On integrating Monte Carlo calculations in and around near-critical configurations – I. Methodology

Mariya Brovchenko, William Kenneth Burn, Patrizio Console Camprini

Annals of Nuclear Energy

ON INTEGRATING MONTE-CARLO CALCULATIONS IN AND AROUND NEAR-CRITICAL CONFIGURATIONS – I. METHODOLOGY

--Manuscript Draft--

Manuscript Number:	ANUCENE-D-21-01092
Article Type:	Full Length Article
Keywords:	Monte-Carlo; variance reduction; Direct Statistical Approach; eigenvalue calculation; reactor core
Corresponding Author:	Kenneth William Burn, Ph.D. ENEA Bologna Research Centre: ENEA Centro Ricerche Bologna Bologna, Bologna ITALY
First Author:	Kenneth William Burn, Ph.D.
Order of Authors:	Kenneth William Burn, Ph.D.
	M Brovchenko
	Patrizio Console Camprini, Ph.D.
Abstract:	Montecarlo with variance reduction (VR) is employed to calculate radiation-induced responses in a variety of fixed source problems. In eigenvalue problems analog Montecarlo is used for in-core integral responses (keff, reactivity coefficients, burn-up effects, etc.). More differential in- or ex-core responses are either treated without VR, limiting the range of solvable problems, or by decoupling (all-Montecarlo or deterministic/Montecarlo). We evaluate an all-Montecarlo approach developed to calculate any differential in- or ex-core response. The ex-core problems involve responses around a GEN III PWR core. Comparison is made with an empirical approach involving analog Montecarlo in-core with VR ex-core. The in-core problems study the neutron flux and 96Zr(n,gamma) rate at the surface of a control rod in the VERA Benchmark. Comparison is made with analog Montecarlo. The paper consists of two parts: this part I tests the methodology; part II focuses on comparison with PWR GEN II and III ex-core decoupled results.
Suggested Reviewers:	Herschel Smith smithnucleareng@gmail.com
	Tom Booth montecarlosimulate@gmail.com
	August Winkelman A.J.M.Winkelman@tudelft.nl

To the editors and reviewers:

This paper ("On Integrating Monte-Carlo Calculations in and around Near-Critical Configurations – I. Methodology") is part I of a two-part work. The second part is currently undergoing internal review. If you wish it can be sent in its present form as a draft although the two parts are independent and it is not necessary to have the second part at hand when you read part I.

There is also a support paper entitled "Monte Carlo with Variance Reduction in and around Near-Critical Configurations" which has been published as an ENEA report. This paper (reference 11 in the submitted paper) is cited frequently and we can supply it as an aid to the reviewers.

To the editors:

Some figures in the paper are identical to some figures in the support paper which is public. These figures are: 1, 2, 14, 15, 18, 21, 22, 23, 24 and 25. These figures with one exception describe the geometrical model. Only Fig. 18 contains results. None of the other results in the paper, in figures or in tables are duplicated elsewhere.

The authors

ON INTEGRATING MONTE-CARLO CALCULATIONS IN AND AROUND NEAR-CRITICAL CONFIGURATIONS – I. METHODOLOGY

Brovchenko M[1], Burn K W[2*] and Console Camprini P[2]

[1] IRSN (Institut de Radioprotection et de Sûreté Nucléaire)
31 avenue de la Division Leclerc 92260 Fontenay aux Roses CEDEX France

[2] ENEA (National Agency for New Technology, Energy and Sustainable Economic Development)
Via M.M.Sole, 4, 40129 Bologna, Italy

mariya.brovchenko@irsn.fr, kennethwilliamburn@gmail.com, patrizio.consolecamprini@enea.it

ABSTRACT

Montecarlo with variance reduction (VR) is employed to calculate radiation-induced responses in a variety of fixed source problems. In eigenvalue problems analog Montecarlo is used for in-core integral responses (k_{eff}, reactivity coefficients, burn-up effects, etc.). More differential in- or ex-core responses are either treated without VR, limiting the range of solvable problems, or by decoupling (all-Montecarlo or deterministic/Montecarlo). We evaluate an all-Montecarlo approach developed to calculate any differential in- or ex-core response. The ex-core problems involve responses around a GEN III PWR core. Comparison is made with an empirical approach involving analog Montecarlo in-core with VR ex-core. The in-core problems study the neutron flux and $^{96}Zr(n,\gamma)$ rate at the surface of a control rod in the VERA Benchmark. Comparison is made with analog Montecarlo. The paper consists of two parts: this part I tests the methodology; part II focuses on comparison with PWR GEN II and III ex-core decoupled results.

KEYWORDS: Monte-Carlo, variance reduction, Direct Statistical Approach, eigenvalue calculation, reactor core

* Corresponding author

1. INTRODUCTION

The increase in computing power in recent years exploiting a high degree of parallelism together with developments in (semi-)automatic variance reduction (VR) techniques in Monte-Carlo (MC) (directed mainly at neutral particle transport), have made the calculation of most radiation-induced effects from fixed radiation sources (radioactive decay, accelerator beams, cosmic rays, etc.) feasible. Instead whilst the role of Monte-Carlo has also expanded in eigenvalue calculations (applied mainly to reactor cores), from the evaluation of reactivity coefficients to burn-up and multi-physics effects, this has concerned more integral, in-core responses. More differential responses within the core (for example localized responses) or at some optical distance outside the core fall outside these areas. The former are usually evaluated with analog MC, as (within the context of the source iteration approach) employment of VR might alter, or unbalance, the source for the next fission generation. Analog MC obviously limits the range of responses that can be calculated. The latter are normally evaluated by decoupling the in-core part of the calculation, run in eigenvalue mode, from the ex-core part, run in fixed source mode, with consequent approximation at the point of decoupling, as well as a possible deterioration in the quality of the statistical error estimate.

A Monte-Carlo approach for calculating radiation responses in and around critical configurations under conditions of relatively high attenuation has been developed [1-3], based on the Direct Statistical Approach (DSA). The DSA optimizes the phase space cell VR parameters with respect to a compound response figure-of-merit FOM_c, defined as the inverse of the product of the sum over all the responses-of-interest of the squares of the fractional standard deviations (fsd's) and the computer time [1], pp. 219-220. (Note that the responses-of-interest include both the ex- or in-core responses that we wish to calculate plus components of the fundamental mode – see §2.2.) The methodology employs a modified version of MCNP [4] as vehicle, the modifications written in "patch" form. For in-core problems the approach provides results of a higher quality compared with analog or alternatively allows to calculate responses previously considered unfeasible. For ex-core problems, the approach avoids decoupling and calculates the ex-core responses within the eigenvalue calculation.

An additional feature, superhistories is introduced with the new approach, while two already existing features, multi-response optimization and fictitious source cells, merit further discussion:

- Superhistories [5], reviewed in §2.1, play a key role in maintaining the stability of the fundamental mode [6]. [The more stable the fundamental mode, the higher the FOM_c (which includes components of the fundamental mode) and also the higher the FOM of the ex-core response(s) (whose source is the fundamental mode).] Their various functions in this context are summarized in [7].

- We perform a multi-response optimization [8] to the responses of interest which we name the "local response(s)" together with what we name the "global responses" which we form from the fundamental mode (§2.2).

- "Fictitious source cells" ("FS cells"), already employed to execute "pseudo source biasing" ([9], p. 740), in DSA fixed source problems to compensate for a poorly sampled source, are also useful in eigenvalue problems and are discussed in §2.3.

In §3 the sample problems are illustrated. An ex-core problem is treated in §3.1: a variety of responses outside the core of a PWR GEN III. In §3.2, an in-core problem is treated based on the VERA benchmark [10].

Firstly in §3.1.1 the differential energy neutron flux at a number of positions around the core, in the pressure vessel (PV) well and in the basemat below the well, is considered. The results, when compared with an

empirical approach involving analog VR in the core, are disappointing. There are some indications from this problem as well as from another one treated in [11] (the dose above the lid of a flooded fuel flask), that the technique produces too few source fission neutrons entering the geometry. A work-around is suggested in §3.1.2 that improves the results. Secondly the fast neutron flux on a limited part of the lower PV is considered in §3.1.3. The results with the new technique, relative to the empirical approach, are better than those in §3.1.1 or §3.1.2, but are substantially improved when applying the same work-around in §3.1.4 as that employed in §3.1.2.

Three VERA problems are discussed in §3.2. The first two problems are treated in [11] and the third, higher penetration problem, involving a single pin with response function in a core, is treated in §3.2.1.

After each sample problem set, a summary is made in §3.1.5 and §3.2.2 outlining the main conclusions that can be drawn. In §4 concluding remarks are made.

2. METHODOLOGY

Our approach was to employ as starting point the VR techniques developed for fixed source problems [8]. These are based on population or weight control and not on biasing. We also limited ourselves to the source-iteration (or power iteration) technique for solving the eigenvalue problem [12]. Its solution by MC requires a number of successive renormalizations of the fission source [5,13]. Each normalization involves maintaining approximately the same number of starting fission neutrons, all with the same weight. However such normalization introduces biases [*ibid.*]. We decided, at least as a first step, that these features should be conserved (although an alternative considered was to allow starting fission neutrons of different weights, with the total weight being approximately maintained at each normalization). In this way we remained as close as possible to well-known codes such as MCNP [4], at least at the point in the Markov chain at which the normalization is made.

As is usually done when estimating errors in the standard MC solution of the source-iteration technique, when generating the VR parameters we treated the fission neutrons issuing from such a normalization point as if they were fixed source neutrons, and thus independent of fission neutrons issuing from previous normalizations (and of each other). Instead when running with the resultant VR parameters, various statistical hypotheses were employed.

2.1. Superhistories

One way of mitigating the biases introduced by the normalizations is to delay the point at which the normalization is made. That is, rather than to normalize after each fission generation (the standard procedure), instead to normalize after n fission generations. These longer fission chains are called *superhistories* [5]. Their employment has the additional advantage of reducing the bias on the error estimates, which as mentioned above are usually made between normalization points.

Maintaining the standard procedure at the normalization point means that any desired distortion of the ensemble of fission points to better sample some parts of phase space is immediately negated at the normalization point – population control alters the weight but renormalization forces the fission neutrons issuing from the normalization to have the same weight. Clearly normalizing at every fission generation means that there is very little time or space to move the particles towards important phase space regions (either through population control or through biasing) before the next fission at which the tracks are normalized and any desired distortion of the neutron population is cancelled.

This lead us to introduce superhistories to delay normalization and allow the distorted population to achieve its desired effect. Typically a superhistory of 10 fission generations was employed. (Furthermore VR was allowed to vary according to the fission generation within the superhistory [1].) This looked to work for in-core problems. Instead when the local detector(s) was ex-core, at first superhistories were not used, because population control within the fissile configuration was considered much less important (as only the outer regions of the fissile configuration tend to contribute outside the configuration) [1-3]. Subsequently we realized that superhistories play an important role in holding the fundamental mode steady when VR is employed in the fissile zone [6,7] so that they should actually be employed when the local responses are ex- as well as in-core.

2.2. Local and global responses

We name the response(s) that we are interested in calculating, "local" (because it often occupies a localized volume in phase space). Then we do not optimize only to the local responses and assume that the fundamental mode is held steady by contributing to the responses through future fission generations. This

is mainly because the normalization at the end of the superhistory means that the development of the adjoint flux is limited to the fission generations within a superhistory. Some further words are required on this:

Restricting our discussion to a single local detector (generalizing to more than one detector does not change its validity), the source of the adjoint flux is the response function of the local detector. Within the logic of the Iterated Fission Probability interpretation of the adjoint flux [14], the development of the adjoint flux is considered as its variation with fission generation. We wish to follow the neutrons (in the forward mode) for a sufficient number of generations to generate adjoint information that has passed the transient and arrived at or near the asymptotic distribution. However at the end of a superhistory, renormalization occurs and the weights of the fission neutrons are reset. (Here we are considering local detectors that require VR and therefore a large weight variation - this is especially true for in-core detectors. Furthermore the structure of our patch is built around superhistories as independent source particle events. We could conceivably pass the adjoint information through the renormalization event and into the next superhistory, but it would be "challenging".) Therefore if we have 1 fission generation per superhistory, clearly large parts of the core will have zero adjoint flux. 10 fission generations will be better but likely insufficient.

We solve this problem by dividing the fissile configuration into a number of segments (typically 10 – 100) and tallying the fission source in each of these segments. We call these the "global" responses. In this way we add the fundamental mode distribution to the responses-of-interest. Our objective is thus to hold the fundamental mode as steady as possible, compatibly with the necessary distortion of the population so as to improve the calculation of the local response(s). To achieve this goal we optimize the VR parameters to all the local and global responses simultaneously, treating each normalization step as a fixed source in an analogous fashion to what was done in [8]. This way we ensure that the fundamental mode is held reasonably steady even with relatively few fission generations per superhistory.

2.3. Fictitious source cells

We consider the following situation: fixed source sub-critical mode with weight-independent splitting/Russian roulette (RR) ([15] and citations therein) and the normal DSA splitting/RR rules [16] in operation (which are the same rules as those employed in MCNP [4] or other codes). Then for any set of phase-space cell VR parameters, neutrons arriving in some phase space cell, born in the same source cell have the same weight independent of the path from the source cell to the cell in question [15,16]. This does not obtain for neutrons born in different source cells (and this is also the case when the DSA is applied to the source iteration solution of eigenvalue problems). These points are illustrated in [11].

In certain fixed source problems, such as with γ activation sources, we found that employing fictitious source cells could partially compensate for a lack of, or an inappropriate, source biasing. FS cells work by executing splitting/RR on source particles immediately they are born and before they enter the real-world source cells [11] (similar to when source particles are born with a weight outside the weight window [4]). Although this technique is not as efficient as well-executed source biasing – if the particle is split at its point of birth, correlations are introduced which are not present with source biasing – it is quite robust.

Instead in eigenvalue problems, the source is usually distributed over a number of spatial cells and with or without superhistories a neutron born from fission at a normalization point must have the same weight, wherever it is born. Furthermore in contrast to some fixed source problems, the energy spectrum does not greatly differ according to the cell. FS cells then provide a straightforward means of implementing a kind of biasing, through sampling, for the starting fission neutron position and energy.

