LA-UR-10-01700

Approved for public release;
distribution is unlimited.

Title:	Theory, Interface, Verification, Validation, and Performance of the Adjoint-Weighted Point Reactor Kinetics Parameter Calculations in MCNP
Author(s):	Brian C. Kiedrowski
Intended for:	Documentation & Release Notes

Form 836 (7/06)

Theory, Interface, Verification, Validation, and Performance of the Adjoint-Weighted Point Reactor Kinetics Parameter Calculations in MCNP

Brian C. Kiedrowski
Los Alamos National Laboratory
XCP-3: Monte Carlo Codes

1 Introduction

MCNP5 [1] v1.60 has, for the first time, the ability to compute adjoint-weighted tallies in criticality calculations using only the existing random walks. This document specifically details the ability to compute point reactor kinetics parameters: neutron generation times, Rossi-α, total and precursor specific effective delayed neutron fractions, and average precursor decay constants. First, the theory of point-reactor kinetics is developed along with a formulation of the tallies for the kinetics parameters. Details on the user interface are given. Following this is a series of verification and validation problems. Finally, a few problems are selected to test the performance of the methods (in sequential and parallel) with and without calculations of kinetics parameters.

2 Theory

The average time-dependent behavior of a nuclear reactor for small transients can be described with the point reactor kinetics equations. From the derivation presented in [2], it is shown that the neutron generation time and effective delayed neutron fractions take the following definitions:

$$\Lambda = \frac{\left\langle \psi^\dagger, v^{-1}\psi \right\rangle}{\left\langle \psi^\dagger, F\psi \right\rangle}, \tag{1}$$

$$\beta_{\text{eff}} = \frac{\left\langle \psi^\dagger, B\psi \right\rangle}{\left\langle \psi^\dagger, F\psi \right\rangle}. \tag{2}$$

The nomenclature is as follows: the brackets denote integration over all space, energy, and direction in the reactor, ψ is the angular forward flux, $\psi^\dagger$ is the angular adjoint flux, v is the neutron speed, F is the operator for total (prompt plus delayed) fission, and B is the operator for the delayed component of the fission source.

Often, a related quantity called Rossi-α defined as $-\beta_{\text{eff}}/\Lambda$ is measured for criticality experiments. By manipulation, Rossi-α has the following definition:

$$\alpha = -\frac{\left\langle \psi^\dagger B\psi \right\rangle}{\left\langle \psi^\dagger v^{-1}\psi \right\rangle}. \tag{3}$$

It is also possible to define an effective delayed neutron fraction for the ith precursor group as follows:

$$\beta_i = \frac{\left\langle \psi^\dagger, B_i\psi \right\rangle}{\left\langle \psi^\dagger, F\psi \right\rangle}. \tag{4}$$

Here B_i is the delayed neutron source only for neutrons emitted from precursors of group i. Note that the sum of the individual effective delayed neutron fraction is the total effective delayed neutron fraction β_{eff}.

It is also possible to find average decay constants $\bar{\lambda}_i$ for each precursor i. Each individual λ_i (which varies depending on the isotope) for each fission event is averaged:

$$\bar{\lambda}_i = \frac{1}{W_i}\sum_{\text{f}\in i} \lambda_i w_0. \tag{5}$$

The summation is only for fissions of precursor i, λ_i is the decay constant for the isotope in the current fission event, w_0 is the source weight of the neutron, and W_i is the sum of source weights of emissions of precursor i.

There is still the issue of performing the adjoint (importance) weighting of each tally. To do this, MCNP uses the iterated fission probability interpretation of the adjoint flux [3]: the importance of a neutron at a point in phase space is proportional to the expected steady state population of a hypothetical neutron introducted into a critical reactor at that same point in phase space. Finding an estimate of the steady state population would, in theory, require following progeny for an infinite number of generations and measuring the population with some tally. Since that is not possible, MCNP assumes the population has achieved a steady state value after some specified number of generations, and measures the asymptotic population in that asymptotic generation.

To do this, a neutron in some original generation is assigned some index p and has tally contributions T_p. When this neutron has any progeny (and those progeny have successive progeny), the index p is inherited. After some number of generations, the asymptotic population R_p of progeny with index p is found by taking a track-length estimate from all tracks of neutrons with index p in the asymptotic generation:

$$R_p = \sum_{\tau \in p} \nu\Sigma_f w\ell. \tag{6}$$

The summation is over all tracks τ that have associated index p, $\nu\Sigma_f$ is the mean neutron production for that track, w is the weight of the track, and ℓ is the length of the track.

The adjoint-weighted tally A is found by taking the sum of all products of R_p and T_p.

$$A = \frac{1}{W_p} \sum_p R_p T_p. \tag{7}$$

Within the original generation, it is possible for the history to branch (such as particle splitting or an n,2n reaction). In this case, different tally contributions for each branch will be accrued after the branching event. It is therefore important to assign new progenitor indices when a branch occurs to ensure neutrons from one branch are not weighted by the importance of another. Since it is possible to have multiple progenitor indices in the same history within the original generation, it is important for calculation of variances to choose the correct meaning of a score. For this, an individual score for history j, A_j, is the sum of all products $R_p T_p$ over all progenitor indicies p that were assigned in history j.

There is one additional consideration with using existing random walks with non-analog particle tracking. In a non-analog simulation, the particle weight w adjusts the simulated frequency to match (on average) the physical frequency of the path. The importance weighting is, however, performed with a hypothetical particle sampled at unit weight from the physical frequency. The number of neutrons produced from a simulated neutron is scaled by the particle weight w. This is correct for the simulated neutron to match the physical result, but is not consistent had a hypothetical neutron at unit weight been used. To correct for this, a factor of $1/w$ must be applied to each tally contribution, leading to tally results that are not multiplied by particle weight.

