Frej Wasastjerna

Using MCNP for fusion neutronics

Using MCNP for fusion neutronics

Frej Wasastjerna

Abstract

Any fusion reactor using tritium-deuterium fusion will be a prolific source of 14 MeV neutrons. In fact, 80% of the fusion energy will be carried away by these neutrons. Thus it is essential to calculate what will happen to them, so that such quantities as the tritium breeding ratio, the neutron wall loading, heat deposition, various kinds of material damage and biological dose rates can be determined. Monte Carlo programs, in particular the widely-used MCNP, are the preferred tools for this. The International Fusion Materials Irradiation Facility (IFMIF), intended to test materials in intense neutron fields with a spectrum similar to that prevailing in fusion reactors, also requires neutronics calculations, with similar methods.

In some cases these calculations can be very difficult. In particular shielding calculations – such as those needed to determine the heating of the superconducting field coils of ITER or the dose rate, during operation or after shutdown, outside ITER or in the space above the test cell of IFMIF – are very challenging. The thick shielding reduces the neutron flux by many orders of magnitude, so that analog calculations are impracticable and heavy variance reduction is needed, mainly importances or weight windows. On the other hand, the shields contain penetrations through which neutrons may stream. If the importances are much higher or the weight windows much lower at the outer end of such a penetration than at the inner end, this may lead to an excessive proliferation of tracks, which may even make the calculation break down.

This dissertation describes the author's work in fusion neutronics, with the main emphasis on attempts to develop improved methods of performing such calculations. Two main approaches are described: trying to determine near-optimal importances or weight windows, and testing the "tally source" method suggested by John Hendricks as a way of biasing the neutron flux in angle.

Sammanfattning

En fusionsreaktor som använder sig av tritium-deuteriumfusion kommer att vara en mycket stark källa till neutroner med en energi av 14 MeV. Dessa kommer i själva verket att bära med sig 80 % av fusionsenergin. Det är alltså väsentligt att beräkna vad som kommer att hända med dem, så att man kan bestämma sådana storheter som tritiumbridkvoten, den av neutronerna förorsakade väggbelastningen, värmedepositionen, olika slags materialskador och biologiska dosrater. Monte Carlo-program, speciellt MCNP, är det mest allmänt använda redskapet för sådana beräkningar. IFMIF (International Fusion Materials Irradiation Facility), avsedd att testa material i ett intensivt neutronflöde med ett spektrum som liknar det i fusionsreaktorer, kräver också neutronikberäkningar, med liknande metoder.

I vissa fall kan dessa beräkningar vara mycket svåra. I synnerhet strålskydds- beräkningar, sådana som behövs för att bestämma värmeutvecklingen i ITERs supraledande fältspolar eller dosraten, i verksamhet eller efter avstängning, utanför ITER eller i utrymmet ovanför IFMIFs testcell, är särdeles krävande. Det tjocka strålskyddet minskar neutronflödet med många storleksordningar, så att analoga beräkningar är ogenomförbara och stark variansreduktion är nödvändig, främst med viktighet eller viktfönster. Å andra sidan innehåller skyddet genom- föringar längs vilka neutroner kan strömma. Om viktigheterna är mycket högre eller viktfönstren mycket lägre i yttre ändan av en sådan genomföring än i den inre ändan, kan en partikelhistoria yngla av sig i ett orimligt antal spår, vilket kan få beräkningen att bryta samman.

Denna avhandling beskriver disputandens arbete i fusionsneutronikbranschen med huvudvikten lagd vid hans försök att utveckla förbättrade beräknings- metoder. Två huvudsakliga sätt att göra detta beskrivs: försök att bestämma närapå optimala viktigheter eller viktfönster, och testning av "tally source"- metoden, föreslagen av John Hendricks som ett sätt att biasera neutronfödet med avseende på vinkel.

Preface

List of publications

This dissertation is based on the following publications:

I Heikinheimo, L., Heikkinen, J., Linden, J., Kaye, A., Orivuori, S., Saarelma, S., Tähtinen, S., Walton, R. and Wasastjerna, F. Dielectric window for reactor like ICRF vacuum transmission line. Fusion Engineering and Design **55**(2001), pp. 419–436.

II Bibet, Ph., Mirizzi, F., Bosia, P., Doceul, L., Kuzikov, S., Rantamäki, K., Tuccillo, A.A. and Wasastjerna, F. Overview of the ITER-FEAT LH system. Fusion Engineering and Design **66–68**(2003), pp. 525–529.

III Fischer, U., Pereslavtsev, P., Kotiluoto, P. and Wasastjerna, F. Nuclear performance analyses for HCPB test blanket modules in ITER. Fusion Engineering and Design **82**(2007), pp. 2140–2146.

IV Norajitra, P., Bühler, L., Buenaventura, A., Diegele, E., Fischer, U. Gordeev, S., Hutter, E., Kruessmann, R., Malang, S., Maisonnier, D., Orden, A., Reimann, G., Reimann, J., Vieider, G., Ward, D. and Wasastjerna, F. Conceptual design of the EU dual-coolant blanket (model C). 20[th] IEEE/NPSS Symposium on Fusion Engineering (SOFE). San Diego, CA, USA 2003.

V Simakov, S.P, Fischer, U., Möslang, A., Vladimirov, P., Wasastjerna, F. and Wilson, P.P.H. Neutronics and activation characteristics of the international fusion material irradiation facility. Fusion Engineering and Design **75–79**(2005), pp. 813–817.

VI Fischer, U., Chen, Y., Simakov, S.P., Wilson, P.P.H., Vladimirov, P. and Wasastjerna, F. Overview of recent progress in IFMIF neutronics. Fusion Engineering and Design **81**(2006), pp. 1195–1202.

VII   Simakov, S.P., Vladimirov, P., Fischer, U., Heinzel, V., Möslang, A. and Wasastjerna, F. Tungsten spectral shifter: neutronics analysis (dpa evaluation, H, He and other impurities generation, recoil spectrum, etc) of different positions and geometries. Final report on the EFDA task TW5-TTMI-003, Deliverable 8. FZK (Forschungszentrum Karlsruhe) report (2006). 11 p.

VIII   Kotiluoto, P. and Wasastjerna, F. Calculating the neutron current emerging through the beam tubes in IFMIF (EFDA Task TW6-TTMI-001-D3a). VTT Research Report (2006). 9 p.

IX   Wasastjerna, F. A study of variance reduction methods in MCNP4C in a slot with a dogleg. 12th Biennial Topical Meeting of the Radiation Protection and Shielding Division of the American Nuclear Society. Paper 14, Santa Fe, NM, USA 2002. 11 p.

X   Wasastjerna, F. Importances for MCNP Calculations on a Shield with a Penetration. 14th Biennial Topical Meeting of the Radiation Protection and Shielding Division of the American Nuclear Society, Carlsbad, NM, USA 2006. Pp. 478–488.

XI   Wasastjerna, F. Application of importances or weight windows in MCNP4C to a geometry similar to an ITER equatorial port. Annals of Nuclear Energy **35**/3 (2008), pp. 425–431.

XII   Wasastjerna, F. Calculating streaming with the "tally source" method, as applied to IFMIF. Annals of Nuclear Energy **35**/3 (2008), pp. 438–445.

In Publications I through III the contribution of the disputant was to model one or more port plugs for an equatorial and/or upper port in ITER, fit this into the general ITER model prepared by a team of other people, and perform the required neutron shielding calculations. In Publication III the modeling was limited to the rear (shield) part of the plug, and even this modeling was shared with one co-author. In Publication IV the disputant modeled the whole reactor, at a more modest level of detail than for ITER or IFMIF, and performed the neutronics calculations. In Publications V through VIII the IFMIF test cell and

its contents were modeled in considerable detail. The author also performed some of the calculations described in Publication VIII.

In more detail, the disputant was responsible for the following:

- Publication I: Section 2
- Publication II: Subsection 3.5
- Publication III: most of Section 5, with a contribution from P. Kotiluoto.

Note that the disputant's share of the modeling in these three publications was mostly limited to the port plug being investigated, and in Publication III only to part of that. For the rest of the reactor, existing standard models were used, though they often had to be modified slightly, e.g. by subdividing the port walls into smaller cells to permit sufficient resolution for importances or weight windows.

- Publication IV: Subsection IIC

- Publication V: The model described in the second paragraph of Section 2 and used in the calculations

- Publication VI: Again, the model used in the Monte Carlo calculations

- Publication VII: All of the several geometry model variants used in the calculations

- Publication VIII: The model and the initial calculations

- All of Publications IX through XII.

Contents

Appendices

*Appendix B: Publications I, II, III, V, VI, XI, XII are not included in the PDF
version. Please order the printed version to get the complete publication
(http://www.vtt.fi/publications/index.jsp)*

1. Introduction

1.1 The need for fusion neutronics

In a fission reactor, approximately one neutron gets released for each 80 MeV of fission energy, and the average energy of fission neutrons is about 2 MeV. In deuterium-tritium fusion, one neutron is released per 17.6 MeV of fusion energy, and the neutrons initially have a kinetic energy of about 14.1 MeV. In other words, a fusion reactor relying on the D-T reaction is a much more prolific source of neutrons than a fission reactor, and in addition the neutrons are much more energetic.

While fusion reactors do not depend on neutrons to maintain a chain reaction the way fission reactors do, neutrons are needed to generate the fuel they burn. Tritium does not exist in significant quantities in nature, so it must be produced by neutrons reacting with lithium (e.g., ^{6}Li(n,t)^{4}He). Since every fusion reaction consumes one tritium atom (ignoring such side reactions as D-D fusion) and provides only one neutron, and since losses of both neutrons and tritium are unavoidable, a certain amount of neutron multiplication in materials such as beryllium or lead is required. The lithium also needs to be enriched in ^{6}Li. Checking that the tritium breeding ratio (TBR) is adequate is one of the problems in fusion neutronics.

The neutrons carry away 80% of the fusion energy, depositing it mostly in the blanket. Calculating the energy deposition rate as a function of position is essential, so that the designers know how much heat the coolant must be able to remove in each location.

The neutrons also have harmful effects. They induce radioactivity in the materials of which a fusion reactor is constructed and in the coolant flowing through it, e.g., through the ^{16}O(n,p)^{16}N reaction. They may damage materials through displacement of atoms, affecting mechanical and electric properties. They produce hydrogen and helium, which make welding difficult in concentrations above about 1 atomic ppm, and they heat superconducting coils in devices using magnetic confinement. These effects also need to be calculated.

1.2 The international magnetic confinement fusion program

The current centerpiece of the international magnetic confinement fusion program is ITER [1, 2], which will be built at Cadarache in France. It is a tokamak with a fusion power of 500 MW, with superconducting coils and a water-cooled blanket, with the coolant temperature limited to 150°C. It will lack equipment to convert the produced power into electricity, since the low temperature and low duty cycle would make such conversion uneconomic. The goal of ITER is to achieve a substantial energy gain[1] and long pulses (at least 400 s), thus demonstrating the practicality of fusion as an energy source. It will provide experience of operating an industrial-scale fusion reactor, enabling scientists and engineers to design reactors suitable for actual power production.

In order to be economic, future fusion reactors will probably need to be more compact than ITER, with a higher wall loading (the power carried as kinetic energy by neutrons per square meter of wall surface) and will also need a high capacity factor. This will mean that, at least near the plasma, the materials will be exposed to high fluences of high-energy neutrons. Developing materials capable of standing such fluences is a major challenge and will require lots of testing. With its modest wall loading and low duty cycle, ITER will be unable to provide the required fluence. Therefore, a dedicated facility for material testing will be needed. This is the reasoning behind IFMIF, the International Fusion Materials Irradiation Facility [3]. IFMIF will use two beams of 40 MeV deuterons, each with an intensity of 125 mA, impinging on a jet of liquid lithium to produce an intense flux of neutrons with energies up to 55 MeV over a volume of several liters and a lower flux over a larger volume. Although the maximum neutron energy of 55 MeV will be substantially higher than the 14.1 MeV maximum in a fusion reactor, most of the neutrons will actually be born at much lower energies. In any case, it has been estimated that the ratios of different types of material damage, such as displacements and hydrogen and helium production, will be similar to what can be expected in a fusion reactor. Thus IFMIF can, in a few years, provide data about the behavior of materials at fluences corresponding to decades of exposure in a fusion reactor. It can also test

---

[1] The ratio Q of the fusion power to the power fed in to heat the plasma and drive the current is intended to be at least 10.

the behavior of materials under simultaneous stress and irradiation or test tritium retention and release under irradiation, all at different temperatures.

In addition to facilities intended to actually be built, such as ITER and IFMIF, the fusion community has also prepared a number of preliminary designs for future reactors. These are not supposed to be built as such; their purpose is to get some idea of what future fusion reactors might be like and thus to clarify what research is needed. One such study is the European Power Plant Conceptual Study (PPCS). In the words of Andreani et al. [4],

> *"Within the European Power Plant Conceptual Study (PPCS) four fusion power plant "models", delivering 1500 MW of electric power to the network, have been developed, named PPCS-A to PPCS-D, which are illustrative of a wider spectrum of possibilities. For the two "near-term" models, PPCS-A and PPCS-B, the plasma physics scenarios are based on parameters slightly better than the design basis of ITER. The high plasma currents, the high current-drive power and the divertor heat load constraints drive these devices to larger size. The technology employed stems from the use of near-term solutions for the blanket, respectively the "water-cooled lithium-lead" concept for model A and the "helium-cooled pebble bed" concept for model B, on which large amount of work has already been performed in Europe. Associated with these are water-cooled and helium-cooled divertors. The water-cooled divertor is an extrapolation of the ITER design and can resist peak heat loads of 15 MW/m^2. A helium-cooled divertor has been chosen for model B to prevent the risk of hydrogen formation resulting from the water-beryllium reaction in case of accident and to simplify the balance of plant by using the same coolant for all internal components. This divertor must resist a peak heat load of at least 10 MW/m^2 and, to achieve a reasonable system efficiency, the high pressure helium pumping power should be limited to no more than 10% of the thermal power. Operating temperature however is limited below 550 °C by the use of EUROFER as structural material so that the benefit of the use of helium on thermodynamic efficiency is limited by the high pumping power.*

> *"PPCS-C and PPCS-D are based on progressive improvements in plasma physics, especially in relation to plasma shaping, stability and density*

limit, and on minimisation of divertor loads without penalisation of the core plasma conditions. In both cases higher fusion power densities and higher operating temperatures of the coolant and, consequently, higher efficiencies in energy conversion are foreseen. The blanket technology stems from a "dual-coolant" blanket concept for model C (helium for the first wall and high temperature self cooled lithium-lead breeder with steel structures protected by silicon carbide inserts acting as thermal insulators) and from a "self-cooled" blanket concept (even higher temperature lithium-lead coolant with a silicon carbide structure) for model D."

The present author performed some neutronics calculations for PPCS-C (Publication IV).

1.3 Typical features of fusion reactors and related facilities

Fusion reactors and related facilities have some characteristic differences from fission reactors, which influence the neutronics calculations. One obvious difference is that (assuming deuterium-tritium fusion and ignoring side reactions) all source neutrons will be born at an energy in the neighborhood of 14.1 MeV, though an exact representation of the source spectrum will require accounting for the thermal motion of the fusing nuclei. This is not a problem, since cross section libraries for neutronics calculations usually extend to energies above 14.1 MeV, typically 17.33 or 20 MeV.

Another, more problematic aspect is that fusion reactors contain lots of voids. To begin with, the plasma where the neutrons are born is essentially a void, as far as the neutrons are concerned. More difficult for the neutronics specialist is the fact that the structures surrounding the plasma contain voids that provide streaming paths for neutrons, with deleterious effects on the shielding capability of these components. For instance, in ITER the innermost structure, the blanket, consists of modules separated by gaps of about 2 cm thickness. Similar gaps separate the plugs to be inserted into the ports from the port walls, and these gaps perforce run right through both the blanket and the vacuum vessel, to locations where humans may have to work during shutdowns. The plugs themselves may contain

waveguides or other ducts. The designers try to insert bends or doglegs into these penetrations, but even so they represent weak spots in the shielding.

The situation is even worse in such locations as the neutral beam injection (NBI) ports, where large, straight void channels are unavoidable. Such cases will, however, not be considered in this thesis, being outside the scope of the author's work.

