Transactions of the Korean Nuclear Society Autumn Meeting
Gyeongju, Korea, October 29-30, 2026

Positive Angular Representations of Multigroup Scattering Kernels from MC^2-3 Libraries

Minjae Lee*, Jong-Hyuck Won, Yeonsang Jung

Korea Atomic Energy Research Institute, 989-111 Daedeok-daero, Yuseong-gu, Daejeon, Korea, 34057
**Corresponding author: lmj@kaeri.re.kr*

***Keywords :** Multigroup Monte Carlo, Anisotropic Scattering, Consistent-P1 Correction, Maximum Entropy, ISOTXS*

1. Introduction

One of the most widely used code systems for fast reactor analysis is the MC²-3 and DIF3D path: MC²-3 [1] condenses continuous-energy data into a multigroup ISOTXS library, and DIF3D [2] solves the core with it. Replacing only the transport solver with a multigroup Monte Carlo (MGMC) code keeps the established group constants while removing the spatial homogenization the deterministic solver needs. The obstacle in that substitution is the scattering kernel. The library stores the angular distribution of each group-to-group transition as truncated Legendre moments, and a Monte Carlo code must sample an actual direction from them, which requires a non-negative density.

The negativity of the reconstructed kernel has two distinct layers. The first is ordinary truncation: a strongly anisotropic but perfectly physical distribution, cut at P3, oscillates below zero. Moments of this kind still belong to some positive distribution, so a positive representation can in principle preserve all of them. The second layer arises from the consistent-PN approximation used to formulate the multigroup transport equation. The consistent-PN approximation preserves every scattering matrix in full and modifies only its diagonal entry, compensating for the multigroup approximation's use of a single scalar-flux-weighted total cross section in place of the moment-weighted total cross section that the exact moment balance requires [3]. The library itself shows this: wherever elastic scattering is isotropic in the CM frame, the group-summed mean cosine equals the kinematic value 2/(3A) exactly, but in self-shielded resonance groups it departs by an order of magnitude and even turns negative (Fe-56 near 3 keV), which is impossible for A = 56 elastic scattering (Fig. 1); the departure above 1 MeV is real forward peaking of the CM distribution, not part of the correction. For a few transitions the resulting moment vector cannot be reproduced by any non-negative distribution at all.

This work measures how accurately the sampling treatments used in production codes reproduce that operator. A per-transition linear-programming test decides in advance whether an exact positive representation exists, using no information about the problem being solved.

2. Methods and Results

Throughout, $a_l = \Sigma_{s,l}/\Sigma_{s,0}$ denotes the normalized moments of a transition $g \to g'$ and

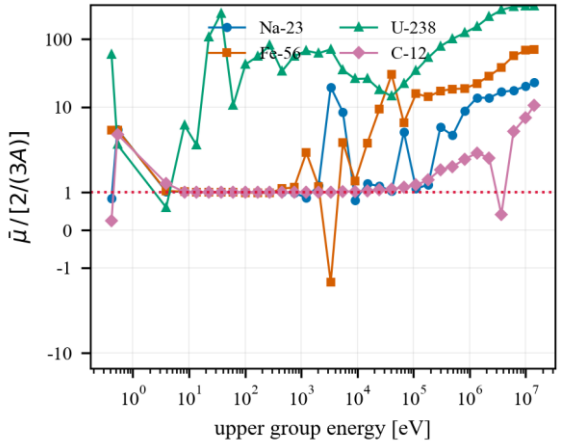


Fig. 1. Group-summed elastic mean cosine over the CM-isotropic value 2/(3A)

$$f(\mu) = 1/2 + \Sigma_\ell(2\ell+1)/2 \cdot a_\ell P_\ell(\mu), \ell = 1 \dots L \quad (1)$$

the density truncated at order L. A P3 library defines a transport operator with $a_l = 0$ for $l > 3$; that operator, not any smooth reconstruction of it, is what the deterministic solver sees, and it is taken as the reference here. A representation is judged by the moments it delivers to the sampler: it should return $a_1 \dots a_3$ unchanged and add as little as possible above them. Power moments $m_n = \int \mu^n f(\mu) d$ are used where the construction is stated in that form.

