Improving Nuclear Data Uncertainty Propagation in Neutronics Calculations - From on-the-fly Doppler broadening to on-the-fly uncertainty propagation -

Pablo Ducru

Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139

Abstract

The future of nuclear energy will in part depend on our capacity to produce high-fidelity predictive simulations enabling the design of optimized advanced systems for enhanced safety and competitiveness. One of the major challenges hindering neutronics calculations towards this goal lies in the uncertainty we have of the nuclear interactions data input parameters. When propagated through a full core calculation, our lack of knowledge of nuclear interaction cross sections results in an uncertainty of the pin power ranging from 1% to 4%. Propagating nuclear data uncertainty is thus crucial if high fidelity calculations are to be achieved.

Nuclear uncertainty data is however seldom propagated in neutronics calculations. This is primarily due to either the limited scope or high computational cost of present uncertainty propagation methods - respectively the Perturbation method or the fast Total Monte Carlo (TMC) one. Indeed, the Perturbation method can only propagate uncertainty on one quantity (for instance the reactivity ρ) per Monte Carlo simulation, and the fast-TMC method requires to run hundreds or thousands of times over a Monte Carlo calculation. In addition, when examining these two methods, accounting for the temperature dependency of nuclear reactions in the resonance region (Doppler-broadening operation) is detected as a computationally determinant step.

To tackle this issue, several prospective methods are introduced, which seek to adapt new, faster, Doppler broadening algorithms to the present state-of-the-art uncertainty propagation methods so as to accelerate them. The here proposed Epistemic MultiPole Accelerated by Windowing and Regression (EMPAWR) method – which introduces nuclear data uncertainty at the heart of the Monte Carlo process by means of a new neutron random walk – stands out as a possible game changer, and a plan for its development and testing is presented.

Indeed, the EMPAWR method would propagate nuclear reactions cross sections uncertainty data to any tallied quantity in only one Monte Carlo calculation. Its implementation into OpenMC could thus yield to an overall nuclear data propagation that would be twice faster and thirty times more statistically accurate than a fast-TMC on a SERPENT-type rejection-sampling Monte Carlo code.

Keywords: Nuclear data, Uncertainty propagation, Sensitivity analysis, Total Monte Carlo, Doppler broadening.

1. Introduction

It is often believed that the greatest source of uncertainty in nuclear reactors power calculations emanates from thermohydraulics (the instantaneous power not being known by more than 2% - that is of the order of 80 MWth for a PWR), while neutronics has achieved an accuracy of the order of a couple hundred pcm. Thus, efforts should focus on thermohydraulics if safety margins are to be reduced by means of high-fidelity simulations, or if predictive computational models are to be achieved for the design and developments of IV$^{\text{th}}$ Generation nuclear reactors.

Email address: p_ducru@mit.edu (Pablo Ducru)

However, the veracity of this statement is seriously compromised if one considers the effects of uncertainty propagation on neutronics calculations. Indeed, nuclear data – be it nuclear reaction cross sections, fission products yields, or decay constants – provided in evaluated nuclear data files (ENDF, JEFF, JENDL, etc.) come with uncertainties due to the limitations of measurements recorded into covariance data. Recent studies have shown that, in the case of the OECD/NEA Martin-Hoogenboom benchmark, the single propagation of ^{235}U, ^{238}U, ^{239}Pu cross sections uncertainties and that of the H and H_2O thermal scattering $S(\alpha, \beta)$ kernels yield local fission pin power uncertainties at mid-hight ranging between 1% (in the center of the core) to 4% (at the periphery), while the statistical uncertainty of the Monte Carlo simulation was always below 1% [1].

