing code to NJOY21 [82], steps in this direction are already underway.

Chapter 5

Concluding remarks

This thesis has explored the causes and effects of several uncertainties in nuclear fusion reactor analysis and operation. It began by describing motivations for developing controlled nuclear fusion power, by situating new power generation technologies within a broader environmental and socio-political context. A growing world population demanding improving material living standards has increased global power demand to double its value of 40 years ago. Human-induced climate change, deteriorating air quality and a desire for greater energy security have helped spur the development of low-carbon electricity production technologies such as renewables, GEN-IV fission and fusion to meet contemporary and future demands for power.

Whilst the body of nuclear fusion knowledge has massively accumulated since the process was first discovered, constructing and operating a working fusion reactor, as opposed to an experiment, is still yet to happen. However we are now in the final stages of research and development before such a reactor is constructed and it is tremendously important to identify and quantify any risks to the realisation of controlled nuclear fusion as an economical, practical electricity generation scheme. Many of these risks are political and organisational, but as the introductory chapter alluded to, many are technical, engineering challenges. Hopefully none are physical impossibilities.

Many of the technical risks are amplified by our inexact knowledge of them. A tremendous amount of work is currently undertaken to model the performance of current and future fusion devices, simulating performance and estimating parameters of interest. However, these parameters are always subject to some sort of uncertainty, whether it is quoted or not. Uncertainty in powers, Mean Time Between Failures (MTBF), particle fluxes, cost of electricity, Tritium Breeding Ratios and more. To decrease these uncertainties, to narrow their distribution around the mean value, is to reduce the maximum potential risk associated

with them and to give room for manoeuvre in the trade-offs that come with engineering a functioning system.

The work which comprises this thesis has investigated several sources of uncertainty of pertinence to the development of nuclear fusion as a power generation scheme. First, a stochastic, sampling technique known as Total Monte Carlo was used to explore the effects of nuclear data uncertainty in the TBR for a future fusion power plant design, DEMO. This technique has never before been used to estimate uncertainty on TBRs. Investigating the contribution of uncertainty from lead nuclear data, many radiation transport simulations of the DEMO device were performed, each tallying the TBR and sampling different lead nuclear data. This work used the TENDL2015 nuclear dataset, with fully correlated cross-channel behaviour for the reaction channels, angular distributions and other variables. The results of the work were to determine the standard deviation of the HCLL DEMO TBR due to lead, 1.2% of the mean value. The simulated TBR distribution was not normally distributed, instead it had a negative skewness–a low-value tail. As a result, 5.8% of the TBR distribution was less than unity. This only serves to reinforce the importance of higher-order moments in parameter probability distributions and the value of TMC style methods. Generally, where parameter mean values are close to the limits of some operational range, one should seek to know the shape of that parameter distribution, not just the extent. Whilst a TBR in a liquid-metal breeder blanket could potentially be tailored through on-line ^{6}Li enrichment, this is not the case for ceramic type breeding concepts. For those, an overestimated TBR could be a costly mistake. After sampling the TBR distribution, the relationships between fundamental nuclear parameters and the TBR were investigated, with a handful of local Optical Model Potential (OMP) parameters responsible for most of the variation in TBR. In terms of future research, the determination of uncertainty contributions to TBR in lead blankets from other nuclides such as ^{16}O, ^{56}Fe and ^{182,183,184,186}W would be a worthwhile effort. Advances in theory or new nuclear physics experimental data could help refine our models of lead nuclei and their behaviour, reducing uncertainties in lead based blanket designs.

The subsequent chapter looked at a particular modelling approximation, spatial homogenisation, where heterogeneity in material composition is artificially reduced by replacing many realistic materials with mass-conserving mixtures of materials. The effects of this approximation on radiation shielding were investigated. First the basic theory of radiation shielding was introduced, before a description of a comparative method for determining the discrepancy between heterogeneous and homogeneous modelling approaches. This method was employed to analyse how the approximation affects calculated dose rates for parameters relevant to the ITER tokamak. In one scenario, the on-load dose rate from D-T fusion was calculated on the far side of a reinforced concrete wall similar to the ITER bio-shield. The

discrepancy induced by the homogeneous approximation was a function of wall thickness, and attained a maximum of a 22% underestimate of the neutron dose. This underestimate in the homogeneous simulation is due to the dispersed absorbing nuclei from the steel material which act as neutron sinks. There was less effect for photons, with a maximum underestimate of 10%. The impacts of spatial homogenisation on the Shut-Down Dose Rate (SDDR) were also investigated. Activation at the internal face of the shield was found to be overestimated by a small amount by spatial homogenisation, approximately 10%, although this did vary as a function of time since last irradiation. The effects of spatial homogenisation in breeding blankets for fusion have been explored by Pelloni et al. amongst others [76]. However, these results pertaining to radiation shielding in nuclear systems are novel, with no similar study having been published before.

