A Hybrid Continuous Energy and Multi-Group Monte Carlo Method

Mikolaj Adam Kowalski, Eugene Shwageraus

Engineering Department, University of Cambridge, Trumpington St., Cambridge, CB2 1PZ, UK

Abstract

This paper investigates mixing multi-group (MG) and continuous energy (CE) representation of cross-sections depending on location of a particle in a Monte Carlo neutronic eigenvalue calculations in 1D and 2D PWR test cases with UOX and MOX fuel. Different population normalisation needs to be applied to CE and MG region to account for the difference in criticality between CE and MG representation. This normalisation procedure requires a neutron production rate ratio between CE and MG region to be known a priori. A resonance correction in energy spectrum during transition of a particle between the MG and CE region was developed based on the equivalence resonance treatment theory. With these, it was shown that it is possible to accelerate total calculation time, while introducing only a moderate error below 1% in the fission rate distribution. The magnitude of acceleration is heavily dependent on the relative size of CE and MG zones.

Keywords: Variable fidelity, Monte Carlo, Multi-Group Monte Carlo

1. Introduction

It has been shown, that Monte Carlo (MC) calculations based on a multi-group (MG) representation of neutron interaction tend to execute significantly faster for the same problems than continuous-energy (CE) calculations [1]. Moreover, the simpler representation of the interaction physics with MG data simplifies the vectorisation of solution algorithm, which makes MG MC well suited for execution on heavily vectorised processor architectures.

Furthermore, very accurate high fidelity results are often not required everywhere in a calculation domain. For example, in a simulation that aims to

verify whether a given reactor configuration violates thermal limits, faithful representation of the physics is needed only in the vicinity of the high power pins. Much higher systematic error could be tolerated in other parts of the domain without significant detriment to the overall calculation. Therefore, the goal of this paper is to investigate the possibility of combining MG and CE representation in MC calculation in different spatial regions of a problem domain to increase the efficiency of the use of computational resources. In the proposed scheme, only a single MC calculation is performed with the representation of a particle changed when it crosses to a new region. The study is limited to eigenvalue problems. An application to dynamic MC calculations is appealing in light of the recent developments in the use of GPUs for the dynamic MC calculations [2], but it is outside the scope of the study in this paper.

The topic of combining different fidelity representations for different spatial regions in MC calculations is relatively unexplored. It was initially suggested in [3] that MC calculations could be potentially accelerated by representing unimportant regions by homogenised blocks with MG data. The effectiveness of MG MC calculation of a domain composed of homogenised blocks was studied in [4] and [5]. However, results of the initial feasibility study [6] suggested that the performance gain from geometry homogenisation is relatively modest in comparison with the gain obtained by replacing CE cross-sections with MG data, while at the same time, the homogenisation introduces significant difficulties with accurate representation of the current spectrum at the MG-CE interface. As a result, this study focuses on interaction of CE and MG regions with heterogeneous representation of geometry.

A number of previous investigations aimed to combine the lower fidelity method in fast and thermal energy range, with CE MC in resonant range. In [7] and [8], method of characteristics was used as a lower fidelity method and the performance of the scheme was evaluated in small 2D problems composed of a few pin-cells. Another study [9] investigated the combination of MG and CE MC in different energy ranges for simple 1D spherical model of International Thermonuclear Experimental Reactor (ITER).

This paper expands on the previous work by investigating the separation of MG and CE representation in space instead of energy, which poses new challenges related to normalisation and self-shielding. In contrast to the previous studies, the performance of the proposed methodology is studied on relatively large problems composed of UOX and MOX fuel assemblies, which

are more representative of a typical reactor calculations. Lastly, the accuracy of power distribution will be evaluated in favour of criticality. This is in line with the motivation for the study of variable fidelity schemes.

It should be noted that throughout this paper the terms "importance", "important" and "unimportant" are used in a qualitative way. The main goal of the paper is to study the practicality of performing combined MG and CE calculations after decomposition of a domain into notional, user defined important and unimportant zones has already been performed based on expert judgment. In principle, a quantitative scheme could be developed based on some threshold in an adjoint flux associated with a response, which is deemed most important by the user. However, no such scheme is studied in this paper.

2. Methods & Infrastructure

2.1. Monte Carlo Transport Code SCONE

To implement algorithms and obtain results presented in this paper, in-house MC Neutron Transport Code SCONE was used. It stands for **S**tochastic **C**alculator **O**f **N**eutron transport **E**quation and is an object-oriented framework for performing particle transport Monte Carlo calculations. It is currently under development in the Nuclear Energy Group at the Cambridge University Engineering Department. Its primary goal is to provide a simple and accessible environment for prototyping new Monte Carlo solution algorithms and to be used for teaching. It is written in Fortran 2008 based on a conjecture that this language provides a reasonable compromise between ease-of-use and efficiency. The object-oriented structure allows to hide implementation details behind a number of APIs, thus limiting the code scope necessary to be known in order to modify the solution algorithm. The object-oriented nature of SCONE may limit its performance, because making low-level optimisations of the code by both the developer and the compiler would be more difficult. For example, the reliance on polymorphism may reduce the inlining of short functions. Furthermore, the use of Fortran makes the control over memory layout more difficult than it might be if a language like C was used instead. However, it is not considered to be a problem, as SCONE does not aim to compete with well-established Monte Carlo codes in terms of computational speed. It attempts to supplement them by allowing inexperienced users, like graduate students, to explore new

techniques without the need to understand the implementation of the sections that are not directly relevant to their work. Thus, SCONE aims at reasonable, not maximum, performance.

SCONE supports both CE and MG cross-section data. In CE, some limitations still exist. Not all of the ENDF energy laws have been implemented yet. Furthermore, unresolved resonances probability tables and bound thermal scattering is not yet supported. Moreover, it is important to note that MG calculation capability of SCONE is achieved by simple substitution of CE nuclear data object with its MG version. Thus, the performance benefit when using MG data is only due to constant time complexity access to cross-sections without the search and interpolation of an energy grid. Parallel calculation capabilities using OpenMP are currently under development.

2.2. Differential Normalisation

Every variable fidelity scheme needs to have a property that as the distance from the boundary between the different fidelity zones increases, any quantity of interest approaches the value, that would be obtained in a single fidelity calculation of the whole domain. This feature is crucial to achieve the intuitive behaviour, that increasing the size of the high fidelity zone improves the quality of results away from the boundary. Without assuring that this is the case, it would be challenging to employ the scheme in practice. For schemes based on separation of fidelity in space, achieving this property in eigenvalue problems requires careful consideration of normalisation.

To illustrate this point let us, for now, move away from the MC calculations and consider a trivial problem of multiplying, weakly absorbing, long homogeneous one dimensional slab modelled with one speed diffusion theory, and reflective boundary conditions at both ends. L shall denote the half-length of the slab (slab extends from $-L$ to L). Now, let us imagine that we prepare the cross-sections to model this system using two different computational routes. One set of data was obtained in a CE MC calculation and will be considered the high fidelity data. The other, lower fidelity data was prepared by a much simpler equivalence treatment. In general, there will be significant differences in the one group data obtained with high and lower fidelity route. For simplicity, let us assume that the only difference is in macroscopic absorption and that the lower fidelity route has predicted higher capture. If we model the slab in a single fidelity calculation utilising only a single data set, in both cases we obtain the correct, uniform flux distribution. The only difference between the systems would be exhibited in

the criticality, with the calculation using lower fidelity data predicting lower value. However, if we try to perform a variable fidelity calculation of the same problem, by replacing higher fidelity cross-sections in the left half of the slab with their lower fidelity equivalent, we will fail to obtain the correct result and the error in flux distribution will not be decreasing with the distance from the boundary between fidelity zones. In fact, using first-order perturbation theory together with the fact that the problem is self-adjoint, it can be shown (Appendix A) that the calculated flux distribution will be:

$$\phi(x) = \phi_0 - \sum_{n=0}^{\infty} \frac{8\Delta(B_m^2)L^3}{\pi^3(2n+1)^3} sin\left(\frac{\pi(n+0.5)x}{L}\right) \tag{1}$$

Where, ϕ_0 denotes the flux magnitude in the high fidelity result, $\Delta(B_m^2)$ stands for the difference between material buckling ($B_m^2 = \frac{\Sigma_a - \nu\Sigma_f/k}{D}$) with lower and high fidelity cross sections. In this case, $\Delta(B_m^2)$ is equal to $\frac{-|\Delta\Sigma_a|}{D}$, where all the symbols have their usual meaning [10]. Clearly, the calculation would result in a significant error with respect to correct uniform solution and the error would reach its maximum close to the ends of the slab. This is exactly the opposite of the behaviour we would like to achieve in a variable fidelity scheme.