Also in [11] it is shown that introducing FS cells into fixed source or eigenvalue problems allows us, with a judicious choice of VR parameters in the FS cells, to achieve the desirable situation of all particles in a particular phase space cell having the same weight, independent of where they were born and how they

arrived in the cell (either via a renormalization or within the same superhistory). It is easy to see that such choice involves any set of importances with the same values in all the FS cells [11].

Among all possible solutions, those mentioned above ensure that neutrons in a particular phase space cell all have the same weight. However because of the further element discussed above of position and energy sampling of the source, the optimum solution is not necessarily one of these sets as indeed many sample problems have shown.

Introducing superhistories means that there are less source events and therefore less improved sampling of the source with FS cells. Thus using FS cells with superhistories may not give the same gain as without, as we are looking at improving the source sampling once every 10 fission generations rather than every generation.

[It should be mentioned that we are dealing with "weight-independent splitting/RR". If the splitting/RR depends on the weight (such as in [17] or the well-known "weight window" [18]), then splitting or RR is implicitly executed immediately a particle is born.]

3. SAMPLE PROBLEMS

As previously mentioned, the DSA optimizes the phase space cell importances with respect to a compound response figure-of-merit FOM_c. Here we modify slightly the definition of FOM_c by multiplying the previously-defined value by the number of responses. This is implicit in all that follows. (Instead when we only consider a single response, $M = 1$ and we employ the normal nomenclature, FOM.) Thus

$$FOM_c = M.[\sum_{i=1}^{M}(fsd_i)^2 \, . \, T]^{-1} \quad (1)$$

where M is the number of responses, fsd_i is the fractional standard deviation of response i, and T is the CPU time.

In [11] we looked at various ways of estimating the statistical error fsd_i: between superhistories, between independent runs, or between steps containing a number of superhistories of a single calculation. We saw that an adequate number of fission generations (up to some thousands) is required to encompass the statistical variation in the fundamental mode, especially when playing VR in the fissile zone. To compare a sufficient number of independent calculations (> 30–40) each with a sufficient number of fission generations was impractical due to time and machine constraints. Thus we opted to calculate statistics between groups, or steps, of superhistories in sequence within the same long calculation. Each step starts from the fission source written at the end of the previous step and each step is assumed independent. Correlations between these steps were a concern, especially when the number of fission generations within each step was small. We verified possible correlations by calculating statistics of the most important responses between groups of steps, as discussed in [11].

Then following [11] we compare the calculation of the responses employing DSA-generated VR parameters ex- and in-core (with superhistories of typically 12 fission generations and FS cells), with an empirical approach with the same DSA VR parameters ex-core, analog VR in-core and superhistories of 1 fission generation (naming these approaches the "DSA case" and the "analog case", respectively).

An identical total number of fission generations was run in each case, such number being sufficiently large to hopefully encompass the variation in the fundamental mode. However the number of fission generations in the phase of generating the VR parameters was substantially lower than the number employed in the final runs with such parameters.

In general and as discussed in [11] the indicated optimum VR parameters in-core in the DSA case tended to involve more RR than splitting. Although FOM's and FOM_c's, which should be independent of the calculation time, are going to be compared, it was desired to expend roughly the same calculational effort in each case. Given that the same ex-core VR parameters were used, this implies the same in-core effort. Thus $NSRCK$, the number of starting fission neutrons in each normalization step or superhistory, was set lower in the analog case.

We assume that we have a sufficiently large set of fission sites representing the fundamental mode of each problem that has been generated by a separate analog calculation and is the starting point of the calculations.

For all sample problems a selection of the results is shown here. Further results are given in [11] as are also some details of how the optimum importances were generated. Also in [11] are shown the cumulative moving averages (CMA's) of selected responses. As well as giving some qualitative idea of possible correlations between steps, they give a good idea of correlations between responses in the part of the core nearest the ex-core response and the ex-core response itself. In fact we found in general high correlations of the latter kind implying that more calculational effort than we had anticipated should be expended in-core compared with ex-core.

3.1 Ex-Fissile Configuration Problem: PWR GEN III with Thick Steel Reflector

We now examine the PWR GEN III model which was investigated in [6]. It has the characteristics of a large core with a thick steel reflector which raises the higher eigenmodes [19]. A consequence is the greater difficulty a Monte-Carlo simulation has of reaching and maintaining the fundamental mode [6].

Equilibrium state UOX assembly-wise fuel compositions, distinguishing the rods with and without gadolinium, were employed with no axial variation. Note the fuel compositions were different from those employed in part II where the results of decoupled and single eigenvalue calculations are compared. It may also be mentioned that here no gamma's are transported, unlike the calculation in part II or in [6].

We will consider a variety of local ex-core responses. Each collection of responses produced its own set of VR parameters, with a differing amount of splitting/RR in the core. Thus the difference in *NSRCK* between the DSA case and the analog case varied for each set of responses.

The fissile zone was divided into 15 segments for tallying the global responses which in this case were fission heating rates: 5 axial (of roughly equal heights) and 3 radial (the outer assembly ring, the next two assembly rings and the inner assemblies). We employed 7 energy groups with upper limits: 1 eV, 1 keV, 20 keV, 200 keV, 2 MeV, 5 MeV and 20 MeV.

3.1.1 Ex-core mapping of the neutron flux with 48 ex-core responses

The vertical section in Fig. 1 shows the surfaces on which the neutron fluxes were tallied. All tallies are on 360°. There are three zones for the flux tallying:

- Radially in the PV well with 3 axial segments and 6 energy groups (upper limits: 0.2 eV, 20 eV, 2 keV, 200 keV, 2 MeV and 20 MeV);

- On the lower surface of the PV with 2 segments and 3 energy groups (upper limits: 10 eV, 500 keV and 20 MeV);

- In the sacrificial concrete and zirconium oxide basemat at 4 depths, with 2 segments and 3 energy groups (upper limits: 10 eV, 500 keV and 20 MeV).

We divide the 48 ex-core responses into three groups, A (radial in the PV well), B (bottom of the PV well) and C (within the basemat) as illustrated in Fig. 1. A horizontal section at the core mid-plane in Fig, 2 shows the radial tally surface, A.

Optimum importances generated with the DSA were run with the two cases for 40 steps. In the DSA case each step had 1.2×10^6 starting neutrons (*NSRCK*) and 12 superhistories *(KCT*) and a superhistory length of 12 fission generations.

It is well known that the outer assembly ring is responsible for most of the ex-core response. We therefore expect RR between the fictitious and real-world source cells in the inner and 2nd and 3rd outer rings. In fact the analog case required around 30% of the starting neutrons of the DSA case for the same computer time. Thus in the analog case, *NSRCK* was set as 3.75×10^5 and *KCT* as 144. The total number of fission generations in each case is $144\times40=5760$.

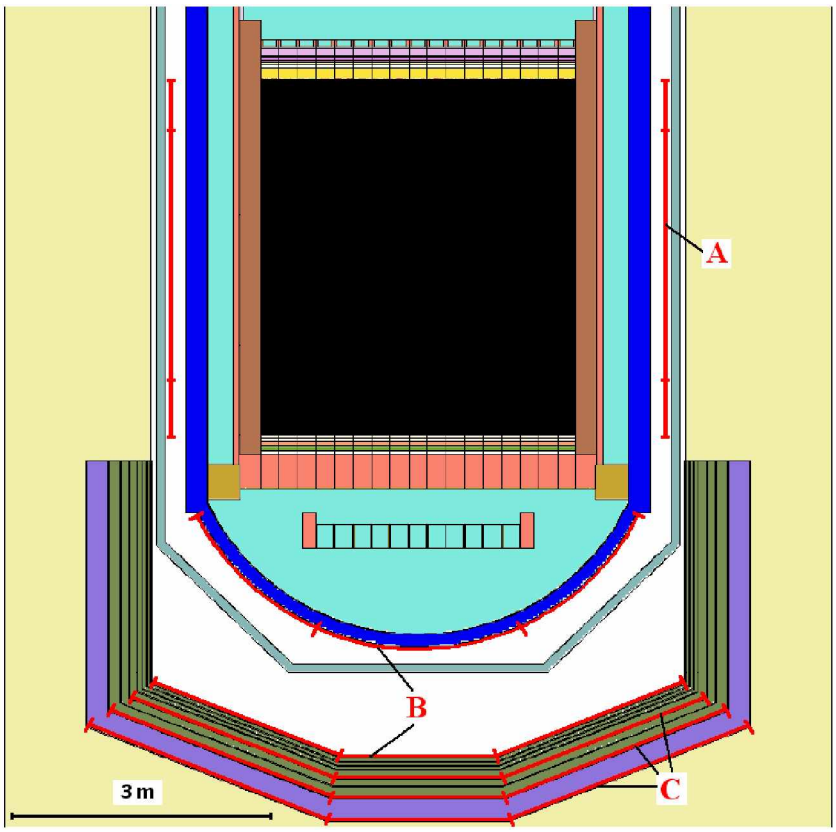

Figure 1. PWR Gen III: ex-core mapping of the neutron flux: vertical section showing surfaces for neutron flux tallying and the three groups of ex-core tallies, A, B and C

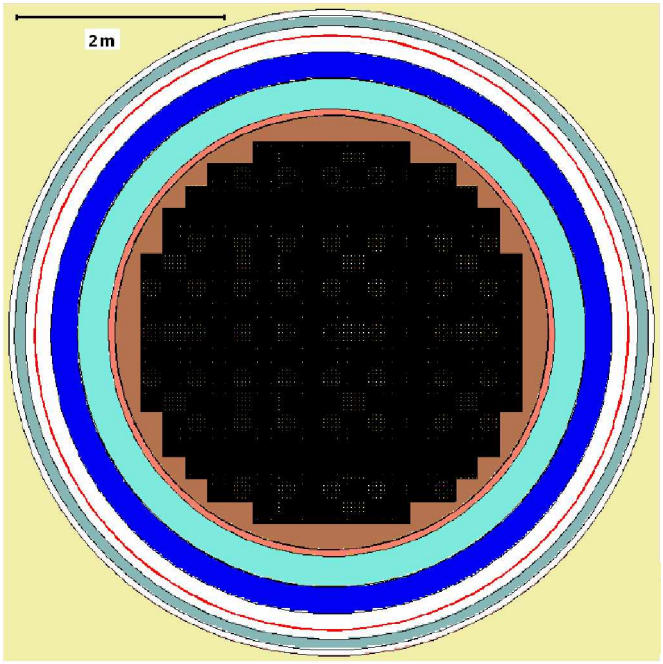

Figure 2. PWR Gen III: ex-core mapping of the neutron flux: horizontal section at core mid-plane showing surface for neutron flux tallying

The FOM_c's of the two cases ("DSA", "analog") and their ratio (DSA/analog) for the ex-core response groups A, B and C are shown in Figs. 3, 4 and 5, respectively. We see from Figs. 3-5 that the analog case is always better than the DSA case. Thus the variation in the fundamental mode is more important than the variation in the ex-core transport, even for the responses in the basemat. The basemat responses (group C) show the least difference between the DSA case and analog case. This is as expected because these responses are at the greatest ex-core penetration. For completeness we see in Fig. 6 the FOM_c's and their ratio for the 15 in-core fission responses. As expected, the ratio is lower than for any of the ex-core response groups.

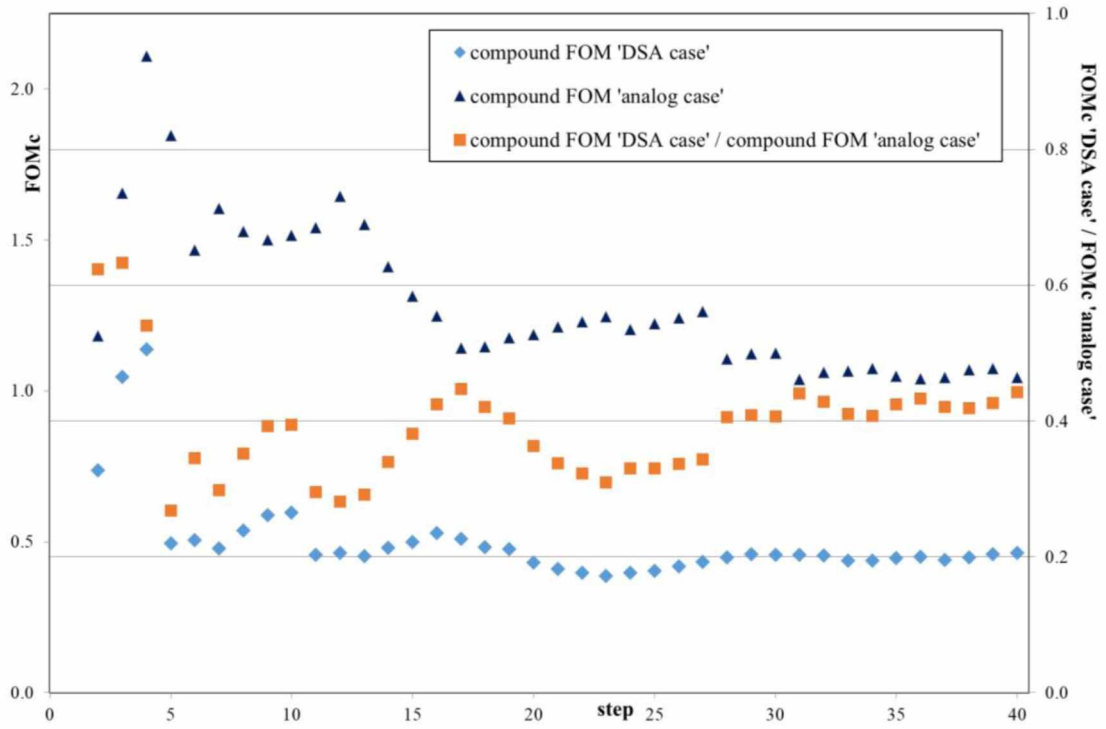

Figure 3. PWR Gen III; ex-core mapping of the neutron flux: ex-core local responses group A, FOM_c's of the two approaches, "DSA" and "analog", and their ratio