From here it is possible to define the tallies necessary to compute the point-reactor kinetics parameters. From equations (1) and (2), it is necessary to define three different tallies. The first is the adjoint-weighted neutron density:

$$\left\langle \psi^\dagger, v^{-1}\psi \right\rangle = \sum_p R_p \sum_{\tau \in p} \frac{\ell}{v}. \tag{8}$$

This is done for every track τ with progenitor index p in the original generation. For every such track, the ratio of the track length to the current neutron speed is summed. The sum of these tracks is weighted by the importance estimate R_p taken from tracks in the asymptotic generation as defined by (6).

The adjoint-weighted fission sources need to be done as absorption estimators and cannot be done with track-length estimators if only the existing random walks are to be used. This is because the collisions must be produced with a frequency based upon the incident energy of the neutron, whereas the importance must be with respect to the energy (and direction) of the neutron after fission emission. A track-length estimator does not actually perform the collision, and therefore the part where a neutron leaves a collision via fission is not sampled.

The adjoint-weighted total fission source takes the following form:

$$\langle \psi^\dagger, F\psi \rangle = \sum_p R_p. \tag{9}$$

This form is simpler than (8) since at fission emission every neutron is assigned a progenitor index p with a weight of unity that remains fixed. Upon any branching event, new indices are selected and the unit weight is, in concept (but not in the simulation to save memory), passed onto all those branches. Therefore, a simple sum of the asymptotic populations will yield the adjoint-weighted fission source.

The adjoint-weighted delayed fission source is similar:

$$\langle \psi^\dagger, B\psi \rangle = \sum_p R_p(1 - \delta_j). \tag{10}$$

The only difference with respect to (9) is the application of Kronecker delta term. Here j corresponds to the emitted precursor index, where $j = 0$ corresponds to a prompt neutron. The Kronecker delta is defined to be one when $j = 0$ and zero otherwise.

Likewise, the delayed neutron source for precursor i can be found by

$$\langle \psi^\dagger, B_i\psi \rangle = \sum_p R_p\delta_{ij}. \tag{11}$$

Again, j corresponds to the emitted precursor index and i corresponds to the index of the precursor group being tallied. The Kronecker delta δ_{ij} is one when $i = j$ and zero otherwise.

By taking the appropriate ratios of these tallies, the kinetics parameters are obtained. Since the three tallies are correlated, a proper uncertainty calculation of a ratio involves finding the correlations between them. Standard techniques in error propogation are used to accomplish this.

3 User Interface

By default, MCNP5 v1.60 does not calculate the kinetics parameters, and therefore must be specified by the user. This is done by use of the KCODE options card, or KOPTS card. In a criticality (KCODE) problem, the user specifies the KOPTS card in the following format:

KOPTS KEYWORD1=ENTRY1 KEYWORD2=ENTRY2 ...

To activate the kinetics parameter calculations, the user must specify the KINETICS keyword on the KOPTS card. There are two valid entries for the KINETICS keyword: YES or NO, with the default being NO. By specifying the following card, MCNP will compute the kinetics parameters:

KOPTS KINETICS=YES

This will only compute Λ, β_{eff}, and α, but not the detailed precursor information (the β_i and $\bar{\lambda}_i$ for each precursor group). To have MCNP compute these as well, the PRECURSOR keyword must be specified. Like with the KINETICS keyword, the valid entries are either YES or NO with the default being NO. Note that it is illegal to request detailed precursor information if the KINETICS keyword is set to NO. When detailed precursor information is requested, MCNP will automatically read the data files to determine how many precursor groups are present (typically six for ENDF and eight for JEFF). The following card will produce a kinetics calculation with detailed precursor information:

KOPTS KINETICS=YES PRECURSOR=YES

By default, MCNP will break the active cycles into blocks of ten cycles for the importance weighting. In this case, tally contributions are taken and stored in the first cycle, progeny is followed through successive cycles, the adjoint weighting of the original tally contributions done in tenth. Following this, a new so-called block begins. Certain problems may require more generations in a block to avoid truncation errors. Conversely, problems that converge quickly may compute the kinetics parameters more efficiently for a smaller progenitor block. The size of the block may be controlled with the BLOCKSIZE keyword on the

KOPTS card. The entry must be an integer greater than or equal to two. A sample KOPTS card with the block size specified is:

KOPTS BLOCKSIZE=10 KINETICS=YES PRECUROSR=YES

This explicitly states that MCNP should use the default block size of ten cycles. Most problems in the validation and verification sections will use this card where applicable unless otherwise specified.

Results of the calculation are placed in the output file below the information on results of k_{eff}. Λ, β_{eff}, and α will be displayed with their absolute standard deviations. In the case where TOTNU NO is specified, Λ will be given and a message will indicate that no calculation of the other two may be performed since no delayed neutrons are present. Λ is given in appropriate units depending on the magnitude of the result: nsec (10^{-9} s), usec (10^{-6} s), or msec (10^{-3} s). Rossi-α is displayed in the inverse of the units of Λ.

Should detailed precursor information be required, the results of those calculations will be printed below the output for the kinetics parameters. This output will contain the effective delayed fractions for each precursor β_i, average emission energies in MeV, average decay constants in s^{-1}, and the corresponding half-lives in s. Absolute standard deviations are given for all of these except the half-lives which can easily be derived from the corresponding decay constant standard deviations.