The reason for this behaviour becomes obvious if we recall the basic idea behind an eigenvalue neutron transport calculation. In practice, no nuclear system is perfectly critical, so in every calculation, the system is modified by the scaling of fission source by the inverse of criticality. In this way, it is possible to model a non-critical system as if it was perfectly critical. In the framework of one-speed diffusion calculation, it corresponds to adjusting the neutron fission term in a material buckling so it matches the geometric buckling of the principal mode (0 in the case of reflected slab). However, after a simple substitution of low fidelity cross-sections in place of their high fidelity equivalent in the left half of the slab, the value of criticality of the combined system changed to a value in-between the values obtained in a single fidelity calculations. As a result, the lower fidelity part of the system was modelled as sub-critical and high fidelity part as supercritical. This accounts for the skewing of the flux solution. In order to obtain the desired variable fidelity scheme, we need to model both parts of the system as perfectly critical. Thus, it is necessary to scale the fission source of the lower and high fidelity regions by different values corresponding to the inverse of lower and high fidelity criticality. If we were to do it in the calculation of the reflected slab, it is

clear that ΔB_m would become 0 and in Eq. 1 we would recover the correct uniform solution.

The observations made for the 1D slab can be extended to a more general case of a MC calculation. Now, low fidelity and high fidelity regions will be modelled with MG and CE cross-sections respectively. The simplified example has demonstrated that when dealing with a variable fidelity eigenvalue calculation, there is a need to scale the fission source in high and low fidelity regions by different values. This is necessary in order to distinguish a variable fidelity system, in which different representations attempt to model the same physical system, from a new physical system in which the cross sections were adjusted (which is also a valid problem). In other words, we need to find an approximation to the physical system eigenfunction instead of looking for an eigenfunction of the adjusted physical system. Based on the physical insight provided by the 1D slab example it may be claimed that a good approximation to the physical system eigenvalue may be obtained if we scale the MG and CE regions by the values of criticality we would obtain in a fully MG or CE calculation. This choice has not been proven to be optimal, so it is possible that a better selection of fission source scaling coefficients may exist. The scaling factor of the fission source ($1/k$) depends only on the fidelity region and not the position of fission event in the phase space. This is in order to preserve the physical justification for the choice of scaling (representing both MG and CE parts of the system as perfectly critical) and to prevent the normalisation procedure from turning into a source fitting scheme. It should be noted that the differential normalisation is independent of the number of tallies required in the problem. It affects only the shape of the fission source and has no effect on the transport of particles and scoring of reaction rates.

As stated above, in order to perform a variable fidelity calculation with MG and CE data, we need to know the criticality of the system represented purely in CE and MG, as a result of the need for differential source normalisation. A simple way to achieve it would be by performing single fidelity calculations prior to the variable fidelity calculations. However, this is unlikely to offer any computational benefit from the variable fidelity approach. Therefore, it is sensible to develop a method to infer the correct normalisation values from some restrictions on the final result. A method proposed in this paper is based on an observation that existing lower fidelity modelling methods (e.g. nodal diffusion) are generally successful at predicting the power distribution averaged over a large mesh. Furthermore, even high fidelity methods, like CE MC will be able to provide accurate results much

faster if the result mesh is coarse. As a result, we will design a normalisation scheme to preserve the ratio of neutron production rate between MG and CE region, which was obtained with other, faster method (e.g. fission matrix model of the domain). Then, it is sufficient to choose such weighting of fission sites, that the ratio is preserved in the combined calculation. If we denote 2x2 fission matrix [11] discretised into MG and CE region by $\bar{F}$, we can calculate the appropriate weighting coefficients by solving the Eq. 2 for α and β to yield Eq. 3 and 4.

$$\begin{bmatrix} s_{mg}^{ref} \\ s_{ce}^{ref} \end{bmatrix} = \bar{F} \begin{bmatrix} \alpha & 0 \\ 0 & \beta \end{bmatrix} \begin{bmatrix} s_{mg}^{ref} \\ s_{ce}^{ref} \end{bmatrix} \tag{2}$$

$$\alpha = \frac{1}{k_{mg}} = \frac{Pr - \frac{k_{12}}{k_{22}}}{(k_{11} - \frac{k_{21}k_{12}}{k_{22}})Pr} \tag{3}$$

$$\beta = \frac{1}{k_{ce}} = \frac{1 - Pr\frac{k_{21}}{k_{11}}}{(k_{22} - \frac{k_{21}k_{12}}{k_{11}})} \tag{4}$$

$$\bar{F} = \begin{bmatrix} k_{11} & k_{12} \\ k_{21} & k_{22} \end{bmatrix} Pr = \frac{s_{mg}^{ref}}{s_{ce}^{ref}} \tag{5}$$

Where Eqs. 5 show the indexing of the fission matrix and definition of a production ratio Pr. In the fission matrix MG region is denoted as 1 and CE as 2. Each k_{ji} term represents the average number of fission neutrons in region j produced by a history of a fission neutron born in region i. s_{mg}^{ref} and s_{ce}^{ref} represent neutron production in CE and MG region in the reference calculation used to obtain Pr. Moreover, k_{mg} and k_{ce} represent the inferred values of whole system MG and CE criticality. It is crucial that in expressions for α and β off-diagonal terms of the fission matrix are not present in the denominator of the nested fractions. In a case of loosely coupled MG and CE zones, these values can become very small and source normalisation terms could become very sensitive to floating point error.

The coefficients of the fission matrix in Eq. 5 used to infer the normalisation coefficients are sampled during the MC calculation. Thus, the only input expected from the user perspective is the decomposition of the domain into MG and CE regions, as well as the value of the production ratio Pr (Eq.5). The normalisation scheme introduces a systematic error into the calculation which is related to the accuracy of the production ratio. Any error in the ratio will cause the neutron production in MG and CE region to be

increased or depressed throughout the region. The effect of the statistical variation in the fission matrix terms will require a further study as this problem is related to the non-linearity in the source convergence introduced by differential normalisation of the CE and MG regions.

There are two possible ways in which the differential normalisation procedure can fail. In the first case, it is possible that the matrix $\bar{F}$ and value of Pr will be such that one of the normalisation coefficients α or β will become negative. Then, it would not be clear how to perform normalisation in the standard Monte Carlo calculation framework. By examining the Eq. 3 and 4 it is possible to show that the allowable values of Pr, required for preventing the negative values of α and β, are given by the following inequality:

$$\frac{k_{12}}{k_{22}} < Pr < \frac{k_{11}}{k_{21}} \tag{6}$$

The second condition under which the procedure may fail is when the fission matrix of the combined system becomes singular. Then, provided that the sparsely discretised fission matrix captures the second eigenvalue without a large error, the dominance ratio of the system must be close to 0. This is unlikely to happen for any system of practical interest.

2.3. MG-CE Interface

In principle, the change of representation of nuclear data for a particle can be performed during transition or collision. In this paper, the change during transition was used in order to ensure clean separation of MG and CE regions, which makes the finding of the neutron production ratio straightforward. Thus, the nuclear data representation was switched when the particle hit a pre-defined boundary during its movement. For simplification, the boundary was chosen in such a way that it lied only in a moderator.

2.4. Resonance Correction

Unlike in a trivial case of transition from CE to MG, the selection of an exact energy value, when only the energy group of a particle is available requires some care. In the thermal range, the variation of cross-sections in energy is smooth and, provided that sufficiently fine group structure is used to represent the thermal spectrum (as is the case in e.g. WIMS-69 group structure), the choice of the exact energy value has little effect on the result. Therefore, we will focus on the selection of energy value in the resonant range.

For the purpose of choosing a specific value of energy when a particle transitions from a region represented in MG to a region with CE treatment, it is necessary to obtain a function that approximates the CE spectrum of partial current leaving the MG region. This function can be used as a probability density function for sampling the exact energy of the particle that enters the CE region from the MG region. Although the methods usually employed for cross-section shielding will be used to obtain it, please note that no re-shielding of MG cross-sections is proposed. It is crucial to capture the depressions of the spectrum caused by the resonances in order to avoid an artificial increase in the capture rate in the pins closest to the MG-CE boundary in the CE region. A reasonable approximation can be made by considering a spectrum of average moderator flux in pin cells close to the boundary. Assuming that the problem consists only of moderator and a single fuel type, and noting that the neutrons energy loss in collisions with the moderator is typically much larger than resonance widths, it is possible to write an expression for the flux in the moderator as a function of energy using the Narrow Resonance Approximation [12]:

$$\begin{aligned} &\phi_m(E)\Sigma_{t,m}(E)V_m = \\ &= \tfrac{1}{E}\left(P_{m\to m}(E)V_m\Sigma_{t,m}(E) + P_{f\to m}(E)V_f\Sigma_{p,f}(E)\right) \end{aligned} \tag{7}$$

Where subscripts m and f stand for the moderator and fuel quantities respectively. $P_{i\to j}$ represents transfer probability between region i and j and V_j is volume of region j. Absorption in moderator was neglected and Σ_p and Σ_t denote the potential scattering and total cross sections respectively. Also, it is assumed that leakage probability is negligible. Thus, it is possible to rearrange Eq. 7 and use that $P_{m\to m}(E) = 1 - P_{m\to f}(E)$ to obtain :

$$\phi_m(E) = \frac{1}{E}\left(1 - P_{m\to f}(E) + P_{f\to m}(E)\frac{V_f}{V_m}\frac{\Sigma_{p,f}(E)}{\Sigma_{t,m}(E)}\right) \tag{8}$$

Using the Reciprocity Relation between the fuel and moderator regions:

$$P_{m\to f}(E) = P_{f\to m}(E)\frac{V_f}{V_m}\frac{\Sigma_{t,f}(E)}{\Sigma_{t,m}(E)} \tag{9}$$

Thus, by substituting Eq. 9 into Eq. 8, we obtain an expression for the flux spectrum in the moderator as a function of energy parametrised by geometry, cross-sections and fuel escape probability.

$$\phi_m(E) = \frac{1}{E}\left(1 + P_{f\to m}(E)\frac{V_f}{V_m}\frac{\Sigma_{p,f}(E) - \Sigma_{t,f}(E)}{\Sigma_{t,m}(E)}\right) \tag{10}$$

The expression for neutron spectrum provided by Eq. 10 will be valid as long as all fuel pins in the neighbourhood of the point of transition have the same composition. This is unlikely to be the case for any practical problem. Thus, the model will need to be extended to account for local lattice heterogeneity. However, the model is sufficient for the initial demonstration. An energy value for the MG to CE transition can be sampled with rejection sampling using Algorithm 1.