1.4 What method to use?

In principle, shielding calculations, like other neutronics calculations, can be performed using either deterministic or Monte Carlo methods. However, deterministic methods have difficulties in coping with the geometrically-complex mixture of shielding materials and voids typical of ITER and presumably other fusion reactors. Typically they use one or the other of two opposite methods of representing the angular dependence of the flux. The S_N method and the method of characteristics use a limited number of discrete flight directions. The P_L method expands the flux in terms of Legendre polynomials of the angular variables. The former method is plagued by ray effects when applied in voids, and the latter introduces a spurious angular spreading. On the other hand, it would presumably be possible to develop methods of calculating neutron fluxes in a geometry consisting mainly of voids, but such methods would be confounded by the fact that there are large amounts of shielding materials present in fusion reactors and by the geometric complexity.

Thus we have to resort to Monte Carlo[2]. Jaakko Leppänen's doctoral thesis [35] contains a very good introduction to this method, so it would be superfluous to provide one here.

Monte Carlo programs are capable of handling both material-filled cells and voids in a correct fashion, subject only to the need to accumulate enough histories to give adequate statistics. The method is also straightforward and intuitive, at least in particle transport problems. That does not necessarily mean

---

[2] Calculation of fluxes and related quantities through the simulation of individual particle histories, using pseudo-random numbers.

that it is easy to use, however. Getting good statistics can sometimes be a very challenging task, as will become apparent below.

In the ITER project, the MCNP program [5–9] has been chosen as a standard. This is a versatile, accurate, well-tested and widely used program. The calculations reported below were performed using MCNP4B, MCNP4C and MCNP5.1.40. In some cases modifications were made (not by the author personally). Specifically, three different modifications were used.

One was intended to model the source in a tokamak more accurately than the standard MCNP source description allows. In a tokamak geometry, this would require modeling the plasma as a number of discrete cells, each with a constant source density. This step-like representation was considered somewhat unsatisfactory and an alternative method was developed at the UKAEA to represent the source in r-θ-z geometry. A mesh of (r,z) points is used, in practice typically some 40*40 points, and each point is assigned an appropriate source intensity, calculated by an auxiliary program on the basis of the plasma density and temperature. The source is considered constant in the θ direction. This method is well adapted to a tokamak geometry.

Another source modification was needed for calculations on IFMIF. Stanislav Simakov at Forschungszentrum Karlsruhe wrote a modified version of MCNP called McDeLicious [10, 11], which calculates the neutron source produced by deuteron beams impinging on lithium. In principle this could presumably also have been calculated using MCNPX [12], but at the time McDeLicious was written, MCNPX did not give satisfactory accuracy. Note that in McDeLicious, the source normalization factor should be the number of deuterons hitting the target per second, not the number of neutrons emitted by the source.

The third modification is required for the Direct 1-Step (D1S) method of activation calculations, described in the next section. In addition to a program, one also needs to decide on a nuclear data library to be used with it. In the ITER project, FENDL, the Fusion Evaluated Nuclear Data Library, was chosen as a standard. The latest version is FENDL-2.1 (Dec. 2004) [13].

1.5 Shutdown dose rate calculations

A prominent issue in neutronics calculations for ITER is the calculation of dose rates with the reactor shut down. The specifications say that it must be possible to carry out hands-on maintenance at the ends of the ports 10^6 seconds (about 11.5 days) after shutdown. Assuming a maximum annual dose of 20 mSv and 200 hours of work per year, this means that the dose rate must be less than 100 μSv/h [14]. This has turned out to be the most difficult requirement to fulfill for the port shielding, much more difficult than limiting the heating in the superconducting coils to an acceptable level.

Calculating shutdown dose rates requires three things. First, one must calculate the neutron flux. Secondly, a calculation of the resulting activation is required. Note that this may be quite sensitive to certain trace materials, such as europium in concrete. Finally, one must then calculate the transport of the delayed gammas and tally the dose rates in the relevant locations.

There are two ways of going about this. One possibility is to do all these calculations separately. This requires two separate MCNP calculations and is accordingly called the Rigorous Two-Step (R2S) method [15]. Between the neutron and gamma calculations, an activation calculation is performed. FISPACT [16] is the preferred program for this purpose. In addition, certain auxiliary programs have been written to facilitate the use of FISPACT on a mass production basis for a large number of geometry cells.

A disadvantage of the R2S method is that the averaging of the flux over each tally cell results in a distortion of the spatial distribution of the activation reaction. The true distribution is approximated by a stepwise distribution. Appendix A provides a rough estimate of this discretization error.

On the other hand, the R2S method permits a full activation calculation, including multi-step reactions. It also allows one to use different geometries in the neutron and gamma calculations.

The other possibility is to handle the neutron and gamma calculations as a single MODE N P (combined neutron-photon) calculation. Then the photons will be born exactly where the neutrons collide, and the spatial discretization error is

eliminated. This method is called the Direct One-Step (D1S) method [17–19]. It requires some changes in MCNP and the data libraries, since the program is intended to handle only prompt gammas. It uses a kind of pseudo-time bins to separate the contributions from different radioactive nuclides. Each such contribution is multiplied by a factor taking into account the irradiation history and subsequent decay.

The D1S method obviously has an advantage in that it avoids the discretization error. It is also simpler to use, at least once the above-mentioned factors have been calculated – and provided one can get the modified program and data libraries to work, which may sometimes be problematic. It requires the activation to be directly proportional to the neutron flux, so multi-step reactions or depletion effects cannot be taken into account. However, one would not expect these to be important at the relatively modest fluences prevailing at the ends of ITER ports.

In principle, combining the neutron and gamma calculations requires the geometry to be the same in both steps. However, the modified version of MCNP currently allows some cells to be void in one step but not the other.

Activity levels after irradiation may be strongly affected by low levels of certain impurities such as europium. This is typically not taken into account in ITER dose rate calculations, which simply use some given composition, not including trace materials, for both transport and activation calculations. However, this is reasonable, since the materials to be used in those parts of ITER subject to significant neutron flux are designed for minimum activation, so the levels of problematic trace materials will presumably be kept low.

1.6 Variance reduction

A few of the calculations described in this work did not require variance reduction. For instance, the wall loading calculation for PPCS-C was easy. However, most of the cases were shielding calculations, and shielding calculations with the Monte Carlo method generally require a lot of variance

reduction[3]. After all, the purpose of shielding is to stop particles from getting where they are not wanted, and, as a result, an analog calculation will have very poor statistics.

The variance reduction techniques most used in the work reported here are implicit capture and splitting / Russian roulette. A few other techniques were occasionally used and will be mentioned in those contexts.

Implicit capture is used by default for neutrons in MCNP. In analog capture, when a particle collides, the track[4] is terminated with a probability equal to the probability of the particle being captured. In implicit capture, the track is not terminated, but its weight[5] is reduced through multiplication by the probability of not being captured. The track will thus continue until it leaks out of the system, by entering a cell with importance zero or a negative sill (weight window lower boundary),[6] or until it is terminated by some other means, typically weight cutoff.

The advantage of implicit capture is that it decreases the risk that a track will be terminated just as it is about to reach a region where it will be tallied. However, there must be some means of terminating a track. For some small systems, it may be feasible to rely on leakage to accomplish this. If an energy or time cutoff is used, they can also terminate tracks effectively. However, in large systems without an energy or time cutoff it is necessary to use some kind of weight cutoff. In this case, the track will be rouletted when the weight drops below a certain limit, so that it will either terminate or continue with an increased weight.

---

[3] Due to its dependence on random sampling, the Monte Carlo method suffers from stochastic uncertainty. This is usually expressed in terms of the **variance** or the **fractional standard deviation** (also called estimated relative error).

[4] In MCNP, a **history** begins when a particle is emitted from the source. It may split into several **tracks**, for instance due to splitting or (n,2n) reactions. Each track is terminated individually. A history ends when all its tracks have been terminated.

[5] In **analog** Monte Carlo, a particle is tracked until it is absorbed or leaves the relevant space, energy or time domain. In non-analog techniques, the number of particles tracked may differ from the physical number, and to compensate for this each track, is assigned a statistical **weight**. The calculation tries to determine the correct total weight in each element of phase space rather than the actual number of particles.

[6] The word "sill" for "weight window lower boundary" was introduced by Hodgdon [25]. Since it is a handy and vivid substitute for a frequently occurring, long and awkward term, it is used in this thesis.

If weight windows are used, the lower boundary of the weight window automatically provides such a cutoff. However, in many of the calculations described here, importances were used instead. Then the weight cutoff defined by the CUT:N card is needed. This card contains the parameters *WC1* and *WC2*. As the MCNP manual [8] (p. 3–135) says,

> "If a neutron's weight *WGT* falls below *WC2* times the ratio *R* of the source cell importance to the current cell importance, then with probability *WGT/(WC1*R)*, the neutron survives and is assigned *WGT = WC1*R*. If negative values are entered for the weight cutoffs, the values

$$|WC1|*W_s \text{ and } |WC2|*W_s$$

> will be used for *WC1* and *WC2*, respectively, where W_s is the minimum weight assigned to a source neutron from an MCNP general source. These negative entries are recommended for most problems."

The default values are *WC1* = -0.50 and *WC2* = -0.25. However, anybody who wishes to use MCNP with a modified source subroutine should bear in mind that the combination of negative *WC1* and *WC2* values with a non-standard source may not work. It has been observed to produce zero or negative weights. Therefore, it is the opinion of the author that when a non-standard source is used, *WC1* and *WC2* should be set to positive values. If the source produces particles with a weight of 1, *WC1* = 0.50 and *WC2* = 0.25 may be appropriate. In McDeLicious, on the other hand, the source particle weight varies but is typically about 0.08 (related to the probability that a deuteron produces a neutron). Then *WC1* = 0.04 and *WC2* = 0.02 may be appropriate.

Probably the most important variance reduction method is geometry splitting; that is, splitting when a particle moves into more important parts of the geometry (closer to a tally cell) and Russian roulette when it moves away. In MCNP this can be implemented using either importances or weight windows.

When importances are used, each cell is assigned an importance. If a particle moves into a cell with a higher importance than the cell it just left, it splits into the number of tracks given by the ratio of the importance in the cell entered to that in the cell it left, while the weight is divided by that same ratio. If the

importance ratio is not an integer, the splitting is done probabilistically so that the average number of tracks will be correct. The weight is still adjusted by the importance ratio, so in individual cases a non-integer ratio of importances will lead to a gain or loss of weight, but on average the figures come out right.

When a particle moves from a more important cell into a less important one, it undergoes Russian roulette. It survives with a probability equal to the ratio of the importance in the cell entered to that in the cell it left, and the weight is adjusted accordingly.

When weight windows are used, the lower boundary of the weight window is specified for each particle type and each cell (or mesh interval, when mesh-based weight windows are used). In addition, weight windows may also be energy or time dependent (not both in MCNP). Angle dependence is not allowed, which is unfortunate in shielding problems such as those considered here, but angle-dependent weight windows would presumably introduce too much additional complexity in the code and input to be practical.

If the weight of a particle drops below the lower boundary, the particle is rouletted. If the weight exceeds the upper bound, the particle is split. The details of the roulette or splitting are determined by parameters on the WWP card, which give the upper bound and the weight after surviving roulette as multiples of the lower bound.

MCNP contains an automatic weight window generator, which calculates weight windows on the basis of the total score from particles in each cell and energy or time bin. However, before the weight window generator can be used, the problem must be run with manually-supplied importances or weight windows. Moreover, the stochastic nature of the weight window generator means that it gives good results only when the variance reduction is already fairly effective. Iterative use of the generator will then produce successively improved results. The manual points out that the generated weight windows should be scrutinized by the user and unreasonable ones adjusted. This is good advice, but unfortunately such scrutiny and adjustment is very time-consuming and laborious when the model consists of thousands of cells, especially if it must be done repeatedly.

A much more extensive description of variance reduction techniques is given in the MCNP manual. In the MCNP5 manual, this description occupies pages 2–134 to 2–163 of Volume I [7].

There is, however, one important point where earlier editions of the MCNP manual are wrong and the MCNP5 manual dated 10/3/05 is somewhat unclear. This concerns the use of weight windows in universes.[7] For importances, the manual says correctly that "an importance in a cell that is in a universe is interpreted as a multiplier of the importance of the filled cell". The MCNP4C manual states that sills are treated similarly, but this is false. Test calculations show that only the weight window of the lowest-level cell is taken into account, any weight windows in filled cells are ignored. (In fact, the manual recommends that mesh-based weight windows be used when the geometry model uses repeated structures, i.e., universes or lattices. This way the issue is avoided. However, mesh-based weight windows may be unable to fit the geometry with sufficient precision.)

So, should one use importances or weight windows?

Weight windows have a number of advantages [7] (p. 2–145). However, entering them manually is even more laborious than doing this for importances, so usually it's best to do the first calculation with importances. Then one can use the weight window generator and recalculate with the resulting weight windows.

Moreover, it's often hard or impossible to avoid the use of universes in modeling complicated geometries. But the fact that only the weight windows for the lowest-level cells are taken into account often means that weight windows will not work as required with universes. Therefore, in such cases one has to use importances.

Fortunately importances and weight windows can be fairly easily converted into each other. There is an inverse relationship between them, so that a high importance means that the weights in that cell will be low, and consequently the sill should also be low. In fact, the WWP card in MCNP provides an option to

---

[7] In MCNP, a cell or a collection of cells can constitute a **universe**, which **fills** another cell.

calculate sills through dividing a constant by the importance for each cell. Similarly, importances can be obtained through dividing a constant by the sills. This is not available as an option in MCNP, but it can be done separately, using a spreadsheet program, so, if necessary, generated weight windows can be translated into importances.

There is also the issue of how large a part of the problem to solve at once. For instance, in shutdown dose rate calculations with the D1S method and in many other applications, combined neutron-photon calculations are necessary. However, it is more practical to start with a calculation covering only fast neutrons (using a CUT:N card to impose an energy cutoff). This can be performed relatively quickly, and the generated weight windows can then be extended to slower neutrons and to photons (suggestions on how to do this are given below). Then one can run a MODE N P calculation and generate improved weight windows, iterating as long as necessary. (If weight windows cannot be used, the generated weight windows must be converted to importances, which must apply to all energies, but one can still have different importances for different particle types.)

One can go further in subdividing the problem. Some analysts prefer to generate a first set of weight windows only for the part of the geometry nearest the source, by using a tally located only a short way into the shielding. It is then possible to get adequate statistics even with very crude starting importances, for instance setting all importances to 1. The generated weight windows can then be used to extend the calculation somewhat farther into the shielding, and so on until good weight windows have been generated for all cells. This stepwise approach is perfectly reasonable, but it involves a large number of calculations. Whether it is preferable to a more direct approach, tackling the whole geometry at once, comes down to a question of how much confidence one has in one's own ability to choose good starting importances. Advice is offered below on how to choose the starting importances, oriented toward the direct approach.

A common problem in shielding calculations with MCNP is that the program "freezes", "locks up" or "hangs", i.e., it keeps running indefinitely without producing any output or dumps. Ctrl-c has no effect, except that if it is repeated a sufficient number of times the run is killed.

It is not known for certain what causes this behavior. However, there are grounds for suspecting that it may be due to excessive splitting. Consider a shielding problem with a long and narrow duct, modeled as a void, running through the shield. The optimum importance will be much higher at the end of the duct furthest from the source than at the near end (or the weight window will be lower). Once in a long while a high-weight particle from the near end will fly straight along the duct, reaching the far end without any collisions on the way. It will then split into a very large number of tracks, so that this history will take a very long time (hours, days or more). Thus we call this the long history problem.

This diagnosis is supported by the fact that such behavior is usually observed in cases where there are voids connecting regions with very different importances. It also appears to be more common in calculations covering the whole energy range for neutrons, and especially in MODE N P calculations, than in calculations considering only fast neutrons, and likewise appears to be more common when absorption in the shielding material is weak. This is consistent with the above diagnosis, since an energy cutoff and high absorption both make for shorter histories.

Possible remedies for the long history problem include minimizing the importance ratio between cells at opposite ends of ducts or other penetrations, but this makes for less effective variance reduction. If only a single duct, of more or less circular cross section, is present, the use of DXTRAN is worth considering, but the geometries encountered in fusion neutronics typically involve numerous or awkwardly-shaped penetrations. John Hendricks has proposed using analog capture, which may help somewhat by terminating tracks sooner.

This dissertation essentially consists of two parts. In the first, the application of MCNP to practical problems in fusion neutronics is described. These calculations often proved to be very difficult, providing an impetus for investigations of how to use variance reduction methods (in particular, importances) as effectively as possible. These studies are described in the second part.

2. Applications

2.1 ITER
(Publications I, II and III)

ITER will contain a number of ports to give access to the plasma and the interior of the machine for various purposes such as auxiliary heating and current drive, diagnostics and maintenance. These ports will be at three different levels: upper, equatorial and divertor (lower) ports, with 18 upper ports and 9 divertor ports. At the equatorial level, there will be 3 neutral beam injection (NBI) ports and 14 non-NBI (regular) ports [2, 20].