4 Verification & Validation

Verification problems are performed for the kinetics parameters. The verification problems are compared against both analytic solutions and with discrete ordinates results obtained from Partisn [4]. Unfortunately, Partisn does not handle delayed neutrons, so only Λ is validated this way. For validation, MCNP computes six values of Rossi-α and these values are compared against experimentally measured values.

Analytic Verification Problems

Four infinte-medium test problems with analytic solutions [5] are used to verify the methods for computing the kinetics parameters within MCNP. Three problems specifically test the calculation of Λ using only prompt neutrons. The first problem is one group, and the second and third have two energy groups. A fourth problem uses delayed neutrons and compares β_{eff} and α.

The analytic solution for Λ in a monoenergetic infinite medium is

$$\Lambda = \frac{1}{v\nu\Sigma_f}. \tag{12}$$

The following data is used: $\Sigma_t = 4.0$ cm^{-1}, $\Sigma_f = 0.5$ cm^{-1}, $\Sigma_c = 0.5$ cm^{-1}, $\Sigma_s = 3.0$ cm^{-1}, $\nu = 2.0$, $v = 1.0$ cm/sh. The problem contains only isotropic scattering. For this problem the value of $k_\infty = 1.0$ and Λ =1 .0 sh or 10 ns.

Because of the infinite nature of the problem, there is no need to iterate on numerous generations, a block size of two (the minimum allowed by MCNP5) will suffice. The problem is run with 100k neutrons per cycle with 40k cycles with the default random number seed and generator with an initial source guess of a point source in the center of the universe. The results yield the expected values of $k_\infty = 1.00001 \pm 0.00001$ and $\Lambda = 10.00014 \pm 0.00030$ ns.

It is also possible to obtain an analytic solution for a multigroup infinite medium problem as well. For simplicity, two groups will be used. Group 1 has the following properties: all fission neutrons are born in this group, the group itself may not cause fission, any scattering event transfers the neutron to group 2 with cross section $\Sigma_s = 0.5$ cm^{-1}, and $\Sigma_{t1} = 1.0$ cm^{-1}. Group 2 has the following features: no scattering is possible, $\nu = 4.0$, $\Sigma_f = 1.0$ cm^{-1}, and $\Sigma_{t2} = 2.0$ cm^{-1}.

For these simplifications and choices of data, it can be shown that,

$$k_\infty = \frac{\nu\Sigma_f}{\Sigma_{t1}}\frac{\Sigma_s}{\Sigma_{t2}} = 1. \tag{13}$$

The analytic solution for the neutron generation time is,

$$\Lambda = \frac{\Sigma_{t1}\Sigma_{t2}}{\nu\Sigma_f\Sigma_s}\left(\frac{1}{v_1\Sigma_{t1}} + \frac{1}{v_2\Sigma_{t2}}\right). \tag{14}$$

For a choice of $v_1 = 1.0$ cm/sh and $v_2 = 0.5$ cm/sh, the analytic solution for Λ is 2 sh or 20 ns. This problem is run with 100k neutrons per cycle with 40k cycles with the default random number seed and generator with an initial source guess of a point source in the center of the universe. MCNP obtains a value for $k_\infty = 1.00000 \pm 0.00000$ and $\Lambda = 19.99963 \pm 0.00045$ ns. For reference the fission lifespan (the mean simulated time between fission events) has a matching value of 19.9995 ns.

Table 1: Cross section data for analytic test problem 3 (in cm^{-1} where applicable).

g	Σ_t	$\nu\Sigma_f$	χ	Σ_{sg1}	Σ_{sg2}
1	2	$\frac{3}{8}$	1	$\frac{1}{2}$	$\frac{1}{2}$
2	3	$\frac{9}{2}$	0	0	1

The third problem uses cross section data given in Table 1. The analytic expression for k_∞ is:

$$k_\infty = \frac{\nu\Sigma_{f1}}{\Sigma_{R1}} + \frac{\nu\Sigma_{f2}}{\Sigma_{R2}}\frac{\Sigma_{s12}}{\Sigma_{R1}} = 1. \tag{15}$$

Note that the removal cross section for group g is $\Sigma_{Rg} = \Sigma_{tg} - \Sigma_{sgg}$.

The neutron generation time has the following analytic expression:

$$\Lambda = \frac{\frac{1}{v_1}\frac{\Sigma_{R2}}{\nu\Sigma_{f2}} + \frac{1}{v_2}\frac{\Sigma_{s12}}{\Sigma_{R2}}}{\frac{\nu\Sigma_{f1}}{\nu\Sigma_{f2}}\Sigma_{R2} + \Sigma_{s12}}. \tag{16}$$

For the cross section data and the choice of $v_1 = 1$ cm/sh and $v_2 = 1/2$ cm/sh, the analytic solution for Λ is 17/12 sh or approximately 14.16667 ns. Like with the others, this problem is run with 100k neutrons per cycle with 40k cycles with the default random number seed and generator with an initial source guess of a point source in the center of the universe. Since the problem is more complicated, a generation block size of five is used. MCNP gives a result of Λ of 14.16669 ± 0.00097 ns, well within the two standard deviation confidence band. The fission lifespan has a result of 12.6660 ns, which shows that the fission lifespan does not always give the same result as the analytic solution.

Checking delayed neutrons and Rossi-α analytically is difficult since MCNP does not natively allow delayed neutrons with multigroup physics. To perform this test, MCNP is modified to read in multigroup delayed neutron data and sample them when required. An analytic two-group solution (with two precursor groups) for β_{eff} is developed for the cross section data given in Table 2. The delayed neutron fractions are independent of incident energy and have values of $\beta_1 = 1/4$ and $\beta_2 = 1/8$. The total delayed neutron fraction is therefore $\beta = 3/8$.