Algorithm 1 Energy selection for MG to CE transition

```
if G ∈ Resonant Groups then
  repeat
    E = draw from pdf uniform in lethargy
    Σ_{t,f} = max(Σ_{t,f}(E), Σ_{p,f})
    P_{f→m} = P_{f→m}(Σ_{t,f})
    Σ_{t,m} = Σ_{t,m}(G)
    T = term in bracket on RHS of Eq. 10
  until rand() < T
else
  E = draw from pdf uniform in lethargy
end if
```

Algorithm 1 requires that the lattice fuel escape probability is provided as a function of fuel total cross-section. A single-term rational approximation was used (Eq. 11) with constant C being a combination of the Bell and Dancoff factors. C was evaluated in a Monte Carlo calculation with SCONE in an infinite fuel lattice by evaluating Eq 12, where the samples of a transition function T were obtained from a random walk of the particles and f and m stand for fuel and moderator regions respectively. Because C is a strong function of the moderator cross-section, ΔE needs to be limited to the resonant energy range where the moderator cross-section is approximately constant ($1eV - 20keV$ for light water).

$$P_{f\to m}(\Sigma_{t,f}(E)) = \frac{C}{\Sigma_{t,f}(E) + C} \tag{11}$$

$$C = \frac{\int_{\Delta E}\int_f\int_m T(r \to r', E)\Sigma_{t,f}(E)dr'drdE}{\int_{\Delta E}\int_f\int_f T(r \to r', E)dr'drdE} \tag{12}$$

2.5. Calculation Scheme

In order to fully realise the acceleration potential of using MG data in MC calculation, the fraction of time spend by the code on obtaining cross-section data needs to be maximised. In order to achieve it, delta-tracking was used [13]. This should not result in a significant reduction of performance of the CE portion of the calculation for a large class of problems. For problems with localised absorbers, it was shown that inclusion of some surface tracking can significantly decrease the penalty due to low collision acceptance threshold [13]. Furthermore other techniques that are based on the implicit treatment of virtual collisions can be used to avoid the reduction in performance [14][15].

In order to implement the differential normalisation described in Section 2.4, an implicit fission site generation rate was adjusted by an appropriate weighting coefficient. SCONE uses the same method for implicit fission site generation as MCNP4 [16], in which at every neutron collision with a fissile nuclide, the following number of secondary fission neutrons is generated:

$$n = \left\lfloor \frac{W\bar{\nu}\sigma_f}{k_{eff}\sigma_t W_{new}} + r \right\rfloor \tag{13}$$

Where k_{eff} is k_{mg} or k_{ce}, defined by Eq. 3 and 4 in MG or CE region respectively, r is a random number between 0 and 1, W is the weight of the incident particle, $W_{new} = 1$ is the weight of the secondary particle and $\lfloor \cdot \rfloor$ is the floor function.

MG cross-sections used in all calculations were prepared with Serpent [17] MC code. Both Serpent and SCONE used the same nuclear data library based on JEFF 3.1.1 for CE calculations. It is important to note that using MC-generated reaction cross-sections in a variable fidelity MC calculation would be impractical in real calculations. When using MG MC calculations, it is preferable to use fine energy group structure, as with the MG data, there is no need to search and interpolate the energy grid. The only section of a MG MC code, in which runtime scales with the number of groups is the sampling of outgoing energy groups in scattering and fission. Thus, the runtime of the MC code has little dependence on the number of energy groups, but the finer representation of cross-sections variation with energy can still significantly improve the accuracy. On the other hand, generation of cross-sections with CE MC code on a fine grid requires very large particle population to accurately capture all low-probability group to group transitions in scattering.

In this paper, Serpent was used only to ensure the data consistency between the CE and MG calculations, as well as because no deterministic code able to provide the required cross-sections was available.

The MG MC calculations were performed using WIMS 69 energy group structure with P1 scattering data and using algorithm for sampling outgoing direction described in [18]. CE calculations were performed without unresolved resonances probability tables and without $S(\alpha, \beta)$ bound scattering both in Serpent and SCONE. Moreover, the results were batched every cycle.

3. Results

The computational efficiency of a variable fidelity MC calculation is heavily dependant on the relative size of CE and MG portion of the calculation domain in terms of number of particles being tracked, as well as relative speed-up of the MG calculation with respect to CE. As a result, we need to define a metric that would allow to compare the performance of the variable fidelity algorithms across a number of different problems and implementations of MG and CE MC particle transport, so that the results obtained with SCONE could be generalised to more advanced, industrial codes. To obtain a consistent reference, we will approximate an ideal speed of variable fidelity calculation by the average of the runtimes of pure MG and CE calculation, weighted by the fraction of neutron production in a given region as shown in Eq 14.

$$T_{ideal} = \frac{T_{mg} N_{prod}^{mg} + T_{ce} N_{prod}^{ce}}{N_{prod}^{mg} + N_{prod}^{ce}} \tag{14}$$

Where T_{mg} and T_{ce} denote the runtimes of MG and CE calculation over the entire domain and N_{prod}^{mg} and N_{prod}^{ce} represent the neutron production rates in MG and CE part of the domain during the variable fidelity calculation.

3.1. Uniform 1D Problem

The first problem was designed to show the effectiveness of the differential normalisation procedure. It consisted of a 34x17 lattice of PWR pins with reflective boundary conditions. The fuel enrichment was 2.25 wt% with Boron-10 in water at the density of $2.0 \cdot 10^{-5} \frac{1}{barn \cdot cm}$ (443 ppm). This problem effectively represents two PWR 17x17 assemblies next to each other. The left assembly (rows 1-17) were represented in MG.

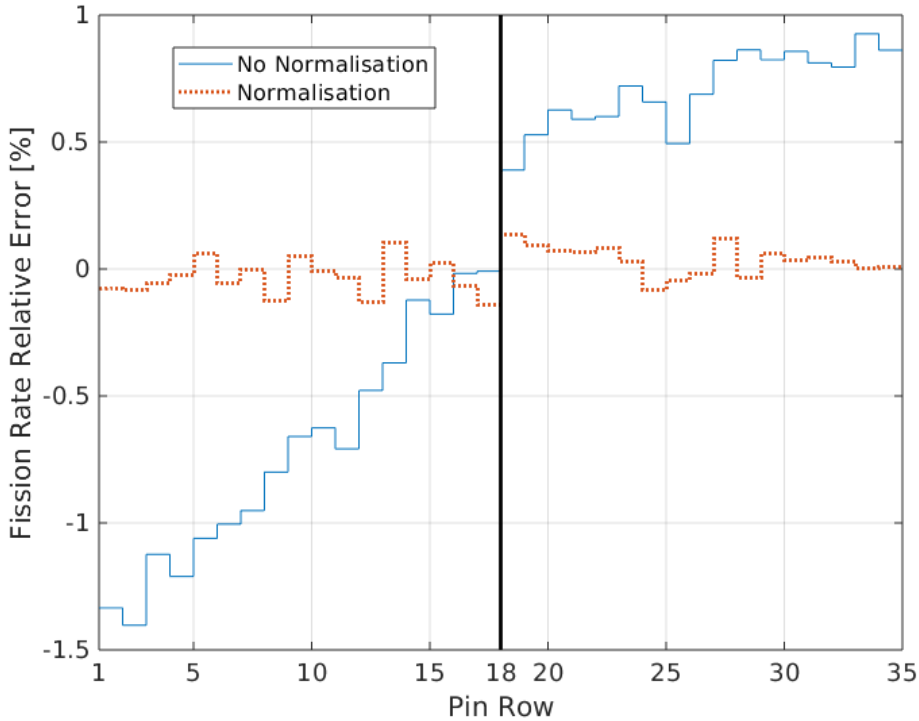

Figure 1: Fission rate distribution in uniform UOX 1D problem. Case with and without differential normalisation is shown. Rows 1-17 were represented in MG. Standard deviation is about 0.082%

Figure 1 shows the distribution of the error in fission rate relative to an analytic, uniform distribution. The importance of differential normalisation even in this simple case is evident. Without it, a large skew in fission rate distribution can be observed, which is consistent with the behaviour suggested by the discussion of a homogeneous slab in Section 2.4. If the differential normalisation procedure is applied, the error is reduced to a level comparable with statistical uncertainty. All calculations of the test case were performed with population of 80000, 200 inactive and 800 active cycles.