Finally, an exploration of energy domain discretisation, or 'group structure optimisation' was undertaken. How the energy domain is subdivided for nuclear analyses can have a significant effect on results. Previous methods for optimisation were mentioned, along with a discussion of nuclear resonant behaviour and the phenomenon of self-shielding. Having developed a framework for the generation of nuclear data on an arbitrary group structure, a method for the targeting of bin density is devised. The method starts with a logarithmically-spaced group structure and determines where self-shielding modifications to cross-sections most impact reaction rates. Distributions of effective self-shielding in energy are computed and summed for chosen nuclides. This distribution is then used as the basis for a bin density distribution and the apportioning of the energy domain. The method is applied to optimise group structures for two cases, a simple tungsten-only example and a more general group, optimised for 9 metals (>40 nuclides) of importance to fusion. Both conferred a significant advantage over traditional group structures in common usage today. The optimised 280 bin group structure was used to determine reaction rates in JET activation foils more accurately than the CCFE 709 bin group structure. Future work could involve the testing of optimised general group structures, to see if there is still an advantage when optimisation is for a very large population of nuclides. Iterative applications of the algorithm were investigated during the course of this work, but did not confer any advantage of single applications. One alternative, related route for optimisation could be using the effective self-shielding distributions as defined in this work, but starting from hyper-fine groups and removing the least necessary bounds until a target is reached. This approach may waste fewer bins than going from relatively coarse to locally fine, as is done in the current implementation.

Nuclear analyses are subject to a variety of sources of uncertainty. These can be from nuclear data, modelling approximations, discretisation of variables, amongst many other factors. Often these sources are simultaneously present in problems. El-Guebaly and

Malang's study of contributions to TBR uncertainty in lithium-lead blankets asserts that 90% of the TBR margin required is to account for uncertainty in its estimation, only the remaining fraction is the required net gain for system losses [38]. The uncertainty comes from both nuclear data (60%) and modelling (30%). Sometimes modelling approximations are synergistic in effect, as Pelloni et al. noted that failing to account for self-shielding effects in breeding calculations is particularly important with homogenised geometries [102].

Minimising uncertainty will become more important as we move towards constructing the first demonstration fusion devices, as unexpected or poor performance may reduce the momentum of these projects. We can continue to reduce sources of uncertainty in nuclear analyses through the development of improved methods, such as the energy group structure optimisation presented here. If approximations must be made for issues of limited computation, as with spatial homogenisation, then further analysis of the effects will be necessary to ensure their effects are adequately understood. Lastly, quantifying parameter distributions through the application of uncertainty propagation techniques like TMC will likely see an increased role–expanding from solely ND to sample other sorts of input distributions, perhaps including manufacturing tolerances and other parameters.

Appendix A

Radiation shielding material definitions

Concrete

Element	% weight
H	0.56
O	49.75
Na	1.71
Mg	0.26
Al	4.69
Si	31.47
S	0.13
K	1.92
Ca	8.28
Fe	1.24

Table A.1 Concrete composition % weight from [63]. The H content may be overestimated in this mixture. Sampling suggests it may be as low as 0.2%wt [7]. The consequences of such an deviation are substantial and were investigated in work not presented here.

Steel (radiation transport)

Element	% weight
Fe	97.24
C	0.22
P	0.05
S	0.05
N	0.012
Mn	0.56
Cr	0.16
Mo	0.16
V	0.16
Ni	0.7
Cu	0.7

Table A.2 Steel composition % weight from [15].

Steel (activation)

Element	% weight
Fe	97.613
C	0.22
P	0.05
S	0.05
N	0.012
Mn	1.08
Cr	0.1
Mo	0.01
V	0.005
Ni	0.05
Cu	0.8
Co	0.01

Table A.3 Steel composition % weight for SDDR calculations from [8]. This composition includes minor constituents which are not important for radiation transport, but may play a significant role in any SDDR.

Nomenclature

Roman Symbols

P_{fus} Fusion power

$t_{\frac{1}{2}}$ Half-life

Acronyms / Abbreviations

ASCII American Standard Code for Information Interchange

CoE Cost of Electricity

EAF European Activation File

EGPT Equivalent Generalised Perturbation Theory

ENDF Evaluated Nuclear Data File

ENDF/B Evaluated Nuclear Data File B

FPY Full Power Year

GHG Green House Gases

GPT Generalised Perturbation Theory

HCLL Helium Cooled Lithium Lead

HWR Heavy Water Reactor

ITER International Thermonuclear Experimental Reactor

JENDL Japanese Evaluated Nuclear Data Library

KDE Kernel Density Estimator

keV kilo electron-Volt

LLN Law of Large Numbers

MeV mega electron-Volt

MG Multi-Group

ND Nuclear Data

ppm parts per million

PT Probability Table

PW Point-Wise

RI Resonance Integral

RR Reaction Rate

RRR Resolved Resonance Range

RSD Relative Standard Deviation

SDDR Shut-Down Dose Rate

TBM Test Blanket Module

TBR Tritium Breeding Ratio

TENDL TALYS-based Evaluated Nuclear Data Library

TF Toroidal Field

TMC Total Monte Carlo

TSUNAMI-3D Tools for Sensitivty and Uncertainty Analysis Methodology Implementation in three Dimensions

URR Unresolved Resonance Range

WWS Wind, Water & Solar

ZETA Zero Energy Thermonuclear Assembly

References

[1] Abdou, M., Morley, N. B., Smolentsev, S., Ying, A., Malang, S., Rowcliffe, A., and Ulrickson, M. (2015). Blanket/first wall challenges and required R&D on the pathway to DEMO. *Fusion Engineering and Design*, 100:2–43.