Table 2: Cross section data for analytic test problem 4 (in cm^{-1} where applicable).

g	Σ_t	$\nu\Sigma_f$	χ_p	χ_1	χ_2	Σ_{sg1}	Σ_{sg2}
1	2	0	1	$\frac{3}{4}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
2	3	$\frac{5}{24}$	0	$\frac{1}{4}$	$\frac{1}{2}$	0	1

Like before, a form of the eigenvalue is derived:

$$k_\infty = \left[(1-\beta) + \xi_1\right]\frac{\nu\Sigma_f\Sigma_{s12}}{\Sigma_{R1}\Sigma_{R2}} + \frac{\xi_2\nu\Sigma_f}{\Sigma_{R2}} = 1. \tag{17}$$

The variable ξ_g is the sum over all precursors i of $\beta_i\chi_{ig}$ for energy group g. Since there are two precursor groups: $\xi_1 = 1/4$ and $\xi_2 = 1/8$.

An analytic solution for the effective delayed neutron fraction is

$$\beta_{\text{eff}} = \frac{\frac{\Sigma_{s12}}{\Sigma_{R1}}\xi_1 + \xi_2}{\frac{\Sigma_{s12}}{\Sigma_{R1}}\left[(1-\beta)+\xi_1\right] + \xi_2}. \tag{18}$$

For the cross section data provided, the analytic value of β_{eff} is 1/2. For the case run like the others, MCNP produces a value of $k_\infty = 1.00002 \pm 0.00001$ and $\beta_{\text{eff}} = 0.50003 \pm 0.00005$.

Rossi-α also has an analytic solution:

$$\alpha = -\frac{\left[\frac{\Sigma_{s12}}{\Sigma_{R1}}\xi_1 + \xi_2\right]\frac{\nu\Sigma_f\Sigma_{s12}}{\Sigma_{R2}-\xi_2\nu\Sigma_f}}{\frac{1}{v_1}\frac{\Sigma_{s12}}{\Sigma_{R2}} + \frac{1}{v_2}\frac{\Sigma_{s12}}{\Sigma_{R2}-\xi_2\nu\Sigma_f}}. \tag{19}$$

For the current nuclear data, the analytic value of Rossi-α is -15/44 sh^{-1} or approximately -3.40909 x 10^{-2} ns^{-1}. MCNP calculates a value of -3.40955 x 10^{-2} $\pm$ 0.00044 x 10^{-2} ns^{-1}. For reference, the analytic solution of Λ is 22/15 sh or approximately 14.66667 ns. MCNP calculates 14.66548 ± 0.00110 ns.

Both the one- and two-group results agree with the respective analytic solutions within two standard deviations. These tests show that the importance weighting routines are being done correctly and handle spectral effects (the forward and adjoint solutions differ in the two group problem for the data chosen). It remains an open question if spatial differences can be accounted, however. To address this, homogeneous and heterogeneous verification problems are performed and compared with equivalent discrete ordinates calculations.

Discrete Ordinates Verification Problems

Multigroup problems with finite geometries are tested using the discrete ordinates method with the code Partisn. Results of Λ are compared between MCNP and Partisn v6.26 (beta release). There are eight multigroup problems: (1) 4-group, bare, fast slab, (2) 4-group fissile slab with thermalizing reflector, (3) 2-group, three region slab problem involving fissile center, strong thermal absorber buffer zone, and moderating reflector, (4) 8-group, bare slab of homogeneous fissile/moderator mixture, (5) 4-group, bare, fast, sphere, (6) 4-group, sphere with reflector, (7) 4-group, bare, subcritical slab, and (8) 4-group, bare, supercritical slab. For each of these problems, a block size of ten generations is used.

For reference, the results of all comparisons are given in Table 3. Also provided are n-σ, the number of MCNP standard deviations that the Partisn result lies outside the MCNP mean result, and C/R which is the ratio of the calculated (MCNP) to reference (Partisn) results. Detailed information on each test problem involving multigroup data, Partisn discretization, MCNP problem specifics, etc. are given below.

Table 3: Λ results of MCNP compared with discrete ordinates calculations from Partisn.

Problem	Partisn	MCNP	n-σ	C/R
1	9.79325 ns	9.79675 $\pm$ 0.00188 ns	+1.86	1.00036
2	135.19020 us	135.22164 $\pm$ 0.03384 us	+0.93	1.00023
3	49.16822 ns	49.20663 $\pm$ 0.01863 ns	+2.06	1.00078
4	112.05232 us	112.29905 $\pm$ 0.13692 us	+1.80	1.00220
5	1.72115 ns	1.72121 $\pm$ 0.00032 ns	+0.19	1.00003
6	10.18997 ns	10.18794 $\pm$ 0.00233 ns	-0.87	0.99980
7	10.17161 ns	10.17110 $\pm$ 0.00230 ns	-0.22	0.99995
8	9.67254 ns	9.67168 $\pm$ 0.00166 ns	-0.81	0.99990

Like with the analytic results, the values of Λ compare well with those computed from the forward and adjoint fluxes of Partisn. A couple of the results (namely problems 1, 3, and 4) are near the 2-σ confidence band, but are not unreasonable. It may be, for this degree of precision, that the default block size of ten generations may be insufficient. As can be seen in the problem descriptions, the value of Λ matches consistently far more frequently than either the prompt removal time (non-adjoint weighted version of Λ) or the fission lifespan (simulated time between fission events).