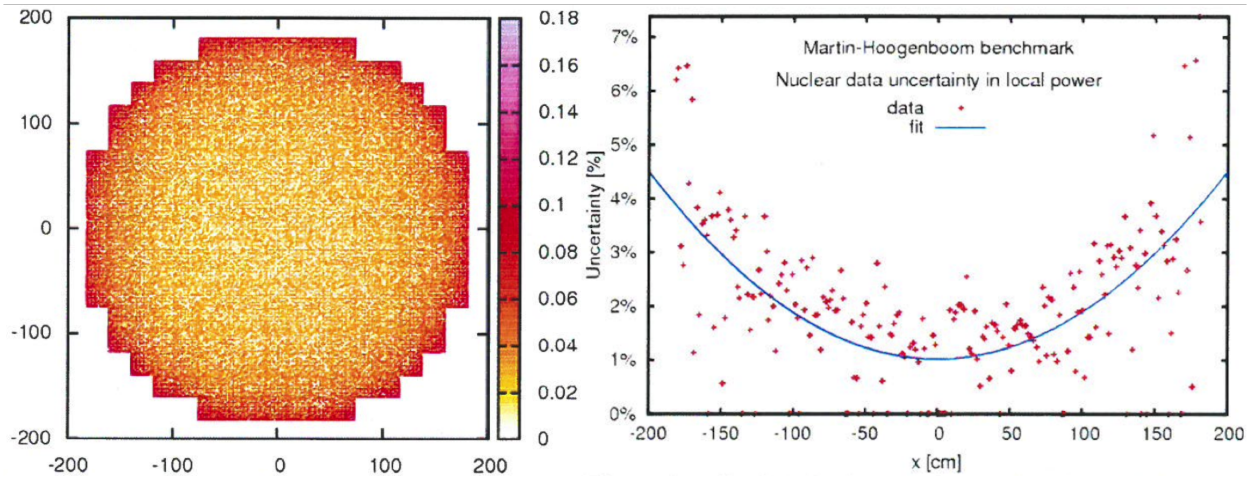

Figure 1: Local power uncertainty at half-hight map in the OECD/NEA Martin-Hoogenboom benchmark, from [1].

This means that, when nuclear data uncertainty propagation is taken into account, neutronics is just as limiting a factor as thermohydraulics when it comes to high-fidelity predictive simulation for reactor design and analysis. Moreover, this stresses the crucial importance of accurately quantifying the nuclear data uncertainty propagation in neutronics calculations, quantification without which no regulatory agency can responsibly diminish safety margins. Important though this may be, present neutronics calculations nonetheless seldom account for nuclear data uncertainty propagation, and this is primarily because of the high computational cost this still represents.

After presenting the current competing two state-of-the-art methods for nuclear data uncertainty propagation in neutronic codes – namely the Sensitivity method and the fast Total Monte Carlo (fast-TMC) – we will point to the Doppler broadening step as an area where improvement could have much impact and introduce tools that could help us progress in them.

Two different methods will then be proposed that could greatly enhance the computational efficiency of nuclear data uncertainty propagation. The Stochastic Accelerated Windowed MultiPole (SWAMP) method proposes to implement the fast-TMC method on the OpenMC Monte Carlo code and make use of its Windowed Multipole Doppler broadening capability. The Epistemic MultiPoles Accelerated by Windowing and Regression (EMPAWR) method proposes to introduce nuclear data uncertainty at every collision by converting the random walk of the neutron into a random walk in a random environment.

The SWAMP method could propagate cross section data uncertainty through a full-core neutronics calculation twice as fast as the alternative on-the-fly rejection sampling method, or with two or three orders of magnitude less memory requirements than the fast-TMC method running on MCNP. Crucially, the EMPAWR method could propagate uncertainty to any output quantity in only one Monte Carlo calculation. It would thus be as fast as the SWAMP method, but 30-times more statistically accurate.

2. Present uncertainty propagation methods and their limitations: Perturbation approach vs. Total Monte Carlo

As of today, two methods are competing in the realm of nuclear data uncertainty propagation: the perturbation method, and the "Total Monte Carlo" one. Both are Monte Carlo based, that is they simulate the random walk of neutrons through the reactor medium and tally the quantities of interest.

2.1. Outlines and limitations of perturbation approach

In Monte Carlo calculations, the Perturbation method looks at the impact of a small deviation of an input parameter (for instance the cross section $\Delta\sigma$) onto some given output quantity of interest (often the k_{eff} multiplication factor) by means of sensitivity coefficients $\left(\frac{\partial k}{\partial \sigma}\right)$ according to the so-called "sandwich rule":

$$\Delta k = \left(\frac{\partial k}{\partial \sigma}\right) \Delta\sigma \left(\frac{\partial k}{\partial \sigma}\right)^{*} \tag{1}$$

The difficulty of this technique lies in the evaluation of the sensitivity gradients $\left(\frac{\partial k}{\partial \sigma}\right)$. The two state-of-the-art Monte Carlo sensitivity methods for this are the Contributon-IFP hybrid method (for Iterated Fission Probability) developed by F. Brown for the Los Alamos MCNP code [2] and the CLUTCH method developed by C. Perfetti for the Oak Ridge National Laboratory TSUNAMI Monte Carlo solver [3]. Both methods rely on some specific algorithm to estimate the "importance" of any given collision.

A main advantage of perturbation/sensitivity methods is that once the sensitivities $\left(\frac{\partial k}{\partial \sigma}\right)$ are calculated, any initial perturbation $\Delta\sigma$ can be propagated. This means that if the effects on k_{eff} of the uncertainty on ^{238}U cross sections is simulated, no additional calculation is required to measure the impact of another isotope's cross section uncertainties.

However, perturbation methods also suffer from major limitations. They restrict themselves to first-order perturbation, and thus provide little insight into what type of shape the output parameter's distribution has. Moreover, it is difficult to propagate quantities other than cross section errors. Yet, most constraining of all is that perturbation method is quantity specific: this means that a calculation can only propagate the error to one quantity (for instance k_{eff}). To estimate the impact of the same perturbation to another quantity (for instance the flux Φ), an entire new calculation has to be run.