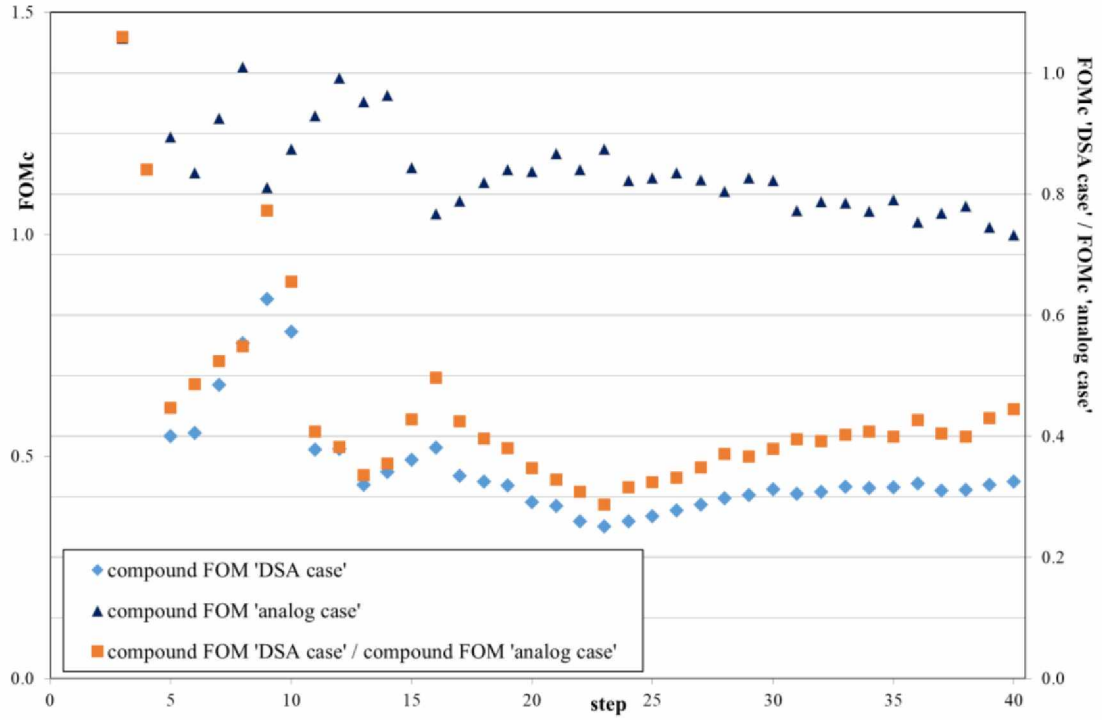

Figure 4. PWR Gen III; ex-core mapping of the neutron flux: ex-core local responses group B, *FOM$_c$'s* of the two approaches, "DSA" and "analog", and their ratio

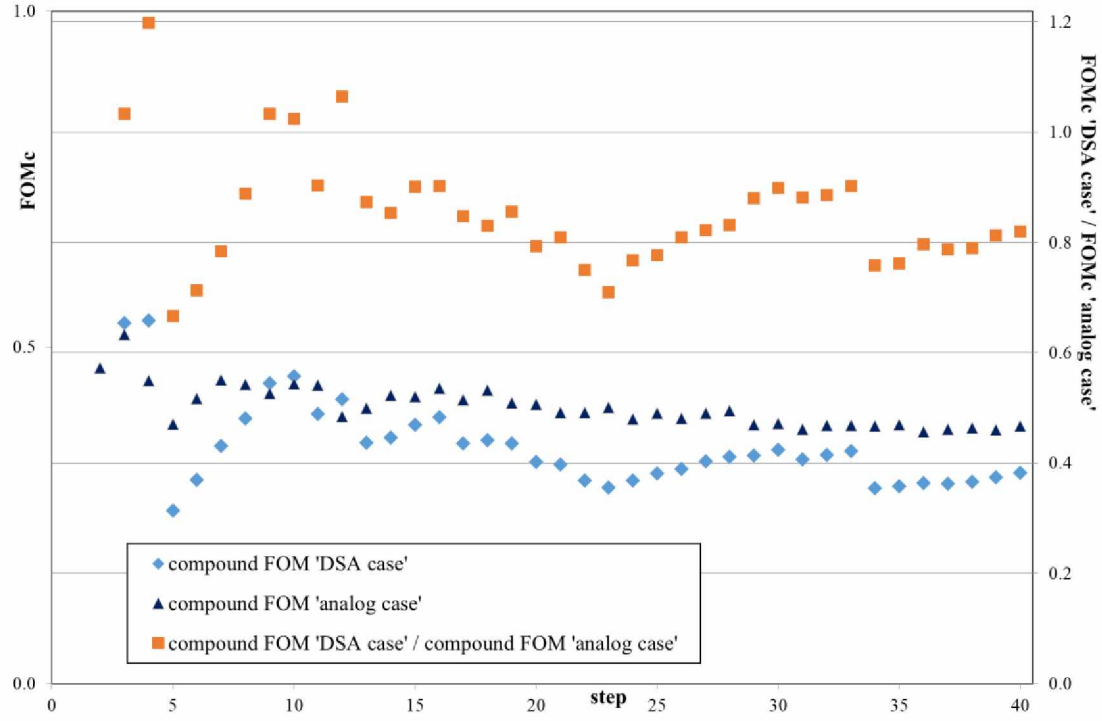

Figure 5. PWR Gen III; ex-core mapping of the neutron flux: ex-core local responses group C, *FOM$_c$'s* of the two approaches, "DSA" and "analog", and their ratio

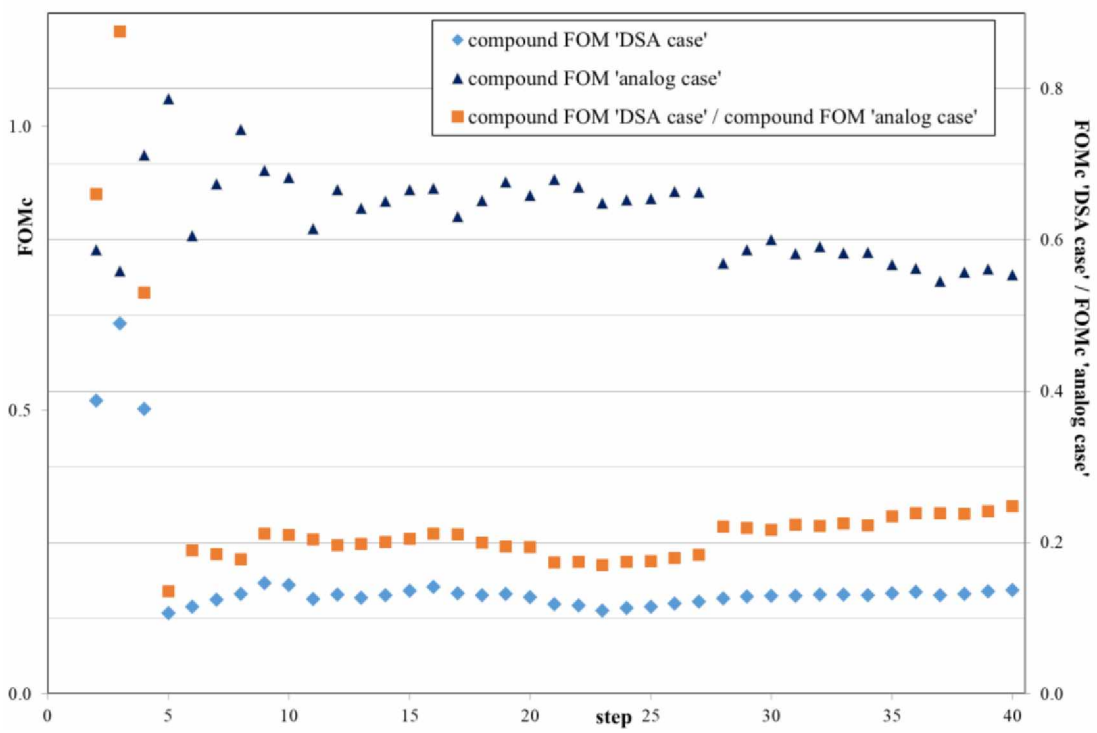

Figure 6. PWR Gen III; ex-core mapping of the neutron flux: *FOM$_c$'s* of the 15 in-core global responses for the two approaches, "DSA" and "analog", and their ratio

3.1.2 Revisiting §3.1.1: increasing the number of fission neutrons passing from the fictitious to the real-world source cells

The results in §3.1.1 were disappointing. Because most of the ex-core response comes from fission neutrons born in the outer assemblies, we had expected that by Russian rouletting fission neutrons in the inner part of the core we would save time and obtain better *FOM$_c$'s* compared with the empirical analog approach. Instead we had indications both from this PWR GEN III problem and from the flooded fuel flask problem in [11] that in the DSA case there are too few fission neutrons entering the real-world source cells from the fictitious source cells. In many of the calculations we saw in [11] a high correlation between the CMA of the fission rate in the part of the fissile configuration that contributes to the ex-core response and that of the ex-core response itself. Thus too little effort seemed to have been spent in-core and too much ex-core. An undersampling of the fissile zone will be more to the detriment of the DSA case than the analog case because the VR performed in the fissile zone in the former case can unbalance the population.

A possible cause of the above imbalance was suspected to be that the DSA was originally developed for fixed source problems [8]. Converting to eigenvalue problems with solution based on source-iteration involved neglecting the fact that the source at each iteration after the first comes from the fission events of the previous iteration. Superhistories only delay or decrease this omission. They do not remove it. To try to acquire a better understanding of how to compensate for this we need to look at how the DSA works.

The DSA operates by generating coefficients of the 2nd moment and (CPU) time functions. If splitting/RR is independent of the weight of the incoming particle, such functions are separable [1,15,16]. If not, an approximation is applied to force the functions to be separable [17]. Once separable into a coefficient part and a part exhibiting a dependence on the splitting/RR parameters, a product called q is made of the 2nd

moment and time functions which is minimized in a standard fashion to generate the optimum splitting/RR parameters. The inverse of this product is proportional to *FOM*.

The 2nd moment function consists of three kinds of terms named "gamma", "delta" and "branching" [1,15-17]. Contributions to the gamma terms occur at every tally score, to the delta terms from correlations at every splitting event at a surface (in space, energy or between fictitious and real-world source cells) and to the branching terms from correlations at every branching (be it natural or due to some other VR technique) that is not a surface splitting event. (In our problem with superhistories, the main branching contribution comes from intra-superhistory fissions.)

For a fixed number of starting neutrons (in the fictitious source cells), splitting decreases the 2nd moment and increases the time whilst Russian roulette increases the 2nd moment and decreases the time. We are however interested in the effect of splitting and RR on q for a fixed amount of (CPU) time.

If we consider the passage from the fictitious to the real-world source cells, we have seen in the test problems here and in [11] that RR is played in most of the fissile zone. Splitting is only played in the very few cells that contribute to the ex-core response. In fixed source rather than eigenvalue problems, splitting source particles is not particularly useful – the 2nd moment could better be reduced by starting a new source particle. At the same time, Russian rouletting source particles is not particularly damaging as we have not wasted much time and we can just start straight away a new source particle that may better reduce the 2nd moment than the particle we have rouletted. So with a fixed source it is not surprising if RR predominates over splitting between the fictitious and the real-world source cells.

We now quantify these ideas by returning to the problem of §3.1.1. Starting from the set of optimum importances that generated the results in Figs. 3-6, we varied the importances in all the fictitious source cells (15 spatial × 7 energy groups) by a constant multiplier that ranged from 1/256 to 256 (with a stride of a factor of 2 making 17 sets of importances). Employing the DSA 2nd moment and time functions, in Fig. 7 we plot the 2nd moment and time values against each set of importances. Note that it is the compound 2nd moment function from 15 global and 48 local responses, normalized by the number of responses, $M=63$.

In Fig. 7 we see that as the multiplier increases, more RR is executed between the fictitious and real-world source cells, the 2nd moment increases and the CPU time decreases. *Vice versa* as the multiplier decreases, more splitting is executed between the fictitious and the real-world source cells, the 2nd moment decreases and the time increases. We see that the delta terms give the major contribution to the 2nd moment, followed by the branching terms. We also note the reduction in gradient in the delta terms at low multiplier values. This is due to correlations between split particles entering real-world source cells and raising the 2nd moment. Finally, present although difficult to see, is a reduction in gradient in the time at high multiplier values, due to time spent resampling source particles after being rouletted between the fictitious and real-world source cells.

In Fig. 8 the inverse of the quality function (equivalent to FOM_c) is plotted against the multiplier. (The optimum values of the fictitious source cell importances are with a multiplier of 1.) We see the degradation at low and high multiplier values due to the above-mentioned effects of correlations and time "wasted" in rouletting source particles respectively. However the main feature of Fig. 8 which we wish to exploit is the wide plateau (with a multiplier between ~0.1 and ~50) either side of the optimum.