Problem 1: This problem is a bare, fast slab of half-thickness 7.7218 cm. The 4-group nuclear data is given in Table 4 and the number density is 0.01 atoms per barn-cm. The upper energy bounds on the 4 groups are 10 MeV, 2.5 MeV, 1.0 MeV, and 0.1 MeV.

In Partisn, 1000 spatial bins are used, each with 512 Gauss-Legendre ordinates. For MCNP, the problem is run with a uniform source with the default random number seed. 50k active cycles are used with 100k neutrons per batch are used.

For this problem $k_{\text{eff}} = 1.00000$ and both MCNP and Partisn match. The non-adjoint weighted lifetime computed by Partisn is 9.89003 ns. Correspondingly, MCNP calculates a prompt removal time of 9.89021 ± 0.00013 ns. There is about a one percent difference between the importance and non-importance weighted lifetimes. The fission lifespan computed by MCNP is 10.1816 ns, which is actually less accurate than the prompt removal lifetime.

Table 4: Nuclear cross section (in barns where applicable) data for test problem 1.

g	σ_a	$\nu\sigma_f$	σ_t	χ	σ_{sg1}	σ_{sg2}	σ_{sg3}	σ_{sg4}
1	2.0	4.5	4.0	0.1	0.5	0.5	0.5	0.5
2	3.5	7.125	6.0	0.6	0.0	1.0	1.0	0.5
3	5.0	10.0	8.0	0.3	0.0	0.0	2.0	1.0
4	8.0	12.5	10.0	0.0	0.0	0.0	0.0	2.0

Problem 2: This problem is a reflected slab reactor with overall half-thickness of 112.5 cm. The inner-core region represents a metallic fuel with a half-thickness of 20 cm. The reflector region is a material containing low-Z isotopes. The 4-group nuclear data for the core and reflector are given in Table 5 and the number density of both materials is 0.001 atoms per barn-cm. The upper energy bounds are: 10 MeV, 1 MeV, 0.1 MeV, and 0.625 eV. Much of the fission (about 75 percent) occurs in either the intermediate or thermal energy range; as such, it is expected that Λ be a few decimal orders of magnitude longer than in Problem 1.

In Partisn, a mesh spacing of 0.1 cm is used in both core and reflector regions. The simulation is run using 512 Gauss-Legendre ordinates. For MCNP, the problem is run with a uniform source with the default random number seed. 50k active cycles are used with 100k neutrons per batch are used.

For this problem $k_{\text{eff}} = 1.00018$ from Partisn and 1.00018 ± 0.00001 for MCNP. The non-adjoint weighted lifetime computed by Partisn is 118.03165 us. Correspondingly, MCNP calculates a prompt removal time of 118.065 ± 0.0002 us. There is about a 12.7 percent difference between the importance and non-importance weighted lifetimes. This fission lifespan computed by MCNP is 136.252 us.

Table 5: Cross section data for fuel and reflector zones of test problem 2.

	g	σ_t	$\nu\sigma_f$	σ_a	χ	σ_{sg1}	σ_{sg2}	σ_{sg3}	σ_{sg4}
Core	1	1.5	2.65	5.0	0.8	1.5	2.0	0.0	0.0
	2	3.0	5.0	8.0	0.2	0.0	2.0	3.0	0.0
	3	10.0	12.5	20.0	0.0	0.0	0.0	10.0	0.0
	4	30.0	70.0	50.0	0.0	0.0	0.0	0.0	20.0
Refl	1	0.0	0.0	20.0	0.0	4.0	6.0	9.0	1.0
	2	0.0	0.0	20.0	0.0	0.0	3.0	14.0	3.0
	3	0.0	0.0	20.0	0.0	0.0	0.0	15.0	5.0
	4	0.1	0.0	20.0	0.0	0.0	0.0	0.0	19.9

Problem 3: This problem is a three-region problem with half-thickness 100 cm. The problem is very similar to problem 2, except for the 1 cm layer of strong thermal absorber between the core and reflector. The core region has a half-thickness of 20 cm. The nuclear data for each material is given in Table 6, and all have number densities of 0.01 atoms per barn-cm. The upper bounds of the two groups are: 1 MeV and 1 eV. This problem is interesting because it is pathological for non-adjoint weighted methods. Neutrons spend a significant amount of time in the reflector, but, once thermalized, are insignificant in terms of the chain reaction because of the strong absorber. For reference, the non-adjoint weighted Λ is about a factor of 200 larger than the correct adjoint-weighted case.

In Partisn, 1000 spatial bins are used, each with 256 Gauss-Legendre ordinates. For MCNP, the problem is run with a uniform source with the default random number seed. 50k active cycles are used with 100k neutrons per batch are used.

For this problem $k_{\text{eff}} = 1.00041$ from Partisn and 1.00042 ± 0.00001 for MCNP. The non-adjoint weighted lifetime computed by Partisn is 10.0678 us. Correspondingly, MCNP calculates a prompt removal time of 10.1060 ± 0.0002 us. There is about a factor of 200 difference between the importance and non-importance weighted lifetimes. The fission lifespan computed by MCNP is 50.0817 ns.