3.2. Heterogeneous 1D Problem

The second problem used the same geometry as uniform UOX problem, but the fuel composition in different rows was varied as shown in Figure 2. Its purpose was to demonstrate that it is possible to recover significant accuracy in CE region of variable fidelity calculation in a simple test case. From

Figure 2: Diagram of the pin rows arrangement in the heterogeneous 1D problem. Atomic composition for diffrent MOX enrichments are taken directly from [19]

Figure 3, which shows the distribution of error in fission rate in variable fidelity and pure MG calculations with respect to CE reference, it is evident that this is indeed the case. The error in fission rate in the CE portion of the domain was reduced to levels comparable to statistical uncertainty. In this case, the ratio of neutron production in MG to CE region, required by the differential normalisation procedure, was obtained from a fission matrix sampled in a CE MC calculation with neutron population of 5000 with 100 active and 300 inactive cycles. All calculations of the test case were performed with neutron population of 80000, 200 inactive and 800 active cycles.

3.3. 2D Problem

In order to investigate the performance of the combined CE-MG MC calculation on a more realistic case, a chessboard configuration of PWR UOX and MOX assemblies was studied. The layout and material compositions of the assemblies was taken from [19]. High fidelity region was square, centred in the middle of the domain and 16 pin cells wide. A Sketch of the configuration is shown in Figure 4. This configuration was chosen to test how the method would behave in the presence of significant thermal flux gradients along the boundary. All calculations were performed in SCONE with neutron population of 500000 particles, 300 inactive and 3000 active cycles. For the combined calculation, the neutron production ratio was obtained from a coarsely discretised fission matrix (9x9) sampled in a CE MC calculation with neutron population of 10000 with 50 inactive cycles and 300 active cycles. MG cross-sections were generated in a single assembly calculations with Serpent.

Figure 5 compares the relative error in fission distribution with respect to CE MC calculation for a case of variable fidelity calculation and pure MG calculation. The effectiveness of combined CE-MG approach is immediately evident. Significant errors of up to 3% were reduced to below 1% within CE region. Although, this imprecision may be considered large, it is important to note that the desired behaviour of decreasing error away from the boundary was achieved. A relatively large 1% error was only observed in the immediate

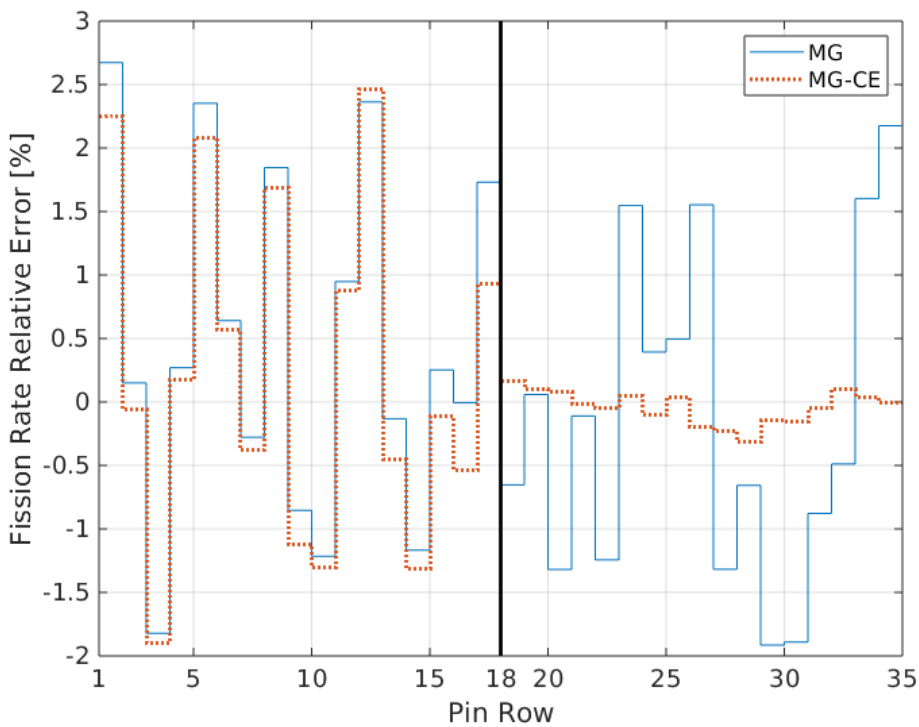

Figure 3: Fission rate distribution in heterogeneous 1D problem. Results are shown for variable fidelity and pure MG case. Rows 1-17 were represented in MG in variable fidelity calculation. Standard deviation is about 0.094%. Error is relative to the full CE solution with SCONE.

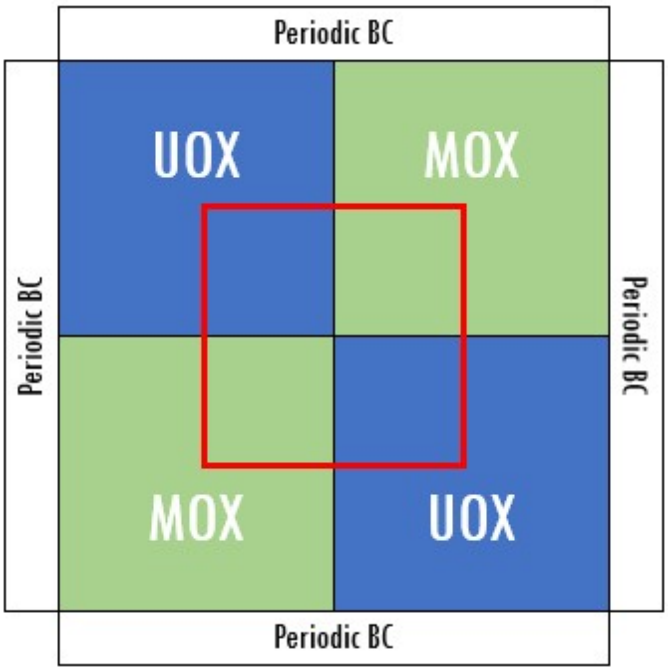

Figure 4: Sketch of the calculation domain of the 2D PWR problem. Chessbord configuration with UOX and MOX assemblies. Red line indicates boundary between CE and MG region.

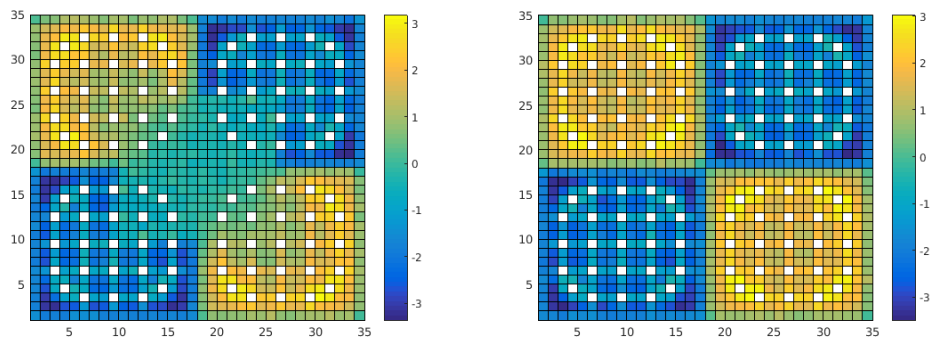

Figure 5: Relative error in fission rate distribution [%] with respect to CE MC calculation. Left: Variable Fidelity Calculation. Right: Pure MG Calculation. Note that the gaps correspond to guide tubes pin cells without fissile material. Standard deviation is about 0.08%.

vicinity of the boundary and it resulted from the large gradient in thermal flux in the MG region. Figures 10 and 11 show the detailed distribution of fission rate error inside the CE region of the variable fidelity calculation. It is clear that in the case of 2D problem, normalisation procedure made little difference in comparison to the effect it had in the 1D cases. The reason for it is a strong coupling between the MG and CE regions. In the 2D case, about 50% of fissions in the CE region were caused by neutrons born in the MG zone. By comparison, in the uniform 1D problem, the same parameter was only about 16%. As a result, the normalisation value in the CE region had little effect on the fission distribution. At the same time, the source in the MG region was also unaffected, because the majority of particles were tracked in MG, so the estimate of criticality without differential normalisation was virtually identical to the pure MG problem. In fact, from Figures 10 and 11 it is clear that the inclusion of the normalisation increased the error in the middle of the CE zone. This was most likely caused by the systematic error in the production ratio.