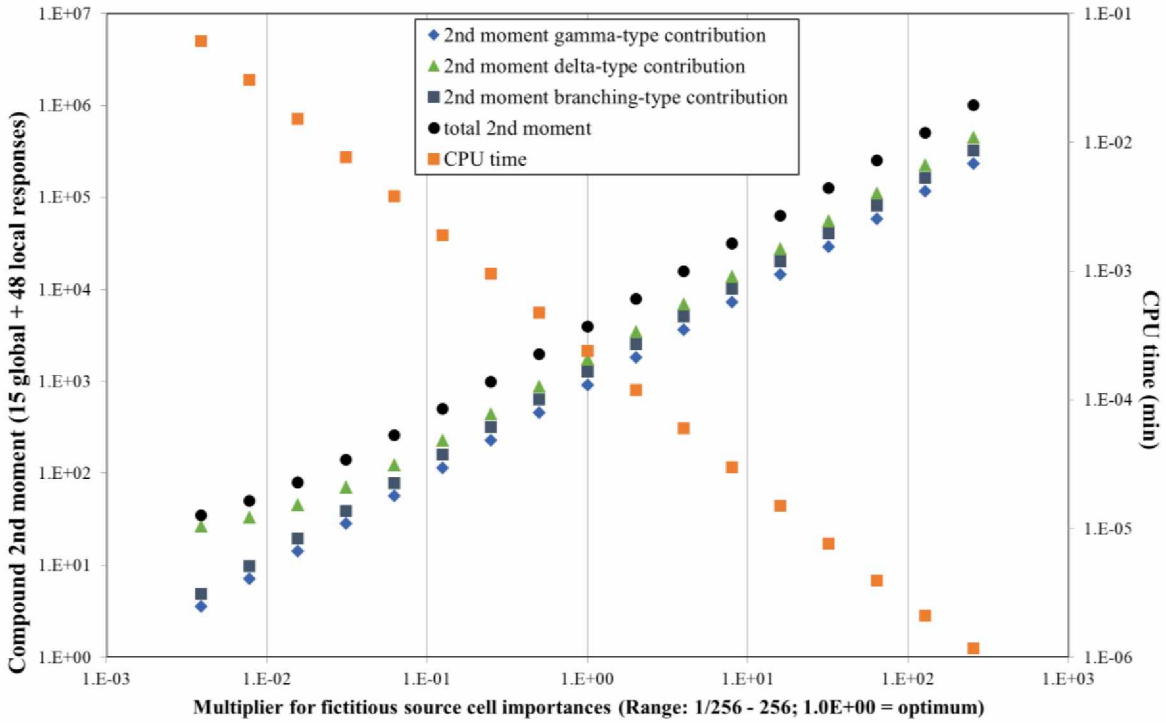

Figure 7. PWR Gen III; ex-core mapping of the neutron flux (increasing no. fission neutrons passing from fictitious to real-world source cells): 2nd moment and time against multiplier of the fictitious source cell importances

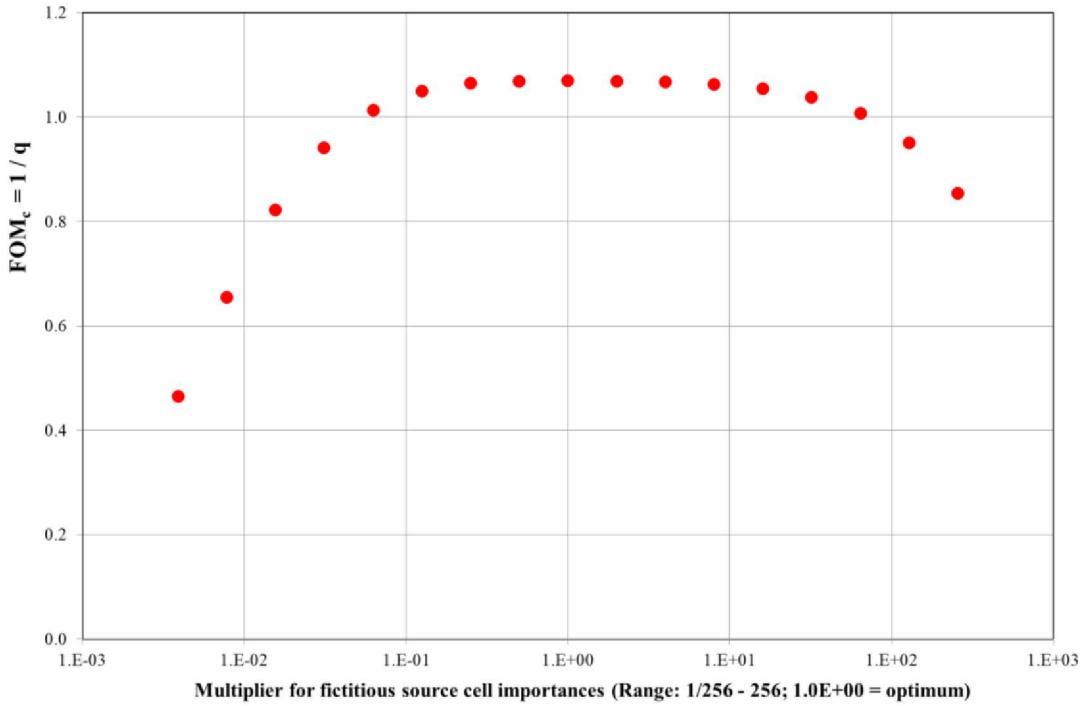

Figure 8. PWR Gen III; ex-core mapping of the neutron flux (increasing no. fission neutrons passing from fictitious to real-world source cells): inverse of the quality (2nd moment × time) against multiplier of the fictitious source cell importances

From Fig. 8 we adopt the multiplication factor 0.125=1/8 to apply to the optimum importances of the FS cells. This should raise the number of neutrons entering the real-world source cells by a factor of 8. We should also like to further raise the number of starting neutrons by raising *NSRCK*. On our system the maximum *NSRCK* is around 9×10^6. However employing this number would violate the constraints listed at the beginning of §3 for the comparison with the analog case as the latter would require a much higher *NSRCK*. Instead in the DSA case we employed a *NSRCK* value of 3.5×10^6 with 3 superhistories (and as usual 12 fission generations per superhistory). Then the analog case with a *NSRCK* value of 9×10^6 and 36 superhistories satisfied our criterion of expending a similar CPU time for the same number of fission generations in the two cases. We ran 60 steps in each case under the usual conditions.

The FOM_c*'s* and their ratio of the three response groups A, B and C are shown in Figs. 9, 10 and 11. Fig. 12 shows the FOM_c*'s* and their ratio for the 15 in-core fission responses.

Comparing Figs. 9-12 with Figs. 3-6 we see that we have increased the FOM_c ratio: DSA case / analog case by around a factor of 2. In particular, the new ratios after 60 steps for the in-core responses and the ex-core responses A, B and C are respectively: 0.44, 0.81, 0.86 and 1.92. Instead the old ratios after 40 steps in §3.1.1 were respectively: 0.25, 0.44, 0.45 and 0.82. Thus we see that now the deeper penetration ex-core responses in group C show a gain with the DSA whilst the ex-core responses in groups A and B still have a slightly higher quality with the analog case.

We have seen here that we should make the effective size of the superhistory in the DSA case as large as possible. We have done this in two ways: by trying to make *NSRCK* as large as possible and by reducing the importance of the fictitious source cells by as much as we can, constrained by the curve in Fig. 8.

We have actually not been able to make *NSRCK* as large as possible because of the constraint of our criterion of the fairest comparison with the analog case. So instead we ran a DSA case with a *NSRCK* value of 9×10^6. (According to our criterion the equivalent value of *NSRCK* in the analog case would then be 23.143×10^6.) We ran this case under the same conditions as above; 12 fission generations per superhistory, the multiplication factor 1/8 for the importances of the FS cells and 60 steps, but now with each step consisting of 1 superhistory. We took the FOM_c*'s* after 60 steps and compared with the results for FOM_c after 60 steps in Figs. 9-12 and after 40 steps in Figs. 3-6. In Fig. 13 we plot these values against the *NSRCK* value of the analog case, with the above estimated value of 23.143×10^6 where the analog case was not run.

We may make two observations from Fig. 13. Firstly between the two *NSRCK* values of 0.375×10^6 and 9×10^6, the DSA case has an appreciably higher gradient then the analog case for all four response groups. This means that it is more important to have a larger size superhistory when VR is played in the fissile zone compared with analog VR. Secondly the DSA case FOM_c values at the highest (equivalent) *NSRCK* value of 23.143×10^6 are appreciably greater than those at 9×10^6.

We thus conclude that the quality of the calculation with VR in the fissile zone can be increased by making *NSRCK* as large as is feasible. We also conclude that we could probably reduce the importances of the FS cells by a greater factor than that of 8 indicated in Fig. 8 by the fixed source DSA model.

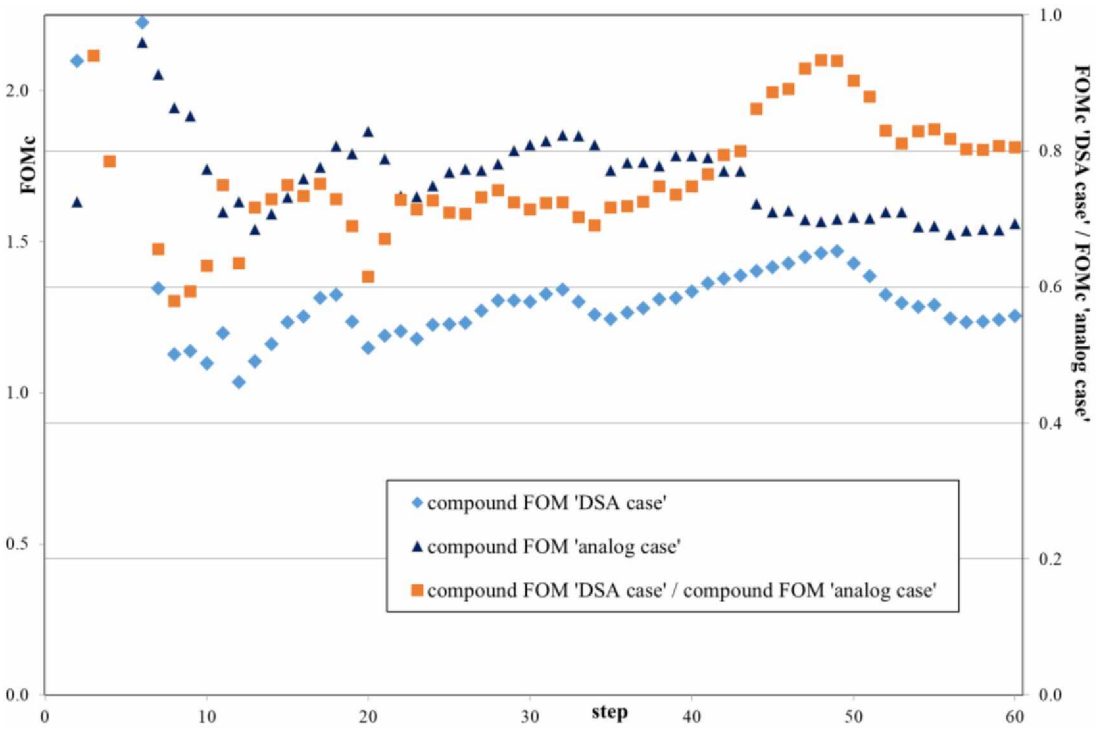

Figure 9. PWR Gen III; ex-core mapping of the neutron flux (increasing no. fission neutrons passing from fictitious to real-world source cells): ex-core local responses group A, *FOM$_c$'s* of the two approaches, "DSA" and "analog", and their ratio

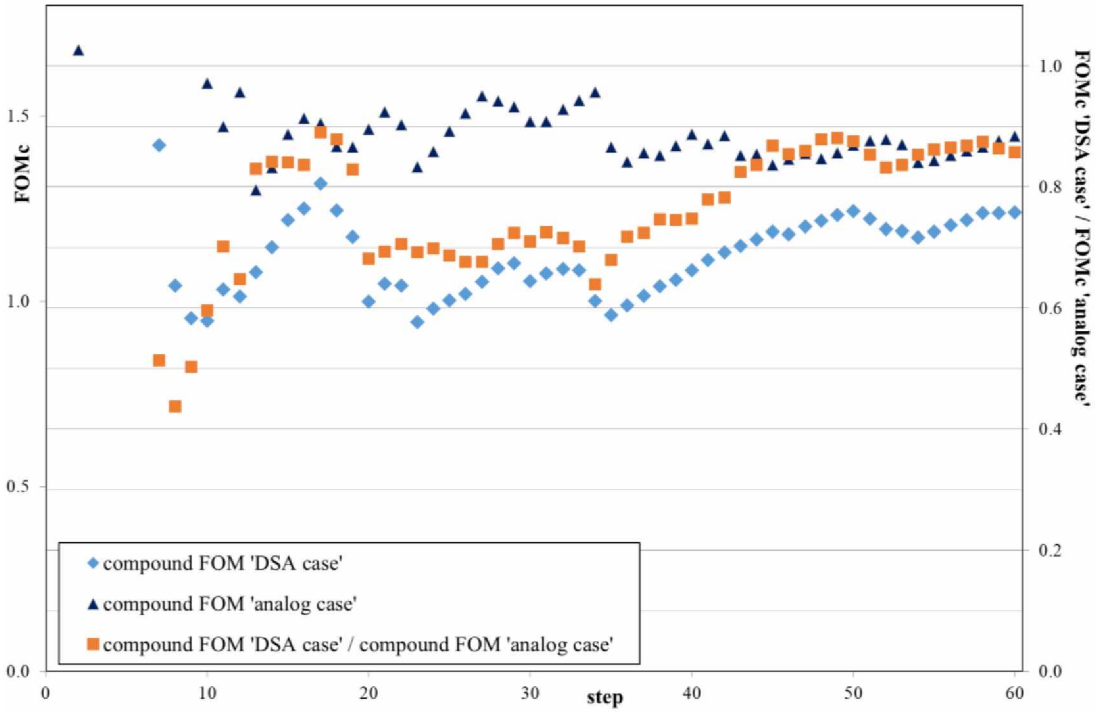

Figure 10. PWR Gen III; ex-core mapping of the neutron flux (increasing no. fission neutrons passing from fictitious to real-world source cells): ex-core local responses group B, *FOM$_c$'s* of the two approaches, "DSA" and "analog", and their ratio