2.2. Total Monte Carlo approach: outer epistemic loop

The other approach that has emerged in the past decade has been dubbed "Total Monte Carlo" (TMC) in the neutronics community.

Two types of randomness are considered: the first one is statistic (called aleatoric) and comes from the random walk of the neutron, where the limited amount of trials entails a discrepancy between the random variables and their estimators. The second comes from our lack of knowledge of nuclear data, and is called epistemic uncertainty. This means we model our lack of knowledge of the nuclear data as a probability distribution of the input parameters, referred to as epistemic. The Total Monte Carlo approach is based on discretizing the epistemic uncertainty into 'M' quantities (about M = 1000) and running a full Monte Carlo calculation for each value of the epistemic uncertainty, as shown in [figure 2].

The Total Monte Carlo method first discretizes the probability distribution functions of the input parameters into M values, as shown in the M curves of [figure 2] in the case of cross sections.

This discretized probability distribution of the nuclear data is then propagated through a full-core neutronics calculation by running M times over a full-core problem, and then using the conditional probability relation for a given random variable X and probabilities P_i.

$$\mathbb{E}[X] = \sum_i \mathbb{P}(P_i)\, \mathbb{E}[X|P_i] \tag{2}$$

At first sight, this process is extremely costly since it covers two loops: one external epistemic (i.e. from nuclear data uncertainty) loop ; and one internal aleatoric (i.e. on statistics) loop. However, unlike perturbation theory, this method is not quantity specific and allows to propagate any type of uncertainty

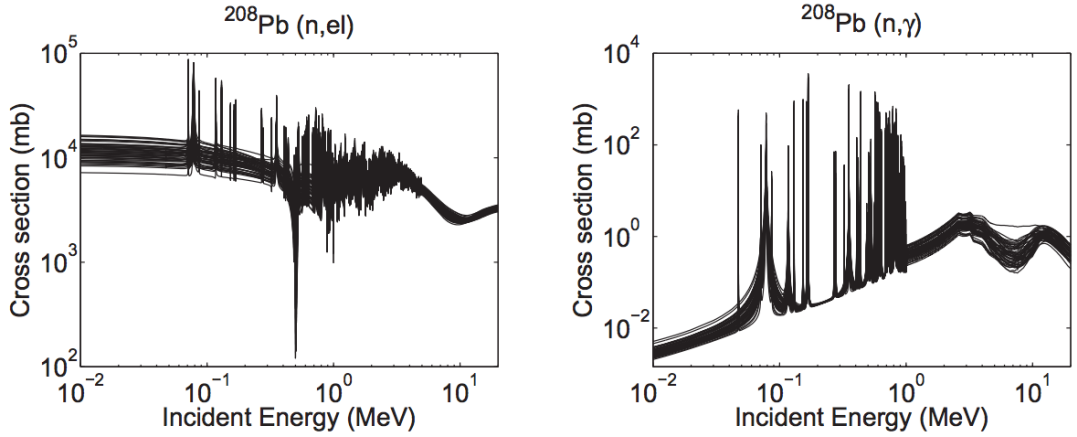

Figure 2: 50 random ACE ^{208} Pb cross sections plotted as a function of incident neutron energy. Left: ^{208} Pb (n,el) and right: ^{208} Pb (n,c). Note that each random ACE files contain a unique set of nuclear data. From [4].

to any type of quantity. Moreover, it enables the propagation of any nuclear data uncertainty distribution, and provides with an accurate sampling of the resulting output distribution.

This however requires to suppose (and then partition) an initial distribution of the resonance parameters. Yet this assumption is already made at the evaluation level since the resonance parameters are fitted with a normal distribution by means of Bayesian inference. Nonetheless, generating resonance parameters normally distributed according to the evaluation files can sometimes lead to unphysical situations, with negative resonance parameters being sampled at tails of the distributions. Moreover, an additional strong drawback of the TMC method is that it faces the insoluble combinatory problem of choosing M initial cross sections. Indeed, any given ℓ-type cross section requires N_λ resonance parameters (Γ, E_λ) to generate. Thus, if one discretizes each resonance paramater's distribution into M values, the total combinatory possibilites of choosing the cross section is M^{N_λ}. In the case of s-wave ^{238}U total cross section, this means 1000^{926} combination possibilities for choosing the $(\Gamma, E_\lambda)_k$ sets, which is far greater than the 10^{80} number of atoms within the observable horizon of the universe. This problem knows no genuine solution. When faced with this issue as they were generating cross sections with their TALYS code, the fast-TMC team had to undertake preliminary studies of the impact of each parameter on the cross sections to determine which ones were significant.