3.4. Performance of Resonance Correction

The need for the resonance correction described in Section 2.4 is clearly evident in Figures 8 and 9, which show the distribution of relative error in resonant capture relative to the high fidelity CE solution for 1D UOX and MOX problems. Resonant capture was defined as any capture reaction in the fuel of neutron with incident energy between 4 eV and 9.118 keV (Resonant Groups (15-27) in WIMS 69 Energy group structure). When the correction is switched-off and the neutrons are distributed uniformly in lethargy during

Error in Fission Rate in CE region - With Normalisation [%]															
1.04	0.96	1.19	0.75	0.65	0.72	0.61	0.04	-0.35	-0.46	-0.31	-0.66	-0.46	0.20	-0.40	-0.55
0.96	0.90	0.82	0.50	0.38	0.58	0.14	-0.11	-0.09	-0.28	-0.27	-0.66	-0.73	-0.71	-0.57	-0.45
1.03	0.75		0.56	0.25		-0.06	-0.24	-0.26	-0.53		-0.44	-0.51		-0.60	-0.06
0.75	0.76	0.49	0.13	0.02	-0.18	-0.04	-0.29	-0.46	-0.65	-0.47	-0.69	-0.60	-0.54	-0.68	-0.45
0.90	0.57	0.18	0.19		-0.18	-0.36	-0.32	-0.31	-0.67	-0.60		-0.69	-0.57	-0.62	-0.60
0.77	0.50		0.03	-0.02	-0.20	-0.48	-0.69	-0.72	-0.59	-0.65	-0.52	-0.51		-0.42	0.13
0.63	0.25	0.00	-0.13	-0.29	-0.45	-0.48	-0.44	-0.47	-0.31	-0.45	-0.45	-0.61	-0.35	-0.45	-0.55
0.08	-0.25	-0.22	-0.12	-0.20	-0.45	-0.50	-0.23	-0.59	-0.46	-0.43	-0.39	-0.37	-0.31	-0.25	-0.27
-0.13	-0.50	-0.24	-0.24	-0.54	-0.61	-0.59	-0.49	-0.62	-0.66	-0.51	-0.71	-0.17	-0.01	0.04	0.22
-0.54	-0.49	-0.64	-0.55	-0.57	-0.58	-0.64	-0.62	-0.15	-0.47	-0.46	-0.45	-0.20	0.02	0.27	0.70
0.12	-0.50		-0.55	-0.69	-0.60	-0.63	-0.35	-0.42	-0.26	-0.21	0.03	0.13		0.52	0.94
-0.41	-0.67	-0.54	-0.78		-0.49	-0.65	-0.32	-0.22	-0.25	-0.21		0.20	0.24	0.53	0.76
-0.49	-0.59	-0.65	-0.61	-0.65	-0.62	-0.50	-0.25	-0.52	-0.10	0.05	0.12	0.42	0.38	0.56	0.63
0.03	-0.68		-0.54	-0.64		-0.46	-0.06	-0.18	0.12		0.24	0.31		0.69	1.03
-0.56	-0.70	-0.47	-0.56	-0.54	-0.45	-0.63	-0.20	0.15	0.29	0.38	0.57	0.68	0.75	1.08	0.92
-0.39	-0.58	0.05	-0.60	-0.57	-0.02	-0.33	-0.37	-0.07	0.85	1.03	0.85	0.92	1.05	1.00	1.16

Figure 6: Relative error in fission rate distribution [%] with respect to CE MC calculation in the CE region of variable fidelity calculation. Case with normalisation correction. Standard deviation is about 0.08%.

Error in Fission Rate in CE region - No Normalisation [%]															
1.05	0.85	1.08	0.97	0.69	0.79	0.82	0.33	-0.35	-0.36	0.00	-0.47	-0.39	0.11	-0.30	-0.46
0.75	0.79	0.67	0.53	0.50	0.61	0.28	0.28	0.09	-0.09	-0.28	-0.55	-0.60	-0.56	-0.71	-0.48
1.18	0.84		0.52	0.34		0.15	-0.02	-0.20	-0.30		-0.40	-0.41		-0.68	0.02
0.83	0.63	0.49	0.28	0.25	-0.03	0.12	0.02	-0.18	-0.35	-0.31	-0.51	-0.40	-0.66	-0.65	-0.27
1.07	0.46	0.46	0.21		-0.02	-0.12	-0.03	-0.15	-0.36	-0.26		-0.51	-0.50	-0.61	-0.51
0.93	0.55		0.20	0.04	-0.05	-0.27	-0.26	-0.21	-0.40	-0.41	-0.37	-0.23		-0.43	0.08
0.80	0.22	0.09	0.04	-0.22	-0.15	-0.10	-0.18	-0.02	-0.01	0.06	-0.07	-0.36	-0.21	-0.15	-0.31
0.27	-0.03	0.12	-0.02	0.20	-0.04	-0.25	0.00	-0.19	-0.43	-0.11	0.04	-0.21	-0.25	-0.22	-0.08
-0.13	-0.41	-0.08	0.01	-0.25	-0.27	-0.40	-0.21	-0.14	-0.09	-0.03	-0.31	-0.03	0.04	0.24	0.32
-0.45	-0.37	-0.56	-0.48	-0.36	-0.07	-0.10	-0.17	-0.01	-0.19	0.13	0.02	0.08	0.23	0.41	0.88
0.13	-0.46		-0.26	-0.20	-0.38	-0.20	-0.08	-0.12	-0.01	0.08	0.17	0.17		0.66	1.17
-0.37	-0.50	-0.47	-0.51		-0.36	-0.20	-0.32	0.05	-0.27	-0.12		0.11	0.45	0.65	1.10
-0.38	-0.48	-0.55	-0.58	-0.44	-0.29	-0.02	-0.19	-0.09	0.08	0.16	0.14	0.42	0.46	0.59	0.92
0.24	-0.79		-0.61	-0.64		-0.31	-0.18	0.02	0.13		0.25	0.34		0.82	1.01
-0.51	-0.62	-0.51	-0.60	-0.57	-0.37	-0.59	0.03	0.02	0.38	0.52	0.61	0.63	0.77	0.96	1.03
-0.48	-0.39	0.03	-0.58	-0.59	0.09	-0.30	-0.13	0.24	1.12	1.13	0.74	0.81	1.23	0.86	1.02

Figure 7: Relative error in fission rate distribution [%] with respect to CE MC calculation in the CE region of variable fidelity calculation. Case without normalisation correction. Standard deviation is about 0.08%.

the MG to CE transitions in all energy groups, a significant artificial peak in capture rate is present in the row of pins closest to the boundary between the regions. It has a magnitude of about 3% in both MOX and UOX case.

The resonance correction makes an assumption of uniform fuel composition in the entire lattice. In the UOX case, this assumption is valid and a single material composition corresponding to fuel can be set without any approximation. In the MOX case, this is no longer the case. For simplicity, the composition of the fuel used to evaluate the macroscopic total cross-section in Eq. 10 was set to the composition of 2.25wt% enriched UOX fuel. This decision was arbitrary. Nevertheless, despite this crude approximation, the artificial capture peak was successfully reduced to a level comparable with statistical uncertainty in both cases. This is a reasonable result because both problems are low enriched. As a result, a majority of heavy metal contribution to the fuel total cross-section originates from U-238, which has almost the same density in all types of fuel. Thus, any resonances of other actinides are diluted and do not contribute to the shape of moderator spectrum in any significant way. Based on these results, it can be concluded that the exact choice of the fuel has little effect on the performance of the resonance correction for the problems with low enriched UOX and MOX fuel.

Figures 10 and 11 show the relative error of the resonant capture rate with and without the resonance correction enabled. The settings for the fuel composition were exactly like in the 1D MOX case. Similar to the 1D cases, the effectiveness of the correction is evident. In the case without correction, an error of up to 7% is observed. It is much higher than the equivalent error in 1D cases. This can be explained by the fact that the 2D case uses different pin cell geometry with volumetric water to fuel ratio reduced from about 0.515 to 0.482, which should increase the importance of self-shielding. When the resonance correction is enabled, the error is reduced to about 1% for the majority of the CE region circumference. The only exception are the pin cells directly opposite the guide tube cells in the MG zone, where the capture rate is underestimated with an error of up to -2.5%. This is reasonable given that in the CE case, the self-shielding in these pins is significantly reduced by the absence of fuel in the guide tubes. The resonance correction does not account for this in any way and uses an estimate of moderator spectrum from an UOX pin cell to CE spectrum of neutron current. It is interesting to observe that this effect is highly localised. Only the pin cells directly opposite the guide tubes are effected. This suggests that this problem could be addressed by simply choosing the fuel composition from the opposite MG