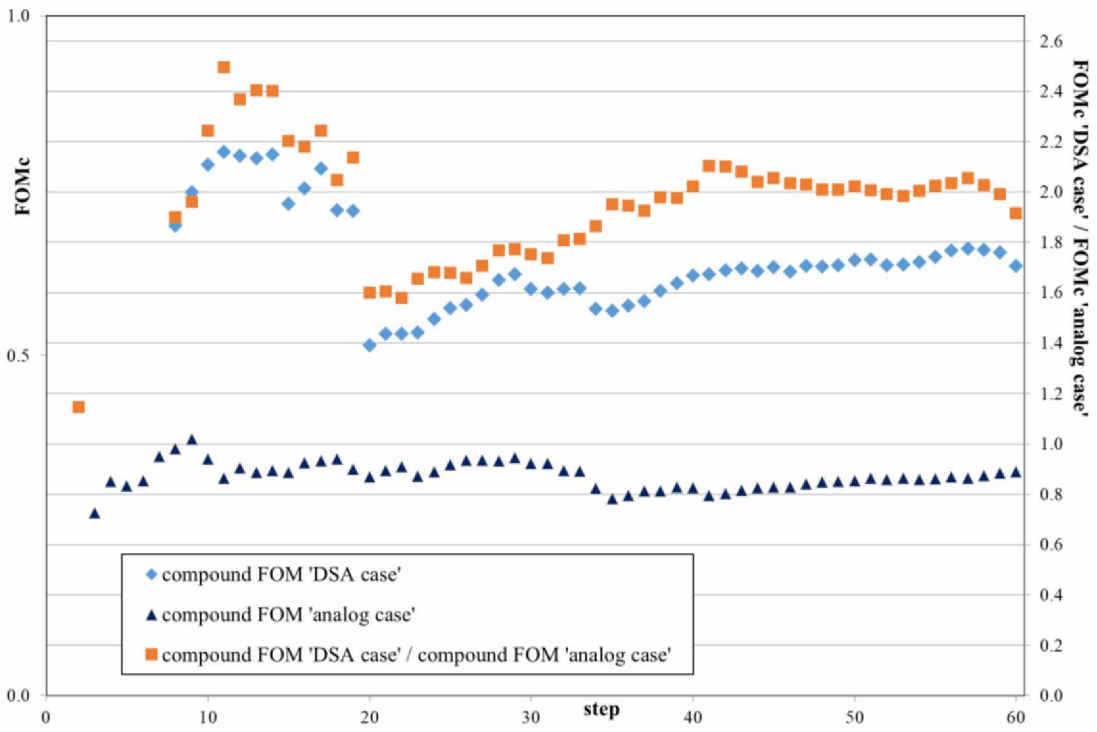

Figure 11. PWR Gen III; ex-core mapping of the neutron flux (increasing no. fission neutrons passing from fictitious to real-world source cells): ex-core local responses group C, *FOM$_c$'s* of the two approaches, "DSA" and "analog", and their ratio

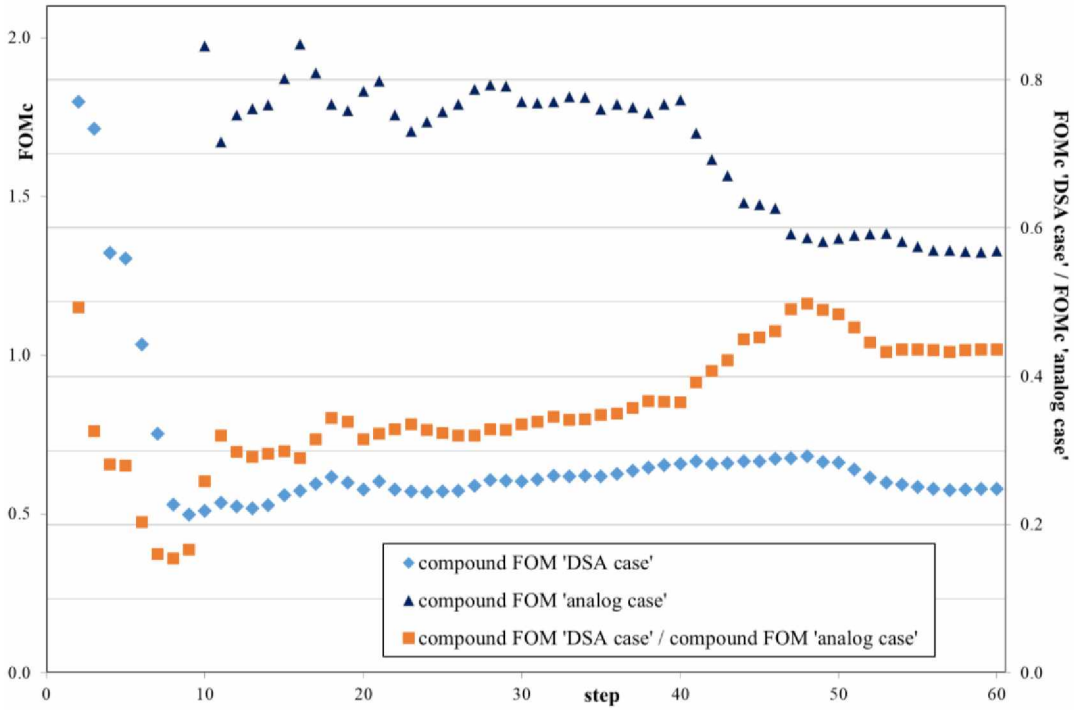

Figure 12. PWR Gen III; ex-core mapping of the neutron flux (increasing no. fission neutrons passing from fictitious to real-world source cells): *FOM$_c$'s* of the 15 in-core global responses for the two approaches, "DSA" and "analog", and their ratio

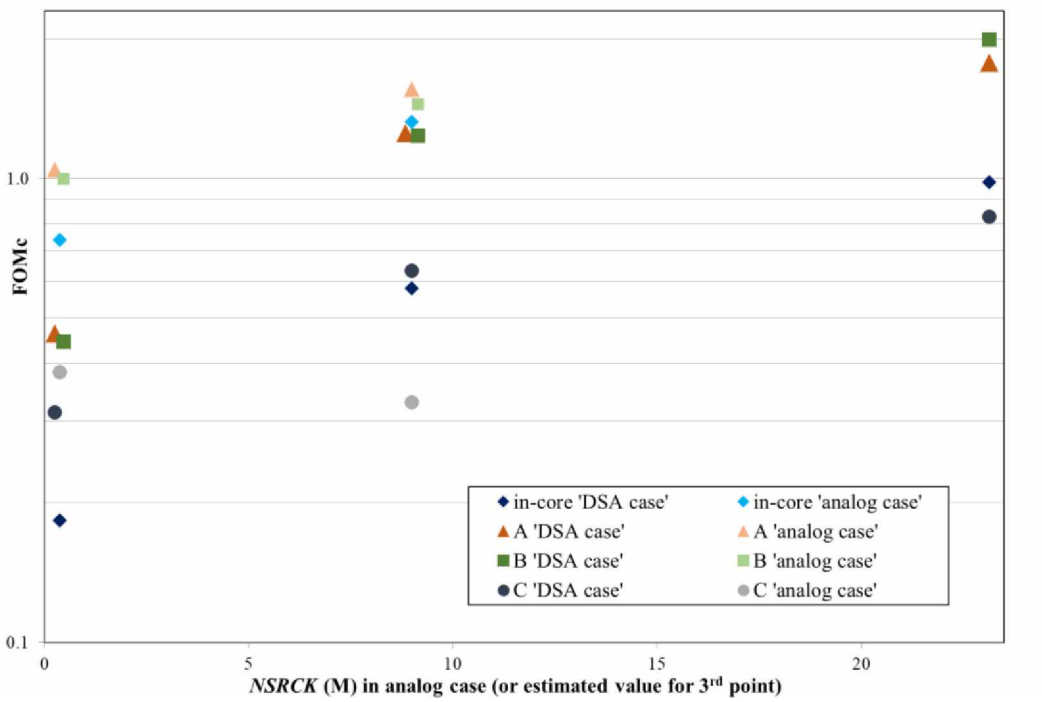

Figure 13. PWR Gen III; ex-core mapping of the neutron flux: final *FOM$_c$'s* of the two approaches, "DSA" and "analog", for the 15 in-core global responses and the three groups of local ex-core responses for the results in Figs. 3-6 and 9-12 and in this section with a *NSRCK* value of 9×10^6 in the DSA approach

3.1.3 Single ex-core response: neutron flux > 1 MeV on PV

The second problem has a single ex-core response: the neutron flux > 1 MeV on the inner surface of the PV. In [11] this detector was placed at three different positions of increasing ex-core penetration: at the level of the bottom of the core on 360°, just below the supporting platforms on 360° and just below the supporting platforms limited to a range of ±2.4° around 0° (directly under a supporting platform). (From [11] for this response at different positions on the PV, the ratio of the *FOM's* DSA/analog looked to be increasing as the ex-core penetration increased.) Here we consider the 3rd position.

A vertical section showing this position is in Fig. 14 and a horizontal section is in Fig. 15. To take into account the limited azimuthal range of the ex-core detector, we needed to add extra spatial cells for VR around 0° from the core out to the PV including in particular the lower plenum, bottom plate and the water between the bottom plate and the distribution plate (see Fig. 14). Also two extra spatial cells for VR were added to the fissile zone, making a total of 17.

Compared with the previous problem, the greater penetration due both to the limited energy range of the detector as well as to the breakdown in azimuthal symmetry, produced a much larger variation in the ratio of importances between the fictitious source cells and their respective real-world cells. This yielded a much smaller fission source size in the analog case to engender a similar computer time to the DSA case. The size of the starting superhistories in the DSA case was therefore increased to 9×10^6. The fission generation length of each superhistory was maintained at 12. To achieve a similar calculational effort, the size in the analog case was then 2.6×10^5. In both cases, each step consisted of 96 fission generations and the total number of steps was 60.

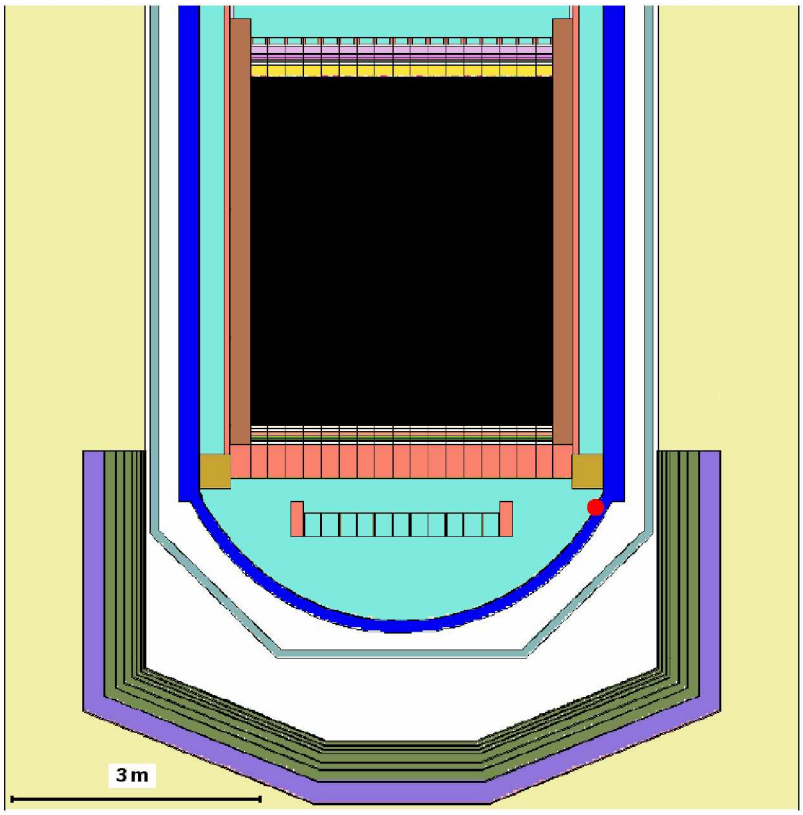

Figure 14. PWR Gen III; fast neutron flux on PV: tally location below supporting platform

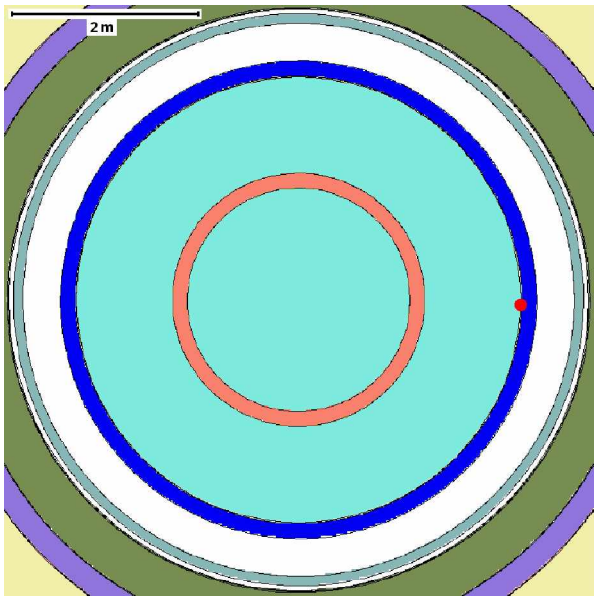

Figure 15. PWR Gen III; fast neutron flux on PV: tally location at 0° below supporting platforms (horizontal section)

The *FOM's* of the two cases ("DSA", "analog") and their ratio for the ex-core response are shown in Fig. 16. After 60 steps (5760 fission generations) the ratio "DSA"/"analog" is around 4.1. Although much better

than the first detector position, and slightly better than the second ([11]), the ratio is not particularly stable and might decrease with further steps. The *FOM$_c$'s* and their ratio for the 17 global responses are shown in Fig. 17. After 60 steps the ratio was 0.54, similar to values in [11] for the first two detector positions.