So as to diminish the computational requirement of the Total Monte Carlo method, the fast-Total Monte Carlo (fast-TMC) method was introduced in the wake of Zwermann's [5] and Rochman's [6] work. The fast-Total Monte Carlo (fast-TMC) methods cuts down the total number of neutrons run back to N by simply running M simulations of N/M neutrons, and noticing that the uncertainties follow:

$$\begin{cases} \sigma^2_{observed} \approx \overline{\sigma_{stat}}^2 + \sigma^2_{data} \\ \overline{\sigma_{stat}}^2 = \dfrac{1}{M}\displaystyle\sum_{k=1}^{M} \sigma^2_{stat,k} \end{cases} \tag{3}$$

where $\sigma^2_{stat,k}$ is the statistical uncertainty for each run. The fast-TMC thus suffers from less precise statistics at two levels: the statistic uncertainty from the aleatoric loop is now reduced to the order of $\frac{1}{\sqrt{(N/M)}}$, and the statistical uncertainty on the outer epistemic loop is itself hindered by the discretization in M different sets of input parameters. However, the fact that the aleatoric uncertainty will now be of the order of $\frac{1}{\sqrt{N/M}}$ is overturned by the fact that as long as $N \gg M$ the statistical uncertainty will be much smaller than the epistemic one.

The fast-TMC method thus poses an optimization problem between the discretization required to model the epistemic uncertainty correctly – M – and the need for enough neutrons per Monte Carlo simulation to get statistically significant results – N/M – while preserving an admissible total runtime - N.

3. Windowed Multipole Doppler broadening method - a tool for accelerating Monte Carlo codes

3.1. Doppler-broadening: a major time-limiting step in Monte Carlo simulations

Accurately accounting for the temperature dependence of the cross sections at microscopic level – through the Doppler broadening phenomenon – is both extremely important for feedback mechanisms and quite cumbersome to do in practice.

Indeed, standard Monte Carlo codes such as MCNP thus relegate this to an entire pre-processing step where the nuclear reaction cross sections are pre-computed at different temperatures and stored into ACE files. These temperature cross section grids represent costly look-up steps that further have to be supplemented with interpolation methods between pre-tabulated temperatures. In practice, this means multiplying the data required to store a cross section at 0K temperature by two orders of magnitude. Newer alternative methods seek to paliate this by Doppler broadening on-the-fly. This can be costly however. For instance, the SERPENT rejection-sampling procedure for Doppler-broadening results in 80% of the neutron transport calculation being spent on cross-section evaluations. An improvement in the Doppler-broadening process can thus have high impact in the overall performance of a Monte Carlo code, hence much work has been undertaken recently to achieving fast on-the-fly methods.

3.2. Windowed Multipole: a fast on-the-fly Doppler broadening technique

One such promising approach is the Multipole representation. The Multipole formalism, which was mostly developed by Hwang [7], looks back at the physical and mathematical properties of the cross section as a function of energy to come with mathematical treatments that make the Doppler broadening step analytically solvable.

From quantum scattering and R-Matrix theory, it can then be shown that applying the elemental theorem of partial fraction decomposition to neutral-channels angle-integrated cross sections yields an expression of the cross section as a sum of $2N_\lambda\left(\ell+1\right)$ poles (p_j) per spin group J^π, and their respective residues(r_j):

$$\sigma(u) = \frac{1}{u^2}\sum_j \Re\Big[\frac{r_j}{p_j - u}\Big] \tag{4}$$

where $u = \sqrt{E}$, $2\sqrt{\xi} = \sqrt{\frac{k_b(T-T_0)}{A}}$, and A is the atomic mass of the target nuclei. This form has been called the multipole representation of a given cross section, and the set of poles and residues pairs the multipoles of the cross section.

If one now considers a multipole representation of a given cross section, the linearity of the Doppler broadening operator can be developed through the multipole sum, each term of which becomes analytically Doppler broadenable. By means of this transformation, and mediating a few approximations, Doppler broadening a cross section in multipole representation comes down to evaluating the Faddeeva function – or scaled complimentary complex error function: $W(z) = e^{-z^2}\left(1-\mathrm{erf}(-iz)\right)$ – at a series of points that depend on the velocity of the incoming neutron, the temperature at which we want to Doppler broaden it, and the multipoles of the cross section:

$$\sigma(u,T) \simeq \frac{1}{2u^2\sqrt{\xi}}\sum_j \Re\left[ir_j\sqrt{\pi}W(z_j^0)\right] \tag{5}$$

$$z_j^0 = \frac{u-p_j}{2\sqrt{\xi}}\ , \tag{6}$$

This process has been further optimized into a "windowed multipole" process where, for each resonance, the poles afar – only contributing as a background term – are fitted with a Laurent development and Doppler-broadened by means of an analytical recurrence formula [8] [9].

The current version of this algorithm for Doppler broadening greatly outperforms any other method. Results from C. Josey's thesis indicate that the Windowed Multipole Doppler broadening method is overall twice as fast as the rejection sampling one encoded into OpenMC (which spends 80% of the neutron transport calculation on cross-section evaluation), and even faster than a ACE-file double lookup interpolation (without counting the preprocessing step). Moreover, the Multipole representation cuts down the memory requirements by two or three orders of magnitude compared to the ACE-file pre-processing steps, as generating a cross section at a given temperature takes 10 times less memory than storing it in point-wise data (150 000 points for U-238 for instance), and that a fine temperature grid multiplies this memory requirement by one or two more orders of magnitude.