Table 6: Cross section data for fuel, absorber, and reflector zones of test problem 3.

	g	σ_t	$\nu\sigma_f$	σ_a	χ	σ_{sg1}	σ_{sg2}
Core	1	5.0	2.425	1.0	1.5	3.5	0.0
	2	50.0	100.0	0.0	45.0	5.0	0.0
Abs	1	0.0	0.0	0.0	0.0	0.0	0.0
	2	1000.0	0.0	1000.0	0.0	0.0	0.0
Refl	1	1.0	0.0	0.0	0.0	0.9	0.1
	2	1.0	0.0	0.0	0.0	0.0	1.0

Problem 4: This problem is a slab consisting of homogenized mixture of fuel and moderator with half-thickness of 77.2 cm. The 8-group cross section data is given in Table 7 and the number density is 0.001 atoms per barn-cm. The upper energy bounds for the eight energy groups are: 10 MeV, 5 MeV, 2 MeV, 1 MeV, 0.1 MeV, 0.1 keV, 1 eV, and 0.1 eV. The cross section data is chosen such that the fast and thermal contributions to the tally for Λ are of the same order. In other words, most of the flux is in the fast region, but each contribution from $1/v$ is small. Contributions from the thermal region are far fewer, but are much larger because $1/v$ is bigger.

In Partisn, a mesh spacing of 0.05 cm is used for each region and using 512 Gauss-Legendre ordinates. For MCNP, the problem is run with a uniform source with the default random number seed. 50k active cycles are used with 100k neutrons per batch are used.

For this problem $k_{\text{eff}} = 1.00039$ from Partisn and 1.00040 ± 0.00001 for MCNP. The non-adjoint weighted lifetime computed by Partisn is 100.42339 us. Correspondingly, MCNP calculates a prompt removal time of 100.507 ± 0.006 us. There is about a 10.4 percent difference between the importance and non-importance weighted lifetimes. The fission lifespan computed by MCNP is 114.462 us.

Table 7: 8-group nuclear cross section (in barns where applicable) data for test problem 4.

g	σ_a	$\nu\sigma_f$	σ_t	χ	σ_{sg1}	σ_{sg2}	σ_{sg3}	σ_{sg4}	σ_{sg5}	σ_{sg6}	σ_{sg7}	σ_{sg8}
1	2.5	6.0	4.0	0.05	0.1	0.2	0.2	0.5	0.5	0.0	0.0	0.0
2	3.0	6.75	5.5	0.20	0.0	0.4	0.6	0.7	0.8	0.0	0.0	0.0
3	3.5	7.8	7.0	0.60	0.0	0.0	1.0	1.0	1.5	0.0	0.0	0.0
4	5.0	10.2	10.0	0.15	0.0	0.0	0.0	1.4	3.5	0.1	0.0	0.0
5	30.0	25.0	40.0	0.00	0.0	0.0	0.0	0.0	9.6	0.4	0.0	0.0
6	40.0	75.0	60.0	0.00	0.0	0.0	0.0	0.0	0.0	12.0	7.0	5.0
7	70.0	150.0	90.0	0.00	0.0	0.0	0.0	0.0	0.0	1.0	14.0	5.0
8	115.0	250.0	140.0	0.00	0.0	0.0	0.0	0.0	0.0	0.5	4.5	20.0

Problem 5: This problem is a bare-fast sphere, similar to problem 1. The sphere has a radius of 4.4608 cm, the same (4-group) nuclear data as problem 1, and a number density of 0.05 atoms per barn-cm.

In Partisn, 1000 mesh intervals are used with 64 Gauss-Legendre ordinates. For MCNP, the problem is run with a central point source with the default random number seed. 50k active cycles are used with 100k neutrons per batch are used.

For this problem $k_{\text{eff}} = 1.00000$ from Partisn and 0.99995 ± 0.00001 for MCNP. The non-adjoint weighted lifetime computed by Partisn is 1.98991 ns. Correspondingly, MCNP calculates a prompt removal time of 1.98982 ± 0.0002 ns. There is about a 13.5 percent difference between the importance and non-importance weighted lifetimes. The fission lifespan computed by MCNP is 1.82160 ns.

Problem 6: This problem is a reflected-fast sphere, similar to problem 2. The overall sphere has a radius of 30 cm. The inner core region has a radius of 21 cm, has the same (4-group) nuclear data as problem 1, and a number density of 0.01 atoms per barn-cm. The outer region is a high-Z reflector with nuclear data given in Table 8. Since the reflector does not thermalize neutrons effectively, the spectrum is, unlike problem 2, expected to be fast.

In Partisn, 300 mesh intervals are used with 64 Gauss-Legendre ordinates. For MCNP, the problem is run with a central point source with the default random number seed. 50k active cycles are used with 100k neutrons per batch are used.

For this problem $k_{\text{eff}} = 1.00086$ from Partisn and 1.00081 ± 0.00001 for MCNP. The non-adjoint weighted lifetime computed by Partisn is 18.04258 ns. Correspondingly, MCNP calculates a prompt removal time of 18.0412 ± 0.0002 ns. There is about a 77 percent difference between the importance and non-importance weighted lifetimes. The fission lifespan computed by MCNP is 11.0911 ns.

Table 8: Nuclear cross section (in barns) data of the reflector region in test problem 6.

g	σ_a	$\nu\sigma_f$	σ_t	σ_{sg1}	σ_{sg2}	σ_{sg3}	σ_{sg4}
1	0.1	0.0	2.0	0.3	0.4	0.6	0.6
2	0.2	0.0	3.0	0.0	0.6	1.0	1.2
3	0.4	0.0	4.0	0.0	0.0	1.6	2.0
4	1.0	0.0	5.0	0.0	0.0	0.0	4.0

Problem 7: This problem is exactly like problem 1 except that the slab has a half-thickness of 5 cm and is therefore in a subcritical configuration.

In Partisn, 1000 mesh intervals are used with 512 Gauss-Legendre ordinates. For MCNP, the problem is run with a uniform source with the default random number seed. 50k active cycles are used with 100k neutrons per batch are used.