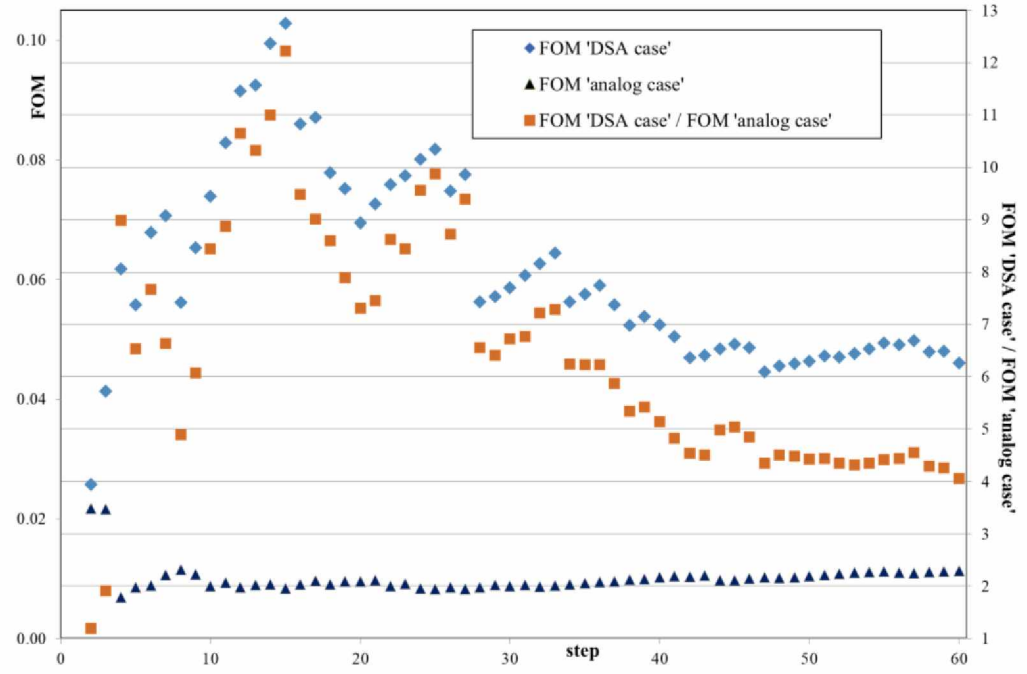

Figure 16. PWR Gen III; fast neutron flux on PV at 0° below supporting platforms: *FOM's* of the ex-core response for the two approaches, "DSA" and "analog", and their ratio

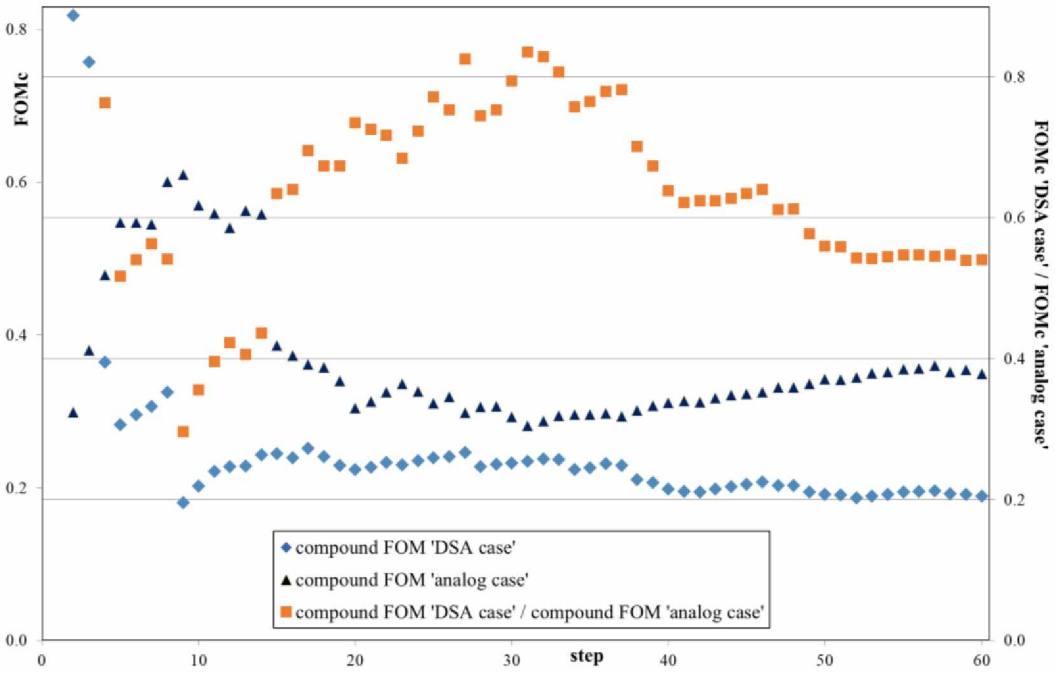

Figure 17. PWR Gen III; fast neutron flux on PV at 0° below supporting platforms: *FOM$_c$'s* of the 17 in-core global responses for the two approaches, "DSA" and "analog", and their ratio

3.1.4 Revisiting §3.1.3: increasing the number of fission neutrons passing from the fictitious to the real-world source cells

We now apply the work-around in §3.1.2 to the problem of §3.1.3. Although the improvement of the *FOM$_c$'s* in the DSA case relative to the analog case in §3.1.2 is limited, there looks to be the possibility of greater gains with the problem of §3.1.3 given the fact that it is a much smaller part of the fissile zone that contributes to the ex-core detector. Following §3.1.2 we varied the importances in all the fictitious source cells (17 spatial × 7 energy groups) by a constant multiplier. In Fig. 18 we plot the inverse of the quality function against the multiplier and we see a similar curve to that shown in Fig. 8. (In [11] there is shown a figure equivalent to Fig. 7 containing the various 2nd moment and time contributions for this problem.)

Given the fact that Figs. 8 and 18 are similar, we adopt the same multiplication factor 1/8 for the importances of the FS cells as in §3.1.2. Now because of the greater difference of *NSRCK* between the DSA case and the analog case in §3.1.3 (factor 35) compared with §3.1.1 (factor 3.2), we were able to maintain *NSRCK* in the DSA case as 9×10^6. The fission generation length of each superhistory was kept at 12. To achieve a similar calculational effort, the size in the analog case was then made 8 times the value of §3.1.3, or 2.08×10^6. In both cases, each step consisted of just 12 fission generations and the total number of steps was now 80.

The *FOM's* of the two cases ("DSA", "analog") and their ratio for the ex-core response are shown in Fig. 19. We firstly note in Fig. 19 that the *FOM* scale is logarithmic. We see that after 80 steps (960 fission generations) the ratio "DSA" / "analog" is around 15. Again the ratio is not particularly stable (although its instability looks more due to that of the analog case rather than the DSA case), but still looks in the range 10–15.

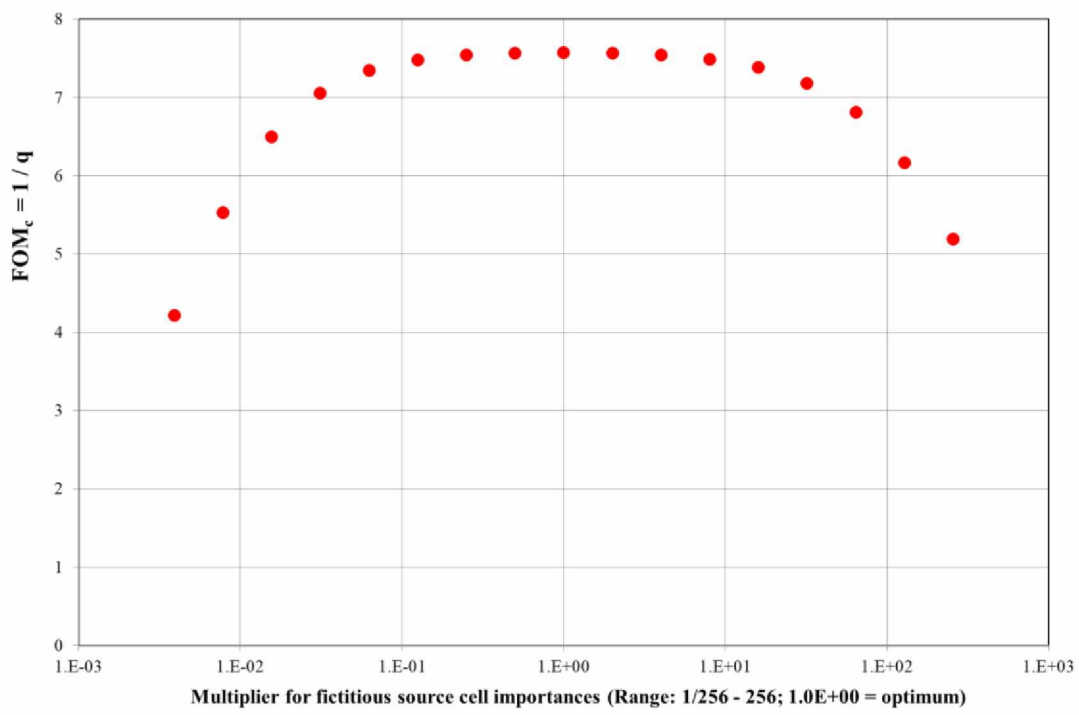

Figure 18. PWR Gen III; fast neutron flux on PV at 0° below supporting platforms (increasing no. fission neutrons passing from fictitious to real-world source cells): inverse of the quality (2nd moment × time) against multiplier of the fictitious source cell importances

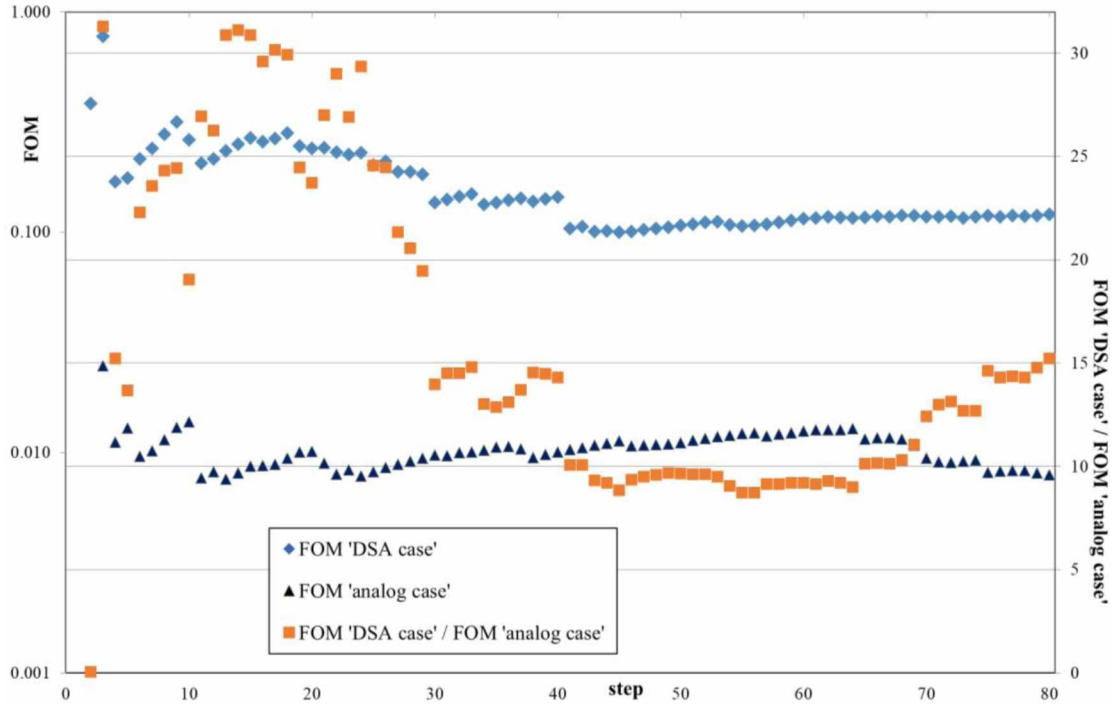

Figure 19. PWR Gen III; fast neutron flux on PV at 0° below supporting platforms (increasing no. fission neutrons passing from fictitious to real-world source cells): *FOM's* of the ex-core response for the two approaches, "DSA" and "analog", and their ratio

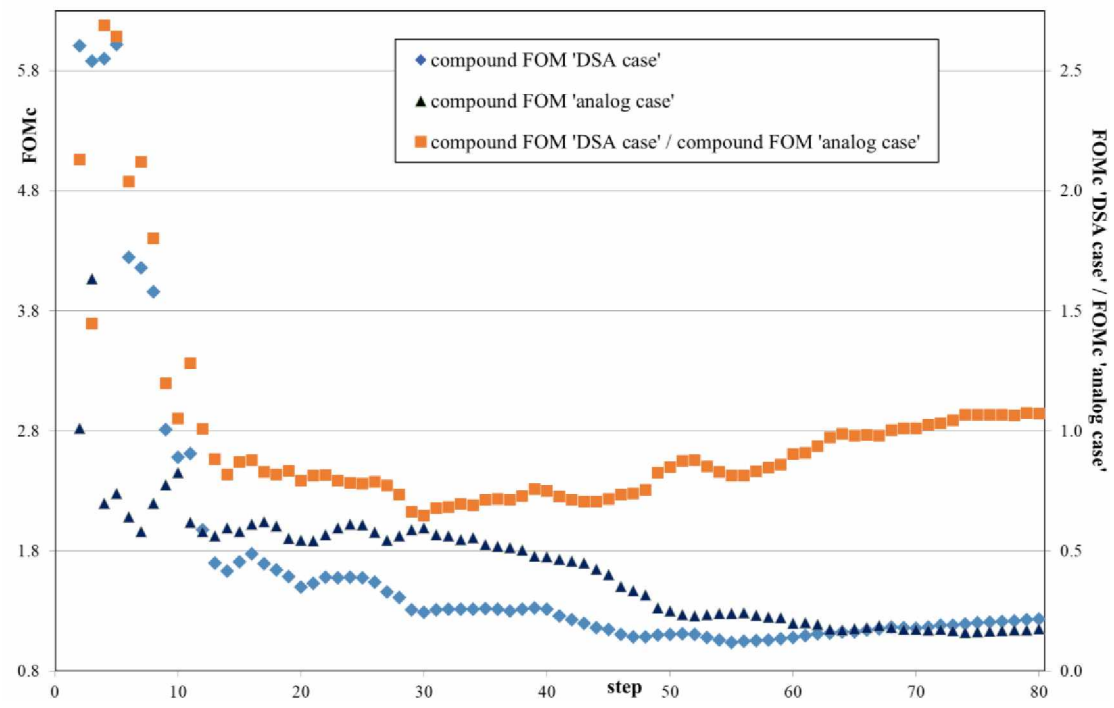

Figure 20. PWR Gen III; fast neutron flux on PV at 0° below supporting platforms (increasing no. fission neutrons passing from fictitious to real-world source cells): *FOM$_c$'s* of the 17 in-core global responses for the two approaches, "DSA" and "analog", and their ratio

The FOM_c's and their ratio for the 17 global responses are shown in Fig. 20. The ratio of the FOM_c's ("DSA" / "analog") after 80 steps was 1.07. In contrast to previous results, the quality of the fundamental mode looks to be, at least, maintained in the DSA case compared with the analog case.