It is thus hoped that by a proper extension of the multipole representation to deal with epistemic randomness, uncertainty propagation could be achieved at even lower computational cost what present methods represent without error propagation.

4. Methods proposal for enhancing uncertainty propagation by means of Doppler broadening acceleration

We have shown that present state-of-the-art uncertainty propagation methods all suffer from the slowness with which Monte Carlo codes deal with the Doppler broadening temperature effect. Now that we have new tools for on-the-fly Doppler broadening at our disposal, we here propose a plan to expand them to handle uncertainty propagation, thus enabling to propagate uncertainty through Monte Carlo codes faster than the speed at which present days non-uncertainty-propagating Monte Carlo codes run.

4.1. SWAMP - Acceleration of fast-TMC method through accelerated Doppler broadening

A first approach to fast-TMC acceleration through efficient on-the-fly Doppler broadening would be to simply adapt the present Windowed Multipole method to be run in a fast-TMC fashion. We call this the Stochastic Windowed-Accelerated Multi-Pole (SWAMP) approach. This would require – like in the fast-TMC method – to discretize the cross-sections probability distribution into M values, so as to run M times a full OpenMC calculation with (N/M) neutrons.

A preliminary step for the development of the SWAMP approach would be to transform the resonance parameters probability distributions (assumed to be normal from the evaluation) into Multipole probability distributions – $\mathbb{P}(\Gamma, E_\lambda) \implies \mathbb{P}(p_j, r_j)$ – and discretize the latter into M sets of multipoles $(r_j, p_j)_{k\in[\![1;M]\!]}$. M full Monte Carlo simulations will then be run (with N/M neutrons each), each run using one of the discretized sets of multipoles $(r_j, p_j)_k$ as a fixed input parameter for the entire simulation. This SWAMP method would thus exactly mimic the fast-TMC approach, but taking advantage of the fact that each Monte Carlo simulation would now be twice as fast thanks to the Multipole on-the-fly Doppler broadening.

However, by straightforwardly adapting the fast-TMC to the Windowed Multipole on-the-fly Doppler broadening scheme, the SWAMP method also carries along all of the fast-TMC drawbacks, among them, poorer statistics. Moreover, when transposed into multipoles, the combinatory conundrum is enhanced as $2N_\lambda(\ell+1)$ multipoles (p_j, r_j) are required to generate the cross section.

Though it does not solve the fast-TMC problems, the SWAMP method would nonetheless be the simplest first approach to implement. By combining the fast-TMC with the Windowed Multipole algorithm, SWAMP would cut back the memory requirements of the current fast-TMC method by two to three orders of magnitude while preserving the same speed, or be twice as fast as an alternative on-the-fly rejection sampling method for Doppler broadening.

4.2. EMPAWR - Internalizing the epistemic loop into the Monte-Carlo code, enabling new statistics and on-the-fly uncertainty propagation

It can be argued that the major limitation of the fast-TMC method is that is imposes an outer epistemic loop (of size M) upon an inner statistic loop (of size (N/M)). This entails both statistical precision diminution

and epistemic discretization error. The main reason why the cross sections are discretized in M values is that it is not possible to "change ACE file" within the Monte Carlo calculation. In other words, cross sections (pre-Doppler broadened or simply at 0K for "on-the-fly" rejection sampling methods) are global input parameters of a Monte Carlo run, provided once and for all at the beginning of the calculations, and it is impossible to change them at will within the Monte Carlo simulation.

However, this might not be true of the Multipole representation. Indeed, generating random cross sections at each collision at reasonable computational cost – presently impossible to achieve – would come within reach of the Windowed Multipole method. This would open the entirely new perspective of inserting the epistemic loop within the Monte Carlo simulation, in other words, of expressing the uncertainty of our knowledge of the physics within each physical reaction during the random motion of a neutron. We call this method EMPAWR (Epistemic Multipoles Accelerated by Windowing and Regression), and hereafter present the logical steps necessary to develop it.