[2] Abdou, M. A., Vold, E. L., Gung, C. Y., Youssef, M. Z., and Shin, K. (1986). Deuterium-Tritium Fuel Self-Sufficiency in Fusion Reactors. *Fusion Technology*, 9(2):250–285.

[3] Akbari, M., Khoshahval, F., Minuchehr, A., and Zolfaghari, A. (2013). A novel approach to find optimized neutron energy group structure in MOX thermal lattices using swarm intelligence. *Nuclear Engineering and Technology*, 45(7):951–960.

[4] Akbari, M., Minuchehr, A., Zolfaghari, A., and Khoshahval, F. (2012). An investigation for an optimized neutron energy-group structure in thermal lattices using Particle Swarm Optimization. *Annals of Nuclear Energy*, 47:53–61.

[5] Alhassan, E. (2014). *Nuclear data uncertainty propagation for a lead-cooled fast reactor: Combining TMC with criticality benchmarks for improved accuracy*. PhD thesis, Uppsala.

[6] Alhassan, E., Sjöstrand, H., Helgesson, P., Koning, A. J., Österlund, M., Pomp, S., and Rochman, D. (2015). Uncertainty and correlation analysis of lead nuclear data on reactor parameters for the European Lead Cooled Training Reactor. *Annals of Nuclear Energy*, 75:26–37.

[7] Aramburu, I. (2016). Concrete Block Characterization 051857C2. Technical report, Tecnalia.

[8] Barabash, V. (2016). Private communication.

[9] Barrett, B. R., Navrátil, P., and Vary, J. P. (2013). Ab initio no core shell model. *Progress in Particle and Nuclear Physics*, 69(1):131–181.

[10] Batistoni, P., Angelone, M., Bettinali, L., Carconi, P., Fischer, U., Kodeli, I., Leichtle, D., Ochiai, K., Perel, R., Pillon, M., Schäfer, I., Seidel, K., Verzilov, Y., Villari, R., and Zappa, G. (2007). Neutronics experiment on a helium cooled pebble bed (HCPB) breeder blanket mock-up. *Fusion Engineering and Design*, 82(15-24):2095–2104.

[11] Batistoni, P., Fischer, U., Ochiai, K., Petrizzi, L., Seidel, K., and Youssef, M. (2008). Neutronics and nuclear data issues in ITER and their validation. *Fusion Engineering and Design*, 83(7-9):834–841.

[12] Bell, G. and Glasstone, S. (1970). *Nuclear Reactor Theory*. Van Nostrand Reinhold Inc., 1st edition.

[13] Bossmann, T. and Staffell, I. (2015). The shape of future electricity demand: Exploring load curves in 2050s Germany and Britain. *Energy*.

[14] Bradshaw, A. M., Hamacher, T., and Fischer, U. (2011). Is nuclear fusion a sustainable energy form? *Fusion Engineering and Design*, 86(9-11):2770–2773.

[15] British Standards (2005). Steel for the reinforcement of concrete. Weldable reinforcing steel. BS EN 10080:2005.

[16] Brown, D. and Kawano, T. (2017). An analytic approach to probability tables for the unresolved resonance region. *EPJ Web of Conferences*, 146:12008.

[17] Brown, D. A., Chadwick, M. B., Capote, R., Kahler, A. C., Trkov, A., Herman, M. W., Sonzogni, A. A., Danon, Y., Carlson, A. D., Dunn, M., Smith, D. L., Hale, G. M., Arbanas, G., Arcilla, R., Bates, C. R., Beck, B., Becker, B., Brown, F., Casperson, R. J., Conlin, J., Cullen, D. E., Descalle, M. A., Firestone, R., Gaines, T., Guber, K. H., Hawari, A. I., Holmes, J., Johnson, T. D., Kawano, T., Kiedrowski, B. C., Koning, A. J., Kopecky, S., Leal, L., Lestone, J. P., Lubitz, C., Márquez Damián, J. I., Mattoon, C. M., McCutchan, E. A., Mughabghab, S., Navratil, P., Neudecker, D., Nobre, G. P., Noguere, G., Paris, M., Pigni, M. T., Plompen, A. J., Pritychenko, B., Pronyaev, V. G., Roubtsov, D., Rochman, D., Romano, P., Schillebeeckx, P., Simakov, S., Sin, M., Sirakov, I., Sleaford, B., Sobes, V., Soukhovitskii, E. S., Stetcu, I., Talou, P., Thompson, I., van der Marck, S., Welser-Sherrill, L., Wiarda, D., White, M., Wormald, J. L., Wright, R. Q., Zerkle, M., Žerovnik, G., and Zhu, Y. (2018). ENDF/B-VIII.0: The 8thMajor Release of the Nuclear Reaction Data Library with CIELO-project Cross Sections, New Standards and Thermal Scattering Data. *Nuclear Data Sheets*, 148:1–142.