For this problem $k_{\text{eff}} = 0.77971$ from Partisn and 0.77972 ± 0.00001 for MCNP. The non-adjoint weighted generation time computed by Partisn is 9.65371 ns. Correspondingly, MCNP calculates a prompt removal time (divided by k_{eff} to get Λ) of 9.65383 ± 0.00001 ns. There is about a five percent difference between the importance and non-importance weighted lifetimes. The fission lifespan computed by MCNP divided by k_{eff} is 10.55002 ns.

Problem 8: This problem is exactly like problem 1 except that the slab has a half-thickness of 10 cm and is therefore in a supercritical configuration.

In Partisn, 1000 mesh intervals are used with 512 Gauss-Legendre ordinates. For MCNP, the problem is run with a uniform source with the default random number seed. 50k active cycles are used with 100k neutrons per batch are used.

For this problem $k_{\text{eff}} = 1.14008$ from Partisn and 1.14007 ± 0.00001 for MCNP. The non-adjoint weighted generation time computed by Partisn is 10.01983 ns. Correspondingly, MCNP calculates a prompt removal time (divided by k_{eff} to get Λ) of 10.0199 ± 0.00001 ns. There is about a 3.6 percent difference between the importance and non-importance weighted lifetimes. The fission lifespan computed by MCNP divided by k_{eff} is 10.06561 ns.

Validation with Experimental Results

Comparisons are made with experimental measurements of six criticality experiments from the OECD/NEA benchmark handbook [6]. These are: Godiva, Jezebel, BIG TEN, Flattop-233, Stacy (run 29), and WINCO (run 5). The corresponding designators are: HEU-MET-FAST-001, PU-MET-FAST-001, IEU-MET-FAST-007, U233-MET-FAST-006, LEU-SOL-THERM-007, HEU-SOL-THERM-038.

The kinetics parameters are computed for each of the experiments and are given in Table 9. Detailed precursor information is also given; Table 10 gives values of β_i and Table 11 gives values of $\bar{\lambda}_i$. All calculations use 50k active cycles with 100k neutrons per cycle, a block size of ten, and ENDF/B-VII.0 data.

There are relatively few measured values of Rossi-α for the OECD/NEA benchmarks. A comparison of those measurements that are available are given in Table 12. All six of the validation tests appear to match within two percent aside from the Flattop-233 benchmark that exhibits about a ten percent difference. One

Table 9: MCNP results for the point-reactor kinetics parameters.

	Λ	β_{eff} (pcm)	α (ms^{-1})
Godiva	5.69465 ± 0.00143 ns	649 ± 1	-1139.57 ± 2.35
Jezebel	2.87273 ± 0.00080 ns	184 ± 1	-640.238 ± 2.374
BIG TEN	62.42441 ± 0.01609 ns	721 ± 1	-115.518 ± 0.219
Flattop-233	12.72031 ± 0.00679 ns	372 ± 1	-292.401 ± 0.808
Stacy-29	59.72607 ± 0.01262 us	730 ± 2	-0.122155 ± 0.00296
WINCO-5	7.36401 ± 0.00248 us	823 ± 2	-1.11723 ± 0.00311

Table 10: MCNP results for the effective delayed fractions for each precursor group (in pcm). The uncertainies are less than 1 pcm.

	β_1	β_2	β_3	β_4	β_5	β_6
Godiva	21	108	104	297	88	31
Jezebel	6	49	34	72	20	3
BIG TEN	16	102	104	330	129	39
Flattop-233	25	83	63	148	44	9
Stacy-29	24	121	118	333	99	35
WINCO-5	27	137	134	375	110	39

Table 11: MCNP results for the average precursor decay constants (s^{-1}). The uncertainies are usually negligible.

	λ_1	λ_2	λ_3	λ_4	λ_5	λ_6
Godiva	0.01249	0.03180	0.10948	0.31745	1.35233	8.67870
Jezebel	0.01252	0.02997	0.10749	0.31681	1.31001	9.52755
BIG TEN	0.01249	0.03122	0.11212	0.32855	1.33222	9.36772
Flattop-233	0.01248	0.03144	0.11057	0.32484	1.30870	9.90878
Stacy-29	0.01249	0.03181	0.10944	0.31727	1.35328	8.65641
WINCO-5	0.01249	0.03182	0.10938	0.31701	1.35393	8.63788

Table 12: Comparison of MCNP calculations of to experimentally measured values of Rossi-α (ms^{-1}).

	Experiment	MCNP	C/R
Godiva	-1100 ± 20	-1139.57 ± 2.35	1.017
Jezebel	-640 ± 10	-640.238 ± 2.374	1.000
BIG TEN	-117 ± 1	-115.518 ± 0.219	0.987
Flattop-233	-267 ± 5	-292.401 ± 0.808	1.095
Stacy-29	-0.122 ± 0.004	-0.122155 ± 0.00296	1.001
WINCO-5	-1.109 ± 0.003	-1.11723 ± 0.00311	1.007

other consistent peculiarity is $\bar{\lambda}_6$ exhibits a value that is consistently different than other published values. A verification of MCNP's ability to caculate $\bar{\lambda}_i$ is performed by comparing the results from a problem containing one fissile isotope directly with the nuclear data in ENDF. Both results are consistent indicating the difference is because of the data itself.

Additions to the Regression Suite

To assist developers and users with ensuring these routines function as expected, two test problems are included in the Regression Suite. Geometrically, both problems are the same: a two concentric region sphere of radius 8 cm (inner radius of 5 cm) containing the same material (90 percent uranium-235 and 10 percent uranium-238).

The first problem tests a calculation of all kinetics parameters with detailed precursor information.

