Windowed Multipole for Cross Section Doppler Broadening

C. Josey^a, P. Ducru^a, B. Forget^a, K. Smith^a

^a*Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139*

Abstract

This paper presents an in-depth analysis on the accuracy and performance of the windowed multipole Doppler broadening method. The basic theory behind cross section data is described, along with the basic multipole formalism followed by the approximations leading to windowed multipole method and the algorithm used to efficiently evaluate Doppler broadened cross sections. The method is tested by simulating the BEAVRS benchmark with a windowed multipole library composed of 70 nuclides. Accuracy of the method is demonstrated on a single assembly case where total neutron production rates and ²³⁸U capture rates compare within 0.1% to ACE format files at the same temperature. With regards to performance, clock cycle counts and cache misses were measured for single temperature ACE table lookup and for windowed multipole. The windowed multipole method was found to require 39.6% more clock cycles to evaluate, translating to a 7.9% performance loss overall. However, the algorithm has significantly better last-level cache performance, with 3 fewer misses per evaluation, or a 65% reduction in last-level misses. This is due to the small memory footprint of the windowed multipole method and better memory access pattern of the algorithm.

Keywords: Monte Carlo, Multipole, Doppler Broadening

---

1. Introduction

Monte Carlo methods for neutron transport calculations have always been synonymous with high fidelity but often limited in use to benchmarking and code verification. The process of randomly tracking particles over large nuclear reactor cores is computationally intensive but is now within reach using leadership class computing platforms and novel communication algorithms [1]. Despite such advancements, high fidelity nuclear reactor simulations are still beyond immediate reach when one considers the need to account for temperature variations during operation and the necessity of tallying reaction rates over millions of regions for tracking core evolution. As part of the CASL project, progress has been made on the last front by using data decomposition [2], overlapping domain decomposition [3] or combinations of non-overlapping domain decomposition and replication [4].

This paper focuses on the former issue of modeling temperature variations in an operating nuclear reactor. Nuclear data evaluations contain parameters and distribution functions for each nuclide representing the interaction probabilities with neutrons at 0K. They are derived by complex fitting procedures relating experimental data with quantum models applicable in different energy ranges. Common practice requires reconstructing the parametric models of cross sections as point-wise values that can be interpolated linearly over 10 orders of magnitude in energy. Such reconstruction is essential to perform Doppler broadening which captures the vibrational effect of temperature on the target material using convolution procedures such as SIGMA1 [5]. A complete data set for a single temperature can require approximately 1GB of data that is then parsed by Monte Carlo codes using a binary search process over each nuclide at each event. This nuclear data lookup is a current bottleneck of nuclear simulations since the memory access pattern is

---

Email addresses: cjosey@mit.edu (C. Josey), p_ducru@mit.edu (P. Ducru), bforget@mit.edu (B. Forget), kord@mit.edu (K. Smith)

somewhat random requiring large data movements. Decomposition of this data is possible but has inherent performance limitations due to the frequency with which it is needed [6, 7].

The simple approach to deal with temperature dependence has been to pre-process nuclear data at many temperatures and interpolate [8, 9, 10]. Such schemes require storing massively large libraries or performing costly pre-processing to encapsulate the temperature range of interest in operating reactors (300K-3000K) while preserving accuracy with sufficiently small temperature spacing. In recent years, novel approaches have also been proposed to treat the Doppler broadening on-the-fly, such as a regression fit using a multi-term expansion at each energy point [11] or directly sampling target motion with an added rejection sampling scheme [12, 13]. While effective, both these methods still require substantial storage of point-wise data and memory access patterns remain largely random.

Recent work, as part of the CESAR project, identified a new approach to efficiently obtain Doppler broadened cross sections based on the multipole formalism [14]. The multipole formalism is a physically and mathematically equivalent formulation of the resonance parameters found in nuclear data evaluations [15]. Its strengths lie in its ability to analytically perform the convolution integral needed to obtain Doppler broadened data without requiring reconstruction of point-wise libraries thus greatly reducing storage needs. Cross sections are constructed at the needed energy point via a sum of consecutive poles and residues thus greatly reducing data movement. To improve efficiency, the windowed multipole method was proposed and optimized [14, 16] on a few key nuclides and thermal-hydraulic coupling studies were also performed as a proof-of-concept of using Monte Carlo methods for multi-physics simulations [17].

In this paper, the windowed multipole method is extended to 70 nuclides and accuracy is assessed on the BEAVRS reactor benchmark. Additionally, performance is assessed by profiling both the single temperature point-wise data lookup approach and the windowed multipole approach. Section 2 presents a background of the R-matrix formalism commonly used to represent nuclear data in the resolved resonance range, as well as the conversion to the multipole formalism. Section 3 introduces the windowed multipole method and a new approach for Doppler broadening high order polynomials used in the fitting process. The contents of the 70 nuclide library are described in Section 4, while Section 5 presents accuracy and performance results on a PWR benchmark. Conclusions and future research directions are presented in Section 6.

2. Background

A common way to model nuclear reactions in the resonance energy region between a projectile and a target at rest is through R-matrix theory and its simplifications [18]. In R-matrix theory, the nuclear potential inside of the target particle is considered a black box, and the unknowns are fitted to experiments through Bayesian statistics. In practice, the R-matrix formalism is often simplified into subsequent, easier to handle forms such as the Reich-Moore formalism and the Multi-Level Breit-Wigner (MLBW) formalism [19]. Once experiments are fit to these formalisms, the results are standardized and recorded in evaluations.

2.1. R-matrix theory

The R-matrix theory is at the heart of most present-day cross section evaluations, and is thoroughly covered by Lane & Thomas [20]. It represents the quantum modeling of the particle-nuclei interaction in the resonance region. It is based on non-relativistic two body interactions forming a compound nucleus in a spherical potential of otherwise unknown form. All the possible reactions are divided into orthogonal entry and exit channels. A reaction channel is defined as a pair of interacting particles (α) and their quantum state, respectively the orbital and intrinsic angular momentum of the two particles, the total angular momentum and its projection on the z axis: $c \equiv \{\alpha, \ell, j, J, m_J\}$. A given reaction can thus be written as: $\alpha \rightarrow A^* \rightarrow \alpha'$ [21]. For each channel, space is divided in two regions: an internal region of radius a_c , corresponding to the inner nucleus potential (considered as a black box), and an external region, where the nuclear potential is null (if relevant, the Coulomb potential exists). A key assumption of the model is that the C^1 continuity condition of the wave function is ensured for each channel by imposing that the radial logarithmic derivative at the channel boundary a_c of the Eigenfunctions χ_λ of the Hamiltonian be a real constant value B_c , that is: $B_c = \left. \frac{\partial_{r_c} \chi_\lambda}{\chi_\lambda / r_c} \right|_{r_c=a_c}$.

From these hypotheses, the R-matrix theory establishes that the angle-integrated cross section for a given entry c channel and an exit c' channel is given by:

$$\sigma_{cc'} = \pi \lambda^2 g_c |\delta_{cc'} - U_{cc'}|^2 \quad (1)$$

where λ is the reduced neutron wavelength, g_c is the spin statistical factor, $\delta_{cc'}$ is the Kronecker delta, and $U_{cc'}$ is the cc' entry of the collision (or scattering) $\mathbf{U}$ matrix [19].

The collision $\mathbf{U}$ matrix, which is unitary and symmetric so as to conserve probability and time reversibility respectively, is then expressed as a function of the $\mathbf{R}$ matrix and the boundary continuity condition $\mathbf{L}$:

$$\begin{aligned} U_{cc'} &= e^{-i(\varphi_c + \varphi_{c'})} P_c^{1/2} \left([\mathbf{I} - \mathbf{R}(\mathbf{L} - \mathbf{B})]^{-1} \times [\mathbf{I} - \mathbf{R}(\mathbf{L}^* - \mathbf{B})] \right)_{cc'} P_{c'}^{1/2} \\ &= e^{-i(\varphi_c + \varphi_{c'})} \left(\delta_{cc'} + 2i P_c^{1/2} \left[(\mathbf{I} - \mathbf{R} \mathbf{L}^0)^{-1} \mathbf{R} \right]_{cc'} P_{c'}^{1/2} \right) \end{aligned} \quad (2)$$

$$R_{cc'} = \sum_{\lambda} \frac{\gamma_{\lambda c} \gamma_{\lambda c'}}{E_{\lambda} - E} \quad (3)$$

$$L_{cc'}^0 = L_{cc'} - B_{cc'} = (L_c - B_c) \delta_{cc'} \equiv (S_c + iP_c - B_c) \delta_{cc'} \quad (4)$$

where the parameters of the R-matrix are evaluated experimentally by the evaluators and are given inputs for the problem [19].

Importantly, the fact that B_c is real and constant entails that the resonance parameters E_{λ} and $\gamma_{\lambda c}$ are also real and constant. All the energy dependence of the collision matrix $\mathbf{U}$ is thus encapsulated into the single poles of the $\mathbf{R}$ matrix terms of Equation (3), φ_c and $L_c = S_c + iP_c$, all functions of $u = \sqrt{E}$, where E is the kinetic energy of the incoming particle hitting the target nucleus.

For neutral projectiles (no Coulomb interactions), the penetration P_c , shift S_c and phase shift φ_c factors only depend on ℓ . Further, both the penetration factor P_{ℓ} and the shift factor S_{ℓ} are rational functions in $\rho_c = \rho_{0c} \sqrt{E} = k_c r_c$, the values of which can be derived from Powell's recurrence formulas and have been tabulated in [20].

The approximations of R-matrix theory (i.e. MLWB and Reich-Moore) do not capture all of the physical phenomena provided by the experimental data. Thus, additional point-wise background terms are often provided in File 3 of the ENDF-6 data format as a correction and must be added to the cross section calculated by the models.

2.2. Multipole Formalism

The fact that $R_{cc'}$, P_{ℓ} and S_{ℓ} are all rational functions of $\rho \propto u = \sqrt{E}$ is essential to construct the multipole formalism. If we neglect the energy dependence of the phase shift factor by treating $\varphi_c(\rho)$ as if it were a constant, the cross section $\sigma_{cc'} = \pi \lambda^2 g_c |\delta_{cc'} - U_{cc'}|^2$ in turn also becomes a rational function in u . Further, according to Hwang [15], all the poles of this proper rational fraction are simple poles, to the exception of zero itself which is a double pole. We can thus apply the elemental theorem of partial fraction decomposition to write the cross section as a sum of poles and their respective residues :

$$\sigma = \frac{1}{u^2} \sum_j \Re \left[\frac{r_j}{p_j - u} \right] \quad (5)$$

We shall henceforth call this form the multipole representation of a given cross section, and the set of poles and residues pairs the multipoles of the cross section.

In order to cast a cross section into its multipole representation, one must thus transform the standard ENDF evaluated nuclear data file parameters for a given formalism into its poles and residues. The methods for converting ENDF parameters into multipoles can be found in [15] for the Reich-Moore formalism and in [22] for the version of the MLBW formalism as described in the ENDF-6 manual. Assuming that we have at our disposal a multipole representation of a nuclear cross section, this formalism becomes a powerful tool since it yields an analytical expression of the Doppler broadened cross section.

2.3. Multipole Doppler Broadening

Though not temperature dependent in the center of mass frame, cross sections become temperature dependent in the laboratory frame due to the thermal agitation of the target nuclei. This agitation is generally assumed to be a Maxwellian-Boltzmann ideal gas distribution. Doppler broadening is a linear operation that transforms the 0K cross section into its TK value through a convolution product of the Maxwellian distribution over the relative velocities of the particle and the target nucleus.

Let $u = \sqrt{E}$ and $\xi = \frac{k_b(T-T_0)}{4A}$, where A is the atomic mass of the target nuclei. Deriving from Maxwell-Boltzmann's free-gas model theory for the target nuclei, Solbrig's Kernel gives the Doppler broadened cross section by [23]:

$$\left\{ \begin{array}{l} u\sigma(u, T) = \int_0^\infty x\sigma(x, T_0)S(u, x)dx \\ \text{where: } S(u, x) = \frac{x}{2u\sqrt{\pi\xi}} \left[e^{-\frac{(u-x)^2}{4\xi}} - e^{-\frac{(u+x)^2}{4\xi}} \right] \end{array} \right. \quad (6)$$

The convolution product that transforms a cross section from temperature T_0 to temperature T defines a Doppler broadening linear operator $\mathfrak{D}\mathfrak{B}_{T_0}^T$:

$$\left\{ \begin{array}{l} \sigma(u, T) = \mathfrak{D}\mathfrak{B}_{T_0}^T [\sigma(u, T_0)] \\ \text{with: } \mathfrak{D}\mathfrak{B}_{T_0}^T [\sigma(u, T_0)] \equiv \int_0^\infty \sigma(x, T_0) \frac{x^2}{2u^2\sqrt{\pi\xi}} \left[e^{-\frac{(u-x)^2}{4\xi}} - e^{-\frac{(u+x)^2}{4\xi}} \right] dx \end{array} \right. \quad (7)$$

If one now considers a multipole representation of a given cross section, the linearity of the Doppler broadening operator yields:

$$\sigma(u, T) = \mathfrak{D}\mathfrak{B}_{T_0}^T [\sigma(u, T_0)] = \sum_j \Re \left[\mathfrak{D}\mathfrak{B}_{T_0}^T \left[\frac{1}{u^2} \frac{r_j}{p_j - u} \right] \right] \quad (8)$$

Each term is then individually and analytically Doppler broadened to a Faddeeva function, or scaled complementary complex error function: $W(z) = e^{-z^2} (1 - \text{erf}(-iz))$.

$$\sigma(u, T) = \mathfrak{D}\mathfrak{B}_{T_0}^T [\sigma(u, T_0)] = \frac{1}{2u^2\sqrt{\xi}} \sum_j \Re \left[ir_j \sqrt{\pi} W(z_j^0) - \frac{r_j}{\sqrt{\pi}} C \left(\frac{p_j}{\sqrt{\xi}}, \frac{u}{2\sqrt{\xi}} \right) \right] \quad (9)$$

where

$$u = \sqrt{E} \quad (10)$$

$$p_j = a_j + ib_j \quad (11)$$

$$z_j^0 = \frac{u - a_j}{2\sqrt{\xi}} - i \frac{b_j}{2\sqrt{\xi}} \quad (12)$$

$$C \left(\frac{p_j}{\sqrt{\xi}}, \frac{u}{2\sqrt{\xi}} \right) = 2p_j \int_0^\infty du' \frac{\exp \left[-\frac{u+u'}{4\xi} \right]^2}{p_j^2 - u'^2}. \quad (13)$$

It can be shown that the function $C \left(\frac{p_j}{\sqrt{\xi}}, \frac{u}{2\sqrt{\xi}} \right)$ is negligible at most energies of interest in reactor analysis [15, 24].

Through this transformation, and neglecting the corrective term C , Doppler broadening a cross section in multipole formalism reduces to evaluating the Faddeeva function at a series of points that depend on the velocity of the incoming neutron, the temperature at which we want to Doppler broaden, and the poles of the multipole representation.

3. Windowed Multipole

There are two major issues with the multipole formalism. The first is one of performance. For every single point, a Faddeeva function for every pole must be evaluated. For some nuclides, the number of poles can be in the tens of thousands. The second one involves the File 3 data in the ENDF-6 specification. For many nuclides, the model chosen may be insufficient to accurately portray the experimental data. For these nuclides, the resonance data is fit to as high accuracy as possible, and then point-wise data from File 3 is added in to correct it. The multipole formalism does not allow for Doppler broadening such data without also using an algorithm like SIGMA1.

By taking advantage of the fact that Equation (5) only fluctuates in proximity to the pole p_j , a lower order approximation can be used far away from p_j . This low order approximation can be a simple curve fit. By selectively choosing which poles need exact evaluation and which poles can be approximated, the windowed multipole formalism was developed [14]. Further, File 3 data can often be approximated with sufficient accuracy with the same form of low order function. This leaves two new problems: how to choose which poles will be evaluated exactly, and how to Doppler broaden the low order function, both of which will be dealt with in the next two sections.

3.1. Windowing Process

Starting with the issue of choosing which poles will be evaluated exactly, a mesh is overlaid on top of the data such that each energy point belongs to a unique cell. These cells will be referred to as "inner windows". This is the blue region in Figure 1. The purpose of this step is to define a domain on which the next set of steps are valid.

Now that this array of windows is constructed, each of them are bracketed by an outer window, shown in red in Figure 1. The purpose of this window is to separate the entire energy range in two. All poles within the outer window will be evaluated exactly at temperature T while all poles outside of this region will be curve fit using $0K$ data. However, the actual "size" of the outer window is irrelevant, and so only the pointer to the first pole and the pointer to the last pole are actually stored. The pointer to the first pole and last pole are fully independent from all other windows. For the inner window, variation in size can have a large penalty. To avoid costly binary searches to locate a given energy value, the inner window is fixed as uniform along a given energy-related variable. Possible choices include equally spaced in energy, momentum, or lethargy. Prior studies have concluded that, of those three, equally spaced in momentum windows are the most beneficial [16].

In this implementation, the inner window size and the curve fit order are nuclide dependent, and the outer window pointers are dependent on which inner window they bracket. This provides many pathways to optimize the problem. If given an inner window size and curve fit order, Algorithm 1 is used.

The actual details of the shrink, expand, and convergence steps can be implementation dependent. In this implementation, a simple bisection algorithm was used. The problem was considered converged once the difference in S_{ow} between cycles only changed the number of poles between $j_{begin,W}$ and $j_{end,W}$ by one, and $\epsilon_{max} < \epsilon_{allowed}$. To find the maximum error in the region requires a global optimizer. The one used is the differential evolution technique [25].

An array of average inner window sizes and curve fit orders were fed into Algorithm 1 and tested on memory and computational efficiency. Outer window sizes were determined by the prescribed desired accuracy. The result is a series of graphs that look like Figure 2 and Figure 3. These were analyzed for each nuclide to find the optimum settings. Computational efficiency was favored for the library generated for this paper.

With these graphs, all three variables have been optimized to yield a highly efficient library. As part of the optimization process, a new algorithm was developed to Doppler broaden the curve fits which is described in the next section.

3.2. Doppler broadening of high order polynomials

During optimization it was noticed that curve fitting was most difficult at lower energies requiring unreasonable amounts of poles or the ability to perform broadening of the fitting function. As such, a series

Figure 1: The construction of windows.

Algorithm 1 Basic Single Run Optimization

```

 $O \leftarrow$  order of curve fit
for Inner Window  $W \in$  Windows do
   $M \leftarrow$  midpoint of  $W$ 
   $S_{iw} \leftarrow$  size of  $W$ 
  Outer Window Size  $S_{ow} \leftarrow$  Initial Guess
  repeat
     $j_{\text{begin}, W} \leftarrow$  index of first pole with real component  $> M - S_{ow}/2$ 
     $j_{\text{end}, W} \leftarrow$  index of last pole with real component  $< M + S_{ow}/2$ 
     $\vec{C} \leftarrow$  least squares curve fit, order  $O$ , of  $\frac{1}{u^2} \sum_{j < j_{\text{begin}} \text{ or } j_{\text{end}} < j} \Re \left[ \frac{r_j}{p_j - u} \right] + \text{Background}$ 
     $\epsilon_{\text{max}} \leftarrow$  maximum error in bounds of  $E \in [M - S_{iw}, M + S_{iw}]$  and  $T \in [T_{\min}, T_{\max}]$ 
    if  $\epsilon_{\text{max}} > \epsilon_{\text{allowed}}$  then
      Expand  $S_{ow}$ 
    else
      Shrink  $S_{ow}$ 
    end if
  until Converged
end for

```

Figure 2: ^{238}U , Memory Requirements

Figure 3: ^{238}U , Time Per Evaluation

expansion was performed at $u = 0$. For this particular function, a Laurent expansion will converge if and only if $|u| < |p_j|$, which is most often the case in thermal energies.

$$\begin{aligned}\frac{1}{u^2} \left[\frac{r_j}{p_j - u} \right] &= \frac{r_j}{p_j u^2} + \frac{r_j}{p_j^2 u} + \frac{r_j}{p_j^3} + \frac{r_j u}{p_j^4} + \mathcal{O}(u^2) \\ \frac{1}{u^2} \left[\frac{r_j}{p_j - u} \right] &\approx \sum_{n \geq -2}^N a_n u^n\end{aligned}\tag{14}$$

Many individual terms of this sum have well-known Doppler broadened forms [26]. The second and third terms correspond to a $1/V$ and a constant cross section, respectively. The rest of this section is dedicated to the development of an algorithm that will Doppler broaden Equation (14), when said equation is of an arbitrary order.

Let the curve fit component of the cross section, $\sigma(u, T_0)$, take the form of the expansion in Equation (14). Applying the convolution product Doppler broadening linear operator to the polynomial yields:

$$\begin{cases} \sigma(u, T) = \mathfrak{D} \mathfrak{B}_{T_0}^T [\sigma(u, T_0)] = \sum_{n \geq -2} a_n \mathfrak{D} \mathfrak{B}_{T_0}^T [u_{T_0}^n] = \sum_{n \geq -2} a_n \mathfrak{D}_{n+2} \\ \text{where :} \\ \mathfrak{D}_{n+2} \equiv \mathfrak{D} \mathfrak{B}_{T_0}^T [u_{T_0}^n] = \int_0^\infty \frac{x^{n+2}}{2u^2 \sqrt{\pi \xi}} \left[e^{-\frac{(u-x)^2}{4\xi}} - e^{-\frac{(u+x)^2}{4\xi}} \right] dx = \left(\frac{e^{-\frac{u^2}{4\xi}}}{u^2 \sqrt{\pi \xi}} \right) \int_0^\infty x^{n+2} \sinh\left(\frac{ux}{2\xi}\right) e^{-\frac{x^2}{4\xi}} dx \end{cases}\tag{15}$$

The following linear recurrence relation enables us to calculate the $\mathfrak{D}_n$ from elemental Gaussian profiles and error functions:

$$\begin{cases} \forall n > 0, \mathfrak{D}_{n+2} = [2\xi(2n+1) + u^2] \mathfrak{D}_n - (2\xi)^2 n(n-1) \mathfrak{D}_{n-2} \\ \mathfrak{D}_2 = [2\xi + u^2] \mathfrak{D}_0 + \frac{2}{u} \sqrt{\frac{\xi}{\pi}} e^{-\frac{u^2}{4\xi}} \\ \mathfrak{D}_1 = \frac{1}{u} \\ \mathfrak{D}_0 = \frac{1}{u^2} \operatorname{erf}\left(\frac{u}{2\sqrt{\xi}}\right) \end{cases}\tag{16}$$

This method constitutes a complete analytical recurrence scheme for exact Doppler broadening of high order polynomials - generalized to include terms in a_{-2} and a_{-1} - with minimal computational cost. The two most expensive operations, the error function and the exponential, only need to be evaluated once. This technique thus provides a highly accurate and efficient Doppler broadening method for curve fits, so long as the curve fits themselves are not only sufficiently accurate in their domain but also in proximity to the domain.

With this component, all of the curve fits formed in Section 3.1 can be Doppler broadened. This also provides a slight advantage to the accuracy of windowed multipole over the general multipole formalism. At very low energies, the function C in Equation (13) is no longer neglected since it is now captured by the Doppler broadened curve fit.

For efficiency reasons, this algorithm is only used when the accuracy requirements warrants it. The optimization algorithm determines if broadening the curve fit is needed which is often only in the thermal energy region.

4. Windowed Multipole Library

A windowed multipole library was generated for most of the nuclides of the BEAVRS fresh core benchmark which utilizes a total of 90 nuclides. In the ENDF/B-VII.1 library that was used in this paper, certain nuclides are only given in a point-wise form thus making it impossible to extract the resonance parameters needed for the multipole formalism. This was historically done for lighter nuclides where very few points were needed to represent the cross sections or for nuclides requiring resonance models not yet supported.

Newer evaluations will have less and less of these point-wise only nuclides but a few of the lighter nuclides will surely remain. Additionally, a few evaluations contain unphysical discontinuities or oscillations (e.g. point-wise resonances) in file 3 causing issues with the windowed multipole optimization. These issues are more common with older evaluations where newer resonance formats that properly account for interference effects were not available. The cause of failure for the two Vanadium nuclides has not yet been determined. Table 1 contains the list of all nuclides excluded from the windowed multipole library. For the results presented in the next section these nuclides were kept in the point-wise format at the closest pre-processed temperature. It should be noted that with a few exceptions (ie. H-1, B-10, B-11) all nuclides on this list are only present as trace elements in the material compositions of the benchmark. Additionally, the few nuclides on this list present in significant quantity have minimal temperature feedback due to the smoothness of their cross section.

Nuclide	Reason
H-1	Point-wise
H-2	Point-wise
He-3	Point-wise
He-4	Point-wise
B-10	Point-wise
B-11	Point-wise
Carbon (natural)	Point-wise
N-14	Point-wise
N-15	Point-wise
O-16	Point-wise
O-17	Point-wise
P-31	Point-wise
S-36	Point-wise
Ar-36	File 3
Ar-38	File 3
Ar-40	File 3
Ca-46	Point-wise
V-50	Incomplete at time of run
V-51	Incomplete at time of run
Nb-93	File 3

Table 1: Nuclides not in windowed multipole format.

In total, 70 nuclides were processed using the windowed multipole method in the resolved resonance range. The complete list is available in Appendix B. Temperature dependent probability table data of the unresolved range and temperature independent data for the fast range was taken from the ENDF library. The library was optimized first and foremost for efficiency using the configuration found in Table 2.

Parameter	Value
Max Relative Error	0.1 %
Threshold Error	10^{-5} b

Table 2: Library Configuration

The threshold error is the minimum absolute cross section value for which errors were checked. There are a few exceptions where the threshold error for the capture cross section was increased slightly in a given window due to sharp unphysical discontinuities in the data for Fe and Cu. The strong discontinuities hampered the low order fitting process since larger errors were observed near the discontinuity. Since step

changes are not expected in nuclear data, the threshold error was relaxed near the discontinuities in order for the algorithm to complete.

5. Results

Two tests were run to analyze the accuracy and efficiency of the windowed multipole method. The first was focused on investigating the accuracy of windowed multipole using a single assembly case. The total fission rate and ^{238}U absorption rate were tallied for both an ACE library and the windowed multipole library described in the previous section. This case used the MIT Faddeeva function to evaluate the Faddeeva function in Equation (9) [27]. The details of this analysis are presented in Section 5.1. The last run was performed on the full core BEAVRS benchmark [28] in OpenMC on an emulated CPU in Callgrind. Callgrind is a tool that allows the measure of cache misses along with other performance metrics [29]. This tool was used to quantify the performance of windowed multipole. These results are described in Section 5.2.

5.1. Accuracy

The first test case consists of a high enrichment assembly from the BEAVRS benchmark performed with the ENDF71X ACE library distributed with the current MCNP release [30], and the aforementioned multipole library. Both runs were done at 600K, the nearest temperature of the ACE library to the hot zero power conditions. The total neutron production rate, $\nu\Sigma_f\phi$, and the ^{238}U capture rate, $\Sigma_{a,U-238}\phi$ were tallied and compared, both in an energy integrated sense and in a volumetrically integrated sense. The bins used for each tally are described in Appendix A.

The single assembly is a 3.2% enriched assembly with no burnable absorber with reflective boundary conditions. The run parameters are listed in Table 3 and the resulting eigenvalues in Table 4.

Parameter	Value
Inactive Batches	200
Active Batches	1000
Neutrons Per Batch	4 million

Table 3: Assembly Accuracy Run Configuration

Data	Eigenvalue	N/s Inactive	N/s Active
ACE	1.21204 ± 0.00001	24138	14841
WMP	1.21207 ± 0.00001	21063	9337

Table 4: Assembly Results

The results in Table 4 show great agreement between the two libraries by being within 3 standard deviations of each other with very tight statistical convergence. Despite being based on the same evaluation, the processing and temperature broadening introduces slight differences which can explain the observed discrepancies. In addition to eigenvalue, Figure 6 and Figure 7 present the volume integrated capture and fission rates in the assembly.

Neither reaction rates, with corresponding 1 sigma statistics, ever exceed the $\pm 0.1\%$ accuracy limit used when generating both libraries for the given energy bins used in this study. Notably, the fission production rate is always within $\pm 0.04\%$, with the most inaccurate bin being the one between 0.625 eV and 1 eV. This provides great confidence in the accuracy of the generated library and the optimization algorithm used.

Additionally, spatial distribution of reaction rates were also analyzed by comparing energy integrated tallies over a pin size mesh. These results are shown in Figure 4 and Figure 5. All pins reaction rates are within $0.12\% \pm 0.05$ between the two libraries. The distribution is quite random indicating no clear bias once again supporting the accuracy of the multipole library. It should also be noted that the standard

Figure 4: ²³⁸U Absorption Rate, Assembly, Volume Integrated

Figure 5: Fission Neutron Production Rate, Assembly, Volume Integrated

deviation provided are from single runs and thus do not properly capture the correlation effects that exist between consecutive fission banks [31].

Figure 6: ^{238}U Absorption Rate, Assembly, Energy Integrated

Figure 7: Fission Neutron Production Rate, Assembly, Energy Integrated

5.2. Performance

In the prior section, the number of neutrons per second wall clock time for the assembly simulation were reported in Table 4. A similar study was performed for performance analysis on the full core BEAVRS benchmark with 4 million neutrons per cycle with results reported in Table 5.

Data	Eigenvalue	N/s Inactive	N/s Active
ACE	0.99733 ± 0.00003	23191	15488
WMP	0.99735 ± 0.00003	20248	14753

Table 5: BEAVRS Results

These tables indicate a decrease in performance anywhere from 4.7% to 37.1% depending greatly on the size of the problem. This performance is quite encouraging considering the large memory reduction and the added capability of calculating cross sections at any temperature. In this section, the performance of the windowed multipole is compared to a single and exact temperature ACE lookup approach in a

more comprehensive way. The BEAVRS full core benchmark was repeated with the configuration shown in Table 6. The goal was not to fully resolve the problem but to provide a realistic problem size-wise that can be emulated in Callgrind in a realistic period of time, so the number of neutrons was kept small.

Parameter	Value
Inactive Batches	200
Active Batches	200
Neutrons Per Batch	500

Table 6: Performance Run Configuration

The entire cross section calculation routine in OpenMC is stored in the `calculate_xs` function. Specific instrumentation of interest was the total estimated clock cycle cost, the first level cache misses, and the last level cache misses. These are shown in Tables 7, 8, and 9. Comparing the clock cycles required, the multipole nuclide library provides a 39.6% performance loss in the `calculate_xs` function. This translates as a 16.3 % decrease in transport operator performance and a 7.9% performance decrease on the overall eigenvalue calculator.

Function	Clock Cycles		
	ACE	WMP	% Diff.
Total	4257	5946	39.6 %
Binary Search	951	741	-22.1 %
Nuclide Cross Section	2255	4148	83.9 %
/ Multipole	—	2330	—
// Faddeeva	—	733	—
// Broaden Curve Fits	—	107	—

Table 7: Cycle Estimation per call of `calculate_xs` (/ and // indicate embedded sub-functions)

For the L1 misses in Table 8, windowed multipole increases the total number of misses by 40.4%, mostly due to the increased access of small objects. The binary search cost decreased due to the decreased data size. Most of L1 misses in the windowed multipole routine itself involve the first access of data from an array (for example, the pointers to the first and last pole, curve fit data) once a window index is known.

Function	L1 Miss		
	ACE	WMP	% Diff.
Total	81.9	115	40.4%
Binary Search	4.93	4.50	-8.7%
Nuclide Cross Section	66.2	99.6	50.4%
/ Multipole	—	40.9	—
// Faddeeva	—	7.55	—
// Broaden Curve Fits	—	1.08	—

Table 8: L1 Misses per call of `calculate_xs` (/ and // indicate embedded sub-functions)

Lastly, the last-level misses in Table 9 are substantially decreased from ACE lookup to windowed multipole, reflecting a significant decrease in the access of large objects during calculation. A majority of the last-level misses in the windowed multipole routine (59.1%) involve access to the rather large curve fit data arrays, which can often be the largest single object in a library. The Faddeeva function and the curve fit broadening function contributed negligible amounts to the number of last-level misses.

Function	LL Miss		% Diff.
	ACE	WMP	
Total	4.72	1.63	-65.5%
Binary Search	0.489	0.061	-87.5%
Nuclide Cross Section	4.23	1.56	-63.1%
/ Multipole	—	0.569	—
// Faddeeva	—	0.00779	—
// Broaden Curve Fits	—	0.000967	—

Table 9: Last-level Misses per call of `calculate_xs` (/ and // indicate embedded sub-functions)

The multipole library shows great benefit in reducing data movement by substantially reducing the library size when compared to a single temperature ACE library. In temperature dependent calculations, the multipole calculations would remain essentially the same, while point-wise ACE libraries would have even move last-level misses to access multiple temperature libraries for the interpolation.

6. Conclusions

With the advent of greater computing power, Monte Carlo methods gain in popularity for reactor design and analysis. To do so, however, currently requires large memory requirements or pre-computing costs for properly accounting of temperature variations everywhere in the core. In this paper, a new approach for approximating the Doppler broadening of nuclear cross sections is presented which requires minimal memory for representing the full range of temperatures required in reactor analysis and with comparable efficiency to a single ACE file temperature lookup.

The windowed multipole method is an approximation of the multipole formalism presented by Hwang [15] that reduces the number of costly Faddeeva function evaluations while preserving the needed accuracy. Additionally, a novel high order polynomial Doppler broadening approach was developed that accurately captures the low-energy free gas thermal scattering effects on curve-fitted background cross section terms.

A windowed multipole library of 70 nuclides was generated and tested on an assembly and full 3D core of the BEAVRS benchmark. Results were compared with an ACE library of similar accuracy on a single assembly benchmark. The total neutron production rates and ^{238}U capture rates were shown to be within library accuracy and statistical error.

In testing the performance, the cross section calculator was found to be slower by 40% compared to a single temperature ACE table lookup. This resulted in a change in the overall run time of 7.9% on the full core BEAVRS simulation. The algorithm also reduces the last-level cache misses by as much as 65%, for a reduction of around 3 last-level cache misses per cross section calculation. Further study on the algorithm in a line-by-line basis may be able to improve the performance further as well as investigations of alternative Faddeeva function libraries.

Acknowledgements

This research was supported in part by the Consortium for Advanced Simulation of Light Water Reactors (CASL), an Energy Innovation Hub for Modeling and Simulation of Nuclear Reactors under U.S. Department of Energy Contract No. DE-AC05-00OR22725, and by DOE's Center for Exascale Simulation of Advanced Reactors (CESAR). The first author is also partially supported under a Department of Energy Nuclear Energy University Programs Graduate Fellowship, while the second author was partially supported by the joint Ecole Polytechnique - EDF "Chaire Énergies Durables" grant.

References

- [1] P. Romano, B. Forget, Parallel fission bank algorithms in Monte Carlo criticality calculations, Nucl. Sci. Eng. 170.
- [2] P. K. Romano, A. R. Siegel, B. Forget, K. Smith, Data decomposition of monte carlo particle transport simulations via tally servers, Journal of Computational Physics 252 (0) (2013) 20 – 36. doi:http://dx.doi.org/10.1016/j.jcp.2013.06.011. URL <http://www.sciencedirect.com/science/article/pii/S002199911300435X>
- [3] J.Wagner, S.Mosher, T.Evans, D.Peplow, J.Turner, Hybrid and parallel domain-decomposition methods development to enable monte carlo for reactor analyses, Progr. Nucl. Sci. Technol. 2.
- [4] N.Horelik, A.Siegel, B.Forget, K.Smith, Improved cache performance in monte carlo transport calculations using energy banding, Parallel Comput. URL <http://dx.doi.org/10.1016/j.parco.2014.10.001>
- [5] D. Cullen, C. Weisbin, Exact doppler broadening of tabulated cross sections., Nucl. Sci. Eng. 60 (3).
- [6] J.Walsh, Computational methods for efficient nuclear data management in Monte Carlo neutron transport simulations, Tech. rep., Massachusetts Institute of Technology (2013).
- [7] A.Siegel, K.Smith, K.Felker, P.Romano, B.Forget, P.Beckman, Improved cache performance in monte carlo transport calculations using energy banding, Comp. Phys. Comm. 185.
- [8] T. Trumbull, Treatment of Nuclear Data for Transport Problems Containing Detailed Temperature Distributions, Nucl. Tech. 156 (2006) 75–86.
- [9] S.W.D.Hart, G.I.Maldonado, C.Celik, L.C.Leal, Problem-dependent doppler broadening of continuous-energy cross sections in the KENO Monte Carlo computer code, in: Proceedings of PHYSOR 2014, 2014.
- [10] J. Conlin, W. Ji, J. Lee, W. Martin, Pseudo Material Construct for Coupled Neutronic-Thermal-Hydraulic Analysis of VHTGR, Trans. Am. Nucl. Soc. 92 (2005) 225–227.
- [11] G. Yesilyurt, W. Martin, F. Brown, On-the-Fly Doppler Broadening for Monte Carlo Codes, Nucl. Sci. Eng. 171 (3) (2012) 239–257.
- [12] T. Viitanen, J. Leppanen, Explicit treatment of thermal motion in continuous-energy Monte Carlo tracking routines, Nucl. Sci. Eng. 171 (2012) 165–173.
- [13] T.Viitanen, J.Leppanen, B.Forget, Target motion sampling temperature treatment technique with track-length estimators in OpenMC - preliminary results, in: Proceedings of PHYSOR 2014, 2014.
- [14] B. Forget, S. Xu, K. Smith, Direct doppler broadening in monte carlo simulations using the multipole representation, Annals of Nuclear Energy 64 (2014) 78–85.
- [15] R. N. Hwang, A rigorous pole representation of multilevel cross sections and its practical applications, Nucl. Sci. Eng. 96 (1987) 192–209.
- [16] C. Josey, B. Forget, K. Smith, Efficiency and accuracy evaluation of the windowed multipole direct doppler broadening method, in: Proceedings of PHYSOR 2014, 2014.
- [17] B.R.Herman, B.Forget, K.Smith, Progress toward Monte Carlo-thermal hydraulic coupling using low-order nonlinear diffusion acceleration methods, Ann. Nucl. Energy (2014) 75–86. URL <http://dx.doi.org/10.1016/j.anucene.2014.10.029>
- [18] M. Chadwick, M. Herman, P. Obložinský, M. E. Dunn, Y. Danon, A. Kahler, D. L. Smith, B. Pritychenko, G. Arbanas, R. Arcilla, et al., ENDF/B-VII.1 nuclear data for science and technology: cross sections, covariances, fission product yields and decay data, Nuclear Data Sheets 112 (12) (2011) 2887–2996.
- [19] F. H. Frohner, Evaluation and analysis of nuclear resonance data, Jeff Report (18).
- [20] A. M. Lane, R. G. Thomas, R-matrix theory of nuclear reactions, Reviews of Modern Physics 30 (2) (1958) 257.
- [21] R. C. Block, Y. Danon, F. Gunsing, R. C. Haight, Neutron cross section measurements, Handbook of Nuclear Engineering 1 (1) (2010) 1–74.
- [22] C. Jammes, R. N. Hwang, Conversion of single- and multilevel breit-wigner resonance parameters to pole representation parameters, Nuclear Science and Engineering 134 (2000) 37–49.
- [23] D. E. Cullen, C. R. Weisbia, Exact doppler broadening of tabulated cross sections, Nuclear Science and Engineering 60 (1976) 199–229.
- [24] R. N. Hwang, An Extension of the Rigorous Pole Representation of Cross Sections for Reactor Applications, Nucl. Sci. Eng. 111 (2) (1992) 113–131.
- [25] R. Storn, K. Price, Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces, Journal of Global Optimization 11 (4) (1997) 341–359. doi:10.1023/A:1008202821328.
- [26] A. Hebert, Applied reactor physics, Presses inter Polytechnique, 2009.
- [27] S. G. Johnson, Faddeeva package - abinitio (December 2013). URL http://ab-initio.mit.edu/wiki/index.php/Faddeeva_Package
- [28] N. Horelik, B. Herman, B. Forget, K. Smith, MIT BEAVRS: Benchmark for evaluation and validation of reactor simulations, International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, Sun Valley, ID.
- [29] J. Weidendorfer, M. Kowarschik, C. Trinitis, A tool suite for simulation based analysis of memory access behavior, in: Computational Science-ICCS 2004, Springer, 2004, pp. 440–447.
- [30] J. T. Goorley, M. R. James, T. E. Booth, F. B. Brown, J. S. Bull, L. J. Cox, J. W. Durkee Jr, J. S. Elson, M. L. Fensin, R. A. Forster III, et al., Initial mcnpp6 release overview-mcnpp6 version 1.0, Tech. rep., Los Alamos National Laboratory (LANL) (2013).
- [31] E.Dumontel, F.Malvagi, A.Zoia, A.Mazzolo, D.Artusio, C.Dieudonne, C. de Mulatier, Particle clustering in Monte Carlo criticality simulations, Ann. Nucl. Energy 63 (2014) 612–618.

A. Energy Bins

Bin Index	Bottom of Bin	Top of Bin
1	10^{-5} eV	3×10^{-2} eV
2	3×10^{-2} eV	0.1 eV
3	0.1 eV	0.3 eV
4	0.3 eV	0.625 eV
5	0.625 eV	1 eV
6	1 eV	4 eV
7	4 eV	6 eV
8	6 eV	10 eV
9	10 eV	25 eV
10	25 eV	50 eV
11	50 eV	100 eV
12	100 eV	1 keV
13	1 keV	10 keV
14	10 keV	100 keV
15	100 keV	1 MeV
16	1 MeV	20 MeV

Table A.10: Energy Bins Used For All Tallies

B. Nuclides Processed

Element	Isotope List
Aluminium	27
Silicon	28, 29, 30
Sulfur	32, 33, 34
Calcium	40, 42, 43, 44, 48
Titanium	46, 47, 48, 49, 50
Chromium	50, 52, 53, 54
Manganese	55
Iron	54, 56, 57, 58
Nickel	58, 60, 61, 62, 64
Copper	63, 65
Zirconium	90, 91, 92, 94, 96
Molybdenum	92, 94, 95, 96, 97, 98, 100
Silver	107, 109
Cadmium	106, 108, 110, 111, 112, 113, 114, 116
Indium	113, 115
Tin	112, 114, 115, 116, 117, 118, 119, 120, 122, 124
Uranium	234, 235, 238

Table B.11: Nuclides in the BEAVRS core that were processed into windowed multipole format

LaTeX Source Files

[Click here to download LaTeX Source Files: JCP Multipole paper 2014 \(1\).zip](#)