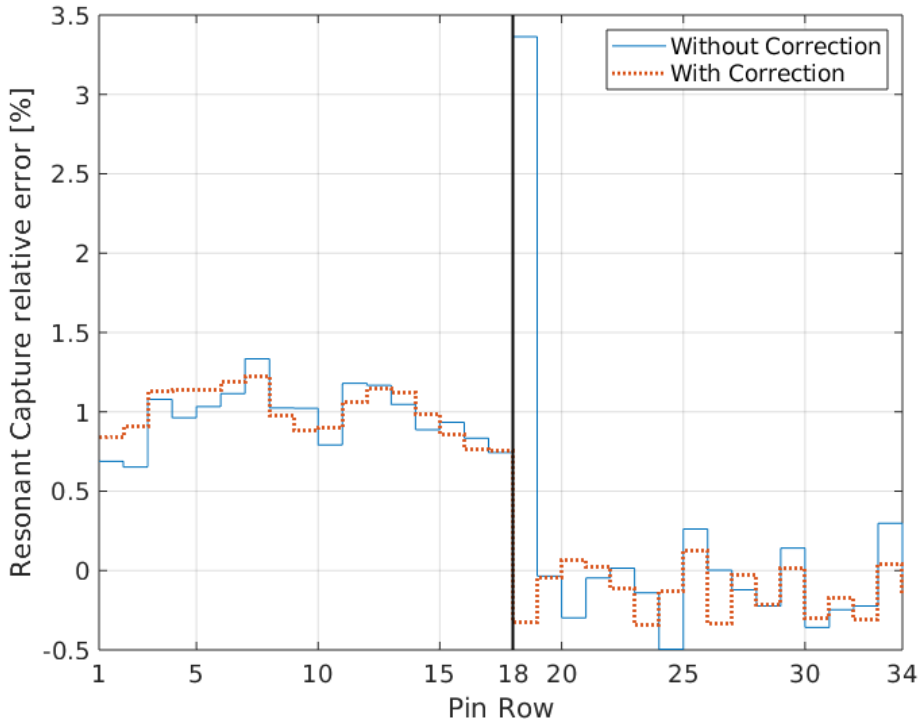

Figure 8: Relative error of resonant capture in fuel with respect to single fidelity CE solution for 1D UOX problem. Standard deviation is about 0.17%. Rows 1-17 were represented in MG.

pin cell in the Eq. 10 (in case of guide tube having total fuel macroscopic cross-section equal to 0). In principle, as long as the MG-CE boundary is constrained to a fuel lattice, such solution could be implemented without too much of programming difficulty into SCONE. However, this places a rather large number of restrictions on the geometry definitions and limits the feasible geometries to fixed-pitch lattices of pin cells. Thus, the choice to use such approximation is a case of a trade-off between accuracy and generality and it is not yet clear which is more preferable in this case.

3.5. Computational Performance

All calculations shown in this paper were performed using a single core on a desktop computer with Intel i7-3770 CPU. The runtime of all types of calculations for the three test problems is shown in Table 1. Figure of merit was based on the average relative standard deviation of the pins fission

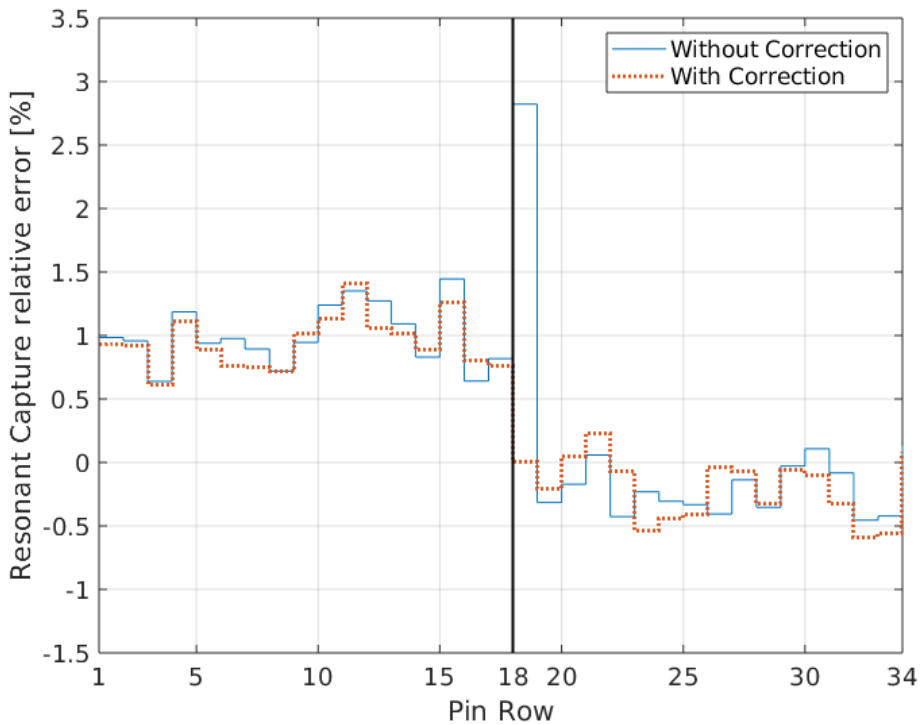

Figure 9: Relative error of resonant capture in fuel with respect to single fidelity CE solution for 1D MOX problem. Standard deviation is about 0.17%. Rows 1-17 were represented in MG.

Error in Resonant Capture in CE region - With Correction [%]															
-0.83	-0.16	-1.76	-0.75	-0.47	-2.89	-1.32	-0.51	-0.55	-0.80	-2.15	-1.41	-1.42	-2.29	-1.07	-1.37
-0.31	0.63	0.27	0.14	0.19	-0.28	-0.43	-0.45	-0.39	-0.67	-0.39	-1.18	-0.59	-0.65	-0.64	-1.42
-1.88	0.34		-0.32	-0.04		-0.46	-0.69	-0.73	-0.79		-0.82	-0.54		-0.77	-2.31
-0.65	0.06	-0.17	0.05	-0.03	0.03	-0.38	-0.52	-0.68	-0.82	-0.59	-0.90	-0.90	-0.64	-0.70	-1.05
-0.99	-0.13	-0.16	-0.58		-0.12	-0.16	-0.45	-0.22	-0.22	-0.63		-0.80	-0.86	-0.35	-1.21
-2.61	-0.25		-0.09	-0.26	-0.61	-0.71	-0.30	-0.52	-0.64	-0.83	-0.32	-0.33		-0.47	-2.53
-1.15	-0.40	-0.14	0.04	-0.53	-0.35	-0.67	-0.55	-0.79	-0.98	-0.46	-0.85	-0.63	-0.77	-0.33	-0.83
-0.96	-0.13	-0.68	-0.80	-0.94	-0.64	-0.43	-0.64	-0.85	-0.73	-0.97	-0.63	-0.65	-0.75	-0.79	-0.59
-0.43	-0.43	-0.56	-0.60	-0.59	-0.51	-0.64	-0.54	-0.99	-0.56	-0.44	-0.80	-0.49	-0.66	-0.56	-0.42
-0.85	-0.95	-0.76	-0.54	-0.53	-0.66	-1.14	-0.72	-0.68	-0.91	-1.09	-0.55	-0.52	-0.59	-0.69	-1.02
-2.01	-0.59		-0.88	-0.39	-0.18	-0.92	-0.43	-0.22	-0.40	-0.38	-0.16	-0.14		-0.19	-2.26
-1.05	-0.73	-0.80	-0.91		-0.56	-0.39	-0.65	-0.83	-1.01	-0.38		-0.57	0.15	-0.27	-0.68
-1.39	-0.66	-0.61	-0.74	-0.53	-0.70	-0.65	-0.40	-0.45	-0.02	-0.24	-0.02	-0.32	-0.39	0.28	-0.53
-2.67	-0.39		-0.39	-0.12		-0.61	-0.35	-0.14	-0.40		0.50	0.09		0.70	-1.99
-1.45	-1.17	-0.44	-1.15	-0.59	-0.26	-0.66	-0.84	-0.50	-0.41	-0.34	0.05	0.09	0.44	0.55	-0.10
-1.16	-1.28	-2.37	-1.00	-1.11	-2.13	-1.00	-0.25	-0.59	-1.00	-2.06	-1.11	-0.67	-1.80	-0.17	-0.36

Figure 10: Relative error [%] of resonant capture rate in fuel with respect to CE single fidelity calculation. Case with resonance correction enabled. Standard deviation is about 0.18%.

Error in Resonant Capture in CE region - No Correction [%]															
7.16	3.93	2.94	3.85	3.55	1.80	3.61	3.92	3.38	2.56	1.19	2.01	2.23	1.03	2.27	4.87
4.13	0.83	0.63	0.25	-0.03	0.14	-0.09	-0.14	-0.18	-0.32	-0.55	-0.89	-0.70	-0.86	-0.89	2.39
2.60	0.32		-0.36	-0.37		-0.75	-0.52	-0.55	-0.82		-0.58	-0.65		-0.42	1.08
3.46	-0.24	-0.36	0.22	-0.24	-0.28	-0.33	-0.62	-0.86	-0.63	-0.42	-0.97	-0.73	-0.69	-0.73	2.43
3.65	0.13	-0.05	-0.37		-0.19	-0.50	-0.33	-0.46	-0.09	-0.58		-0.94	-0.90	-0.58	2.01
1.96	-0.01		-0.39	0.01	-0.97	-0.47	-0.66	-0.40	-0.64	-0.97	-0.10	-0.66		-0.40	1.18
3.22	-0.19	-0.74	-0.05	-0.56	-0.15	-0.70	-0.75	-0.72	-0.68	-0.82	-0.42	-0.70	-0.98	-0.53	2.41
3.57	-0.08	-0.59	-0.70	-0.69	-0.82	-0.71	-0.67	-0.48	-0.85	-0.72	-0.48	-0.54	-0.83	-0.52	3.51
3.48	-0.27	-0.45	-0.72	-0.42	-0.55	-0.69	-0.64	-1.22	-0.60	-0.54	-0.78	-0.37	-0.56	-0.12	3.82
2.69	-0.87	-0.55	-0.42	-0.42	-0.59	-0.83	-0.88	-0.64	-0.48	-0.56	-0.36	-0.39	-0.30	-0.08	3.45
1.58	-0.72		-0.65	-0.60	-0.59	-0.90	-0.81	-0.07	-0.47	-0.29	-0.12	-0.50		-0.24	2.12
2.23	-0.88	-1.06	-0.95		-0.82	-0.86	-0.45	-0.70	-0.73	-0.20		-0.59	-0.01	0.04	3.58
2.15	-0.50	-0.66	-0.83	-0.90	-0.84	-0.60	-0.56	-0.50	-0.60	-0.17	-0.13	-0.02	-0.21	0.39	4.05
0.54	-0.54		-0.65	-0.52		-0.77	-0.62	-0.22	-0.17		0.13	-0.03		1.01	2.12
2.09	-0.93	-0.51	-0.86	-0.55	-0.26	-0.80	-0.42	-0.38	0.20	-0.30	0.07	0.16	0.52	0.76	4.28
4.96	1.97	1.26	2.46	2.45	1.63	2.53	3.64	3.74	3.54	2.52	3.40	3.37	2.47	4.43	7.34

Figure 11: Relative error [%] of resonant capture rate in fuel with respect to CE single fidelity calculation. Case with resonance correction disabled. Standard deviation is about 0.18%.

Test Case	Type	Run Time [s]	$\sigma[10^{-4}]$	FoM [-]
UOX	CE	4449	8.05	347
	MG	1118	8.36	1281
	MG-CE	3190	8.20	466
	MG-CE*	3153	8.20	471
	Ideal	2784	N/A	N/A
1D MOX	CE	6387	8.42	221
	MG	1249	8.50	1106
	MG-CE	4320	8.44	325
	MG-CE*	4300	8.44	332
	Ideal	3827	N/A	N/A
	PR	135	N/A	N/A
2D MOX	CE	151797	8.44	9.25
	MG	35531	8.43	39.65
	MG-CE	76013	8.43	18.52
	MG-CE*	74880	8.43	18.80
	Ideal	61836	N/A	N/A
	PR	305	N/A	N/A

Table 1: Runtime comparison of CE, MG and variable fidelity MG-CE calculations of three test problems. Ideal performance was calculated with Eq. 14. MG-CE* case has the resonance correction disabled. Figure of Merit (FoM) is based on average relative standard deviation of fission rate. PR case shows runtime of the precalculation to get value of production ratio from Eq.5.

rate. One interesting feature is the non-ideal convergence of the 2D MOX results, which despite significantly higher population settings failed to achieve better standard deviation on fission powers. However, it is easy to check that relative performance of CE and MG calculations is consistent between the 1D and 2D MOX problem. Thus, the non-ideal convergence will not be investigated in detail as this is not the focus of this paper.

Based on Table 1, it is evident that in all the considered problems, the variable fidelity calculation was noticeably faster than the equivalent pure CE calculation. However, it is also clear that the magnitude of this speed-up varies across different test cases. The fastest acceleration by a factor of about 2 was obtained for 2D MOX case. This is reasonable given that in this case, the CE region occupied the smallest proportion of the domain. This result was also improved by a better relative performance of the MG calculation with respect to the CE calculation in problems with MOX fuel.

This is related to the deficiency of CE nuclear data in current implementation of SCONE, which lacks the energy grid search acceleration mechanism, so the runtime scales linearly with the number of nuclides in the problem. It is also important to note that the computational performance of the combined MG-CE calculations was relatively modest in all test cases due to the large size of the CE region compared to MG region. In a full core calculation, it is conceivable that a CE region would span no more than few assemblies. Thus, in such cases, the computational performance of the whole calculation should be substantially improved and follow the pure MG speed more closely. It should also be noted that the time required for the calculation of the production ratio was small in comparison with the runtime of the combined calculation for both the 1D and 2D MOX case. It constituted only 3% and 0.4% of total calculation time (pre-calculation + combined calculation) in 1D and 2D MOX case respectively.

Different perspective can be obtained by comparing the relative difference of the actual variable fidelity runtime to the theoretical ideal one (Eq. 14). Then, it is clear that the 2D MOX case was actually the least efficient with runtime of about 19% above the ideal. The same metric was 12.8% and 11% for 1D MOX and UOX case respectively. Clearly, the overhead of the variable fidelity calculation is non-negligible and cannot be ignored. Further study is required to estimate the exact reason for the overhead. However, based on the results in Table 1, it is clear that the extra runtime as a result of rejection scheme in resonance correction is not the main contribution. Only in the 2D MOX case, it seems to have had a notable effect on the runtime and it increased the calculation time by only 1.5%.

4. Conclusions & Further Work

This paper investigated the possibility of combining MG and CE nuclear data representation in different spatial regions of MC calculation domain. It was shown that this approach is viable and can provide a noticeable acceleration with respect to CE calculation, which is heavy dependent on the relative size of the MG and CE zones. However, some error is introduced into a high fidelity CE region. Most of the inaccuracies are contained close to the boundary and the error in fission rate distribution was contained below 1% in the majority of the CE zone in all considered test cases, which were based on MOX-PWR geometry and compositions. The current implementation of the scheme in the in-house MC neutron transport code SCONE, suffers from

relativity large overhead of about 20% and 10% over the theoretical performance in 2D and 1D cases respectively. However, promising acceleration was achieved in all problems in terms of absolute runtime (up to the factor of 2).

Furthermore, it was demonstrated that a differential normalisation procedure is required for MG and CE zones for problems where different fidelity regions are loosely coupled. It requires an additional constraint in the form of a ratio of neutron production in MG to CE region to be introduced into the calculation. The ratio needs to be known a priori and in the cases presented in this paper was obtained from a coarsely discretised fission matrix sampled in a CE MC calculation with much lower population settings. In the 1D MOX case, the initial calculation constituted only 3% of total runtime. In the 2D MOX case, it was 0.4%. The results for the 2D MOX test case suggest that in a majority of realistic cases, in which the CE region is a single assembly sized region in a larger MG zone, there is a strong coupling between the CE and MG region and the differential normalisation makes little difference to the results. The magnitude of the maximum error at the boundary did not change significantly. However, the error increased in the middle of the CE region, possibly as a result of a systematic error in the value of the production ratio.

The results for the studied test cases demonstrate that it is necessary to account for resonance shielding effects in the energy spectrum when switching the particle representation from MG to CE. Otherwise a strong artificial peak in resonant capture can occur in the row of pins closest to the MG-CE boundary. In the 1D test cases, the resulting error in resonant capture rate reached about 3%. In the 2D case, the error increased up to 7%. To mitigate this error a "resonance correction" based on slowing down equations was developed. In the 1D cases it successfully eliminated the artificial peak. In the 2D case it reduced the error to within 3%.

The test cases in this paper were constrained only to PWR-type problems with reflective boundary conditions. Thus, further effort will be focused on extending the current methodology to other systems. In particular, BWR problems should prove challenging to the current resonance correction due to the large variation in moderator density across different fuel assemblies. In addition, various contributions to the computational overhead must be quantified and reduced. Furthermore, the sensitivity of the fission and capture rate distribution to the parameters in resonance correction and normalisation procedure needs to be quantified. In particular the effect of the non-linearity introduced by the normalisation procedure on the source convergence and

true variance of reaction rate estimates must be investigated in more detail.

5. Acknowledgements

This research was supported by a David and Susan Hibbit Scholarship from Jesus College, Cambridge and by EPSRC Doctoral Training Partnership with University of Cambridge.

6. Declaration of Interest

None

7. References

[1] S. Goluoglu, L. M. Petrie, M. E. Dunn, D. F. Hollenbach, B. T. Rearden, Monte Carlo criticality methods and analysis capabilities in SCALE, Nuclear Technology 174 (2011) 214–235. doi:10.13182/NT10-124.

[2] B. Molnar, G. Tolnai, D. Legrady, A GPU-based direct Monte Carlo simulation of time dependence in nuclear reactors, Annals of Nuclear Energy 132 (2019) 46–63. doi:10.1016/j.anucene.2019.03.024.

[3] J. C. Kuijper, S. C. Van Der Marck, A. Hogenbirk, Using homogenized macroscopic group cross sections in continuous-energy Monte Carlo neutron transport calculations with MCNP, in: Proc. M&C+ SNA 2007, Monterey, CA, USA, 2007.

[4] J. Leppänen, On the feasibility of a homogenised multi-group Monte Carlo method in reactor analysis, in: Proc. PHYSOR 2008, Casino-Kursaal Conference Center, Interlaken, Switzerland, 2008.

[5] Z. Peng, W. Kan, Y. Ganglin, A multi-group Monte Carlo core analysis method and its application in SCWR design, in: Proc. PHYSOR 2012, Knoxville Convention Center, Knoxville, TN, USA, 2012.

[6] M. A. Kowalski, E. Shwageraus, Accelerating Monte Carlo by variable fidelity: First results, in: Proc. PHYSOR 2018, Cancun, Mexico, 2018.

[7] D. Lee, H. Lee, S. Choi, A hybrid Monte Carlo/Method-of-Characteristics method for efficient neutron transport analysis, Nuclear Science and Engineering 180 (2015) 69–85. doi:10.13182/NSE13-102.

[8] Q. Zheng, W. Shen, Y. Li, H. Wu, L. Cao, A deterministic-stochastic energy-hybrid method for neutron-transport calculation, Annals of Nuclear Energy 128 (2019) 292–299. doi:10.1016/j.anucene.2019.01.019.

[9] L. Deng, Z. Hu, R. Li, T. Cheng, C. Yang, Y. Fu, D. Shi, G. Li, B. Zhang, L. Zhang, L. Wen, The coupled neutron transport calculation of Monte Carlo multi-group and continuous cross section, Annals of Nuclear Energy 127 (2019) 433–436. doi:10.1016/j.anucene.2018.12.032.

[10] W. M. Stacey, Nuclear reactor physics, Wiley, New York, 2001.

[11] S. Carney, F. Brown, B. Kiedrowski, W. Martin, Theory and applications of the fission matrix method for continuous-energy Monte Carlo, Annals of Nuclear Energy 73 (2014) 423–431. doi:10.1016/j.anucene.2014.07.020.

[12] W. Rothenstein, Resonance absorption calculations in thermal reactors, Progress in Nuclear Energy 5 (1980) 95–144. doi:10.1016/0149-1970(80)90003-7.

[13] J. Leppänen, Performance of Woodcock delta-tracking in lattice physics applications using the Serpent Monte Carlo reactor physics burnup calculation code, Annals of Nuclear Energy 37 (2010) 715–722. doi:10.1016/j.anucene.2010.01.011.

[14] L. Morgan, D. Kotlyar, Weighted-delta-tracking for Monte Carlo particle transport, Annals of Nuclear Energy 85 (2015) 1184–1188. doi:10.1016/j.anucene.2015.07.038.

[15] D. Legrady, B. Molnar, M. Klausz, T. Major, Woodcock tracking with arbitrary sampling cross section using negative weights, Annals of Nuclear Energy 102 (2017) 116–123. doi:10.1016/j.anucene.2016.12.003.

[16] J. F. Briesmeister, MCNP- A General Monte Carlo N-Particle Transport Code, Tech. Rep. LA-13709-M, Los Alamos National Laboratory (2000).

[17] J. Leppänen, M. Pusa, T. Viitanen, V. Valtavirta, T. Kaltiaisenaho, The Serpent Monte Carlo code: Status, development and applications in 2013, Annals of Nuclear Energy 82 (2015) 142–150. doi:10.1016/j.anucene.2014.08.024.

[18] I. Lux, L. Koblinger, Monte Carlo particle transport methods: neutron and photon calculations, CRC Press, Boca Raton, FL, USA, 1991.

[19] Benchmark calculations of power distribution within fuel assemblies, phase II: Comparison of data reduction and power reconstruction methods in production codes, Tech. Rep. NEA/NSC/DOC(2000)3, Nuclear Energy Agency (2000).

[20] D. J. Griffiths, Introduction To Quantum Machanics, Cambridge University Press, Cambridge, UK, 2017.

Appendix A. Variable Fidelity Calculation of 1D Slab Problem

This section contains the derivation of the Equation 1. We consider a one-dimensional multiplying slab with reflective boundary conditions at both ends. The slab has the length of $2L$ and the middle of the slab is at co-ordinate 0. We model the slab with one speed diffusion equation. D denotes the diffusion coefficient and Σ_a and $\nu\Sigma_f$ stand for the absorption cross-section and neutron production cross-section respectively. Please note that absorption includes both fission and radiative capture. Using the usual $\phi(x)$ to represent neutron flux distribution we can write:

$$-D\nabla^2\phi(x) + \Sigma_a\phi(x) = \frac{\nu\Sigma_f}{k}\phi(x) \tag{A.1}$$

$$\left.\frac{d\phi(x)}{dx}\right|_{\pm L} = 0 \tag{A.2}$$

In order to derive the effect of change of representation in the left half of the slab to lower fidelity, we will use 1st order perturbation theory. Since the problem is self adjoint, we can rely on standard approach from quantum mechanics [20]. For convenience, we rewrite Eq. A.1 and define transport operator $\hat{A}$ as follows.

$$\hat{A}\phi(x) = \left[-\nabla^2 + \frac{\Sigma_a - \frac{\nu\Sigma_f}{k}}{D}\right]\phi(x) = [-\nabla^2 + B_m^2]\phi(x) \tag{A.3}$$

Using the new notation, we can solve for the eigenstates of the operator $\hat{A}$ and obtain the following solutions:

$$\hat{A}\phi_l(x) = \lambda_l\phi_l(x) = (-B_{gl}^2 + B_m^2)\phi_l(x) \tag{A.4}$$

$$\begin{array}{ll} \phi_0(x) = 1 & B_{g0} = 0 \\ \phi_n(x) = sin(B_{gn}x) & B_{gn} = \frac{\pi}{2L} + n\frac{\pi}{L} \\ \phi_m(x) = cos(B_{gm}x) & B_{gm} = \frac{\pi}{L} + m\frac{\pi}{L} \end{array} \tag{A.5}$$

The eigenfunctions shown in Eq. A.5 are all orthogonal, but have not been normalised and are not orthonormal. Both n and m are natural numbers. Now, we can consider the perturbed eigenvalue problem and retain only the first order terms to obtain:

$$(\hat{A} + \Delta\hat{A})(\phi_l(x) + \Delta\phi_l(x)) = (\lambda_l + \Delta\lambda_l)(\phi_l(x) + \Delta\phi_l(x)) \tag{A.6}$$

$$(\Delta\hat{A} - \Delta\lambda_l)\phi_l(x) = (\lambda_l - \hat{A})\Delta\phi_l(x) \tag{A.7}$$

Then, we take the inner product with some other eigenfunction n, such that $n \neq l$. The inner product is defined as $\langle f(x), g(x)\rangle = \int_{-L}^{L} f(x)g(x)dx$. We also use the self-adjoint property of the operator $\hat{A}$, which gives that $\langle f(x), \hat{A}g(x)\rangle = \langle \hat{A}f(x), g(x)\rangle$. Thus, we can obtain the projection of the perturbation $\Delta\phi_l$ on the eigenfunction. If we expand the perturbation into all the eigenvalues we obtain:

$$\Delta\phi_l(x) = \sum_{n\neq l} \frac{\langle \phi_n(x), \Delta\hat{A}\phi_l(x)\rangle}{\lambda_l - \lambda_n} \tag{A.8}$$

We are interested in the effect of changing the representation of the cross-sections in the left half of the slab. Thus, the perturbation in the transport operator $\hat{A}$ is non-negative only in the interval $\langle -L; 0\rangle$. In Section 2.4, we have assumed that the only difference between low and higher fidelity cross-section set is in capture. Thus, $\Delta\hat{A} = \Delta(B_m^2) = -\frac{|\Delta\Sigma_a|}{D}$ for $x \in \langle -L; 0\rangle$ and is zero elsewhere. We are interested in the perturbation to the principal eigenfunction, thus, $l = 0$. For the cosine and sine eigenfunctions, we obtain respectively:

$$\int_{-L}^{0} cos(B_{gm}x)dx = 0 \tag{A.9}$$

$$\int_{-L}^{0} sin(B_{gn}x)dx = -\frac{1}{B_{gn}} \tag{A.10}$$

In Eq. A.10 we use the fact that $cos(-B_{gn}L) = 0$ due to boundary condition in Eq. A.2. Using these results, we can obtain the final expression.

$$\begin{array}{l} \phi(x) = \phi_0 - \sum_{n=0}^{\infty} \frac{\Delta(B_m^2)}{B_{gn}^3} sin\left(B_{gn}x\right) \\ \phi(x) = \phi_0 - \sum_{n=0}^{\infty} \frac{8\Delta(B_m^2)L^3}{\pi^3(2n+1)^3} sin\left(\frac{\pi(n+0.5)x}{L}\right) \end{array} \tag{A.11}$$