With the exception of the NBI ports, the equatorial and upper ports will be filled by plugs which contain equipment for the purpose in question. Calculations are necessary to determine the heat production in the inner part of each plug and the adequacy or otherwise of the shielding properties of the plug. The latter involves two main criteria: the heating in the superconducting coils, especially the toroidal field coils, must be limited, and so must the shutdown dose rate around the end of each port (see above). It may also be necessary to determine the fluence to which certain components, such as vacuum windows, will be exposed.

The author has been involved mainly in calculations on plugs containing non-NBI heating and current drive equipment. This heating is done by electromagnetic waves in three different frequency ranges: the ion cyclotron (IC) resonance, lower hybrid (LH) and electron cyclotron (EC) resonance ranges. Launchers for all three kinds of heating will be present in some equatorial ports, and some upper ports will contain EC launchers. The author has performed calculations on all four kinds of plugs, as described in Publications I through III. His contribution has been to model the plugs themselves, insert these models into existing models of a 10- or 20-degree sector of ITER, modifying the general ITER model as necessary (some cells had to be subdivided to allow sufficient resolution for importances or weight windows) and to perform the calculations.

These plugs contain voids, e.g., waveguides, running through them, with bends or doglegs to limit neutron and gamma streaming. To make it possible to insert and remove them, they are separated from the vacuum vessel and blanket by a

gap, which is 2 cm thick in most places according to present designs. This gap contains a dogleg at the boundary between the blanket and the vacuum vessel.

These calculations were very challenging. On one hand, the system is designed to limit the radiation reaching the outer end of the ports, and the distances and material thicknesses involved are considerable. Thus a lot of variance reduction is required, such as use of large importances. On the other hand, there are streaming paths that make it possible for a particle from a low-importance region to reach a high-importance region without passing through material-filled cells on the way. This is a recipe for "long histories" (see above) and the calculations were in fact plagued by this problem, especially in the full MODE N P (combined neutron+photon) calculations required by the D1S method. To avoid long histories, it was necessary to limit the importance variation, which meant that the variance reduction could not be as effective as desired. Thus fractional standard deviations (FSD) below 0.1 could generally not be achieved with reasonable runtimes, say a week or less. Typical FSDs were about 0.2. This, combined with the simplifications and approximations (for instance, geometric simplifications or killing neutrons and, especially, photons in the less important parts of the system) meant that the calculated results are not likely to be very accurate. They could be in error by a factor of 2 or so either way. This is not a problem where the heating of the superconducting field coils is concerned, since this was in all cases found to be well below the stipulated limits. The shutdown dose rate, on the other hand, was more problematic. The calculated dose rates were generally below 100 μSv/h, as required, except for locations near the roots of the ports, where hands-on maintenance will probably not be needed. However, the safety margin to the 100 μSv/h limit was frequently less than a factor of 2.

Most of the shutdown dose rate is attributable to neutron streaming in the gap between the port plug and port walls. Flagging neutrons in this gap shows that they cause about 90% of the dose rate, whatever the interior of the plug may contain. This is therefore a generic problem, common to all non-NBI equatorial ports. Moreover, survey calculations have shown that the streaming is strongly affected by small changes in the shape and dimensions of the gap (though not by changes in the offset of the dogleg, so long as it is sufficient to prevent straight lines of sight along the gap). Thus, careful attention needs to be paid to design details affecting this gap.

In addition to heating and current drive plugs, the author has also been involved in calculations on one plug containing test blanket modules, specifically for the helium-cooled pebble bed blanket design; see Publication III. This work was shared with Petri Kotiluoto. The results were consistent with what was said above regarding equatorial ports in general.

2.2 PPCS, model C
(Publication IV)

Design studies have been performed for actual fusion power plants. One such study was the Power Plant Conceptual Study (PPCS), which involved four different models of 1500 MWe power plants with progressively more advanced technology, from the "near-term" models A and B to the very advanced model D:

- Model A uses a water-cooled lithium-lead (WCLL) blanket and, in general, technology close to the present state of the art.

- Model B uses a helium-cooled ceramic breeder blanket and slightly more advanced materials, technologies and plasma physics.

- Model C is a dual-coolant lithium-lead (DCLL) design, with the lithium-lead breeder material pumped through the blanket and heat exchangers, but with helium cooling for the steel structures. The technology level can be described as intermediate.

- Model D is a self-cooled lithium-lead design with SiC as the structural material.

The author was involved in neutronics calculations on an early, incomplete design for model C. Ports and other details were not represented, only the breeding blanket, hot shield, cold shield, vacuum vessel and toroidal field coils. The blanket used a "banana segment" design, unlike that described in Publication IV.

These calculations covered the tritium breeding ratio, the neutron wall loading, the nuclear heating, and calculation near the inboard mid-plane of the fast and total neutron flux, the helium production in Eurofer steel, and the epoxy dose, copper dpa, heating and fast and total neutron fluence at the front of the toroidal field coil.

2.2.1 Geometry modeling

The geometry modeled in this work was a 10-degree sector, one half of the spacing between toroidal field coils, assuming 18 such coils which was the value used at that time. No ports were included. The divertor was modeled only very roughly, since no information about it was available, and the calculations concentrated mainly on the midplane. On the other hand, the breeding blanket was modeled in some detail, based on the drawings then available. Beyond this was a shield, divided into hot and cold parts, and then the vacuum vessel and toroidal field coils (TFCs).

The blanket, shield and vacuum vessel were modeled using cylindrical rather than toroidal surfaces; see Fig. 1.
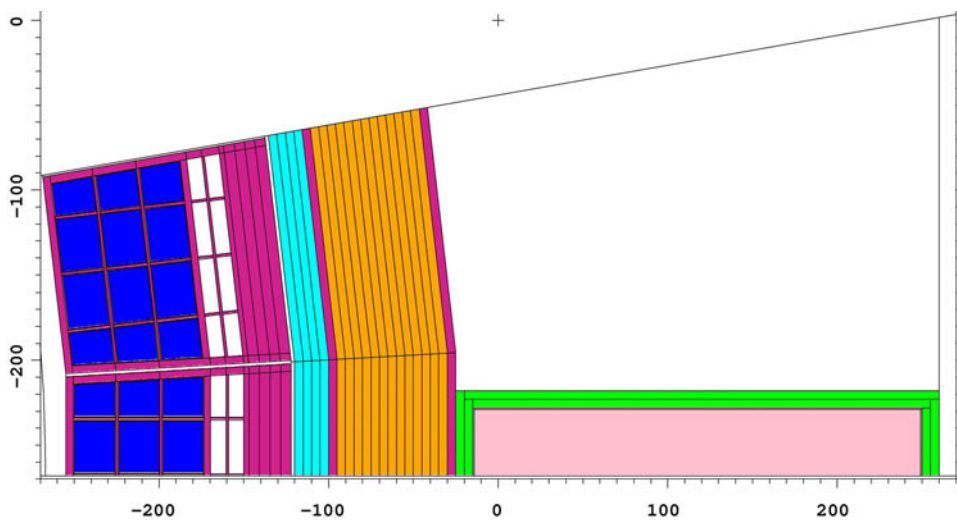
Figure 1. Plan view of the outboard part of the PPCS C model (cut near the equatorial plane). Dimensions in cm, relative to the picture origin (located at x = 1270 cm, y = 268 cm, z = 0.1 cm in the coordinate system centered on the machine axis). Colors: Purple = Eurofer low-activation steel, Blue = Pb-17Li (83 atomic % Pb, 17% Li enriched to 90% Li-6), Yellow = SiC (not visible in picture, present as a liner between steel and Pb-17Li), Cyan = 80 volume % Eurofer, 20% water, Orange = 60 volume % borated (2%) stainless steel ASTM-A887-89, 40% water, Green = 316-LN cryogenic stainless steel, Pink = homogenized toroidal field coil.

The curvature and locations of these surfaces were determined using the method described below (quoted from an informal memo by the author) [21]:

"*If the plasma dimensions in a tokamak are given in terms of the major radius R_o, the minor radius a, the elongation κ and the triangularity δ, the separatrix has the shape*

$$r = R_o + a\, cos\, (\theta + \delta\, sin\, \theta)$$

$$z = \kappa\, a\, sin\, \theta$$

in a cylindrical coordinate system, where the poloidal angle θ runs from 0 to 2π. (In a tokamak, the poloidal direction is the one in which the toroidal field coils run.) This shape can easily be plotted using Excel.

"*The above formulas are not exact, since there needs to be at least one X-point where helium and impurities can be bled off into the divertor. However, they do give at least a rough idea of the plasma shape. If no information is available on the location of the X-point, one can guess at it by drawing a pair of tangents to the separatrix defined above and calling the place where they cross the X-point, but this is obviously wildly inexact if there is no information where one should place these tangents.*

"*Sometimes it may be necessary to estimate the position of the first wall before the designers have decided where to locate it. Even at this stage a decision on the distance between the separatrix and the first wall will probably have been made, however. Then one can fit a series of circles to the separatrix and increase their radius by the given amount. The first wall can then be modeled as a series of toroidal or cylindrical surfaces.*"

2.2.2 Neutron wall loading

One quantity that is of interest in fusion neutronics is the neutron wall loading, the rate at which neutrons transfer kinetic energy through the first wall. In addition to the average neutron wall loading, the poloidal distribution is also required. This was easy to calculate. Once the source region was modeled, it was

surrounded with the surface at which the loading should be calculated, then it was just a matter of tallying and killing all neutrons crossing that surface.

The wall loading is customarily expressed in MW/m^2. Therefore the FM multiplier (typically used to normalize tallies to a given source strength [8], p. 3–95) should be the *neutron* power (80% of the fusion power) in MW, and the SD divider (used for areas, volumes or masses) for each segment of the tally surface should be the area in m^2. In a complicated geometry it is most practical to determine these areas stochastically, as described in the MCNP manual. The same often applies to the volumes needed in calculating many other quantities.

In the downward direction (around bin 27 in Fig. 2) the tallying surface was not the divertor surface but a horizontal surface spanning the mouth of the divertor and touching the plasma surface.

No variance reduction was needed for this calculation.
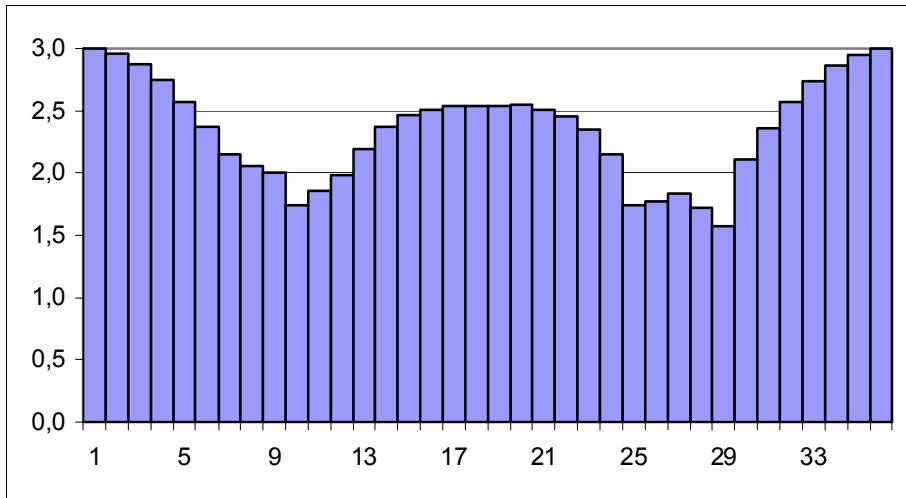
Figure 2. Poloidal distribution of neutron wall loading, unit MW/m^2. The numbers along the x axis are the poloidal angles in units of 10 degrees at the end of each 10-degree angular bin, counting up from the outward direction, then down along the inboard wall.

2.2.3 Tritium breeding ratio

As mentioned in the introduction, large-scale use of fusion energy will require fusion reactors to produce their own tritium. For this purpose a breeding blanket will be required, containing lithium, probably enriched with respect to Li-6, and a neutron multiplier. The most famous neutron multiplier is beryllium, but lead can also be used for this purpose. In fact, the (n,2n) cross section of lead is greater than that of beryllium above ca. 2 MeV.

In PPCS model C, the breeding material was chosen to be Pb-17Li, a mixture of lead and 17 atomic percent lithium enriched to 90 percent Li-6. Note that the shielding properties of this material are poor; the neutron flux declines much more slowly with distance than in a steel-water mixture. The tritium breeding ratio (TBR, the ratio of tritium atoms produced to those consumed) was calculated using data for total tritium production (reaction number MT = 205) in the FENDL-2 library. With the geometry model believed to be most realistic, a TBR of 1.156 was obtained. A value of at least 1.06 was considered necessary to ensure tritium self-sufficiency, since in a real machine such details as ports would lower the TBR.

Several variations of the model were tried, providing a kind of sensitivity study. The main result was that the TBR is highly sensitive to any water (or, presumably, other hydrogenous material) in the hot shield. 10 volume percent of water in the hot shield is enough to lower the TBR by more than 0.05. The value of 1.156 was obtained based on the assumption that no water can be present in the hot shield, since the temperature is too high.

The total thickness of lithium-lead is also important. Increasing the inboard PbLi thickness from 37.5 cm to 67.5 cm was found to increase the TBR by 0.069.

The influence of the water content in the cold shield was not studied. A water content of 20 volume percent was assumed. Subsequent shielding calculations suggested that a higher water content may be desirable, and perhaps this will lower the TBR. However, it is expected that the composition of the cold shield will have much less influence on the TBR than that of the hot shield, and with a calculated value of 1.156 there is enough margin to accommodate some lowering of the TBR.

In later versions of PPCS model C (not covered by the calculations), a more triangular plasma shape was chosen. This might raise the TBR still further, if it means that the outer blanket, with its greater PbLi thickness, covers a greater part of the poloidal circumference. On the other hand, if the blanket is broken up into modules smaller than the large ones assumed, that will lower the TBR somewhat, since some neutrons will disappear into the gaps between the modules and into the module walls.

The TBR calculations also required no variance reduction.

2.2.4 Shielding calculations

Calculating responses beyond the shield was somewhat more demanding, but the required variance reduction was still easy due to the absence of penetrations in the model. The following quantities were calculated for the inboard midplane region (from 50 cm below to 50 cm above the midplane):

- The dose rate [Gy] to the epoxy at the front of the winding pack, tallied in F2 (surface flux) tallies multiplied by the total cross section and the average heating number (MeV/collision). These quantities are available with the FM numbers -1 and -4 for neutrons, -5 and -6 for photons.

- The fast fluence (E > 0.1MeV)[cm^{-2}] at the front of the winding pack. This was calculated using a simple F4 (volume flux) tally with 2 energy bins.

- The displacement damage to copper [dpa] at the front of the winding pack. An F2 tally was used, multiplied by a dpa cross section entered using E (energy) and EM (energy multiplier) cards.

- The nuclear heating [Wcm^{-3}] at the front of the winding pack, calculated the same way as the epoxy dose rate.

- The helium production [atomic fraction] at the rear of the cold shield. This was calculated for Eurofer steel and for the elements Cr, Fe, Ni and B separately using the reaction cross sections for MT = 22 (n,n'α) and 107 (n,α) or, in the case of B, MT = 207 (total α production), in F2 (surface flux) or F4 (volume flux) tallies.

Calculations were performed for three water volume fractions in the cold shield: 20, 40 and 60%.

The results indicated that the design was adequate to keep the fast fluence and the heating density in the winding pack below the appropriate limits. The helium production, unlike the other quantities, increased monotonically with increasing water fraction, and was slightly above the limit at the higher water fractions. The epoxy dose and copper dpa had a minimum at a water volume fraction of about 40 to 50%, but even then they were well above the limits, by factors of 10 and 6, respectively. To solve this problem, the thickness of the cold shield and the vacuum vessel were increased inboard by 10 and 5 cm, respectively. Even this was not quite enough to bring the epoxy dose down to the desired level. It was thus necessary to change the material composition of the hot and cold shields so that both included 70 volume percent tungsten carbide. This finally brought down the epoxy dose to an acceptable value.

2.3 IFMIF
(Publications V–VIII)

2.3.1 Background and geometry

As mentioned in the introduction, the purpose of IFMIF is to test materials exposed to neutron irradiation in a spectrum sufficiently similar to that of a fusion reactor (in terms of the primary recoil spectrum and relative transmutation rates) and in a flux sufficient to provide relevant fluences (equivalent to 150 dpa$_{NRT}$) in a few years. The author was involved in modeling the test cell of IFMIF and carrying out some calculations.