2.1. Moment-preserving discrete angles (MCNP)

The multigroup mode of MCNP represents the kernel by discrete angles [4],

$$f(\mu) = \Sigma_i \omega_i \delta(\mu - \mu_i),\ \Sigma_i \omega_i \mu_i^n = m_n \quad (2)$$

with i = 1...N and 2N conditions n = 0...2N−1 fixing the nodes and weights. A P3 library gives N = 2, a two-point Gauss quadrature on $m_0 \dots m_3$. This is the extremal measure of the truncated moment problem, so its a_4 always sits at one end of the attainable interval defined in Section 2.3 (confirmed to four decimal places in every realizable case tested). When a node falls outside [−1, 1] the construction fails and has no solution.

2.2. Tabular representation with rejection (OpenMC)

OpenMC converts the expansion to a 33-point equally spaced table and rejects the negative part [5]. Mathematically this is the renormalized positive part

$$f_+(\mu) = max(f(\mu), 0) / \int_{-1}^{1} max(f(\mu'), 0) d\mu' \quad (3)$$

called clipping below; the two implementations agree within 9 pcm in every case tested. Clipping preserves no moment, and its error grows in proportion to the discarded negative area (correlation 0.97 over the eight transitions of the data set, six of which are reported here).

2.3. Maximum entropy with moment constraints

Zhu et al. [6] generate ACE-format multigroup libraries with the maximum-entropy density in exponential form,

$$f(\mu) = exp\left(\Sigma_\ell \lambda_\ell P_\ell(\mu)\right), \ell = 0 \ldots L \quad (4)$$

whose multipliers λ_ℓ follow from L+1 nonlinear equations. The form is positive by construction and, for the mixtures examined there, tracks the P5 Legendre curve closely. Two properties matter for the present problem. The moments of Eq. (4) are integrals of an exponential, so what the sampler receives carries the residual of the nonlinear solve, and the solve itself has a solution only when the moment vector lies in the interior of the set reachable by non-negative densities. Neither point is a limitation for kernels that are merely truncated, but the corrected library of Section 1 contains vectors outside that set.

The representation used here is therefore a histogram of K = 32 equal bins with masses w_k, maximizing the entropy

$$S = -\Sigma_k w_k ln(w_k / \Delta\mu) \quad (5)$$

subject to $w_k \geq 0$, $\Sigma_k w_k = 1$ and

$$\Sigma_k w_k \bar{P}_{\ell,k} = a_\ell, \ell = 1,2,3 \quad (6)$$

where $\bar{P}_{\ell,k} = (1/\Delta\mu) \int_{bin\ k} P_l(\mu) d\mu$ is the bin-averaged Legendre polynomial. Because the constraints are written on the same basis the sampler uses, the moments of the sampled histogram are the constrained values to machine precision, not to the tolerance of a solve. The same linearity settles realizability, which for moments on [−1, 1] is a truncated Hausdorff problem: a vector is realizable when Eq. (6) admits a non-negative solution, and the attainable interval of any higher moment is

$$[a_4^-, a_4^+] = \{ \min, \max \} \Sigma_k w_k \bar{P}_{4,k} \quad (7)$$

over the same feasible set. Both are linear programs of 32 variables and cost a few milliseconds per transition. When the library vector is unrealizable it is projected to the nearest realizable point, holding a_1 as an equality and conceding the highest orders first.

Because the reference operator has $a_4 = 0$ while any positive distribution carries a non-zero a_4, the injected fourth moment is the leading bias of every positive representation. The variant proposed here therefore adds one constraint: pin a_4 to the point of the interval of Eq. (7) nearest zero. The interval follows from a_1..a_3 alone, so the rule uses no information about the geometry and can be evaluated per transition at library-generation time; the intervals of the test cases are reported in Section 3.1. Extending the constraint to a_5 was tested and gives no further gain, because suppressing the fifth moment enlarges the sixth and higher.

2.4. Negative-weight sampling

The remaining option is to sample the cosine from the normalized absolute density and carry the sign in the particle weight [7],

$$p(\mu) = |f(\mu)| / C, \quad w' = w \cdot C \cdot \operatorname{sgn} f(\mu), \quad C = \int |f| d\mu \quad (8)$$

which reproduces the operator exactly, negative lobes included, at the cost of a weight magnitude that grows by $C > 1$ at every collision. The divergence of the power iteration is suppressed by cancelling fission-source weights on a spatial mesh [8]. This treatment is unaffected by the negativity of the kernel and serves as the Monte Carlo check of the deterministic reference.

3. Numerical Results

3.1. Test problems

Test moments were taken from the ISOTXS library used in PGSFR core design [9] (ANL-33 group structure, P5). Macroscopic kernels were assembled for a fuel region (H01) and a sodium region (H10), and six transitions were selected to span the range of negative-lobe fraction and moment structure (Table I). Five are realizable; A6 is not, and it belongs to the resonance-group down-scatter family identified in Fig. 1.

Table I: Test cases from the PGSFR ISOTXS library (fuel)

ID	$g \to g'$	$\Sigma_{s,0}$	neg. [%]	a_1	a_2	a_3
A1	18→18	1.0652	20.7	0.589	0.535	0.566
A2	18→18	0.7609	3.3	0.298	0.345	0.323
A3	1→1	0.0721	24.5	0.829	0.684	0.574
A4	30→30	0.3346	1.7	0.007	−0.146	−0.148
A5	18→19	0.1888	2.8	−0.393	−0.102	0.052
A6	18→19	0.0911	4.3	−0.399	−0.133	0.056

Because the modified kernel destroys any analytical criticality condition, k = 1 was restored as a design condition, following the benchmark construction of the negative-weight verification study. The host is a one-group homogeneous slab with vacuum boundaries using the cross sections of Problem 37 of Ref. [10], and for each case the half-thickness was searched until a discrete-ordinates (S_N) solution of the same moments gave k = 1; the critical half-thickness is 0.88 to 1.07 transport mean free paths. The S_N reference uses Richardson extrapolation of 800- and 1600-cell diamond-difference solutions at S_{1024}, with a verified discretization residual below 0.1 pcm. All Monte Carlo runs used the in-house MGMC code written for this study; the negative-weight treatment reproduced k = 1 within 1.1σ in all runnable cases, which closes the loop between the two solvers.

The two ends of that range are drawn in Fig. 2. In A1 the P3 expansion is forward peaked and dips below zero

over about a fifth of the cosine range, once around μ = 0.25 and again near μ = −1. In A5 the mean cosine is negative, the peak sits in the backward hemisphere and the negative part is small. The four positive representations handle these differently. Clipping follows the expansion wherever it is positive and simply removes the rest, so it keeps the shape but loses the moments. The maximum-entropy histogram is the flattest curve that still returns $a_1..a_3$, which in A1 means a broad spike against μ = 1 rather than the smooth rise of the expansion. Constraining a_4 splits the mass into two narrower groups, moving the representation toward the two-point quadrature drawn on the right axis; the pair of discrete angles is the limit of that trend, all mass at two cosines. None of the four looks like the P3 curve, and Section 3.2 shows this distinction is less significant than it initially appears: the solver sees only the moments.

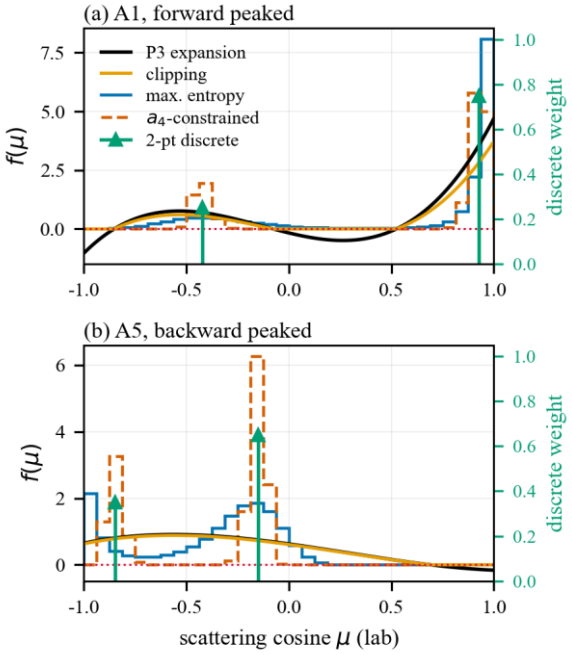


Fig. 2. Angular density of the P3 expansion and of its positive representations, cases A1 and A5

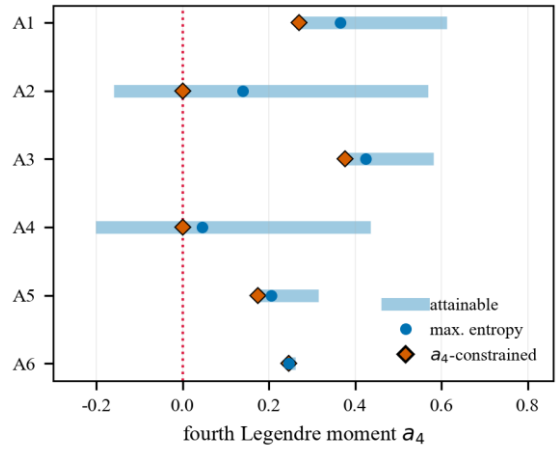


Fig. 3. Attainable interval of a_4 with $a_1 \ldots a_3$ fixed

Fig. 3 gives the attainable interval of a_4 from Eq. (7) for each case, with the fourth moment carried by the plain and the constrained histograms. In A2 and A4 the interval contains the operator value a_4 = 0 and the constraint is met exactly. In the other four it does not, and the distance from zero to the near edge is a floor on the bias of any positive representation, no matter how it is built. For A6 the interval is 0.025 wide, so the choice of a_4 inside it changes almost nothing.

3.2. Multiplication factor

In A6, one node of Eq. (2) falls at μ = −1.12, outside the physical range. The two-point set is then built from $m_0..m_2$ instead, which leaves +236 pcm; collapsing to a single angle would give the one-point value, −865 pcm. Table II and Fig. 4 give the bias of each representation at 5×10^4 particles per generation (1σ of 15 to 27 pcm). Clipping is the worst treatment where the negative lobe is large, reaching +2353 pcm for A3. The single discrete angle is worse still in five of the six cases, because it preserves nothing beyond the mean cosine. All representations that preserve $a_1..a_3$ agree within about 100 pcm of each other, reflecting the fall of the sensitivity dk/da_l by a factor of three to four per order. Only the negative-weight treatment brackets zero within its error bars in every case (Fig. 4).

Table II: Bias against the S_N reference [pcm], 1σ = 15–27 pcm

ID	neg. wt.	max. ent.	a_4-constr.	2-pt	1-pt	clip
A1	−24	+218	+191	+228	−1709	+1622
A2	−12	+53	−20	−43	−2335	+49
A3	+56	+308	+326	+334	−428	+2353
A4	+9	+8	+6	−104	−1205	−165
A5	−30	+111	+37	+53	−1061	−356
A6	+29	+155	+144	+236	−865	−557

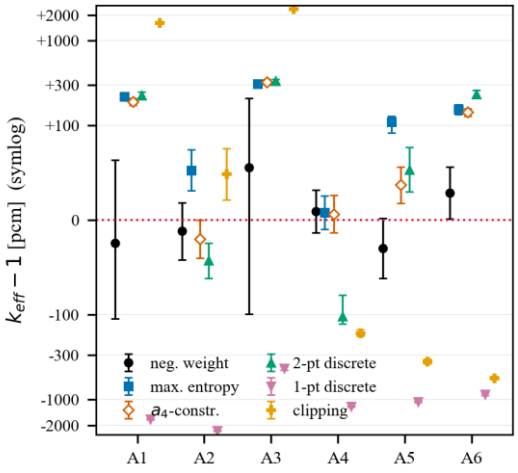


Fig. 4. Bias of each representation against the S_N reference

Differences of a few pcm were resolved by repeating the cases where a_4 = 0 is attainable with forty times the statistics (5×10^5 particles, 2,000 active generations, 1σ ≈ 3 pcm). For A2 the maximum-entropy bias is +55.6 ± 3.1

pcm, an 18σ effect matching the S_N prediction of +54.8 pcm; constraining the fourth moment removes it, +0.7 ± 3.2 pcm, indistinguishable from the exact negative-weight result of +0.3 ± 4.7 pcm.

3.3. Flux shape

Fig. 5 compares the S_N scalar flux of each representation against the reference for the worst case A3, both solutions normalized to unit sum so that only the shape is compared. The shape error peaks at the boundary, where leakage forms, and reaches 6.7% for clipping against 0.6% for the constrained histogram. Across all cases the shape error tracks the k error (correlation 0.81 to 0.99) for every treatment except the single discrete angle, whose correlation is −0.32: its k error alone misjudges the quality of the solution.

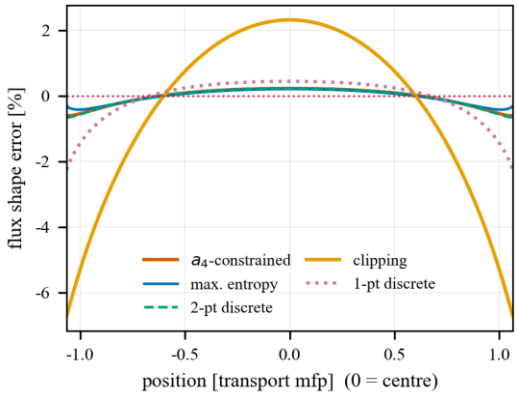


Fig. 5. Scalar-flux shape error, case A3

3.4. Library-wide classification

The linear-programming test was applied to every transition of the library above a $\Sigma_{s,0}$ cut of 10^{-6} (Table III). Most of the scattering admits an exact positive representation; the part that does not is concentrated in the sodium region, where the residual floor is +49 to +225 pcm in the cases measured here. Transitions whose moments are unrealizable, the only ones requiring projection, carry 1.6% of the fuel-region scattering and none of the sodium region.

Table III: Classification of all transitions, $\Sigma_{s,0}$-weighted [%]

Class	fuel	sodium
exact (a_4–a_6 = 0 attainable)	90.7	70.8
a_4 = 0 only	4.4	4.5
realizable, a_4 = 0 not attainable	3.3	24.7
a_1–a_3 not realizable	1.6	0.0

4. Conclusions

Multigroup Monte Carlo handles corrected scattering kernels with negative lobes reasonably well: every representation that preserves the first three moments stays within a few hundred pcm of the P3 operator, and constraining the fourth moment to zero, a rule that needs no geometry information, reduces the residual to the statistical level (+55.6 ± 3.1 to +0.7 ± 3.2 pcm). Clipping, the current default of production codes, is a factor of 2 to 26 worse than the constrained histogram and should be replaced.

In the PGSFR library the problematic kernels are few. 90.7% of the fuel-region and 70.8% of the sodium-region scattering admits an exact positive representation, and the transitions that no positive distribution can reproduce, a consequence of the consistent-P1 correction concentrated in self-shielded resonance groups, carry at most 1.6% of the scattering. The effect of the negative kernels on a full-core calculation is therefore expected to be limited, with the leakage-dominated sodium regions as the place to watch.

ACKNOWLEDGEMENT

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIT) [grant numbers RS-2025-16012975].

REFERENCES

[1] C. H. Lee and W. S. Yang, MC^2-3: Multigroup Cross Section Generation Code for Fast Reactor Analysis, ANL/NE-11/41 Rev. 4, Argonne National Laboratory, 2019.

[2] M. A. Smith, DIF3D-VARIANT Usage Guide, ANL/NSE-22/33, Argonne National Laboratory, 2022.

[3] E. E. Lewis and W. F. Miller, Jr., Computational Methods of Neutron Transport, John Wiley & Sons, New York, 1984.

[4] J. C. Wagner, E. L. Redmond II, S. P. Palmtag, and J. S. Hendricks, MCNP: Multigroup/Adjoint Capabilities, LA-12704, Los Alamos National Laboratory, 1994.

[5] P. K. Romano, N. E. Horelik, B. R. Herman, A. G. Nelson, B. Forget, and K. Smith, OpenMC: A State-of-the-Art Monte Carlo Code for Research and Development, Annals of Nuclear Energy, Vol. 82, pp. 90-97, 2015.

[6] S. Zhu, X. Ma, H. Tang, L. Yang, and B. Cao, Analyzing the Neutron Scattering Angular Distribution in an ACE-formatted Multigroup Cross Section Library Using the Maximum Entropy Method, EPJ Web of Conferences, Vol. 239, 10003, 2020.

[7] H. Brockmann, Treatment of Anisotropic Scattering in Numerical Neutron Transport Theory, Nuclear Science and Engineering, Vol. 77, pp. 377-414, 1981.

[8] P. Singh and H. Belanger, Negative Weights and Weight Cancellation to Treat Anisotropic Scattering in Multigroup Monte Carlo Simulations, Nuclear Science and Engineering, Vol. 199, pp. 1273-1295, 2025.

[9] J. Yoo, J. Chang, J.-Y. Lim, J.-S. Cheon, T.-H. Lee, S. K. Kim, K. L. Lee, and H.-K. Joo, Overall System Description and Safety Characteristics of Prototype Gen IV Sodium Cooled Fast Reactor in Korea, Nuclear Engineering and Technology, Vol. 48, pp. 1059-1070, 2016.

[10] A. Sood, R. A. Forster, and D. K. Parsons, Analytical Benchmark Test Set for Criticality Code Verification, LA-13511, Los Alamos National Laboratory, 1999.