3.1.5 Remarks on the PWR GEN III with Thick Steel Reflector problems

This problem has a large core. We know that the external assemblies give the predominant contribution ex-core so we expect to save time by playing RR in the internal assemblies. One would therefore expect the DSA case with VR parameters in the core to display a more efficient calculation compared with an empirical approach for which analog VR is performed in the core. Instead in §3.1.1 for a variety of neutron responses in the PV well and basemat, we saw that even for neutron responses which are at quite deep penetration in the basemat (Fig. 1), playing VR in the core is worse than doing nothing (in the core).

The reason that the results were poorer than expected was suspected to be that too few fission neutrons were entering the real-world source cells from the fictitious source cells. In §3.1.2 we tested a hypothesis that might explain why our approach produced this result. This was that a basic assumption of the DSA approach (its fixed source nature) was responsible for the lack of fission neutrons. A work-around was made that did not involve any empiricism. Now we saw some improvement in the results of the DSA approach, with a gain of around a factor 2 over the previous results.

We then looked at a fast neutron response nearer the core that is "seen" by only a limited portion of the core. The response was the flux > 1 MeV on the inner surface of the PV. From a number of positions of this response considered in [11], we chose the highest penetration one: below the supporting platforms subtended by an azimuthal angle of ±2.4°. In §3.1.3 we saw that playing VR in the core produced a benefit of a factor of around 4 compared with the empirical approach although the *FOM* ratio DSA case to analog case was not particularly stable.

The same work-around employed in §3.1.2 was used in §3.1.4 on the fast neutron response in §3.1.3. Now we saw a definite improvement in the results of the DSA approach, with a gain of at least 10 over the analog approach in the ex-core result and a fundamental mode that seems to be maintained at least as well as in the analog approach. The results in Figs. 19 and 20 look to go some way towards supporting the hypothesis outlined in the opening paragraphs of §3.1.2. That is, to back the idea that if we perform VR in the fissile zone, it is important that the size of the fission generation should be adequate, and that this is more important than for the case that no VR is performed in the fissile zone.

3.2 In-Core Problem: Neutron Flux and Response at Control Rod Surface in VERA

As far as in-core environments are concerned, we tested the technique firstly on a single assembly model (work in [20] had revealed the need for better statistics, which prompted this analysis), then a full-core model of the VERA (Virtual Environment for Reactor Applications) computational benchmark [10].

The results for the neutron flux crossing the lower surface of a single control rod (CR) in an assembly in an infinite lattice and in a full core are reported in [11]. The gains over analog were ~2.5 and ~50 respectively (in both cases without using FS cells and estimating errors between superhistories). It may be remarked that in the latter case the results are consistent with results over a short axial segment of a single pin found previously [1].

3.2.1 Assembly in a reactor core with a ^{96}Zr(n,γ) response

We take the problem of the single control rod (CR) in an assembly in a full core reported in [11] and add a ^{96}Zr(n,γ) response function at the lower surface of the CR in question. (This was inspired by an earlier attempt to calculate the same response in a wire in the core of a WWRSM–type research reactor [21].)

Fig. 21 shows a horizontal cross-section of an assembly with CR's. The CR's are in the nearly-out position (~11% in – see Fig. 22). This assembly is placed in a full core (see Fig. 23) with all the other assemblies having all 25 guide tube positions empty (i.e. without Pyrex burnable neutron absorber rods, just with water). The assemblies have three different enrichments: 2.11, 2.619 and 3.1% as shown in Fig. 24.

The local tallies were the total neutron flux and the ^{96}Zr(n,γ) reaction rate at the lower surface of the single CR in Fig. 21 at position [3 -6] in the assembly at position [5 -3]. This assembly is shown in Fig. 25 together with the 8 radial VR subdivisions. An axial VR subdivision of 7 segments was adopted. The global fission heating responses in the core were defined as identical to the spatial VR cells, thus 8×7=56 global responses. The ^{96}Zr(n,γ) cross section together with the chosen 7 energy group structure for VR is shown in Fig. 26. We employed a superhistory of 10 fission generations with different VR in each of 5 fission generation sub-groups and included FS cells.

We again use the names "DSA case" and "analog case", although unlike §3.1 the latter employed analog VR both in- and ex-core.

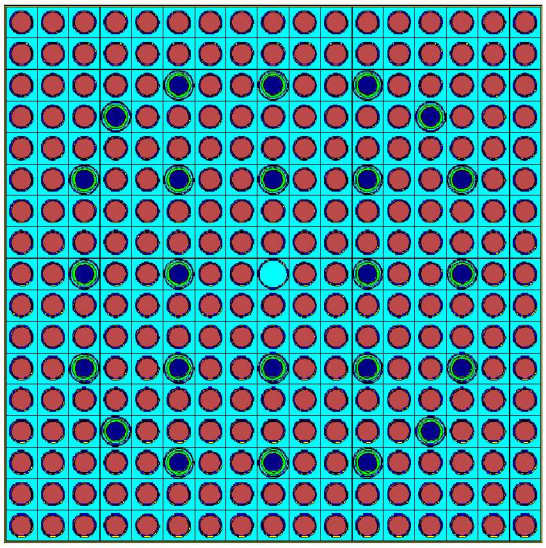

Figure 21. VERA assembly horizontal cross-section with CR's (scale: 21.5×21.5 cm^2)

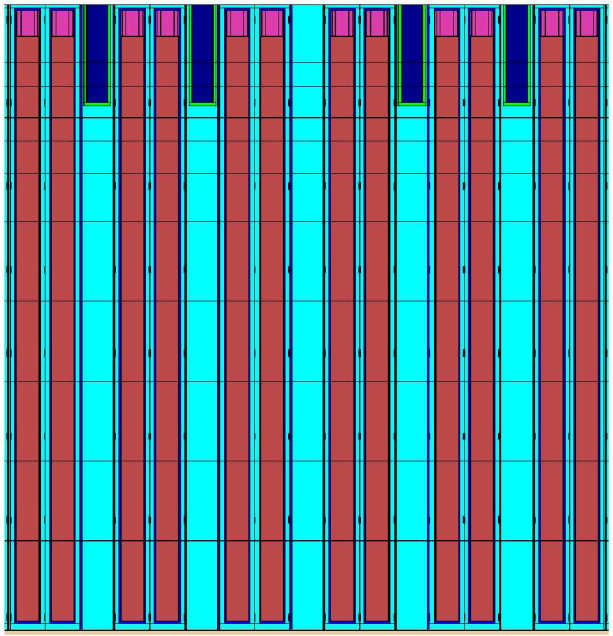

Figure 22. VERA assembly vertical cross-section with CR's ~11% in (scale: 21.5×391.5 cm^2)

We firstly estimated the error between superhistories; that is in the DSA case between groups of 10 fission generations and in the analog case between fission generations. In both cases the size of each superhistory was 2×10^6 and the number of superhistories was 25000 in the analog case and 2500 in the DSA case. A summary of the results is shown in Table I. As far as the ^{96}Zr(n,γ) results are concerned, the analog results are noisy even after a long run and no confidence may be attached to them. With the DSA functions (§3.1.2

and [1,15-17]) the estimation of the gain in FOM_c over analog is around 50000. This value should be more reliable than that in Table I of 5300.

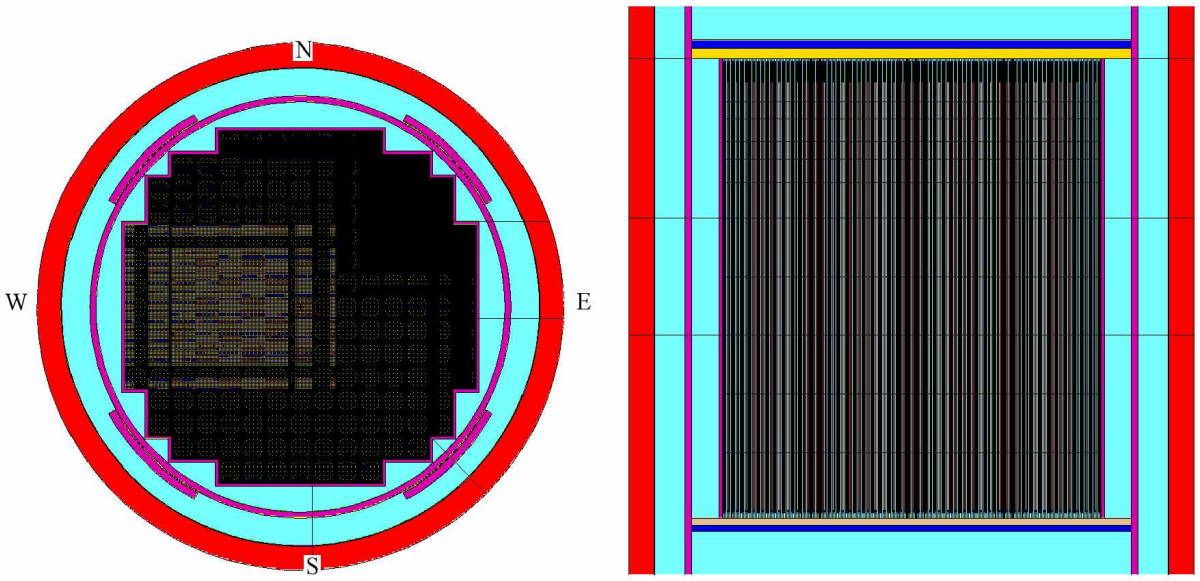

Figure 23. VERA full core horizontal (left) and vertical (right) cross-sections

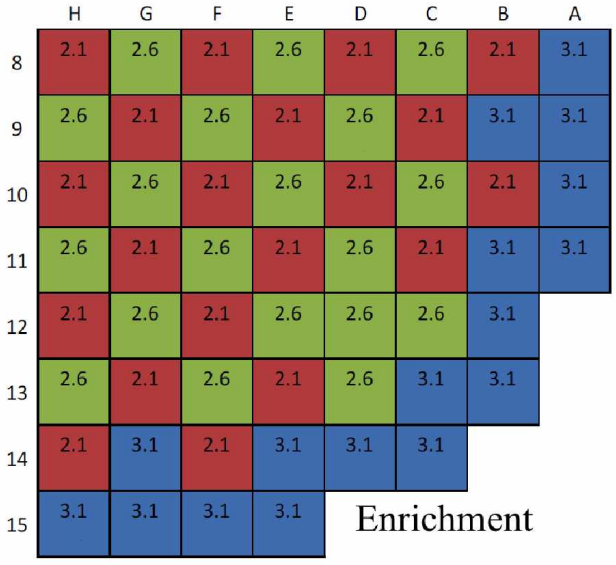

Figure 24. VERA core assembly configuration (from Fig. 9 of [10]) (1/4 core, 1/8 core symmetry)

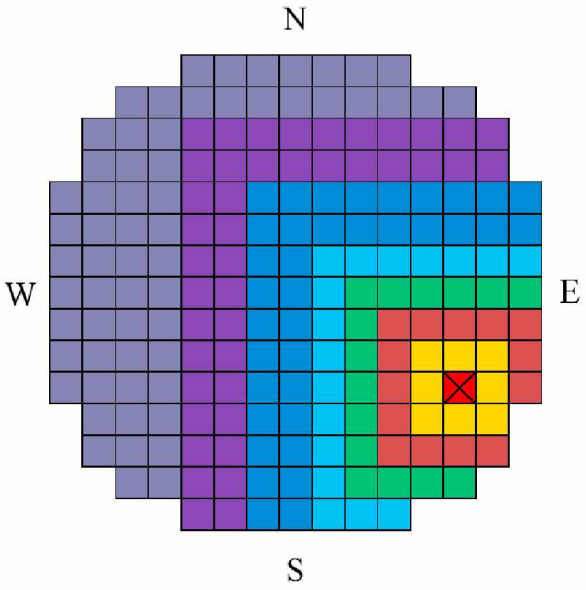

Figure 25. VERA core: assembly containing the tally-of-interest and radial VR division

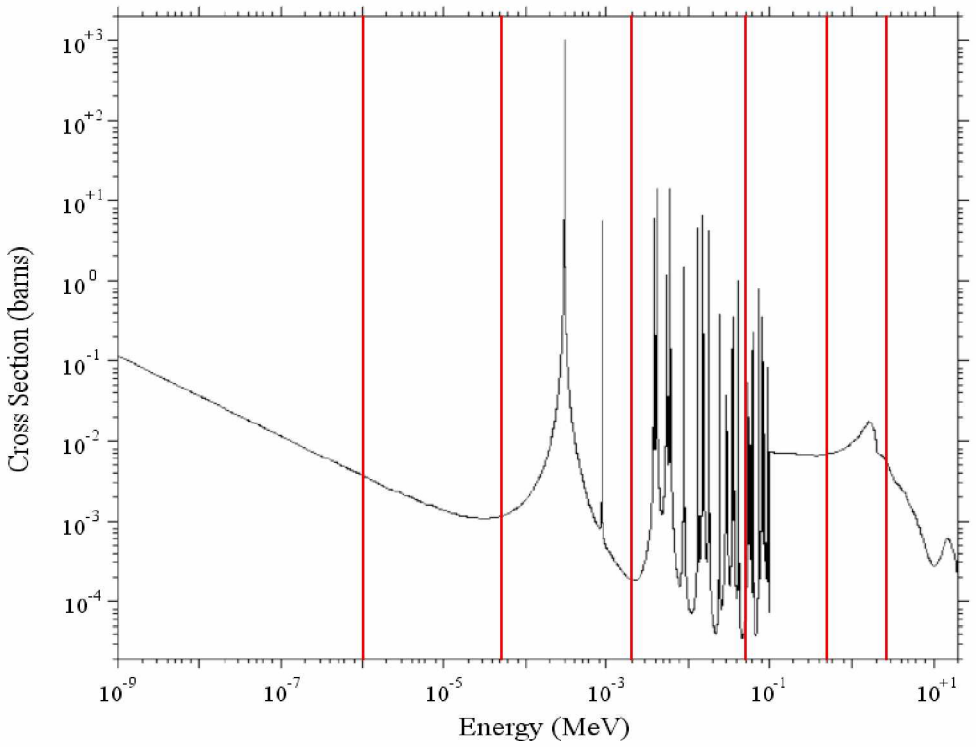

Figure 26. VERA: ^{96}Zr(n,γ) cross section together with the chosen 7 energy group structure for VR

Table I. VERA full core revisited: Results for neutron flux and ^{96}Zr(n,γ) rate at lower surface of a single control rod with statistics between superhistories (*shist* = no. fission generations / superhistory)

statistics	Details of each run	CTM min	Φ_n (unnormalized) at lower surface of CR (fsd)	*FOM*	*FOM* / FOM_{an}	^{96}Zr(n,γ) (unnormalized) at lower surface of CR (fsd)	*FOM*	*FOM* / FOM_{an}	FOM_c (58 responses)	FOM_c/ $FOM_{c,an}$
between super-histories	*shist*=1 with analog VR in fissile zone (*NSRCK*=2M; *KCT*=25000)	2749026	3.84E-7 (0.0157)	0.00148	1	1.07E-7 (0.5721)	1.11E-6	1	6.6E-5	1
between super-histories	*shist*=10 with VR & FS cells (*NSRCK*=2M; *KCT*=2500)	219149	3.99E-7 (0.0051)	0.175	118	7.52E-8 (0.0256)	0.00696	6273	0.350	5300

We then estimated the error in a similar fashion to what was done in §3.1. That is, a calculation was run over a relatively large number of fission generations, splitting it into a number of steps. The fundamental mode, written at the end of a step is used as source for the subsequent step. Each step consisted of 800 fission generations and is assumed independent of the previous step (this assumption is examined in [11]). The DSA case (step of 80 superhistories) had a superhistory size of 4×10^6 neutrons. Because with VR, RR is mainly played between fictitious and real-world source cells, in the analog case this size became 2×10^5

neutrons (step of 800 generations) so as to take roughly the same computer time. Results after 36 steps are shown in Table II.