In the model, the test cell of IFMIF, shown in Fig. 3, consists of a space 3 m long, 2.5 m wide and 6 m deep, surrounded by walls of heavy concrete, 3 m thick. The cover, also composed of heavy concrete, has a thickness of 2.4 m except for the removable plugs where the thickness is 2.7 m. The deuterons enter through two slits with a rectangular cross section in one wall and strike a jet of liquid lithium. Beyond the backplate shaping the lithium flow lies the high-flux test module (Fig. 4), in which small metal specimens suitable for embrittlement tests are irradiated at controlled temperatures, which may be different in each of

the 12 cassettes. Next comes the medium-flux test module, which in the basic model contained a creep fatigue testing machine, also referred to as a universal testing machine, which subjected samples to mechanical stress while they were being irradiated; a tungsten spectral shifter; and an in-situ tritium release experiment. Finally the low-flux / very-low-flux module contains some vertical tubes for insertion of relatively large specimens.

Each module is connected to the space above the test cell by several ducts penetrating the cover. The cover contains four removable plugs: VTA1 (Vertical Test Assembly 1) connects to the high-flux test module, VTA2 to the medium flux module, and the VIT (Vertical Insertion Tubes) plug to the low-flux / very-low-flux module. There is also the CSP (Central Shielding Plug) with no connections to anything in the test cell.

The model has undergone many modifications through the years. Fig. 5 shows the geometry according to md39, the 39[th] generation of the model.

The geometry is fairly complicated (though still much simpler than ITER), and the model is correspondingly complicated. (The input file md39 is 8193 lines long without tallies.) In particular the cover is penetrated by three sets of ducts, each with a different geometry, and in addition there are also large, gas-filled coolant ducts to cool the concrete in each of the four removable shield blocks. No information was available about coolant ducts in the fixed part of the cover, so none were modeled. Even without them, there are thus seven different sets of ducts, each with its own geometry and streaming properties.
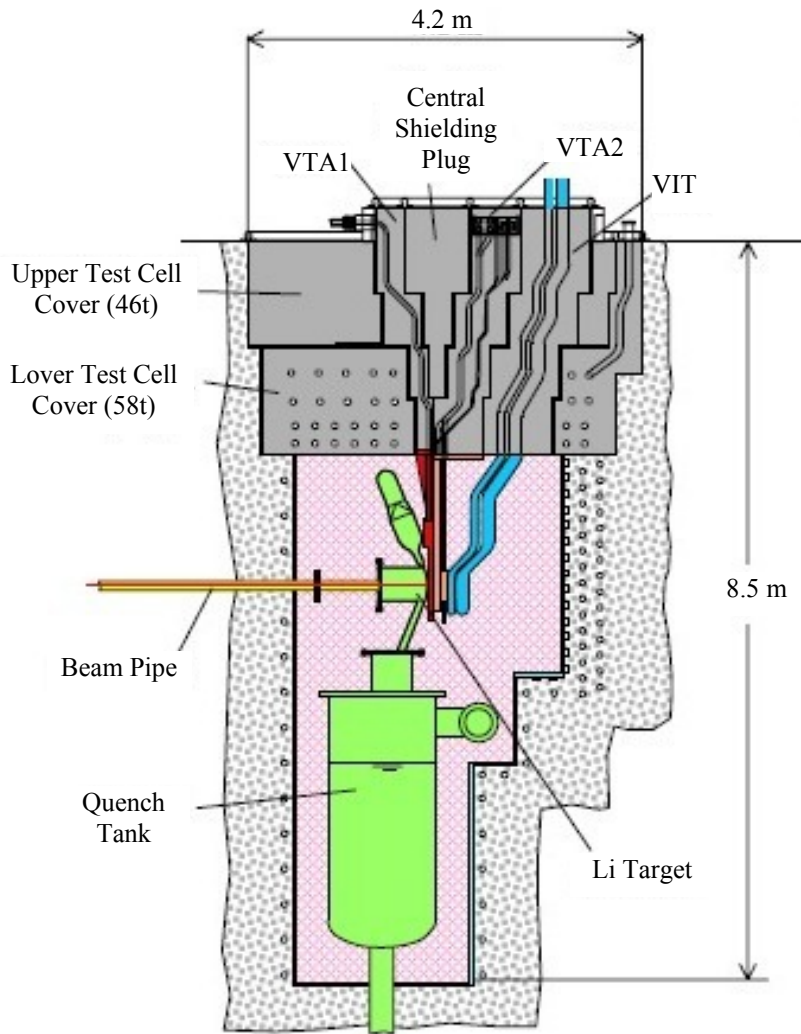
Figure 3. The test cell of IFMIF, according to A. Möslang, 7[th] *May 2002.*
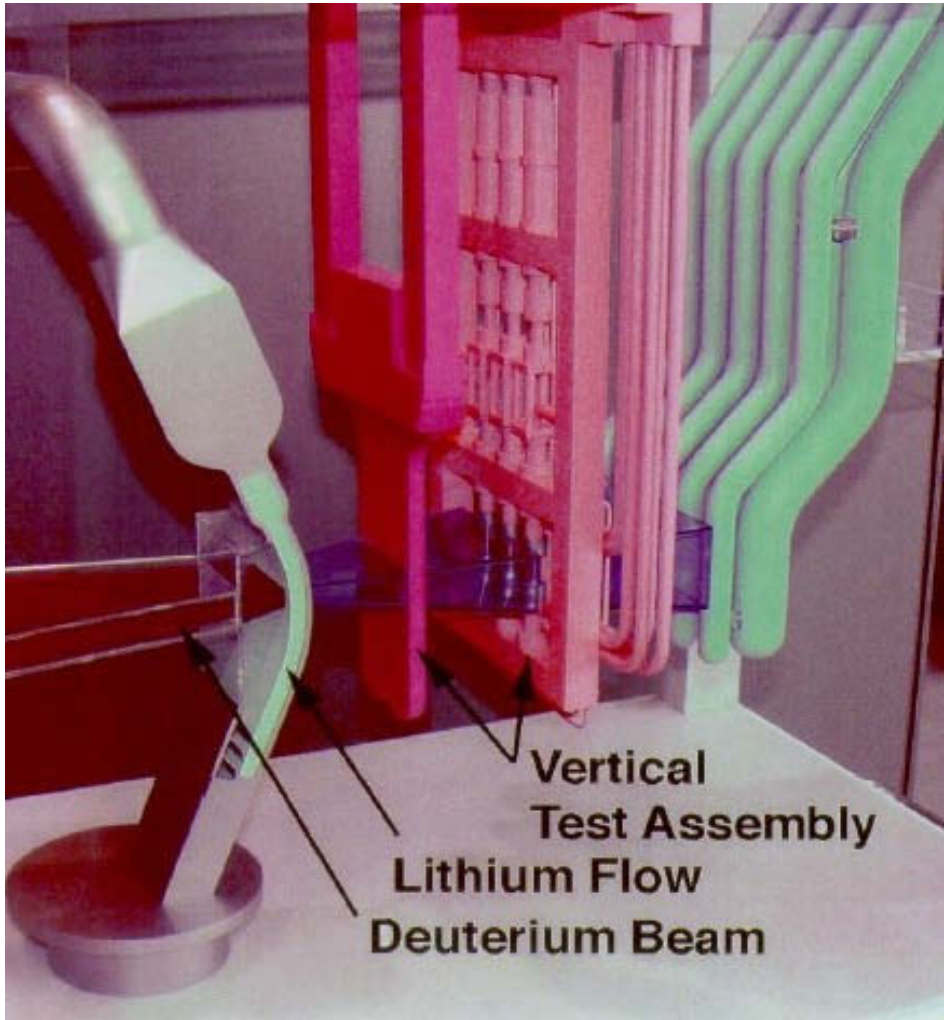
Figure 4. Exploded detail of IFMIF test cell, according to A. Möslang, 7th May 2002.
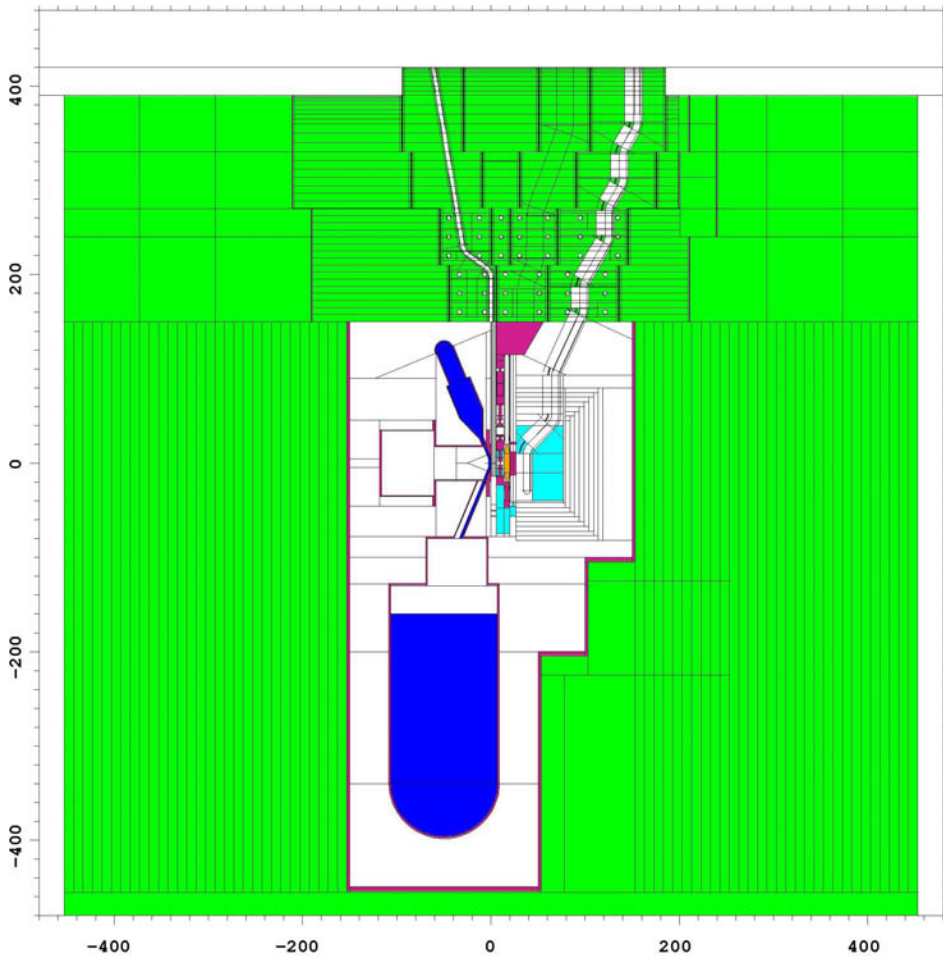
Figure 5. Elevation view of model md39 of IFMIF. Green = heavy concrete, Purple = steel (with very low density in the adapter connecting the medium flux test module to shield block VTA2), Blue = lithium, Cyan = graphite, Orange = tungsten. Certain volumes in the low-flux module and the cover are shown incorrectly as voids because the program has difficulties in plotting lattices. The cells surrounding the graphite reflector of the low-flux module are shown correctly as voids. These cells were added to model the horseshoe reflector but were voided in model md39.

The cover was modeled using universes, in such a way that the ducts were modeled in universes which were then used to fill cells that were stacked vertically to form the shield blocks. This was done before the author discovered

that the MCNP4C manual was misleading as explained in Section 1.6. It turned out to be an unfortunate choice, as described below in connection with the shielding calculations for the cover, since it precluded the use of cell-based weight windows in these calculations. However, modeling the cover in any other way would probably have been much more difficult.

Universes were also used to model the test equipment, at the request of one of the users of the model (P. Vladimirov). This makes it easier to move different parts with respect to each other. In particular, in some variants of the model the tungsten spectral shifter was moved from its original location between the creep fatigue testing machine and the in-situ tritium release experiment, to a location between the high-flux test module and the creep fatigue testing machine. In other variants, the thickness of the spectral shifter was reduced, and in one case thinner spectral shifters were placed in both locations. All this was much easier to do with universes. In this instance there was also no trouble with weight windows, because calculations on the test modules themselves do not require much variance reduction.

Most of the actual calculations were done by other people at FZK, but the author was involved in some calculations. The two most demanding sets of these calculations are described below.

2.3.2 Calculating neutron transport through the beam tubes

This work was part of the effort to determine the activation of the accelerators and adjacent structures in IFMIF. One source of activation will be the neutrons emerging from the test cell, mainly along the beam tubes. The task of the author was to calculate this neutron current in cooperation with Petri Kotiluoto, Stanislav Simakov and Ulrich Fischer. This current was then treated as a surface source and provided to Michael Loughlin at the UKAEA, who was responsible for the activation calculations. A detailed description of the calculations is given in Publication VIII, so only a summary will be provided here.

The calculations were performed using the program McDeLicious and a geometry model that was not quite up-to-date concerning the details of the cover but had a cell structure in the walls of the test cell that was more appropriate for

these calculations than that in later models. The main objective was to let the program store the neutrons emerging from the outer surface of the wall containing the beam tubes as a surface source. This could be done with the SSW (Surface Source Write) option, which records the location, direction, energy and weight of each particle crossing a specified surface. This surface source was the most important output of the calculations, but we included some tallies to check what was going on.

The neutrons emerging from the wall surface could be divided into three components:

- Neutrons coming directly from the source or scattered from the adjacent components and flying through the beam tubes without traversing any material.

- Neutrons colliding in the beam tube walls and scattering back into the beam tubes.

- Neutrons traveling, at least part of the way, through the test cell wall itself.

Since every point in the source, where the deuterons hit the lithium target, has a direct line of sight to the ends of both beam tubes without any intervening material, the first of these components is the most important. (In fact, it could be calculated analytically if the angular distribution of the source were known. It was not known to the author, however.)

The small solid angle subtended by the ends of the beam tubes means that only a small fraction of the neutrons produced will be included in this first component. This contribution cannot be emphasized through the use of importances or weight windows, nor is angular biasing of the source available in McDeLicious. DXTRAN could perhaps have been used, but would have been problematic due to the rectangular rather than circular cross section of the beam tubes, so it was not employed. That meant that the only way of reducing the variance was to use low importances to cut short the histories wandering into unimportant parts of the geometry.

The first calculation, with 7,500,000 histories, gave a quite respectable number of tracks in the surface source, namely 127,338. However, it was clear that about

99% of these were very low-weight tracks, belonging to the second and third of the components described above, and that including them in the calculation of neutron transport beyond the test cell wall (the responsibility of the UKAEA) would have been a waste of time. Thus the importances used in this calculation were poorly chosen and had to be modified in order to de-emphasize these components.

Improved importances did solve the problem of low-weight contributions. The next calculation, with 7,104,165 histories, gave 3,073 tracks in the surface source. This indicated that the low-weight tracks were much less numerous and had a higher weight, so this source was much more practical. However, the number of tracks was insufficient to get decent statistics in the next step, so longer runs were required.

This led to difficulties. Extending the number of histories beyond about 7,500,000 made the program freeze, similarly to the behavior described in Section 1.6. However, in this case it was clear that the "long history" explanation previously assumed for such behavior couldn't be correct, since it happened even when all importances were set to 1 or 0, so that no splitting could occur. Instead, the reason was found to be a programming error in McDeLicious, which was diagnosed and corrected by Dr Simakov.

Thus, in this case the "long history" hypothesis was not a correct explanation for the hanging run problem. However, the programming error involved was specific to McDeLicious, and hanging runs occur with standard MCNP as well. Considering how widely MCNP is used, it is extremely unlikely that a similar bug would have escaped detection. Moreover, while it is not feasible to examine in detail a history that takes hours, days or longer, limited attempts to study them by dumping event logs for them have been made, and the results are consistent with the hypothesis that they are due to excessive proliferation of tracks, often combined with long survival of each track before it is terminated.

While this work was going on, Dr Fischer made a proposal to ensure that adequate numbers of tracks would be made available in the surface source without incurring excessive run times. In the full geometry calculations, many histories waste a lot of time wandering around in the complicated geometry of the test modules, and these make a relatively modest contribution to the

emerging current, while the most important contribution comes from the source neutrons emerging directly through the beam tubes. Therefore, one could save time by simply killing all neutrons except in the target, the beam tubes and a few cells between and around these. This could be considered a more radical application of the same principle embodied in setting the importances to less than 1 in the test modules. An input file based on this proposal was written by the author and then run by Dr Simakov, since the author was otherwise occupied at the time. A run with 10^9 histories created a surface source file with 122,836 tracks.