[18] Cacuci, D. G. (2010). *Handbook of Nuclear Engineering*.

[19] Cambi, G., Cepraga, D. G., Di Pace, L., Druyts, F., and Massaut, V. (2010). The potential presence and minimisation of plutonium within the irradiated beryllium in fusion power plants. *Fusion Engineering and Design*, 85(7-9):1139–1142.

[20] Cantor, R. and Hewlett, J. (1988). The economics of nuclear power. *Resources and Energy*, 10:315–335.

[21] Capote, R., Herman, M., Obložinský, P., Young, P. G., Goriely, S., Belgya, T., Ignatyuk, A. V., Koning, A. J., Hilaire, S., Plujko, V. A., Avrigeanu, M., Bersillon, O., Chadwick, M. B., Fukahori, T., Ge, Z., Han, Y., Kailas, S., Kopecky, J., Maslov, V. M., Reffo, G., Sin, M., Soukhovitskii, E. S., and Talou, P. (2009). RIPL - Reference Input Parameter Library for Calculation of Nuclear Reactions and Nuclear Data Evaluations. *Nuclear Data Sheets*, 110(12):3107–3214.

[22] Capote, R., Smith, D., and Trkov, a. (2010). Nuclear data evaluation methodology including estimates of covariances. *EPJ Web of Conferences*, 8:04001.

[23] Chadwick, M. B., Herman, M., Obložinský, P., Dunn, M. E., Danon, Y., Kahler, A. C., Smith, D. L., Pritychenko, B., Arbanas, G., Arcilla, R., Brewer, R., Brown, D. A., Capote, R., Carlson, A. D., Cho, Y. S., Derrien, H., Guber, K., Hale, G. M., Hoblit, S., Holloway, S., Johnson, T. D., Kawano, T., Kiedrowski, B. C., Kim, H., Kunieda, S., Larson, N. M., Leal, L., Lestone, J. P., Little, R. C., McCutchan, E. A., MacFarlane, R. E., MacInnes,

M., Mattoon, C. M., McKnight, R. D., Mughabghab, S. F., Nobre, G. P., Palmiotti, G., Palumbo, A., Pigni, M. T., Pronyaev, V. G., Sayer, R. O., Sonzogni, A. A., Summers, N. C., Talou, P., Thompson, I. J., Trkov, A., Vogt, R. L., van der Marck, S. C., Wallner, A., White, M. C., Wiarda, D., and Young, P. G. (2011). ENDF/B-VII.1 nuclear data for science and technology: Cross sections, covariances, fission product yields and decay data. *Nuclear Data Sheets*, 112(12):2887–2996.

[24] Church, J., Clark, P., Cazenave, a., Gregory, J., Jevrejeva, S., Levermann, a., Merrifield, M., Milne, G., Nerem, R., Nunn, P., a.J. Payne, Pfeffer, W., Stammer, D., and a.S. Unnikrishnan (2013). Sea level change. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 1137–1216.

[25] Chuyanov, V. a., Campbell, D. J., and Giancarli, L. M. (2010). TBM Program implementation in ITER. *Fusion Engineering and Design*, 85(10-12):2005–2011.

[26] Clack, C. T. M., Qvist, S. A., Apt, J., Bazilian, M., Brandt, A. R., Caldeira, K., Davis, S. J., Diakov, V., Handschy, M. A., Hines, P. D. H., Jaramillo, P., Kammen, D. M., Long, J. C. S., Morgan, M. G., Reed, A., Sivaram, V., Sweeney, J., Tynan, G. R., Victor, D. G., Weyant, J. P., and Whitacre, J. F. (2017). Evaluation of a proposal for reliable low-cost grid power with 100% wind, water, and solar. *Proceedings of the National Academy of Sciences*, 114(26):6722–6727.

[27] Colling, B. R. and Monk, S. D. (2012). Development of fusion blanket technology for the DEMO reactor. *Applied Radiation and Isotopes*, 70(7):1370–1372.

[28] Community, T. S. (2018). numpy.corrcoef.

[29] Cross Sections Evaluation Working Group (2011). ENDF-6 Formats Manual. Technical report, Brookhaven National Laboratory.

[30] Crutzen, P. J. (2006). *The "Anthropocene"*, pages 13–18. Springer Berlin Heidelberg, Berlin, Heidelberg.

[31] Cullen, D. E. (1988). The Accuracy of Processed Nuclear Data. *Nuclear Science and Engineering*, 99(2):172–181.

[32] Cullen, D. E. (2017). PREPRO 2017 ENDF/B Pre-processing Codes. Technical report, IAEA.

[33] Davis, A. (2010). *Radiation shielding of fusion systems*. PhD thesis, Unversity of Birmingham.

[34] Davis, A. and Pampin, R. (2010). Benchmarking the MCR2S system for high-resolution activation dose analysis in ITER. 85:87–92.

[35] Dembia, C. L., Recktenwald, G. D., and Deinert, M. R. (2013). Bondarenko method for obtaining group cross sections in a multi-region collision probability model. *Progress in Nuclear Energy*, 67:124–131.

[36] Dunford, C. L. (2004). ENDF Utility Codes Release 7. Technical report.

[37] Eade, T., Stonell, D., and Turner, A. (2015). MCR2S unstructured mesh capabilities for use in shutdown dose rate analysis. *Fusion Engineering and Design*, 100:321–333.

[38] El-Guebaly, L. a. and Malang, S. (2009). Toward the ultimate goal of tritium self-sufficiency: Technical issues and requirements imposed on ARIES advanced power plants. *Fusion Engineering and Design*, 84(12):2072–2083.

[39] Entler, S., Horacek, J., Dlouhy, T., and Dostal, V. (2018). Approximation of the economy of fusion energy. *Energy*, 152:489–497.

[40] Environmental Justice Foundation (2017). Beyond Borders: Our changing climate - its role in conflict and displacement. Technical report.

[41] Fernbach, S., Serber, R., and Taylor, T. B. (1949). The scattering of high energy neutrons by nuclei. *Physical Review*, 75(9):1352–1355.

[42] Fiorito, L., Žerovnik, G., Stankovskiy, A., Van den Eynde, G., and Labeau, P. E. (2017). Nuclear data uncertainty propagation to integral responses using SANDY. *Annals of Nuclear Energy*, 101:359–366.

[43] Fischer, U., Bachmann, C., Palermo, I., Pereslavtsev, P., and Villari, R. (2015). Neutronics requirements for a DEMO fusion power plant. *Fusion Engineering and Design*, pages 2–5.

[44] Fleming, M. J., Morgan, L. W. G., and Shwageraus, E. (2016). Optimization Algorithms for Multigroup Energy Structures. *Nuclear Science and Engineering*, 183:173–184.

[45] Forrest, R. A. (2011). The role of nuclear data for fusion technology studies. *Nuclear Engineering and Design*, 241(10):4326–4330.

[46] Gandini, A. (1967). A generalized perturbation method for bi-linear functionals of the real and adjoint neutron fluxes. *Journal of Nuclear Energy*, 21(10):755–765.

[47] Gandini, A., Palmiotti, G., and Salvatores, M. (1986). Equivalent generalized perturbation theory (EGPT). *Annals of Nuclear Energy*, 13(3):109–114.

[48] Gilbert, M. R. and Sublet, J. C. (2011). Neutron-induced transmutation effects in W and W-alloys in a fusion environment. *Nuclear Fusion*, 51(4).

[49] Gilbert, M. R., Sublet, J. C., and Dudarev, S. L. (2016). Spatial heterogeneity of W transmutation in a fusion device. *Nuclear Fusion*, 57.

[50] Goorley, T., James, M., Booth, T., Brown, F., and Bull, J. (2012). Initial MCNP6 Release Overview. *Nuclear Technology*, 180(3):298–315.

[51] Hansen, J., Ruedy, R., Sato, M., and Lo, K. (2010). Global surface temperature change. *Rev. Geophys.*, 48(4):RG4004.

[52] Hanus, M. (2014). *Mathematical Modeling of Neutron Transport*. PhD thesis, University of West Bohemia.

[53] Harms, A. A. (1975). *An introduction to the CANDU nuclear energy conversion system*. McMaster University, Hamilton.

[54] Hauser, W. and Feshbach, H. (1952). The inelastic scattering of neutrons. *Physical Review*, 87(2):366–373.

[55] Helgesson, P. and Sj, H. (2017). Combining Total Monte Carlo and Unified Monte Carlo: Bayesian nuclear data uncertainty quantification from auto-generated experimental covariances. 96:76–96.

[56] Herman, M. and Trkov, A. (2010). ENDF-6 Formats Manual. *Brookhaven National Laboratory*, page 392.

[57] Hodgson, P. E. (1971). The nuclear optical model. *Reports on Progress in Physics*, 34:765–819.

[58] IAEA NDS (2017). Experimental Nuclear Reaction Data (EXFOR).

[59] Ihli, T., Basu, T. K., Giancarli, L. M., Konishi, S., Malang, S., Najmabadi, F., Nishio, S., Raffray, a. R., Rao, C. V. S., Sagara, a., and Wu, Y. (2008). Review of blanket designs for advanced fusion reactors. *Fusion Engineering and Design*, 83(7-9):912–919.

[60] International Renewable Energy Agency (IRENA) (2018). Renewable Power Generation Costs in 2017. Technical report, IRENA.

[61] IPCC Core Writing Team (2007). Climate Change 2007: An Assessment of the Intergovernmental Panel on Climate Change. Technical report, IPCC.

[62] Jacobson, M. Z., Cameron, M. A., Hennessy, E. M., Petkov, I., Meyer, C. B., Gambhir, T. K., Maki, A. T., Pfleeger, K., Clonts, H., McEvoy, A. L., Miccioli, M. L., von Krauland, A. K., Fang, R. W., and Delucchi, M. A. (2018). 100% clean and renewable Wind, Water, and Sunlight (WWS) all-sector energy roadmaps for 53 towns and cities in North America. *Sustainable Cities and Society*, 42(June):22–37.

[63] Jakhar, S. (2016). Shutdown dose rate analysis of the bio-shield plug design at equatorial level RJLPH8. Technical report, ITER Organisation.

[64] Kodeli, I. (1995). Neutron and gamma field characteristics after shutdown and a possible application to determine the coolant inventory. *Proceedings of the Meeting on Nuclear Energy: Central Europe: Present and Perspectives*, (September).

[65] Kodeli, I. (2001). Multidimensional Deterministic Nuclear Data Sensitivity and Uncertainty Code System: Method and Application. *Nuclear Science and Engineering*, 138(1):45–66.

[66] Kolbasov, B. N., Khripunov, V. I., and Biryukov, A. Y. (2016). On use of beryllium in fusion reactors: Resources, impurities and necessity of detritiation after irradiation. *Fusion Engineering and Design*, 109-111:480–484.

[67] Koning, A. (2015). Bayesian Monte Carlo Method for Nuclear Data Evaluation. *Nuclear Data Sheets*, 123:207–213.

[68] Koning, A., Hilaire, S., and Goriely, S. (2017). TALYS-1.9 Manual. Technical report.

[69] Koning, A. and Rochman, D. (2013). Nuclear data uncertainty propagation using a Total Monte Carlo approach. *Workshop on Uncertainty Propagation in the Nuclear Fuel Cycle*.

[70] Koning, A. J. and Delaroche, J. P. (2003). Local and global nucleon optical models from 1 keV to 200 MeV. *Nuclear Physics A*, 713(3-4):231–310.

[71] Koning, A. J., Hilaire, S., and Duijvestijn, M. C. (2005). TALYS: Comprehensive nuclear reaction modeling. *AIP Conference Proceedings*, 769:1154–1159.

[72] Koning, A. J. and Rochman, D. (2008). Towards sustainable nuclear energy: Putting nuclear physics to work. *Annals of Nuclear Energy*, 35:2024–2030.

[73] Koning, A. J. and Rochman, D. (2012). Modern Nuclear Data Evaluation with the TALYS Code System. *Nuclear Data Sheets*, 113(12):2927–2934.

[74] Kovari, M., Coleman, M., Cristescu, I., and Smith, R. (2018). Tritium resources available for fusion reactors. *Nuclear Fusion*, 58(2).

[75] Krane, K. S. (1987). *Introductory Nuclear Physics*. Wiley.

[76] Kumar, A., Watanabe, Y., Youssef, M. Z., and Abdou, M. A. (1989). Analysis for the Selection of Experimental Configurations for Heterogeneity and Be Multi-Layered Experiments of United-States Doe Jaeri Collaborative Program on Blanket Neutronics. *Fusion Technology*, 15(2):1309–1314.

[77] Lawson, J. D. (1955). Some criteria for a useful thermonuclear reactor. Technical report, Atomic Energy Research Establishment.

[78] Leichtle, D., Fischer, U., Perel, R. L., and Serikov, A. (2011). Sensitivity and uncertainty analysis of nuclear responses in the EU HCLL TBM of ITER. *Fusion Engineering and Design*, 86(9-11):2156–2159.

[79] Lengar, I. (2017). Private communication.

[80] Libby, J. (2005). Breit-Wigner cross section and angular momentum.

[81] Lobel, R. (2008). A Brief History of the Way Neutron Activation has Affected the Construction, Maintenance and Operation of JET. Technical report, EFDA.

[82] Los Alamos National Laboratory (2018). NJOY21.

[83] Loughlin, M. (2009). Recommendation on plasma scenarios 2V3V8G. Technical report, ITER Organisation.

[84] Lovering, J. R., Yip, A., and Nordhaus, T. (2016). Historical construction costs of global nuclear power reactors. *Energy Policy*, 91:371–382.

[85] MacFarlane, R. E. and Kahler, A. C. (2010). Methods for Processing ENDF/B-VII with NJOY. *Nuclear Data Sheets*, 111(12):2739–2890.

[86] MacFarlane, R. E. and Muir, D. W. (2016). The NJOY Nuclear Data Processing System. Technical report, Los Alamos National Laboratory.

[87] MacKay, D. (2009). *Sustainable Energy — without the hot air*. UIT Cambridge Ltd., Cambridge, 1st edition.

[88] Mann, D. H., Groves, P., Reanier, R. E., Gaglioti, B. V., Kunz, M. L., and Shapiro, B. (2015). Life and extinction of megafauna in the ice-age Arctic. *Proceedings of the National Academy of Sciences*, 112(46):14301–14306.

[89] Markandya, A. and Wilkinson, P. (2007). Electricity generation and health. *Lancet*, 370(9591):979–990.

[90] Morgan, L., Sublet, J.-C., Haeck, W., and Pasley, J. (2013). Optimising the energy group structure used for fusion systems. *Annals of Nuclear Energy*, 55:108–115.

[91] National Audit Office (NAO) (2017). Hinkley Point C. Technical Report June, National Audit Office.

[92] National Institute of Standards and Technology (NIST) (2018). The NIST reference on constants, units and uncertainty.

[93] National Snow and Ice Data Centre (2018). Sea Ice Index.

[94] Nazarewicz, W. (2016). Challenges in Nuclear Structure Theory. *Journal of Physics G: Nuclear and Particle Physics*, 43.

[95] NOAA Global Monitoring Division (2018). Trends in Atmospheric Carbon Dioxide.

[96] OECD and NEA (2013). Transition Towards a Sustainable Nuclear Fuel Cycle. Technical report, OECD.

[97] Oliphant, M. L., Harteck, P., and Rutherford, E. (1934). Transmutation Effects observed with Heavy Hydrogen. *Nature*, page 413.

[98] Packer, L., Sublet, J., Kopecky, J., and Forrest, R. (2011). Recent Progress in Neutron-, Proton- and Deuteron-induced Reaction Nuclear Data for EAF-2010 and the European Activation System. *Journal Of The Korean Physical Society*, 59(2):1100–1103.

[99] Pampin, R. (2005). Tungsten transmutation and resonance self-shielding in PPCS models for the study of sigma-phase formation. Technical report, UKAEA.

[100] Pampin, R., Davis, A., and James, D. (2007). Systematic study of neutron shielding options for ITER generic diagnostic port plug design. Technical report, UKAEA.

[101] Pattie, R. W., Callahan, N. B., Cude-Woods, C., Adamek, E. R., Broussard, L. J., Clayton, S. M., Currie, S. A., Dees, E. B., Ding, X., Engel, E. M., Fellers, D. E., Fox, W., Hickerson, K. P., Hoffbauer, M. A., Holley, A. T., Komives, A., Liu, C. Y., MacDonald, S. W. T., Makela, M., Morris, C. L., Ortiz, J. D., Ramsey, J., Salvat, D. J., Saunders, A., Seestrom, S. J., Sharapov, E. I., Sjue, S. K., Tang, Z., Vanderwerp, J., Vogelaar, B., Walstrom, P. L., Wang, Z., Wei, W., Weaver, H. L., Wexler, J. W., Womack, T. L., Young, A. R., and Zeck, B. A. (2018). Measurement of the neutron lifetime using an asymmetric magneto-gravitational trap and in situ detection. *Science*, 360(May):627–632.

[102] Pelloni, S., Cheng, E. T., and Embrechts, M. J. (1989). Self-Shielding Characteristics of Aqueous Lithium Salt Blankets for Next-Generation Fusion Devices. *Fusion Technology*, 16(1):53–64.

[103] Perez, A. (2014). Construction Design – Tokamak Complex PBS 62.11, PBS 62.14 & PBS 62.74 Detailing calculation – Level L1 – Hypotheses & methodology. ENG-50-CR-110093-CW-v02.0. Technical report, ITER Organisation.

[104] Plechaty, E. F., Cullen, D. E., Howerton, R. J., and Kimlinge, J. R. (1978). Tabular and graphical presentation of 175 neutron-group constants derived from the LLL evaluated-nuclear-data library (ENDL). Technical report, Lawrence Livermore Laboratory.

[105] Price, B. T. (1959). *Radiation Shielding*. Pergamon Press.

[106] Pushkina, D. and Raia, P. (2008). Human influence on distribution and extinctions of the late Pleistocene Eurasian megafauna. *Journal of Human Evolution*, 54(6):769–782.

[107] Python Software Foundation (2018). multiprocessing — Process-based parallelism.

[108] Rearden, B. T. (2004). Perturbation Theory Eigenvalue Sensitivity Analysis with Monte Carlo Techniques. *Nuclear Science and Engineering*, 146(3):367–382.

[109] Rhodes, J., Smith, K., and Lee, D. (2006). CASMO-5 development and applications. In *Proc. ANS Topical Meeting on Reactor Physics*.

[110] Ricke, K. L. and Caldeira, K. (2014). Maximum warming occurs about one decade after a carbon dioxide emission. *Environmental Research Letters*, 9(12).

[111] Rider, T. H. (1995). A general critique of inertial-electrostatic confinement fusion systems. *Physics of Plasmas*, 2(6):1–56.

[112] Rising, M. E. (2012). *Quantification and Propagation of Nuclear Data Uncertainties*. PhD thesis, University of New Mexico.

[113] Rochman, D., Koning, A., van der Marck, S., Hogenbirk, A., and Sciolla, C. (2011a). Nuclear data uncertainty propagation: Perturbation vs. Monte Carlo. *Annals of Nuclear Energy*, 38(5):942–952.

[114] Rochman, D., Koning, A. J., Sublet, J. C., Fleming, M., Bauge, E., Hilaire, S., Romain, P., and Morillon, B. (2015). Neutron sub-library for Pb (Z=82).

[115] Rochman, D., Koning, A. J., Sublet, J. C., Fleming, M., Bauge, E., Hilaire, S., Romain, P., and Morillon, B. (2016). The TENDL library: hope, reality and future. In *ND 2016 International Conference on Nuclear Data for Science and Technology*.

[116] Rochman, D., Koning, A. J., and van der Marck, S. C. (2011b). Exact Nuclear Data Uncertainty Propagation for Fusion Design. *Journal of the Korean Physical Society*, 59(2):1386–1389.

[117] Rochman, D., Marck, S. V. D., and Hogenbirk, A. (2010). Nuclear data uncertainty propagation (adjustment procedure). Technical report, NRG.

[118] Rochman, D., Zwermann, W., van der Marck, S., Koning, A. J., Sjostrand, H., Helgesson, P., and Krzykacz-Hausmann, B. (2014). Efficient Use of Monte Carlo: Uncertainty Propagation. *Nuclear Science and Engineering*, 177:337–349.

[119] Sabouri, P. (2013). Application of Perturbation Theory Methods to Nuclear Data Uncertainty Propagation using the Collision.

[120] Sanz, J. (2014). Effect of Bioshield Plug Modelling Features on Port Cell Shutdown Dose Rate Computations. Technical report.

[121] Schuster, E. (2017). Power Balance in Fusion Plasmas.

[122] Shapiro, S. S. and Wilk, M. B. (1965). An analysis of variance test for normality (complete samples). *Biometrika*, 52(3 & 4):591–611.

[123] Shibata, K., Iwamoto, O., Nakagawa, T., Iwamoto, N., Ichihara, A., Kunieda, S., Chiba, S., Furutaka, K., Otuka, N., Ohsawa, T., Murata, T., Matsunobu, H., Zukeran, A., Kamada, S., and Katakura, J. I. (2011). JENDL-4.0: A new library for nuclear science and engineering. *Journal of Nuclear Science and Technology*, 48(1):1–30.

[124] Shimwell, J. (2018). ShimPlotWell.

[125] Shimwell, J., Lilley, S., Kovari, M., Zheng, S., Morgan, L., and McMillan, J. (2014). Reducing beryllium content in solid-type breeder blankets. *Fusion Engineering and Design*.

[126] Shuman, E. K. (2010). Global climate change and infectious diseases. *The New England Journal of Medicine*, 362(1):1061–1063.

[127] Sjöstrand, H., Conroy, S., Helgesson, P., Hernandez, S. A., Koning, A., Pomp, S., and Rochman, D. (2017). Propagation of nuclear data uncertainties for fusion power measurements. *EPJ Web of Conferences*, 146:0–3.

[128] Stone, R. (1997). An Element of Stability. *Science*, 278(October):571–572.

[129] Sublet, J.-C. (2011). CALENDF-2010: USER MANUAL. Technical report, CEA Saclay.

[130] Sublet, J.-C., Eastwood, J. W., Morgan, J. G., Fleming, M., and Gilbert, M. R. (2015). The FISPACT-II User Manual. Technical Report 6, UK Atomic Energy Authority.

[131] Sublet, J. C., Eastwood, J. W., Morgan, J. G., Gilbert, M. R., Fleming, M., and Arter, W. (2017a). FISPACT-II: An Advanced Simulation System for Activation, Transmutation and Material Modelling. *Nuclear Data Sheets*, 139:77–137.

[132] Sublet, J.-C., Fleming, M., and Gilbert, M. R. (2017b). CALENDF Probability Tables Advanced Self-shielding Factors Usage in the Inventory Code FISPACT-II. In *M&C 2017*.

[133] TAE (Tri Alpha Energy) (2018). Technology Overview.

[134] The PyNE Development Team (2018). PyNE - The Nuclear Engineering Toolkit.

[135] Ward, D. J., Cook, I., Lechon, Y., and Saez, R. (2005). The economic viability of fusion power. *Fusion Engineering and Design*, 75-79(SUPPL.):1221–1227.

[136] Woods, R. D. and Saxon, D. S. (1954). Diffuse surface optical model for nucleon-nuclei scattering. *Physical Review*, 95(2):577–578.

[137] Yamada, H. (2012). Fusion Energy. In Chen, W.-Y., Seiner, J., Suzuki, T., and Lackner, M., editors, *Handbook of Climate Change Mitigation*, pages 1183–1215. Springer US, New York, NY.

[138] Yi, C. and Sjoden, G. (2013). Energy group structure determination using particle swarm optimization. *Annals of Nuclear Energy*, 56:53–56.

[139] Youssef, M. Z. and Abdou, M. A. (1986). Uncertainties in Prediction of Tritium Breeding in Candidate Blanket Designs Due to Present Uncertainties in Nuclear-Data Base. *Fusion Technology*, 9(2):286–307.

[140] Zhang, X., Wan, H., Zwiers, F. W., Hegerl, G. C., and Min, S. K. (2013). Attributing intensification of precipitation extremes to human influence. *Geophysical Research Letters*, 40(19):5252–5257.

[141] Zheng, S., King, D. B., Garzotti, L., Surrey, E., and Todd, T. N. (2016). Fusion reactor start-up without an external tritium source. *Fusion Engineering and Design*, 103:13–20.

[142] Zhu, T. (2015). *Sampling-Based Nuclear Data Uncertainty Quantification for Continuous Energy Monte Carlo Codes*. PhD thesis, ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE.

[143] Zickfeld, K. and Herrington, T. (2015). The time lag between a carbon dioxide emission and maximum warming increases with the size of the emission. *Environmental Research Letters*, 10(3).