Table II. VERA full core revisited: Results for neutron flux and ^{96}Zr(n,γ) rate at lower surface of a single control rod with statistics between steps (each step of 800 fission generations) (*shist* = no. fission generations / superhistory)

statistics	Details of each individual step	Φ_n (unnormalized) at lower surface of CR (fsd)	*FOM*	*FOM* / FOM_{an}	^{96}Zr(n,γ) (unnormalized) at lower surface of CR (fsd)	*FOM*	*FOM* / FOM_{an}	FOM_c (58 responses)	FOM_c/ $FOM_{c,an}$	FOM_c (56 fission neutron responses)	FOM_c/ $FOM_{c,an}$
between 36 steps	*shist*=1 with analog VR in fissile zone (*NSRCK*=200K; *KCT*=800)	4.23E-7 (0.0452)	0.00320	1	2.89E-7 (0.9798)	6.8E-6	1	3.9E-4	1	0.539	1
between 36 steps	*shist*=10 with VR & FS cells (*NSRCK*=4M; *KCT*=80)	3.96E-7 (0.0074)	0.122	38.1	8.40E-8 (0.0348)	0.00557	816	0.104	264	0.152	0.28

We see in Table II that, bearing in mind that the statistics are calculated in a totally different manner, the results are as consistent as can be expected with those in Table I. (We remind ourselves that the analog ^{96}Zr(n,γ) rate has an error that renders the *FOM* meaningless.) The FOM_c in Table II for the 56 fission responses excludes the two local responses and thus the analog value is better than the VR one.

In Fig. 27 we show the running *FOM's* for the two cases for the total flux at the lower surface of the CR and their ratio. The *FOM's* of both cases, especially the DSA case, remain reasonably constant over the last two thirds or so of the problem. The *FOM* for the ^{96}Zr(n,γ) rate at the lower surface of the CR for the DSA case is shown in [11] as are also the running FOM_c*'s* of the 56 global fission responses for the two cases and their ratio.

3.2.2 Remarks on the VERA problems

The gain over analog in calculational efficiency for the neutron flux on the lower surface of a single CR in an assembly in an infinite lattice was relatively low (~2.5). Instead when the assembly is placed in a full core it becomes ~50. When the ^{96}Zr(n,γ) reaction rate is added, the gain runs into the thousands or tens of thousands.

The gains in the full core case were verified by both calculating the error between superhistories (of 1 and 10 fission generations) and between batches of superhistories (each batch, or step, made up of 800 fission generations and employing as source the fission distribution written at the end of the previous step).

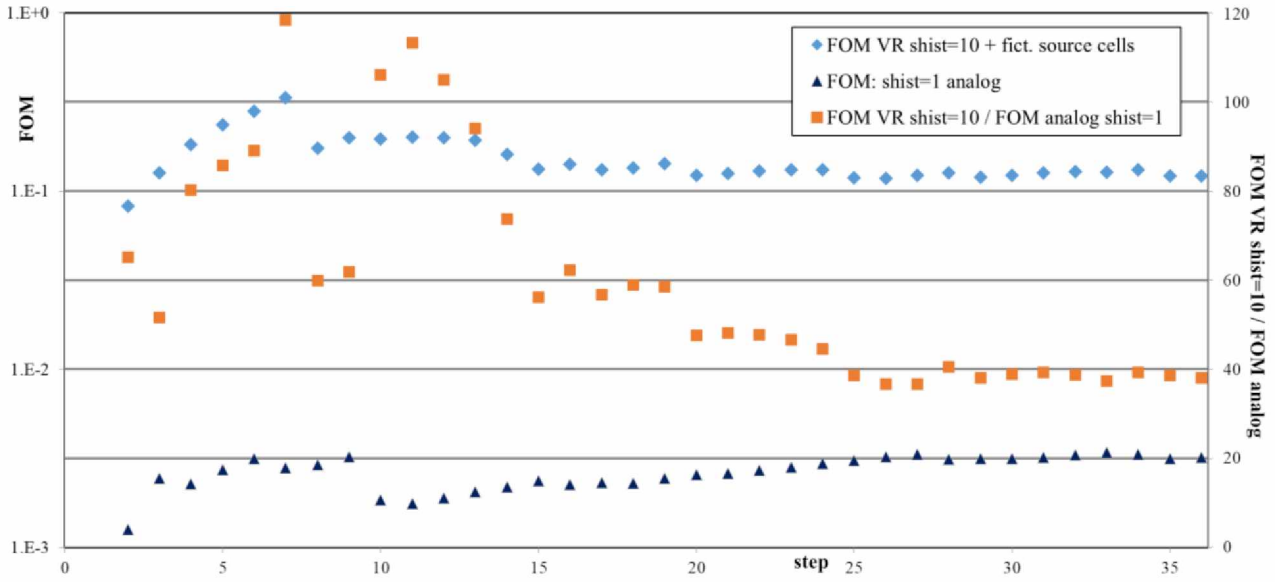

Figure 27. VERA full core with ^{96}Zr(n,γ) response: total flux at lower surface of CR; *FOM's* for the two approaches and their ratio

4. CONCLUDING REMARKS

We have compared the results of our technique for calculating radiation responses in and around critical configurations with an empirical/analog approach. We saw that at first our technique seemed to work less well than we had expected for some responses outside the configuration.

The reason for the underperformance with the ex-core responses looked to be that too few fission neutrons were created in the geometry due to a too strong Russian roulette at their birth point. It was concluded that the cause was weakness in the basic assumptions of the technique. A straightforward work-around was introduced that improved the gain in the sample problems by a factor of 2-3 so that in one of the sample problems there was a factor of 10-15 improvement over the empirical/analog approach.

There seems to be some space for further improvement of the quality of the calculation of ex-core responses by increasing the size of the superhistories either directly, or indirectly through a manipulation of the variance reduction parameters. Further investigation is indicated here to provide analytical support.

When we moved to differential responses within critical configurations, we obtained substantial gains of the order of a factor of 50 in *FOM* for a single pin response in a full core, that were in line with previous results. Introducing a response function – the ^{96}Zr(n,γ) rate – delivered a gain of a factor of some thousands (and solved a problem that cropped up some 25 years ago).

In practice whilst the need to calculate ex-core responses is clear, that of differential in-core ones is less so. Perhaps the existence of a tool will create a market. Also more loosely packed configurations such as those found in criticality safety pose questions of local responses that reactor cores do not.

The evaluation of the technique described in this paper, made through the sample problems, may be considered as a survey of the landscape or topography of the approach. Analytical support is given in the references. The sample problems then pinpoint areas for further analysis.

ACKNOWLEDGMENTS

Dr. Hao Li of the Dept. of Engineering Physics, Tsinghua University, Beijing, China, suggested the VERA single assembly problem and provided the MCNP input deck.

The computing resources and the related technical support used for this work have been provided by CRESCO/ENEAGRID High Performance Computing infrastructure and its staff. CRESCO/ENEAGRID High Performance Computing infrastructure is funded by ENEA, the Italian National Agency for New Technologies, Energy and Sustainable Economic Development and by Italian and European research programmes, see http://www.cresco.enea.it/english for information.

Dedication
This paper is dedicated to the memory of Arie Dubi (1944–2015), the originator of the DSA and a pioneering Monte-Carlo developer.

REFERENCES

1. K.W. Burn, "Optimizing variance reduction in Monte Carlo eigenvalue calculations that employ the source iteration approach," *Ann. Nucl. Energy*, **73**, pp. 218–240 (2014)
2. K.W. Burn, "A correction and a clarification to 'Optimizing variance reduction in Monte Carlo eigenvalue calculations that employ the source iteration approach'," *Ann. Nucl. Energy*, **85**, pp. 776–777 (2015)
3. K.W. Burn, "Estimating Local In- and Ex-Core Responses within Monte Carlo Source Iteration Eigenvalue Calculations," *Proc. PHYSOR-2014*, Kyoto, Sept. 28 – Oct. 3, 2014, JAEA-Conf 2014-003 paper no. 1228305 (2014)
4. Denise B. Pelowitz (Ed.), "MCNP6™ User's Manual – Version 1.0", Los Alamos National Laboratory, LA-CP-13-00634, Rev. 0 (2013)
5. R.J. Brissenden and A.R. Garlick, "Biases in the Estimation of Keff and Its Error by Monte Carlo Methods," *Ann. Nucl. Energy* **13**(2), pp. 63–83 (1986)
6. K.W. Burn and P. Console Camprini, "Radiation transport out from the reactor core: to decouple or not to decouple," *Proc. ICRS-13*, Paris, Oct. 3–6, 2016, EPJ Web of Conferences **153**, 05007 (2017)
7. K.W. Burn and P. Console Camprini, "A Consistent Monte Carlo Treatment of Radiation Responses in and around Critical Configurations," *Proc. PHYSOR-2018*, Cancun, April 22–26, pp. 576–586 (2018)
8. K. W. Burn, "Optimizing Monte Carlo to Multiple Responses: the Direct Statistical Approach, 10 Years On," *Nucl. Technol.* **175**, pp.138–145 (2011)
9. T. E. Booth and K. W. Burn, "Some Sample Problem Comparisons Between the DSA Cell Model and the Quasi–Deterministic Method," *Ann. nucl. Energy* **20–11** 733 (1993)
10. A.T. Godfrey, "VERA Core Physics Benchmark Progression Problem Specifications," Consortium for Advanced Simulation of LWRs, CASL-U-2012-0131-004, Revision 4, Aug. 29, 2014
11. P. Console Camprini, K.W. Burn, M. Brovchenko "Monte Carlo with Variance Reduction in and around Near-Critical Configurations," ENEA NK-N-R-570 (2021)
12. R. von Mises and H. Pollaczek-Geiringer, "Praktische Verfahren der Gleichungsauflösung," *ZAMM - Zeitschrift für Angewandte Mathematik und Mechanik*, **9**, pp. 152–164 (1929)
13. J. Lieberoth, "A Monte Carlo Technique to Solve the Static Eigenvalue Problem of the Boltzmann Transport Equation," *Nukleonik* **11**, pp. 213–219 (1968)
14. H. Hurwitz, "Physical Interpretation of the Adjoint Flux: Iterated Fission Probability" in *Naval Reactor Physics Handbook*, A. Radkowsky (Ed.), **1** 864-869, Naval Reactors, U.S. Atomic Energy Commission (1964)
15. K. W. Burn, "Extending the Direct Statistical Approach to Include Particle Bifurcation between the Splitting Surfaces," *Nucl. Sci. Eng.*, **119**, 44 (1995)
16. K.W. Burn, "Complete Optimization of Space/Energy Cell Importances with the DSA Cell Importance Model," *Ann. Nucl. Energy*, **19–2**, 65 (1992)
17. K. W. Burn, "A New Weight–Dependent DSA Model," *Nucl. Sci. Eng.*, **125**, 128 (1997)
18. T. E. Booth, "Genesis of the Weight Window and the Weight Window Generator in MCNP – A Personal History," LA-UR-06-5807, Los Alamos National Laboratory (2006)
19. A. Sargeni, K.W. Burn and G.B. Bruna, "Coupling Effects in Large PWR Reactor Cores: the Impact of Heavy and Conventional Reflectors on Power Distribution Perturbations," *Proc. PHYSOR-2014*, Kyoto, Sept. 28 – Oct. 3, JAEA-Conf 2014-003 paper no. 1099602 (2014)
20. Hao Li, Ganglin Yu, Shanfang Huang, Gang Wang and Kan Wang, "Mathematics Derivation of the Adjoint-Weighted Tally Value for Geometric Perturbation of K-Eigenvalue based on Continuous-Energy Monte Carlo Method," *Proc. PHYSOR-2018*, Cancun, April 22–26, pp. 3938–3947 (2018)
21. K. W. Burn, "Learning Aspects of the Direct Statistical Approach to the Optimization of Monte Carlo Radiation Transport Calculation," Workshop on Adaptive Monte Carlo Methods, Los Alamos, Aug. 1996, ENEA-IRIS Open Archive: http://hdl.handle.net/20.500.12079/3659

Declaration of interests

☒ The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

☐The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: