Journal of Computational Physics ••• (••••) •••–•••

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Windowed multipole for cross section Doppler broadening

C. Josey, P. Ducru, B. Forget *, K. Smith

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

ARTICLE INFO

Article history:
Received 18 May 2015
Received in revised form 4 August 2015
Accepted 8 August 2015
Available online xxxx

Keywords:
Monte Carlo
Multipole
Doppler broadening

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 ^{238}U 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.

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 0 K. 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

* Corresponding author.
E-mail addresses: cjosey@mit.edu (C. Josey), p_ducru@mit.edu (P. Ducru), bforget@mit.edu (B. Forget), kord@mit.edu (K. Smith).

http://dx.doi.org/10.1016/j.jcp.2015.08.013
0021-9991/

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 1 GB 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 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 [8–10]. Such schemes require storing massively large libraries or performing costly pre-processing to encapsulate the temperature range of interest in operating reactors (300–3000 K) 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 and 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 $\mathcal{C}^1$ 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:

$$\sigma_{cc'} = \pi \lambda\!\!\!^{-2} g_c \left|\delta_{cc'} - U_{cc'}\right|^2 \tag{1}$$

where $\lambda\!\!\!^{-}$ 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) $\boldsymbol{U}$ matrix [19].

Please cite this article in press as: C. Josey et al., Windowed multipole for cross section Doppler broadening, J. Comput. Phys. (2015), http://dx.doi.org/10.1016/j.jcp.2015.08.013