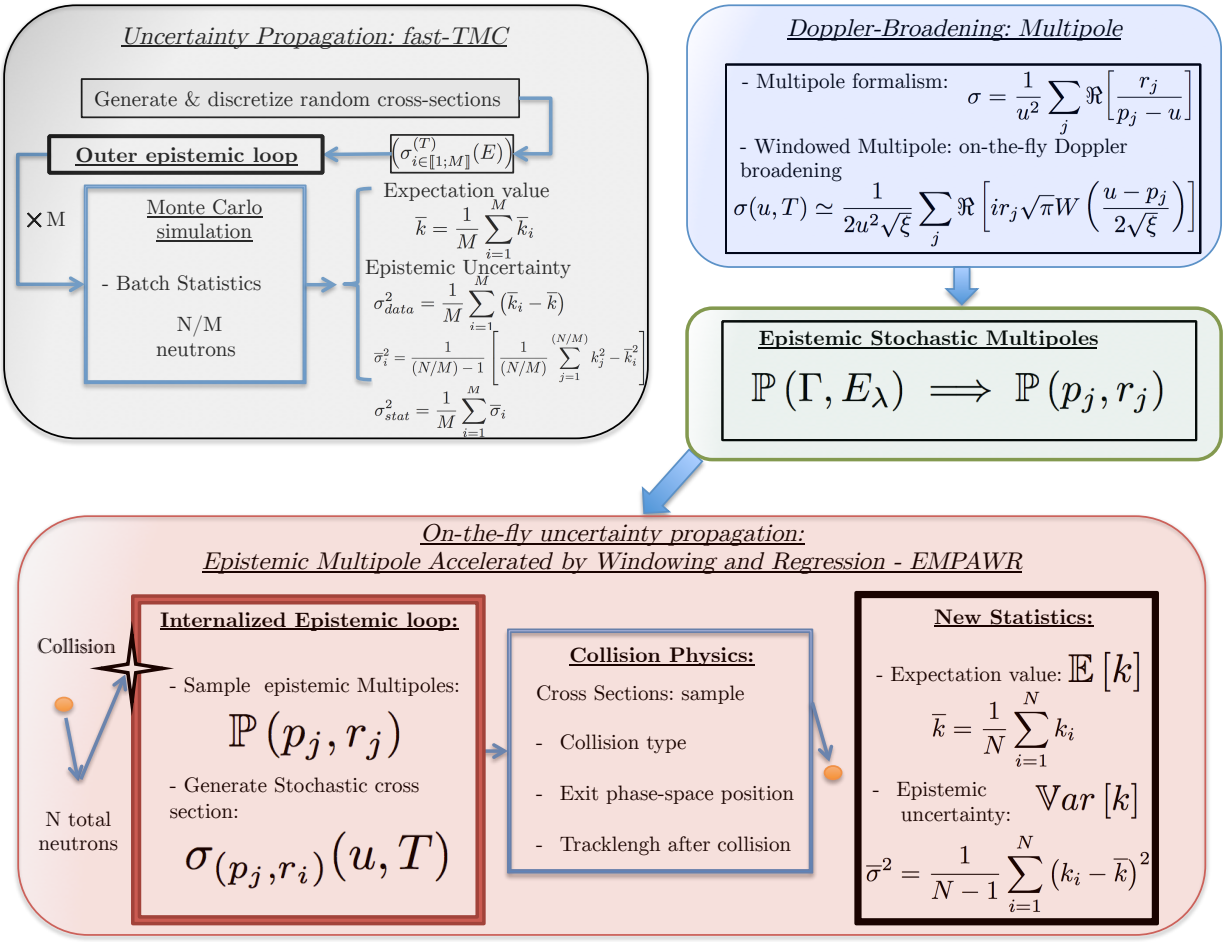

Figure 3: EMPAWR methodology flow chart

First, probability distributions have to be built for the resonance parameters: $\mathbb{P}(\Gamma, E_\lambda)$. Though they are fitted to be Normal distributions during the evaluation, this step runs into the same difficulties previously mentioned for the Total Monte Carlo method of negative widths at the tails of the distribution.

Second, epistemic stochastic Multipoles have to be built by converting the resonance parameters probability distributions $\mathbb{P}(\Gamma, E_\lambda)$ into poles and residues probability distributions $\mathbb{P}(p_j, r_j)$, as represented in the green block of the EMPAWR flow chart [figure 3]. This step will probably be the greatest challenge in

the development of EMPAWR. Indeed, the residues (r_j) are non-trivial functions of the poles (p_j) and the resonance parameters (Γ, E_λ). Moreover, the poles (p_j) of a given isotope's cross sections are the roots of a polynomial the coefficients of which are functions of the resonance parameters (Γ, E_λ).

$$\begin{cases} \forall j \in [\![1; 2N_\lambda(\ell+1)]\!] & P(p_j) = 0 \\ P(X) = \sum_{n=0}^{2N_\lambda(\ell+1)} a_n X^n & a_n = f(\Gamma, E_\lambda) \end{cases} \tag{7}$$

The problem is thus: "given a polynomial the coefficients of which have known probability distribution functions, what are the probability distributions of its roots?". It is however known since Niels Abel and Évariste Galois' theorems that the roots of a polynomial of degree five or more are not in general soluble in terms of radicals. The problem is thus mathematically non-trivial.

Assuming the probability distributions of the poles and residues $\mathbb{P}(p_j, r_j)$ have been obtained for each isotope and partial reaction, this opens the possibility of generating the epistemically stochastic cross sections on-the-fly, at each collision. For each isotope, the probability distribution functions of the poles and residues will have to be stored into a new Multipole library, represented by the arrow connecting the green and red boxes of the EMPAWR flow chart [figure 3].

With such an epistemic Multipole library at our disposal, a new Monte Carlo procedure is possible in which our uncertainty of the physics is embedded. This is the EMPAWR random walk, represented in the red box in the EMPAWR flow chart [figure 3]. At each neutron collision, the cross sections that the neutron "sees" are now generated in a random way that describes our epistemic uncertainty about the nuclear data. In practice, this is achieved by sampling the poles and residues according to their probability distributions, and then building the cross section on-the-fly at the given local temperature through the Windowed Multipole process. It is interesting to further note that if the contribution of the poles afar from the incoming neutron's energy (outside the outer window in the Windowed Multipole process) is assumed unchanged by the epistemic stochasticity, this new internal epistemic perturbations step will only require to sample the poles evaluated within the outer windows. This means this EMPAWR process adds almost no additional cost.

Moreover, this new, doubly random, walk of the neutrons will entail new statistics. Indeed, in current Monte Carlo calculations, the variance of the random walk has no clear meaning and is only used in batch statistics to calculate the deviation of the mean estimator from the searched-for expected value. However, in the new EMPAWR random walk, the variance could actually represent the deviation from the mean due to the internal epistemic loop: in other words the nuclear data related uncertainty on the final result. Moreover, the statistical uncertainty would now be of the order of $1/\sqrt{N}$, compared to $1/\sqrt{N/M}$ in the fast-TMC case, thereby circumventing all the issues the fast-TMC method was faced with with its outer epistemic loop. The previous results are however not trivial to prove: indeed, the EMPAWR method will be replacing a random walk in a given environment with a random walk in a random environment, and more work will have to be undertaken on the mathematical aspect to prove the consistency of the methods.

Thus, the EMPAWR approach could perform on-the-fly uncertainty propagation and Doppler broadening in only one Monte Carlo run, going as fast as the SWAMP method, but thirty times outperforming it in statistical precision. This, of course, would be possible only at the cost of solving the mathematically challenging issue of determining the probability distributions of the epistemic stochastic Multipoles.

4.3. Expanding to other cases not handled by the Multipole representation method

Unfortunately, not all nuclear cross sections data provided by the evaluators are in the form of resonance parameters, nor is it believes that the multipole representation can be extended to any type of cross-section formalism. For instance, angular dependency seems to be presently out-of-reach of the multipole representation. Moreover, crucial isotopes such as ^{1}H and ^{16}O are only provided in pointwise data. It is therefore necessary to also come about with a method to handle these cases.

Fortunately, a strong advance has been very recently made by G. Ferrand to deal with the Doppler-broadening operation applied to an arbitrary function of positive energies [10]. The EMPAWR method

Table 1: Performance comparison of uncertainty propagation and Doppler broadening methods: fixed total neutrons N.

	fast-TMC/MCNP	fast-TMC/SERPENT	fast-TMC/OpenMC = SWAMP	EMPAWR
Doppler broadening	- input data: ACE-file storing $\sigma(E)$ at reference temperatures - process: Lookup and interpolation	- input data: 0K pointwise cross sections - process: Rejection sampling	- input data: Multipole file - process: Windowed multipole	- input data: Epistemic Multipole file - process: Windowed multipole
Memory	ref.	10^1 - 10^2 $\times$ less	10^2 - 10^3 $\times$ less	10^2 - 10^3 $\times$ less
Speed	ref.	2$\times$ slower	same	same
Statistical precision	$\mathcal{O}\left(\frac{1}{\sqrt{N/M}}\right)$	$\mathcal{O}\left(\frac{1}{\sqrt{N/M}}\right)$	$\mathcal{O}\left(\frac{1}{\sqrt{N/M}}\right)$	$\mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$

could therefore find an alter-ego to treat those cases the Multipole formalism cannot treat in the new Fourier Transform method for Doppler broadening. Thought slower, this method could handle more general situations and might also circumvent some of the difficult mathematical problem raised by the EMPAWR approach since the probability transformation would be reduced to taking the characteristic function of the cross sections' probability distributions.

4.4. Roadmap for development of the EMPAWR method

The complete development of the EMPAWR method would have to result from the gradual validation and testing of the following steps and deliverables:

1. *Proof of concept of the double random walk on toy problem*: The preliminary step will be to validate the mathematical consistency of the EMPAWR random walk in a random medium on a toy problem: an analytically solvable 1D, 1 group case where the theoretical 1^{st}-order perturbation can be calculated will serve as a reference case against which test the results of an EMPAWR Monte Carlo code where the cross sections will be assumed normally distributed. A crucial step will be to observe if the mean and variance converge towards the expectation value of the non-EMPAWR random walk and the epistemic uncertainty respectively.
2. *Transform the resonance parameters probability distributions into that of poles and residues*:$\mathbb{P}\left(\Gamma, E_\lambda\right) \implies \mathbb{P}\left(p_j, r_j\right)$ This step represents a mathematically challenging problem. If after having recourse to Galois theory and Polynomial Chaos expansion no theoretical solution is found, it might be necessary to solve the problem numerically through a Monte Carlo approach. The various moments of the output distribution will then have to be fitted to reconstruct an approximate poles and residues probability distribution.
3. *Building an Epistemic Multipoles library*: Once the proof-of-concept made and the poles and residues distribution functions obtained, the EMPAWR method would have to be implemented into OpenMC. In practice, this would first require the building of a new Multipole library, where, for each isotope, the probability distribution functions of the poles and residues are recorded. Considering that for each spin-group of angular momentum ℓ their are $2N_\lambda(\ell+1)$ poles (and as many residues), it is crucial that the probability distributions be stored in a compact way if memory escalation is to be avoided. This could be achieved by storing only the relevant moments of the distributions. Only a detailed study of them will solve this issue.
4. *Implementation into OpenMC*: Once the Multipole library established with the probability distribution function of the poles and residues, the EMPAWR method will then be implemented into OpenMC. For more flexibility, this could be achieved be means of a MP_data_update function that will update

the Multipole data to be read and used according to their probability distribution functions. Thus, the random sampling of the nuclear data could be made at every collision, or at history or batch level, and their results compared.

5. *Validation and testing on reference benchmarks*: It is crucial that once implemented, the EMPAWR method be validated and tested against a real reference case. The most robust reference method would be to run a Total Monte Carlo calculation on the BEAVRS PWR benchmark, or, for cross-comparison with previous works, with other specified OECD/NEA benchmarks. The TMC reference case would have to be performed with isotope specific perturbations on all important isotopes, and then again with all the isotopes "mixed" together. Such a run would take months. So a first step could be to only perturb the cross-sections of some important isotope, for instance ^{239}Pu, and use that as a reference case against which to test the EMPAWR method.

5. Conclusion

Two methods have here been proposed to to accelerate the present fast-TMC state-of-the-art uncertainty propagation method by improving the process by which the temperature dependency of cross sections is handled in the resolved resonance region. Both rely on the Multipole representation to perform the Doppler broadening operation on-the-fly while evaluating cross sections at every collision. Though the Multipole representation is only achievable yet for neutral channels angle-integrated cross-sections provided with resonance parameters, the Fourier transform method for Doppler broadening could also be adapted to perform on-the-fly uncertainty propagation on those cases. While the SWAMP method would be the easiest to implement at first and still cut by two-to-three orders of magnitude the memory requirements of the present fast-TMC method – or run twice faster than a fast-TMC ran on a rejection-sampling code – it would still face all the obstacles the fast-TMC is confronted with. On the other hand, the EMPAWR method could represent a genuine breakthrough by changing the random walk of the neutron in the reactor into a random walk in a random environment. If the considerable mathematical hurdles it poses are overcome, the EMPAWR method would enable both on-the-fly Doppler broadening and uncertainty propagation in a single Monte Carlo run, achieving 30 times better statistical precision than SWAMP at equivalent speed.

References

[1] D. Rochman, S. van der Marck, Nuclear data uncertainties for local power densities in the martin-hoogenboom benchmark, Tech. rep., Nuclear Research and Consultancy Group NGR (2013).

[2] F. Brown, B. Kiedrowski, Continuous-energy sensitivity coefficient capability in mcnp6, American Nuclear Society Winter Meeting.

[3] C. M. Perfetti, Advanced Monte Carlo Methods for Eigenvalue Sensitivity Coefficient Calculation, PhD Thesis, University of Michigan, 2012.

[4] E. Alhassan, H. Sjöstrand, P. Helgesson, A. Koning, M. Österlund, S. Pomp, D. Rochman, Uncertainty and correlation analysis of lead nuclear data on reactor parameters for the european lead cooled training reactor, Annals of Nuclear Energy 75 (2015) 26–37.

[5] W. Zwermann, al., Aleatoric and epistemic uncertainties in sampling based nuclear data uncertainty and sensitivity analyses, in: Proceedings of PHYSOR 2012, 2012.

[6] D. Rochman, W. Zwermann, A. Koning, al., Nuclear data uncertainty propagation: Perturbation vs. monte carlo, Nuclear Science and Engineering 177 (2014) 337–349.

[7] R. N. Hwang, A rigorous pole representation of multilevel cross sections and its practical applications, Nucl. Sci. Eng. 96 (1987) 192–209.

[8] C. Josey, B. Forget, K. Smith, Efficiency and accuracy evaluation of the windowed multipole direct doppler broadening method, in: Proceedings of PHYSOR 2014, 2014.

[9] C. Josey, P. Ducru, B. Forget, K. Smith, Windowed multipole for cross section doppler broadening (2014 - still awaiting publication) Submitted to Journal of Computational Physics.

[10] G. Ferran, M. Gonin, A new method for the doppler broadening of the solbrig's kernel using a fourier transform, Nucl. Sci. Eng. 179 (2015) 285–301.