This "reduced geometry" approach did involve an error, killing some neutrons that would otherwise have made it through the beam tubes or even through the wall. By comparing surface flux tallies for the full geometry and reduced geometry calculations, the author calculated a correction factor of 2.02263, by which the results obtained with the reduced geometry source should be multiplied. This source and the correction factor were then used by Dr Loughlin at the UKAEA to calculate the contact dose rates after shutdown at the raster scanner and other accelerator components near the test cell [22].

This method is not exact. It ignores the fact that the killed neutrons would have emerged with another (probably softer) spectrum than those included in the reduced geometry source and with a different angular distribution. Nonetheless, it is able to provide a rough estimate, sufficient for a dose rate calculation.

2.3.3 Cover shielding calculations

The other main calculation activity for IFMIF in which the author was involved concerned the dose rate in the access cell above the test cell, both during operation and after shutdown. This meant evaluating the shielding effectiveness of the cover, consisting of 2.7 m (2.4 m outside the central area) of heavy concrete, penetrated by about 40 ducts with a circular cross-section and a diameter of several cm. All of these ducts contain bends to limit the streaming, but these bends are rather gentle and in some cases barely sufficient to prevent unobstructed lines of sight. Thus this is a very challenging shielding problem, for the same reasons as those mentioned in connection with ITER calculations. In fact, it turned out to be even more difficult.

The first set of calculations attempted to estimate the dose rate during operation. Originally the goal was to perform a full MODE N P (neutrons + photons) calculation, but this proved very difficult, even using the techniques described in Publication X. Runs with sufficient variance reduction to produce useful results were plagued by a consistent tendency to hang. Event logs showed that this was not only due to track proliferation in some histories; a single track could last very long before termination, due to weak absorption in concrete.

One possible remedy was to use an energy cutoff, ensuring that tracks would be terminated reasonably rapidly by dropping below the cutoff. In this way it was possible to get an estimate of the fast flux (above 0.1 MeV). This showed that, during operation, the dose rate from fast neutrons alone – ignoring the contribution from epithermal and thermal neutrons and from gammas – could approach 2 millisieverts per hour, even when averaging over a substantial area (the top of the VTA2 shield block) and ignoring local hot spots. This suggests that hands-on work in the access cell during operation cannot be recommended.

The problem of determining the dose rate after shutdown remained. This required calculating the delayed gammas resulting from neutron irradiation of the cover. It was decided that the R2S method would be used.

The first step was to calculate the total neutron flux in the cover, especially near the upper surface. This was the same task that had proven intractable when trying to calculate the dose rate during operation, and continued attempts to accomplish it were no more successful, being consistently plagued by hanging runs. Finally it was decided that a calculation of the fast flux would again have to suffice, and such a calculation was performed, with the intent of following this up with activation calculations and calculations of photon transport and dose rates.

Shortly after this decision, the author attended the RPS2006 meeting, the American Nuclear Society's 14th Biennial Topical Meeting of the Radiation Protection and Shielding Division, in Carlsbad, New Mexico, USA. At one session of this meeting he brought up the subject of how to perform shielding calculations in a geometry like this. John Hendricks, one of the developers of MCNP, agreed that this was a very difficult problem. Nonetheless, he had a suggestion. This suggestion, which the author named the "tally source" method, is described in Section 3.3 below. This method was tried in a test calculation on

one of the seven sets of ducts. It worked, although it is very laborious and involves some loss of accuracy.

By the time this test calculation had been finished, it was clear that Association Euratom-Tekes would no longer be involved in IFMIF work, so the involvement of the author in IFMIF calculations was terminated. The reports describing the work done so far and data files constituting the results of that work are available to Forschungszentrum Karlsruhe, which will presumably continue the work.

3. Method development

3.1 Background

The preceding chapter shows that there is a clear need to develop efficient variance reduction methods. The most important variance reduction method in MCNP, as described in Section 1.6, is geometry splitting / Russian roulette. Implementing this requires the user to choose importances or weight windows.

There is a weight window generator in MCNP. However, it has weaknesses, some of which are described in Section 1.6. Besides these, there are cases where weight windows are inappropriate, for instance if the use of universes precludes the use of cell-based weight windows and mesh-based weight windows will not fit the geometry adequately.

In any case, when using the weight window generator in a deep penetration problem, it is difficult to get enough particles to score to provide usable weight windows. One possible approach is to start with importances chosen to be as close to optimal as the user's engineering judgment can provide. This is the approach espoused in this work, but there are other alternatives.

For instance, Sato, Iida and Nishitani [23] recommended the stepwise method mentioned in Section 1.6. The starting importances can be set to 1, but the first calculation involves a target tally close to the source, so one gets good statistics for it. The resulting weight windows are then used in the next step, making it feasible to use a target tally located farther into the shield. The procedure is repeated, until one reaches the location where one really wants to tally the flux.

The same article also recommended a backward transport calculation to set weight windows for gammas. One puts a gamma source in the tally cell(s) and follows the gammas towards the real source. If no gammas reach a given cell, this is taken to mean that it is so far from the tally region that gammas born in that cell have a negligible probability of scoring, so they are killed in the final calculation. Where the backward calculation gives a non-zero flux, this flux can be used to calculate a weight window lower bound. Strictly speaking, this backward calculation should be an adjoint calculation, but for gammas the forward and adjoint fluxes are likely to be sufficiently similar that the forward

flux can be considered a sufficient approximation for the adjoint flux, keeping in mind that the weight windows do not need to be very accurate to be useful.

Another article by Sato and Nishitani [24] (in Japanese) took a different stepwise approach. The target tally is set at its final location, but the densities of the shield cells are reduced. Once weight windows are generated for low densities, the densities may be increased, and so on until the true densities are reached.

Hodgdon [25] applied a backward method – similar to the one Sato, Iida and Nishitani used for gammas – to the neutrons themselves, also getting useful results in this way.

Many articles have been written about the use of deterministic adjoint calculations to determine sills for Monte Carlo. Here we will consider only a few of them.

Sweezy et al. [26] used PARTISN to provide the sills, with the intention of incorporating it into MCNP. The transfer of geometry data from MCNP to PARTISN was automated, based on a weight window mesh, but the fact that the material was taken to be that at the center of each mesh cell underscores the problem that a complicated geometry cannot be represented exactly in a regular mesh. Good results were obtained for an oil-well logging problem, but for a problem involving a dogleg duct in a shield, MCNP using the standard weight window generator performed much better.

Probably the best known combination of a Monte Carlo program with a deterministic program to calculate weight windows is A^3MCNP [27]. Omura et al. [28] described a version of this program, altered to model volume sources rather than just distributions of point sources. This version was tested in a cask problem and a reactor dosimetry problem and performed well even when calculating streaming through a ventilation duct in the cask problem. However, the width of this duct was 10 cm, which is rather wide compared with those described in Chapter 2.

Shahdatullah et al. [29] described calculations using the finite element code EVENT to find sills for MCNP. Good results were obtained even in a case with a dogleg duct through a shield, but again this was a duct with a width of 10 cm.

On the question of cell-based vs. mesh-based weight windows, Hendricks and Culbertson [30] wrote: *"Mesh-based windows can outperform cell-based windows. Cell-based windows are superior only if the geometry is optimally subdivided for variance reduction."* However, the problems they treated were not much like the streaming-dominated problems common in fusion neutronics.

Murata et al. [31] described a method of biasing the scattering direction in a Monte Carlo calculation. This can be very useful in a streaming calculation. Unfortunately, MCNP does not provide such a facility. In some cases, DXTRAN [7], p. 2–156 et seq., or the tally source method described in Subsection 2.3.3 and Section 3.3 can substitute.

Freiesleben, Richter, Seidel and Unholzer [32] described a comparison between a deep penetration experiment using steel and perspex (simulating water) as shielding materials and MCNP4A calculations. The results agreed within 30 percent at the most critical position.

Serikov et al. [33] provided a recent example of ITER shielding calculations, specifically for an upper port ECRH launcher.

The brief literature review above is, of course, very far from exhaustive. But it shows that most methods of finding importances or weight windows for deep penetration Monte Carlo calculations can be grouped into 3 broad categories:

- Direct: Start with the best importances your engineering judgment can provide for the whole geometry, then use the weight window generator to calculate improved weight windows, iteratively if necessary. The main weakness of this approach is that it requires a lot of skill and knowledge to choose the starting importances well. The advantages are that, if done right, it requires only a few intermediate calculations for weight window generation and no extra programs.

- Stepwise: As in the two articles by Sato et al. mentioned above, either start with a target (for the weight window generator) near the source, or reduce the densities. Then gradually move the target toward the final location of interest or increase the densities. The disadvantage of this approach is that it requires several calculation steps, which may be quite laborious if one follows the recommendation of the MCNP5 manual to

review and correct the weight windows at each step. The advantage is that it requires little skill in setting the starting importances. Even setting them all to 1 will do the job. In this case, too, no extra programs are required.

- Deterministic: Perform an adjoint calculation with a deterministic program, often based on the discrete ordinates method, to find optimal weight windows. This method requires a suitable deterministic program, either a stand-alone program, which will necessitate the creation of two separate geometry models, or combined with the Monte Carlo program into one program that requires only one input file. It is likely to be quite effective for reasonably homogeneous shields. For shields containing ducts and gaps it may be less effective, for two reasons: in the first place, a deterministic program using a regular mesh may not be able to model the geometry accurately. In the second place, even if the geometry is modeled accurately, deterministic methods tend to do poorly in voids; for instance, the discrete ordinates method is notorious for ray effects.

Other methods are also conceivable. For instance, the use of backward but non-adjoint Monte Carlo calculations, as in Refs. [23] and [25], may be considered a fourth class of methods. Although theoretically unsound, in that adjoint calculations would be preferable, it may work adequately, especially for gammas. Some Monte Carlo programs have the capability for adjoint calculations, but they tend to have limitations: for instance, MCNP can perform adjoint calculations only on a multigroup basis, and not with continuous energy.

For a variety of reasons, including the unavailability of combined deterministic-Monte Carlo programs at the institutions where the work was performed, the author chose the direct approach. Therefore, the main thrust of the research described below was to develop ways of choosing appropriate starting importances and a good cell structure, though the tally source method was also tested.

In choosing importances, the user faces two contradictory requirements. On one hand, he must use sufficient biasing (in this case, importances that increase fast enough, going through the shield) to ensure that enough particles penetrate to give adequate statistics. On the other hand, heavy biasing will make the individual histories longer, slowing down the calculation. In a simple solid shield (i.e., without penetrations and with all optical path lengths through the shield reasonably similar) it is not very difficult to strike a reasonable balance.

The MCNP manual suggests choosing the importances so that the number of tracks stays roughly constant, proceeding through the shield. The work described below applies the same principle to shields with penetrations, but one must keep in mind that this may in some cases lead to high-weight particles streaming through voids to cells with very high importances, leading to extreme splitting and to the feared "long history" phenomenon. Thus, when in doubt, it is probably better to underbias rather than overbias. (The MCNP manual agrees with this.)

One point worth keeping in mind is that, fortunately, it is not necessary to determine the exact optimum importances. Doing that would be as difficult as solving the transport problem itself, but any reasonable importances are likely to improve the performance of MCNP compared with doing nothing (keeping all importances at 1), so long as one avoids serious overbiasing.

3.2 Choosing importances
(Publications IX, X and XI)

Publications IX, X and XI all deal with the problem of how to choose importances or weight windows for a neutron or neutron-photon calculation on a shield with penetrations, in particular a shield penetrated by a planar slit with a dogleg, similar to the gap between a port plug and the port walls in an ITER non-NBI equatorial port. It is likely that the same principles can be applied to many different kinds of penetrations, though experience from the IFMIF cover shielding calculations suggest that having many penetrations near each other, so that particles may move from one penetration to another, will require some modification of the method presented here. Calculations of photon transport alone have not been investigated, but most likely some similar method would work for those, too.

The three publications deal with different stages of the process. Publication IX is essentially a first exploration of the problem, but it outlines a method of determining the longitudinal distribution of importances, i.e., in the direction of particle transport. It also finds that the weight window generator can give significantly improved performance, though this may partly be due to the fact that the next question had not been properly addressed. This question, the transverse importance distribution, is addressed in Publication X. No dogleg was present in

this study, but that should not detract from the applicability of the results to geometries with doglegs. In Publication XI the results of the two preceding publications are drawn together in a geometry with a dogleg and the resulting method is tested.

3.2.1 The longitudinal importance distribution: the Handy Method

In determining the longitudinal importance distribution, we follow the advice that the track density should be kept roughly constant. First we consider transport through a bulk shield. In such a case the flux and current will decline in a roughly exponential manner, so we divide the shield into layers, each with a thickness roughly similar to, or somewhat greater than, the halving distance h (the shield thickness required to decrease the flux by an additional factor of 2). The value of this halving distance will vary from case to case, but in these investigations we are concerned mainly with fusion reactor applications. In a typical fusion reactor spectrum, experience has shown that an optimum mixture of water and stainless steel (about 25 and 75 volume %, respectively) gives a halving length of about 7.5 cm. (Actually the value for a pure bulk shield might be closer to 4.5 cm, but we are not dealing with such shields here.) Thus we get importances doubling at steps of 7.5 cm.

On the other hand, we also need to consider streaming along the penetration. In a duct with a circular cross-section, we assume that the streaming component of the current, at a certain distance into the duct, is proportional to the solid angle subtended by the opening of the duct at this distance. It is not obvious exactly what the constant of proportionality should be, but it is also not very important. If the radius of the duct is r and the distance from its beginning is x, we can take the flux to be $(r/x)^2$ times the flux at the beginning of the duct (obviously only when $x > r$, otherwise use the factor 1).

If we have a planar slit of thickness d instead of a cylindrical duct, the corresponding attenuation factor could be set to d/x, again only for $x > d$. Alternatively a slightly different constant of proportionality could be used; it doesn't matter much.

In any case, the inverse of this streaming component of the current will give an alternative set of importances. Now, for each layer of the shield we simply

choose the smaller of the two importances. Usually, when r or d is small compared to h, the bulk importances will be lower near the beginning of the shield (when x is small), but the streaming importances will be smaller farther on. If the latter is not the case, this suggests that the shield is not thick enough to reduce the flux or current to the point where streaming dominates.

After a bend or dogleg, one can repeat the process, starting with the importance derived for the cells just before the bend. This issue will be considered in somewhat more detail in Subsection 3.2.3.

As a final step, it's useful to change the importance to the highest power of 2 that does not exceed the importance estimated by the above method. This will ensure that the ratios between importances are always integers, avoiding the non-conservation of weight that may arise in individual histories if this is not the case.

This method has been named the "Handy Method", because it is similar, though not identical, to a method with this name of obtaining rough flux estimates behind shields with penetrations, proposed by Iida et al [34].

Publication IX introduced and tested this method in a simple, two-dimensional test case, a shield penetrated by a slit with a dogleg. The resulting importances were found to work adequately as starting values, but generated weight windows significantly improved the performance. It is likely that much of this improvement was due to the fact that the issue of the transverse variation of the importances was almost ignored when applying the Handy Method. Only a slight transverse variation was used (reducing the importances by a factor of 2 at some distance from the slit). This issue will be considered in the next subsection.

The Handy Method is a direct method in the sense of Section 1.6, in that it determines a set of starting importances for the whole geometry. In Publication IX, the stepwise or progressive method of determining weight windows (see Section 1.6) was also tested and found to give performance comparable to the direct method. The choice between these methods is to some extent a matter of taste, but for users who lack experience at choosing importances, the stepwise method may be better.

For calculations covering the whole neutron energy range, and even more for coupled neutron-photon calculations, it is usually advisable to start with a calculation covering only fast neutrons. One can then generate weight windows, which may give improved performance. This can be iterated if necessary. Once the fast neutron weight windows (or importances) are satisfactory, they can be extended to lower energies and, if required, to photons. In Publication IX, the alternative approach of starting with a coupled neutron-photon calculation right away was also tested, and in this simple case it worked, but it is likely to work less well in more complicated cases.

Mesh-based weight windows were also tried, as were the exponential transform and forced collisions. The results obtained with mesh-based weight windows were disappointing, but it is not clear why. One cannot conclude from a single series of test calculations that mesh-based weight windows are unsatisfactory. Unfortunately, schedule and funding limits precluded an investigation of their poor performance. Both the exponential transform and forced collisions gave small improvements in performance. For this particular kind of problem it seems questionable whether this small improvement is sufficient to justify the extra work these methods require, though of course in some problems the benefit may be greater.

3.2.2 Transverse importance distribution

In Publication IX, the improvement obtained using the weight window generator did not consist of reducing the fractional standard deviations for a given number of histories. On the contrary, the fractional standard deviations increased in some cases, but the runtime decreased substantially, improving the figure of merit (defined as $1/R^2T$, where R is the fractional standard deviation and T the runtime). This means that, within a given time, one can run more histories and thus obtain better statistics.

This behavior is a sign that the generated weight windows were higher – corresponding to lower importances – than the manually determined importances, at least in some cells. This is confirmed by Figure 3 of Publication IX, reproduced below as Fig. 6.
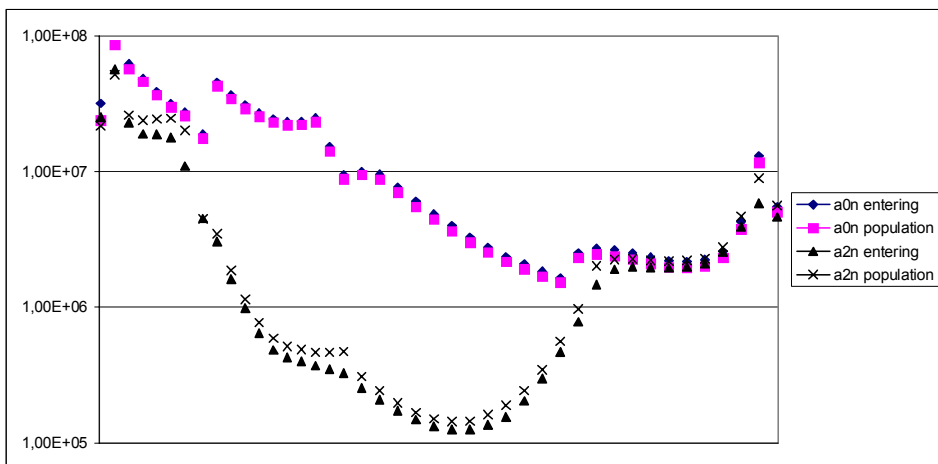
Figure 6. Fast neutron populations in two runs of the test problem described in Publication IX. a0n = run with initial importances, a2n = run using generated weight windows. "Entering" and "population" refer to the "tracks entering" and "population" columns of print table 126 in the MCNP output, the particle activity table. All values are summed over each transverse slice of the geometry, i.e., for all cells in a particular x interval. The horizontal axis shows x, the distance into the shield, from 0 to 302 cm.

We can see that the starting importances did a good job of keeping the number of tracks roughly constant, as recommended in the manual. (Note that the actual flux was reduced by about 7 or 8 orders of magnitude from one end of the geometry to the other.) The generated weight windows did not do this to anything like the same extent. The dip in the middle part shows that here, the weight windows were higher than what would have corresponded to the starting importances. This appears strange at first sight, but there are two things that explain this behavior.

In the first place, the importances derived using the Handy Method as described above are actually appropriate only for the slit itself. However, the cells in the slit were modeled as voids, and MCNP does not split in a void. (Doing so would, in fact, be pointless, since all the progeny of a particle would follow the same path until it entered a non-void cell.) Thus, however high the importance or low the weight window in a void cell may be, it has no effect.

Secondly, the optimal importance obviously declines as one moves away from the slit. As pointed out above, this was almost neglected in the manually-chosen importances used in Publication IX, since the author had no idea of how sharply the importances should decline. This was the subject chosen for investigation in Publication X.

In this work the author used a simple slab-type shield, 240 cm thick, with a straight penetration. The shield was divided into a number of cells, in both the longitudinal and transverse directions. The neutron current on the far side of the shield was tallied. The weight window generator was used, and the weight windows were then translated back into importances by dividing the lower boundary of the weight window in the source cell by the lower boundary for the cell in question. The resulting importance distribution was investigated and conclusions drawn.

First some sensitivity studies were performed to check which features of the model had a major influence on the results. The geometry of the penetration did not make a big difference, so a planar slit 2 cm thick was used. This is similar to the gaps between plugs and port walls in ITER. A monoenergetic 14 MeV source and a fission spectrum source were also found to give similar results, so the former was used. The energy cutoff was also not very important, so 0.1 MeV was used.

On the other hand, changing the shield material had a large influence, so calculations were made for seven different materials: water, quartz (a simplified version of concrete), light concrete, pure iron, stainless steel 316 L(N) ITER grade, stainless steel with 20 volume % water at a density of 0.9 g/cm^3, and tungsten. The calculations showed pure iron to be a very poor shielding material, presumably due to the presence of cross section windows. The best shield materials were tungsten and the 80% steel + 20% water mixture.

The hopefully optimal importances calculated by the above method agreed well with the results of the Handy Method in the slit itself. However, that is not in important in itself, since MCNP does not split in voids. Thus these importances matter only as a scaffolding, upon which one can construct the importances to be used in non-void cells, and the results of interest concern the transverse importance distribution. The decline of this distribution with increasing distance from the slit was striking.

Two aspects of the transverse importance distribution were considered: the ratio of the importance in the cells immediately adjacent to the slit (called "wall cells" in Publication X) to that in the slit itself, and the rate at which the importance declined further as a function of distance from the slit.

In both respects, the transverse importance distribution is relatively flat at the surfaces of the shield, particularly the far surface (where the current is tallied) and much less flat in the middle. It is, in fact, obvious that the transverse importance distribution must be flat near the far side. A neutron that has made it that far has about the same chance of scoring whether it is near the slit or far from it. On the other hand, in the middle of the shield, any neutron that strays far from the slit has only a very small chance of reaching the far side.

Even a neutron in a wall cell has a much smaller chance of scoring than one that is streaming through the slit, which is illustrated by Fig. 7. With increasing distance from the slit, this chance quite rapidly declines further in good shield materials; see Fig. 8.

From these figures, one can conclude that the following recipe should give reasonably good importances in a thick shield of some good shielding material (unlike pure iron) with a straight penetration:

1) The ratio of the importances adjacent to the penetration to those in the penetration itself, derived from the Handy Method, should be about 0.1 at the beginning of the shield, declining to about 0.01 one quarter of the way through the shield, remaining close to that value up to three quarters of the thickness and finally rising to 1 at the end.

2) In the best shielding materials, it is appropriate to make the cells adjacent to the penetration about 1 mm thick or less in the transverse direction. In the next layer of cells, the importances should be lower by a factor of 4 (except near the end, and perhaps the beginning) of the shield, and these cells should extend some millimeters further from the penetration. Beyond that, a further reduction by a factor of 4 or more is appropriate. For shielding materials that are not so good, a somewhat flatter transverse importance distribution is appropriate, maybe increasing the thickness of the second cell layer to several centimeters.

Of course, a more detailed subdivision would give a better fit to Fig. 8, but it is doubtful whether the benefits would be worth the added complexity in the geometry modeling.

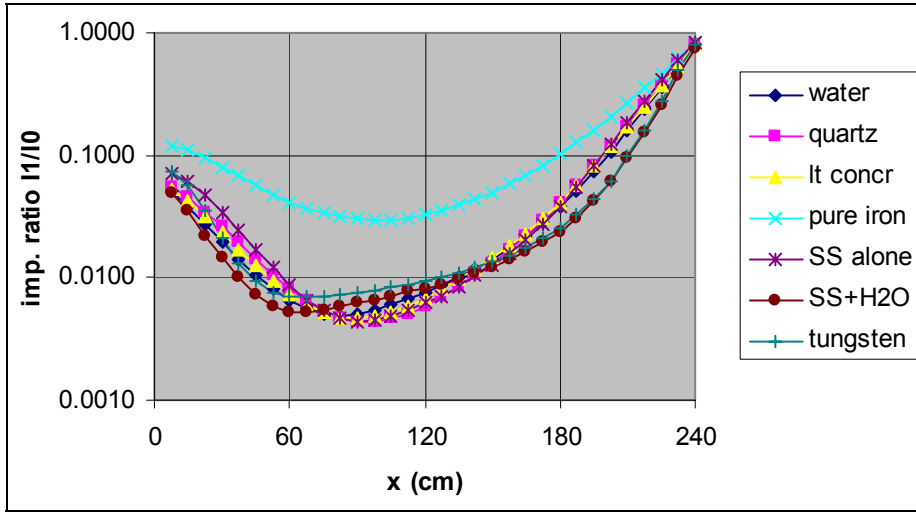*Figure 7. Ratios of importance in the first 0.001 cm of the wall (I1) to that in the slit (I0) as functions of the longitudinal coordinate x.*
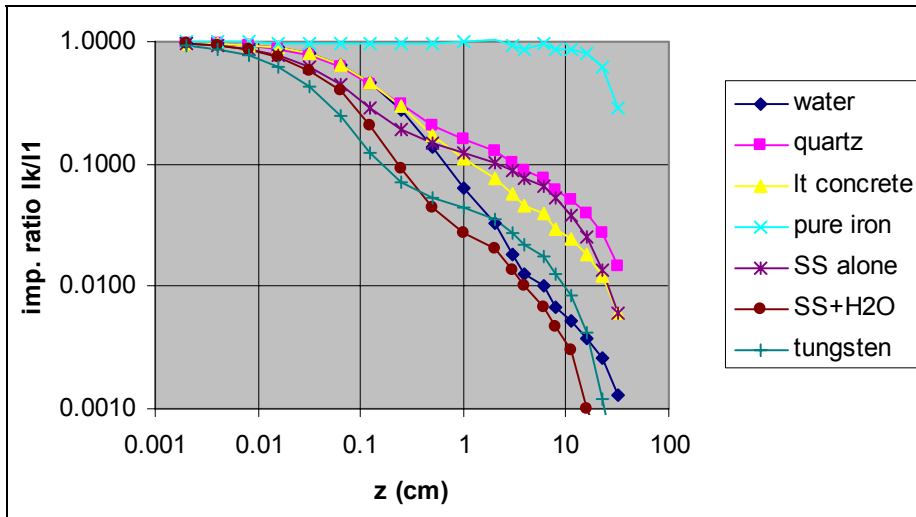

Figure 8. Ratios of importance in the k'th cell (Ik) to that in the first 0.001 cm of the wall (I1) as functions of the lateral coordinate z in the interval 112.5 cm < x < 120 cm (in the middle of the shield).

3.2.3 Testing the methods

In Publication XI the author tested the methods described above in a geometry representing a very simplified version of an ITER non-NBI equatorial port, shown in Fig. 9. The calculations simulated a shutdown dose rate calculation, though they were performed for prompt rather than delayed gammas.
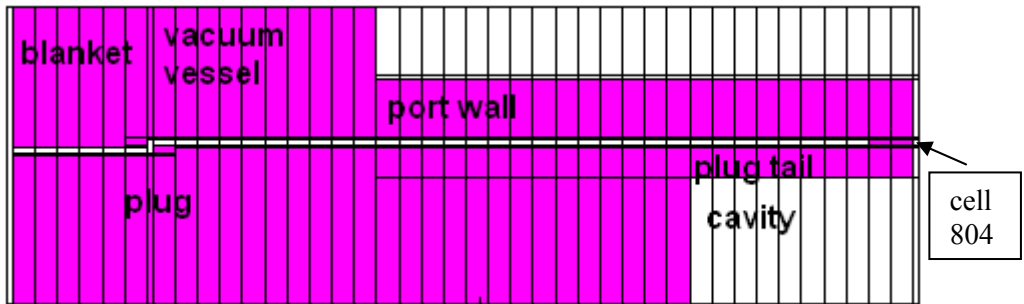

Figure 9. Geometry used in Publication XI (homogeneous version).

Starting importances for the fast flux were estimated using the Handy Method and transverse distributions based on the work described in the preceding subsection. These gave good results for the fast flux at the end of the port wall, but the fact that one also needs to calculate the dose rate around the port walls complicated the issue. It necessitated the use of quite different transverse importance distributions in the outer parts of the geometry, where particles traveling upward through the port wall or downward through the plug tail could contribute to the score. Once this was done and the tally used for the weight generator was modified correspondingly, good results for the fast flux and good weight windows were obtained.

The next step was then to extend the calculation to epithermal and thermal neutrons and to gammas. This necessitated extending the weight windows to these particle types, which was done by multiplying the fast neutron weight windows by appropriate factors, called weight factors in Publication XI. Since it was not known what values might be suitable for these weight factors, they were initially mostly set to 1, and in some cases to 0.1. A calculation using the resulting weight windows was then performed and new weight windows were generated.

The generated weight windows were then analyzed to find improved values for the weight factors. To reduce the stochastic noise, the cells in the model were divided into classes. In each cell, the weight factors needed to obtain the weight windows for epithermal and thermal neutrons and for gammas from the fast neutrons were calculated, and then these factors were averaged over each class. Geometric rather than arithmetic averages were used to limit the influence of excessively large values caused by undersampling. The resulting averages were then rounded, and the rounded weight factors were used to calculate new weight windows, which were then tested in a new calculation.

This calculation gave reasonably good results for the photon flux on the top surface of the port wall, but at the end of the wall the results were very poor. The fractional standard deviation in cell 804, the tally cell in the footprint of the gap, was about 5.5 times as high as in the previous calculation in spite of a doubling of the number of histories.

It is plain that the composite tally cell used in generating the weight windows, consisting of cells both on top of the port wall and at the end of the wall, was not well suited for this application. For fast neutrons it worked well, presumably because the collimated flux at the end of the gap ensured that enough fast neutrons reached cell 804 so that the tally was not excessively dominated by the cells on the top surface. The gammas, on the other hand, were emitted isotropically from wherever they were born, so that many more of them reached the top surface of the port wall than the end. Thus, the weight windows generated with the composite tally cell emphasized upward transport through the tally wall to such a degree that these weight windows did not work well for the gamma flux at the end of the port.

Some minor adjustments of the weight factors were tried first but proved inadequate. Then a new set of weight windows was generated using a tally cell covering only the end of the port wall[8]. The weight factors obtained in this way were compared with the previous ones, and a set of compromise weight factors was worked out. These factors are shown in a table in Publication XI. When used to determine weight windows for the next calculation, they worked well.

---

[8] The weight windows were generated on the basis of a photon energy tally. They were derived for three neutron energy groups, with boundaries at 1 eV and 0.1 MeV, and one photon group.

Usually one can expect further improvements from using the weight window generator to obtain new and better weight windows, iterating the process as far as necessary. In this problem, however, it would then be necessary to use two different sets of weight windows, which one could perhaps combine somehow, for instance by using the lower weight window for each cell. It would probably be more practical to perform separate calculations for the top and end surfaces of the port wall if one wishes to iterate the weight windows.

In the last calculation described in Publication XI, a lower limit was imposed on the weight windows so that extremely low values for the sills were not permitted. The motivation for this was that such a measure may help avoid hanging runs due to long histories. No hanging runs had occurred in this work, but they are a common problem, so testing a measure intended to protect against them was considered justified. It turned out that this actually improved the performance. There was some slight increase in the fractional standard deviations, but the run went more than twice as fast, so the figure of merit improved. This means that the procedure outlined above seems to lead to some degree of overbiasing, so imposition of such a limit appears to be a good idea.

3.3 The "tally source" method (Publication XII)

Even with the methods described above, some problems remain intractable. One example, mentioned in Subsection 2.3.3, was the cover shielding calculation for the IFMIF test cell. As noted in that subsection, John Hendricks made a suggestion, which can be described as follows:

> Tally the neutron current at a suitable location, specifically at the beginning of a long straight stretch of duct, taking care to ensure sufficient resolution in space, angle and energy. Then define a surface source based on this tally, biasing it in angle to favor neutrons streaming along the duct. Use this source to calculate the neutron transport to the next such location and repeat until the end of the shield is reached.

This should not be confused with writing and reading a surface source using the SSW and SSR options in MCNP, since that provides no opportunity for angular

biasing. The advantage of this method is specifically that it provides a way of biasing the flux in angle. This ensures that the neutrons that matter most, those that stream along the ducts, will be better sampled.

If weight windows rather than importances were used, it would also alleviate the risk of long histories, since the neutrons most likely to reach cells with low weight windows would already have low weights and would thus undergo less splitting. Unfortunately, the use of cell-based weight windows was impractical in the cover shielding problem due to the way universes were used in modeling the cover, and revising the model was also considered unfeasible. Modeling the cover in a different way would probably have required much more work than the original modeling. Mesh-based weight windows might have been usable, but they would have fit the geometry poorly.

A test calculation was performed for one of the seven sets of ducts in the test cell cover, as described in Publication XII. The through-going VTA1 ducts were chosen as a test case, because they are relatively simple. These 8 ducts consist of three straight stretches of duct, two of them relatively short and the third long, connected by bends. One such duct can be seen in Fig. 10, where it is the leftmost duct in the cover.

The first tally source was deployed at the entrance of the ducts, at the bottom surface of the cover. The neutrons striking the cover were tallied in space, energy and angle bins. The tallies were analyzed to find a way to represent the current as a surface source, within the limitations imposed by the SDEF source description in MCNP4C. It turned out that the angular distribution depended on the spatial location and the energy spectrum depended on the angle. However, MCNP4C did not allow more than one level of dependence. Since the most important angle bins of the source in this location are the ones closest to the upward vertical, and since the spectrum was reasonably constant for these bins, this spectrum was used, ignoring the softer spectrum in the tally bins farther from the upward vertical. The resulting source was biased in space, angle and energy, to emphasize the most important particles. One obvious drawback of this is a large weight variation among the particles.

The second tally source was deployed on the inclined plane delimiting the bottom end of the topmost straight segment of the ducts (surface 4773 in the

model). In this case the angular distribution was independent of location. Thus the location and angle could be treated as independent variables, while the spectrum was modeled as dependent on the angle bin. The need to use areas of different shapes in modeling the location dependence of the source caused some difficulties, which were resolved by treating a rectangle as a circle, from which the excess parts were trimmed by using cookie cutter cell rejection. The resulting source was then used in a calculation extending to the upper surface of the cover.

Following this, a sequence of three calculations was needed: from the source in the lithium jet to the lower surface of the cover, from that surface to surface 4773, and finally from there to the upper surface of the cover. In each of the first two calculations the parts of the geometry beyond the tally surface had their importances set to zero, killing all neutrons passing the tally surface. Thus, neutrons reflected back and forth across this surface were not included in the tally source. On the other hand, in each subsequent calculation, the full geometry of the preceding calculation was retained, so reflected neutrons were taken into account at this stage.
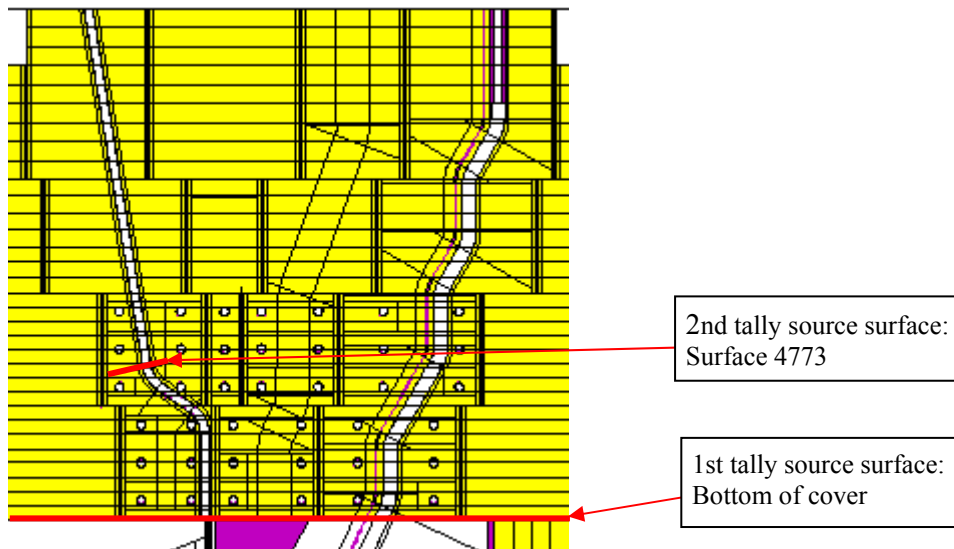

Figure 10. Tally source surfaces.

The result desired from this sequence of calculations was the neutron flux near the ducts, which is required for activation calculations. This was tallied using an F4 tally. In the part of the cover below surface 4773 it was necessary to tally the neutrons in both the second and third calculations. The second gave the contribution from neutrons which had not passed surface 4773, the third the contribution from those which had passed this surface and then been reflected. These two contributions had to be added to get the total flux. Since the auxiliary programs used to prepare the flux data for a FISPACT calculation take this data from MCTAL files, another auxiliary program was written to add fluxes from two or more such files.

As we see, this method involves far more work than an ordinary MCNP shielding calculation. Each duct type has to be treated separately, and several calculation steps are needed, each requiring a substantial amount of work. Moreover, the discretization involved in tallying the current and then defining a source on the basis of this tally involves some loss of accuracy. However, against this one must set the fact that the tally source method can solve otherwise intractable problems.

4. Summary and conclusions

This dissertation describes some experience of using MCNP, mostly MCNP4C, for fusion neutronics calculations, especially shielding calculations for ITER and IFMIF, as detailed in Publications I through VIII. There are good reasons why the Monte Carlo method is used for such problems. In particular, deterministic programs have difficulties in dealing with the mixture of voids and thick shields. On the other hand, it is not easy for Monte Carlo calculations to handle either. Since the whole point of shielding is to reduce the neutron and gamma flux radically, heavy splitting is required in order to get useful results. Conversely, the presence of narrow penetrations means that, once in a while, a particle emanating from a low-importance cell will reach a high-importance cell without splitting on the way. It may then split into an excessive number of tracks, so that a prohibitive amount of time is spent on a single history.

It is therefore essential to choose the importances or weight windows well. The weight window generator in MCNP is useful, but it has limitations. There is consequently a need for some way of choosing good starting importances. In addition, it is also necessary to choose the cell structure of the geometry model in a way that provides sufficient resolution for importances or weight windows, without making the cells so small that the weight window generator is hampered by unacceptably poor statistics. (If one uses mesh-based weight windows, a similar issue arises for the mesh.)

Much of this dissertation addresses this issue. The method espoused here is first to determine an importance distribution in the longitudinal direction (the direction of streaming) in whatever penetration is the most important streaming path. The Handy Method is used for this. This importance distribution is not important in itself, since MCNP does not split in voids, but it constitutes a scaffolding on which one can hang the importances in the non-void cells. This is done using the transverse distributions determined in Publication X. It is notable that they drop off very rapidly with increasing distance from the penetration. Thus, unless the shielding material is very poor, the non-void cells immediately adjacent to the penetration should be quite thin, on the order of 1 mm. The next layer of cells should be somewhat thicker, but only about 2.5 mm for good shielding materials.

Of course this depends on the geometry and on what quantities one is interested in tallying. For instance, the example of Publication XI demonstrates that thicker cells are appropriate when particle transport at right angles to a penetration is of interest.

Sometimes such, relatively conventional, techniques are not enough, and the absence of provisions in MCNP for biasing the flux not only in space and energy but also in angle becomes a serious problem. Then one should consider the tally source method proposed by Hendricks and described and tested in Publication XII. This does provide a kind of angle biasing capability. This method has its drawbacks, being laborious and inexact, but it can be used to solve problems that do not yield to more direct methods.

One issue that has not been explicitly addressed in this text is the accuracy of the calculations. It should be noted that they were generally best estimate calculations rather than aiming for conservatism, so the errors may go in either direction.

MCNP provides estimates not only of fluxes and reaction rates, but also of the stochastic error in the form of fractional standard deviations. These are themselves evaluated stochastically, so the larger the stochastic uncertainty is, the less reliably is it estimated. In addition there are other sources of error, which may often be greater, such as the uncertainties in the cross sections and simplifications in the modeling.

In the easier cases, requiring only calculation of the flux near the plasma, the errors are likely to be small, of the order of a few percent. Examples of such cases are calculations of the wall loading or the tritium breeding ratio. However, as one goes deeper into the materials surrounding the plasma, it becomes more difficult to get good statistics. When estimating the shutdown dose rate at the end of an ITER port, calculated fractional standard deviations of about 0.2 were commonplace. In addition, with increasing depth of penetration the uncertainties of the total cross sections also become more significant. In the strictly intuitive opinion of the author, a factor of 2 might be a reasonable estimate of the uncertainty in such calculations. For the IFMIF test cell cover described in Publication XII, an even greater uncertainty factor, maybe 10, should be used, due to the approximations involved in the tally source method. On the other hand, high accuracy is generally not needed in shielding calculations. The

uncertainties estimated here are reasonably acceptable in such applications, considering that the shutdown dose rate in the IFMIF access cell (above the test cell) are likely to be well below the limit of what can be allowed.

References

1. Shimomura, Y. The present status and future prospects of the ITER project. Journal of Nuclear Materials, 2004, Vol. 5–11, pp. 329–333.

2. Shimomura, Y. Preparation of ITER construction and operation. Fusion Engineering and Design, 2006, 2, Vol. 81, No. 1–7, pp. 3–11.

3. Cozzani, F. Status and possible prospects of an international fusion materials irradiation facility. Fusion Engineering and Design, 1999, 11, Vol. 46, No. 2–4, pp. 413–421.

4. Andreani, R., Diegele, E., Gulden, W., Lässer, R., Maisonnier, D., Murdoch, D., Pick, M. & Poitevin, Y. Overview of the European Union fusion nuclear technologies development and essential elements on the way to DEMO. Fusion Engineering and Design, 2006, 2, Vol. 81, No. 1–7, pp. 25–32.

5. LA-12625-M, Version 4B. Briesmeister, J.F. (Ed.) MCNP – A General Monte Carlo N-Particle Transport Code, Version 4B. 1997.

6. LA-13709-M. Briesmeister, J.F. (Ed.) MCNP – A General Monte Carlo N-Particle Transport Code, Version 4C. 2000.

7. LA-UR-03-1987. X-5 Monte Carlo Team MCNP – A General Monte Carlo N-Particle Transport Code, Version 5. Volume I: Overview and Theory. 2003, revised 2005.

8. LA-CP-03-0245. X-5 Monte Carlo Team MCNP – A General Monte Carlo N-Particle Transport Code, Version 5. Volume II: User's Guide. 2003, revised 2005.

9. X-5 Monte Carlo Team MCNP – A General Monte Carlo N-Particle Transport Code, Version 5. Volume III: Developer's Guide. 2003, revised 2005.

10. Fischer, U., Simakov, S.P., von Möllendorff, U. & Pereslavtsev, P. Validated computational tools and data for IFMIF neutronic calculations. Proceedings of the ANS Annual Meeting AccApp'03. San Diego, California, USA, June 2003.

11. Simakov, S.P., Fischer, U., von Möllendorff, U., Schmuck, I., Konobeev, A.Y., Korovin, Y.A. & Pereslavtsev, P. Advanced Monte Carlo procedure for the IFMIF d-Li neutron source term based on evaluated cross section data. Journal of Nuclear Materials, 2002, 12, Vol. 307–311, Part 2, pp. 1710–1714.

12. Waters, L.S. (Ed.) MCNPX User's Manual, version 2.1.5.

13. FENDL-2.1. Fusion Evaluated Nuclear Data Library (Dec. 2004) [Homepage of IAEA], [Online] Available: http://www-nds.iaea.org/fendl21/.

14. Sato, S., Iida, H., Plenteda, R., Valenza, D. & Santoro, R.T. Evaluation of biological dose rates around the ITER NBI ports by 2-D S/activation and 3-D Monte Carlo analyses. Fusion Engineering and Design, 2000, 1, Vol. 47, No. 4, pp. 425–435.

15. Chen, Y. & Fischer, U. Rigorous MCNP based shutdown dose rate calculations: Computational scheme, verification calculations and applications to ITER. 6[th] Int. Symp. Fusion Nuclear Technology. San Diego, California, USA, 2002.

16. UKAEA FUS 514. Forrest, R.A. & Gilbert, M.R. FISPACT-2005: User manual. 2004.

17. Petrizzi, L., Batistoni, P., Chen, Y., Fischer, U., Iida, H. & Morimoto, Y. Advanced Methodology for Dose rate Calculation of ITER-FEAT. 12th Biennial Topical Meeting of the Radiation Protection and Shielding Division of the American Nuclear Society. Santa Fe, New Mexico, USA, 2002.

18. Iida, H., Valenza, D., Plenteda, R., Santoro, R.T. & Dietz, J. Radiation Shielding for ITER to allow for Hands-on Maintenance inside the Cryostat. Journal of Nuclear Science and Technology 1999, 19 Oct. 1999.

19. Valenza, D., Iida, H., Plenteda, R. & Santoro, R.T. Proposal of shutdown dose estimation method by Monte Carlo code. Fusion Engineering and Design, 2001, 9, Vol. 55, No. 4, pp. 411–418.

20. ITER Technical Basis document G A0 FDR 1 01-07-13 R1.0.

21. Wasastjerna, F. Experience of Fusion Neutronics Calculations, version 3. VTT memo ed.

22. Loughlin, M.J. Final Report, TW6-TTMI-001: IFMIF accelerator facilities. Dose rate evaluation in the IFMIF accelerator. 2007.

23. Sato, S., Iida, H. & Nishitani, T. Evaluation of shutdown gamma-ray dose rates around the duct penetration by three-dimensional Monte Carlo decay gamma-ray transport calculation with variance reduction method. 2002. 1237 p.

24. Sato, S. & Nishitani, T. A study on variance reduction of Monte Carlo calculation with weight window generated by density reduction method. 2007. 5 p.

25. Hodgdon, A.D. A Variance Reduction Management Algorithm for MCNP. Transactions 2003, Iss. 89, pp. 373–374. ISSN 0003-018X.

26. Sweezy, J., Brown, F., Booth, T., Chiaramonte, J. & Preeg, B. Automated variance reduction for MCNP using deterministic methods. Radiation Protection Dosimetry, 2005, Vol. 116, No. 1–4, pp. 508–512.

27. Haghighat, A. & Wagner, J.C. Monte Carlo variance reduction with deterministic importance functions. Progress in Nuclear Energy, 2003, Vol. 42, No. 1, pp. 25–53.

28. Omura, M., Miyake, Y.1., Hasegawa, T., Ueki, K., Sato, O., Haghighat, A. & Sjoden, G.E. Performance of the improved version of Monte Carlo code A^3MCNP for large-scale shielding problems. Radiation protection dosimetry, 2005, Vol. 116 (2), No. 1–4, pp. 493–497.

29. Shahdatullah, M.S., Ziver, K., Eaton, M.D., Pain, C.C. & Goddard, A.J.H. Monte Carlo variance reduction using finite element adjoint weight windows. Vol. 2006. 2006.

30. Hendricks, J.S. & Culbertson, C.N. Assessment of MCNP Weight Windows. 2000.

31. Murata, I., Yamamoto, H., Miyamaru, H., Goldenbaum, F. & Filges, D. Scattering direction biasing for Monte Carlo transport calculation. Nuclear Instruments and Methods in Physics Research, 2006, A 562, pp. 845–848.

32. Freiesleben, H., Richter, D., Seidel, K. & Unholzer, S. Fusion neutronics-streaming, shielding, heating, activation. The CAARI 2000: Sixteenth international conference on the application of accelerators in research and industry. AIP Conference Proceedings, Volume 576. 2001. Pp. 763–767. ISSN 0094-243X.

33. Serikov, A., Fischer, U., Heidinger, R., Henderson, M.A., Spaeh, P. & Tsige-Tamirat, H. Radiation shield analyses in support of the FS design for the ITER ECRH launcher. SOFT-24 Symposium on Fusion Technology, No. 24, Warsaw, POLOGNE (11/09/2006), 2007, Vol. 82, No. 5–14. Pp. 736–743.

34. Iida, H., Santoro, R.T., Valenza, D. & Khripunov, V. A handy method for estimating radiation streaming through holes in shield assemblies. Fusion Engineering and Design, 1998, 10, Vol. 43, No. 1, pp. 1–13.

35. Leppänen, J. Development of a new Monte Carlo reactor physics code. VTT Publications 640, Espoo 2007. 228 p. + app. 8 p. http://www.vtt.fi/inf/pdf/publications/2007/P640.pdf.

Appendix B: Publications I, II, III, V, VI, XI, XII are not included in the PDF version. Please order the printed version to get the complete publication (http://www.vtt.fi/publications/index.jsp)

Appendix A: An Analytic Estimate of the Spatial Discretization Error in the Rigorous Two-Step Method for Calculating Shutdown Dose Rates

Two competing methods of calculating shutdown dose rates with Monte Carlo programs such as MCNP are known as the Direct One-Step (D1S) and Rigorous Two-Step (R2S) methods, respectively. Each has its own advantages and disadvantages. Perhaps the most obvious disadvantage of the R2S method is the spatial discretization error. Instead of each delayed gamma being born at the exact location where the neutron producing it was absorbed, as in the D1S method, the neutron flux is averaged over each cell and then the resulting activation is calculated in a separate program, such as FISPACT. In the gamma transport calculation, the gamma source is taken to be constant in each cell, losing the detailed spatial information within the cell. The aim of the work described here was to obtain at least a rough estimate of the magnitude of this discretization error and its dependence on the cell size, with a view to choosing cell sizes appropriately, for instance in IFMIF shielding calculations.

To make the problem tractable, we make some simplifying assumptions. Take a purely one-dimensional shielding problem in slab geometry, with a shield extending from x = 0 to x = a, embedded in a void on both sides. Neutrons impinge on the shield on the left side (x = 0), and the quantity of interest is the delayed dose rate on the right side. We assume that the neutron flux in the shield follows an exponential distribution,

$$\phi(x) = Ce^{-\kappa x}. \tag{1}$$

The delayed gamma source in the shield is proportional to the neutron flux,

$$Q(x) = S\phi(x). \tag{2}$$

A gamma photon born at x will reach the far side of the shield with a probability of

$$e^{-\lambda(a-x)}. \tag{3}$$

Thus the gamma flux (or dose rate) beyond the shield is

$$\psi = \int_0^a KQ(x)e^{-\lambda(a-x)}dx = \int_0^a KSCe^{-\kappa x}e^{-\lambda(a-x)}dx = \frac{KSCe^{-\lambda a}}{\lambda - \kappa}\left(e^{(\lambda-\kappa)a} - 1\right). \quad (4)$$

We now divide the shield into layers, $i = 1,...,I$, each layer extending from x_{i-1} to x_i $(x_0 = 0, \quad x_I = a)$. Let the thickness of each layer be Δx_i. The exact contribution to the dose rate from each layer (within the assumptions we have made here) is

$$\psi_i = \int_{x_{i-1}}^{x_i} KSCe^{-\kappa x}e^{-\lambda(a-x)}dx = \frac{KSC}{\lambda - \kappa}e^{-\lambda a}e^{(\lambda-\kappa)x_{i-1}}\left(e^{(\lambda-\kappa)\Delta x_i} - 1\right). \quad (5)$$

In the R2S method, however, the neutron flux is treated as constant within each layer,

$$\tilde{\phi}_i = \frac{C}{\Delta x_i}\int_{x_{i-1}}^{x_i} e^{-\kappa x}dx = \frac{C}{\kappa\Delta x_i}e^{-\kappa x_{i-1}}\left(1 - e^{-\kappa\Delta x_i}\right). \quad (6)$$

Accordingly, the gamma flux at the shield surface caused by activation in layer i will be

$$\tilde{\psi}_i = \frac{KSC}{\kappa\Delta x_i}\left(e^{-\kappa x_{i-1}} - e^{-\kappa x_i}\right)\int_{x_{i-1}}^{x_i} e^{-\lambda(a-x)}dx = \frac{KSC}{\kappa\lambda\Delta x_i}e^{-\lambda a}e^{(\lambda-\kappa)x_{i-1}}\left(e^{\lambda\Delta x_i} - 1\right)\left(1 - e^{-\kappa\Delta x_i}\right). \quad (7)$$

For the sake of legibility, let us introduce the symbols K = $\kappa\Delta x_i$ and L = $\lambda\Delta x_i$. Then the ratio between the discretized and non-discretized gamma flux contributions from layer i is

$$\frac{\tilde{\psi}_i}{\psi_i} = \frac{L-K}{LK}\frac{\left(e^L - 1\right)\left(e^K - 1\right)}{e^L - e^K}. \quad (8)$$

When $\kappa = \lambda$, the above formula breaks down, but then it reduces to

$$\frac{\tilde{\psi}_i}{\psi_i} = \frac{\left(e^L - 1\right)\left(e^K - 1\right)}{LKe^K}. \quad (9)$$

Suppose we set as an objective that the discretization error of the most important layer should be less than 10%. This is a fairly tight criterion, in that the total effect of other sources of error in an ITER or IFMIF shielding calculation will usually exceed this by a large margin. Then Table A.1 (see the yellow cells) shows that this is achieved if

$$\Delta x_i < \frac{1}{\sqrt{\kappa\lambda}}. \tag{10}$$

Moreover, when κ and λ are very different, the allowable layer thickness becomes greater than that allowed by the above inequality.

In practice, the thickness allowed by this reasoning may be about 1 cm. However, we do not need to use this cell thickness throughout the system. If we assume that photons are attenuated more sharply than neutrons, i.e., $\lambda > \kappa$, the layer thickness limit need only apply to the cell layer at the end of the shield ($i = I$). For cell layers farther back, their relative contribution will be decreased by the factor

$$e^{-(\lambda-\kappa)(a-x_i)}. \tag{11}$$

It is also conceivable that the neutron attenuation may exceed that of photons, so that $\lambda < \kappa$. In that case the dominant cell layer will be the first one, not the last, and the following layers will have their contributions attenuated by a factor

$$e^{-(\kappa-\lambda)x_{i-1}} \tag{12}$$

caused by the attenuation of the neutron rather than photon flux.

Tables A.2 and A.3 show how one can use the above results in practice. Table A.2 shows the relative error, $|\tilde{\psi}/\psi| - 1$, caused by discretization as a function of κ, λ and Δx, using a value of κ corresponding to a halving length of 7.5 cm and values of λ appropriate for photons with energies of 2.4, 1.2, 0.6 and 0.3 MeV in heavy concrete. In Table A.3 a discretization scheme is shown, together with the resulting contributions to the dose rate error, relative to the dose rate from the last x interval (-1 to 0 cm), taking into account the attenuation correction. We see that in this example the discretization error remains small. Of course, the theory on which this result is based is in itself inexact, but it probably gives an idea of the order of magnitude of the discretization error.

Table A.1. $\tilde{\psi}/\psi$ values for selected $K = \kappa\Delta x_i$ and $L = \lambda\Delta x_i$. The largest K and L values for which $\tilde{\psi}/\psi < 1.1$ are highlighted in yellow, the cells where K = L in green.

K\L	0.100	0.126	0.158	0.200	0.251	0.316	0.398	0.501	0.631	0.794	1.000	1.259	1.585	1.995	2.512	3.162	3.981	5.012	6.310	7.943	10.000
0.100	1.001	1.001	1.001	1.002	1.002	1.003	1.003	1.004	1.005	1.007	1.008	1.010	1.013	1.016	1.019	1.023	1.027	1.031	1.035	1.039	1.041
0.126	1.001	1.001	1.002	1.002	1.003	1.003	1.004	1.005	1.007	1.008	1.010	1.013	1.016	1.020	1.024	1.029	1.035	1.040	1.045	1.049	1.052
0.158	1.001	1.002	1.002	1.003	1.003	1.004	1.005	1.007	1.008	1.010	1.013	1.016	1.020	1.025	1.031	1.037	1.044	1.051	1.057	1.062	1.066
0.200	1.002	1.002	1.003	1.003	1.004	1.005	1.007	1.008	1.011	1.013	1.017	1.021	1.026	1.032	1.039	1.047	1.056	1.064	1.072	1.079	1.085
0.251	1.002	1.003	1.003	1.004	1.005	1.007	1.008	1.011	1.013	1.017	1.021	1.026	1.033	1.040	1.050	1.060	1.071	1.082	1.092	1.101	1.108
0.316	1.003	1.003	1.004	1.005	1.007	1.008	1.011	1.013	1.017	1.021	1.026	1.033	1.041	1.051	1.063	1.076	1.090	1.105	1.118	1.130	1.139
0.398	1.003	1.004	1.005	1.007	1.008	1.011	1.013	1.017	1.021	1.027	1.033	1.042	1.052	1.065	1.080	1.097	1.116	1.134	1.152	1.167	1.179
0.501	1.004	1.005	1.007	1.008	1.011	1.013	1.017	1.021	1.027	1.034	1.042	1.053	1.067	1.083	1.103	1.125	1.149	1.174	1.197	1.217	1.233
0.631	1.005	1.007	1.008	1.011	1.013	1.017	1.021	1.027	1.034	1.042	1.054	1.068	1.085	1.106	1.132	1.161	1.193	1.226	1.256	1.283	1.306
0.794	1.007	1.008	1.010	1.013	1.017	1.021	1.027	1.034	1.042	1.054	1.068	1.086	1.108	1.136	1.169	1.208	1.251	1.296	1.338	1.375	1.406
1.000	1.008	1.010	1.013	1.017	1.021	1.026	1.033	1.042	1.054	1.068	1.086	1.109	1.138	1.175	1.219	1.271	1.330	1.391	1.450	1.503	1.547
1.259	1.010	1.013	1.016	1.021	1.026	1.033	1.042	1.053	1.068	1.086	1.109	1.139	1.177	1.226	1.285	1.357	1.439	1.526	1.611	1.687	1.751
1.585	1.013	1.016	1.020	1.026	1.033	1.041	1.052	1.067	1.085	1.108	1.138	1.177	1.228	1.292	1.373	1.473	1.590	1.718	1.846	1.962	2.060
1.995	1.016	1.020	1.025	1.032	1.040	1.051	1.065	1.083	1.106	1.136	1.175	1.226	1.292	1.379	1.492	1.634	1.807	2.002	2.203	2.390	2.550
2.512	1.019	1.024	1.031	1.039	1.050	1.063	1.080	1.103	1.132	1.169	1.219	1.285	1.373	1.492	1.650	1.858	2.121	2.434	2.772	3.096	3.379
3.162	1.023	1.029	1.037	1.047	1.060	1.076	1.097	1.125	1.161	1.208	1.271	1.357	1.473	1.634	1.858	2.167	2.583	3.112	3.722	4.341	4.897
3.981	1.027	1.035	1.044	1.056	1.071	1.090	1.116	1.149	1.193	1.251	1.330	1.439	1.590	1.807	2.121	2.583	3.255	4.194	5.390	6.713	7.968
5.012	1.031	1.040	1.051	1.064	1.082	1.105	1.134	1.174	1.226	1.296	1.391	1.526	1.718	2.002	2.434	3.112	4.194	5.900	8.408	11.600	14.949
6.310	1.035	1.045	1.057	1.072	1.092	1.118	1.152	1.197	1.256	1.338	1.450	1.611	1.846	2.203	2.772	3.722	5.390	8.408	13.760	22.221	32.920
7.943	1.039	1.049	1.062	1.079	1.101	1.130	1.167	1.217	1.283	1.375	1.503	1.687	1.962	2.390	3.096	4.341	6.713	11.600	22.221	44.608	83.588
10.000	1.041	1.052	1.066	1.085	1.108	1.139	1.179	1.233	1.306	1.406	1.547	1.751	2.060	2.550	3.379	4.897	7.968	14.949	32.920	83.588	220.245

Table A.2. The relative error, $\left|\widetilde{\psi}/\psi\right|-1$, for a given κ and selected values of λ and Δx.

kappa	lambda			
0.092	0.114	0.165	0.235	0.331
delta-x (cm)				
1	0.000874	0.001266	0.001802	0.002538
2	0.003501	0.005068	0.007218	0.010153
3	0.007890	0.011428	0.016271	0.022853
4	0.014061	0.020376	0.029003	0.040648
5	0.022037	0.031957	0.045472	0.063560
10	0.090439	0.131885	0.187179	0.256424

Table A.3. A cell subdivision scheme (with $x = 0$ at the end of the shield), showing the contribution to the total error (as a multiple of the dose rate contribution from the last layer), including the attenuation factor.

	kappa	lambda			
	0.092	0.114	0.165	0.235	0.331
x	delta-x				
0	1	0.000874	0.001266	0.001802	0.002538
-1	2	0.003425	0.004712	0.006256	0.007995
-3	3	0.007386	0.009180	0.010595	0.011157
-6	4	0.012322	0.013149	0.012297	0.009689
-10	5	0.017685	0.015400	0.010882	0.005824
-15	5	0.015843	0.010691	0.005323	0.001763
-20	10	0.058246	0.030629	0.010720	0.002153
-30					

Appendix B: Publications I–XII

Author(s)
Wasastjerna, Frej

Title
Using MCNP for fusion neutronics

Abstract

Any fusion reactor using tritium-deuterium fusion will be a prolific source of 14 MeV neutrons. In fact, 80% of the fusion energy will be carried away by these neutrons. Thus it is essential to calculate what will happen to them, so that such quantities as the tritium breeding ratio, the neutron wall loading, heat deposition, various kinds of material damage and biological dose rates can be determined. Monte Carlo programs, in particular the widely-used MCNP, are the preferred tools for this. The International Fusion Materials Irradiation Facility (IFMIF), intended to test materials in intense neutron fields with a spectrum similar to that prevailing in fusion reactors, also requires neutronics calculations, with similar methods.

In some cases these calculations can be very difficult. In particular shielding calculations – such as those needed to determine the heating of the superconducting field coils of ITER or the dose rate, during operation or after shutdown, outside ITER or in the space above the test cell of IFMIF – are very challenging. The thick shielding reduces the neutron flux by many orders of magnitude, so that analog calculations are impracticable and heavy variance reduction is needed, mainly importances or weight windows. On the other hand, the shields contain penetrations through which neutrons may stream. If the importances are much higher or the weight windows much lower at the outer end of such a penetration than at the inner end, this may lead to an excessive proliferation of tracks, which may even make the calculation break down.

This dissertation describes the author's work in fusion neutronics, with the main emphasis on attempts to develop improved methods of performing such calculations. Two main approaches are described: trying to determine near-optimal importances or weight windows, and testing the "tally source" method suggested by John Hendricks as a way of biasing the neutron flux in angle.

Författarna
Wasastjerna, Frej

Namn
Fusionsneutronik med MCNP

Referat

En fusionsreaktor som använder sig av tritium-deuteriumfusion kommer att vara en mycket stark källa till neutroner med en energi av 14 MeV. Dessa kommer i själva verket att bära med sig 80 % av fusionsenergin. Det är alltså väsentligt att beräkna vad som kommer att hända med dem, så att man kan bestämma sådana storheter som tritiumbridkvoten, den av neutronerna förorsakade väggbelastningen, värmedepositionen, olika slags materialskador och biologiska dosrater. Monte Carlo-program, speciellt MCNP, är det mest allmänt använda redskapet för sådana beräkningar. IFMIF (International Fusion Materials Irradiation Facility), avsedd att testa material i ett intensivt neutronflöde med ett spektrum som liknar det i fusionsreaktorer, kräver också neutronikberäkningar, med liknande metoder.

I vissa fall kan dessa beräkningar vara mycket svåra. I synnerhet strålskyddsberäkningar, sådana som behövs för att bestämma värmeutvecklingen i ITERs supraledande fältspolar eller dosraten, i verksamhet eller efter avstängning, utanför ITER eller i utrymmet ovanför IFMIFs testcell, är särdeles krävande. Det tjocka strålskyddet minskar neutronflödet med många storleksordningar, så att analoga beräkningar är ogenomförbara och stark variansreduktion är nödvändig, främst med viktighet eller viktfönster. Å andra sidan innehåller skyddet genomföringar längs vilka neutroner kan strömma. Om viktigheterna är mycket högre eller viktfönstren mycket lägre i yttre ändan av en sådan genomföring än i den inre ändan, kan en partikelhistoria yngla av sig i ett orimligt antal spår, vilket kan få beräkningen att bryta samman.

Denna avhandling beskriver disputandens arbete i fusionsneutronikbranschen med huvudvikten lagd vid hans försök att utveckla förbättrade beräkningsmetoder. Två huvudsakliga sätt att göra detta beskrivs: försök att bestämma närapå optimala viktigheter eller viktfönster, och testning av "tally source"-metoden, föreslagen av John Hendricks som ett sätt att biasera neutronfödet med avseende på vinkel.

Any fusion reactor using tritium-deuterium fusion will be a prolific source of 14 MeV neutrons. It is essential to calculate what will happen to them, so that such quantities as the tritium breeding ratio, the neutron wall loading, heat deposition, material damage and biological dose rates can be determined. Monte Carlo programs, in particular MCNP, are the preferred tools for these calculations.

   In some cases these calculations can be very difficult. In particular shielding calculations - such as those needed to determine the heating of the superconducting field coils of ITER or the dose rate, during operation or after shutdown, outside ITER - are very challenging. In an analog Monte Carlo calculation very few neutrons will make it through a thick shield, so heavy reliance on variance reduction is necessary. The most important such technique is geometry splitting. Neutrons that have made some progress through the shield are split so that each history branches into several tracks. This may be repeated many times. This technique ensures that a greater number of tracks will penetrate the shield, but it is problematic when there are slits or ducts through which neutrons may stream.

   This thesis describes the author's work in fusion neutronics, with the main emphasis on attempts to develop improved methods for such calculations. Two main approaches are described: trying to determine near-optimal importances or weight windows to control splitting, and testing the "tally source" method suggested by John Hendricks as a way of preferentially following those neutrons whose flight direction gives them the best chances of penetrating a shield with ducts.

Frej Wasastjerna