The problem includes coupled neutron-photon transport and importance splitting of both the neutrons and photons. The tallies should not be affected, statistically speaking, by secondary photons, and this tests this. Also, splitting is performed to ensure that the answers are not affected by that as well.

The second problem uses solely neutrons, no importance splitting, the TOTNU NO option, but still tries to calculate detailed precursor information. In this case, the user is warned that this is impossible and the program proceeds without those calculations. Also, an invalid block size of one is specified; MCNP sets this to the minimum value of two and warns the user. The primary purpose of this test is to catch input errors.

5 Performance Testing

Tests of performance on CPU time/multiprocess scaling and memory usage are performed. MCNP5 v1.60 is compiled using Intel 10.0 on an AMD Operton machine running Red Hat Linux. The test problems run for the CPU time/multiprocess scaling tests are: Flattop-233, BIG TEN, and Stacy-29. The memory usage test uses the Godiva benchmark.

For each benchmark, 10k neutrons per cycle are used, 10 inactive and 500 active cycles are run, and a block size of 10 generations is used. The CPU time is measured with the Linux 'time' command and the setup time is assumed to be small compared to the overall running time. Each benchmark is run with sequentially (1 process) and with 2, 4, 8, 16, 24, 32, 48, and 63 MPI slave processes. For the three benchmarks, the average slowdown experienced for toggling the kinetics parameter calculations (with detailed precursor information tabulated) is around 5-6%. The slowdown for the different number of processors is somewhat erratic because of the non-linear effects associated with caching.

For each benchmark (Flattop-233, BIG TEN, and Stacy-29), the speedup for each with the kinetics parameter calculations is compared with the speedup from a reference case with no such calculation. This comparison is given in Figure 1. For all cases, the speedup is greater for the reference case, as expected, since more data needs to be passed between processors, boosting the non-parallelizable fraction of the program

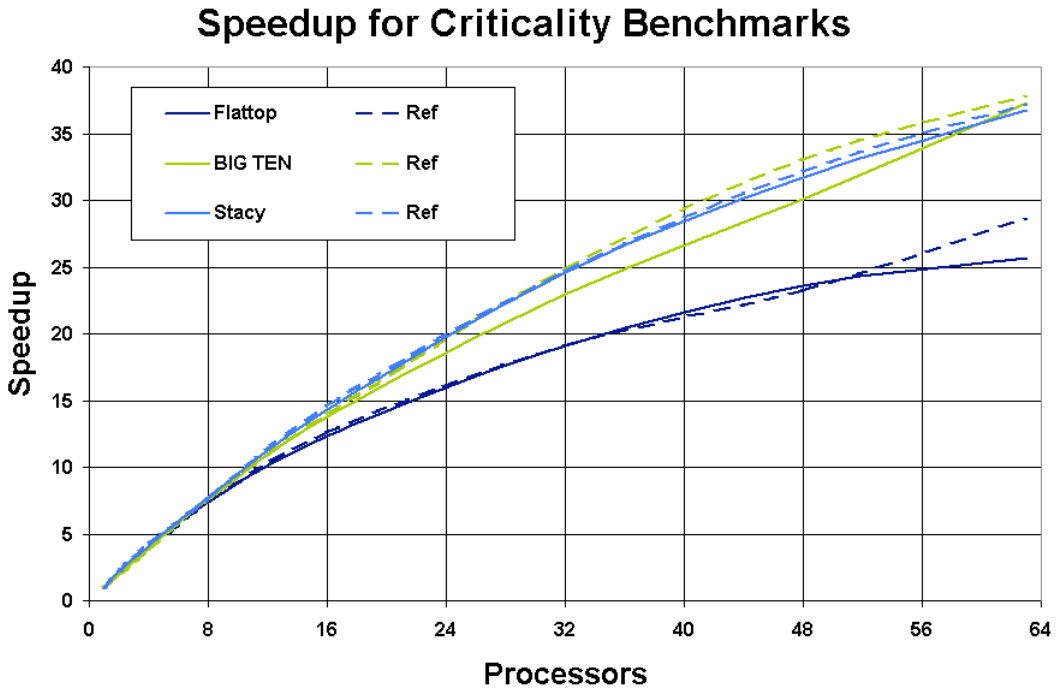

Figure 1: Comparison of speedup using varied numbers of MPI processes for the Flattop-233, BIG TEN, and Stacy-29 criticality benchmarks. The reference cases denote running an equivalent calculation without kinetics parameters.

execution. The penalty in speedup is typically less than five percent, but exceptions are observed depending on the problem and number of processors used.

The Godiva benchmark is run sequentially with varied numbers of neutrons per cycle. The memory of the cases with and without (which is called the reference case) the kinetics parameter calculations (with detailed precursor information tabulated) is measured by way of the Linux 'top' command.

The results of the tests with the Godiva problem are displayed in Figure 2. The increased memory usage is from two components: fixed sized data for each tally and arrays that vary with the number of progenitors allowed. Initially, MCNP selects the size of the latter based on the number of neutrons per cycle specified by the user (the base memory storage does not necessarily grow linearly with the user selection as seen in increasing the batch size from 75k to 100k); however, during problem execution MCNP will automatically adjust the number of progenitors allowed should more be required. This leads to the non-linear growth in memory usage seen in this test.

As a general trend, the memory usage typically increases with batch size at a faster rate than for the other arrays found in a simple criticality problem. To illustrate, the kinetics calculations require about 5% of the memory usage for the 10k batch size case, whereas the the kinetics calculations consume over 50% of the memory for a batch size of 250k. However, the growth is not always monotonic as the percentage of memory used by the kinetics calculations actually decreases between the 100k and 150k batch sizes.

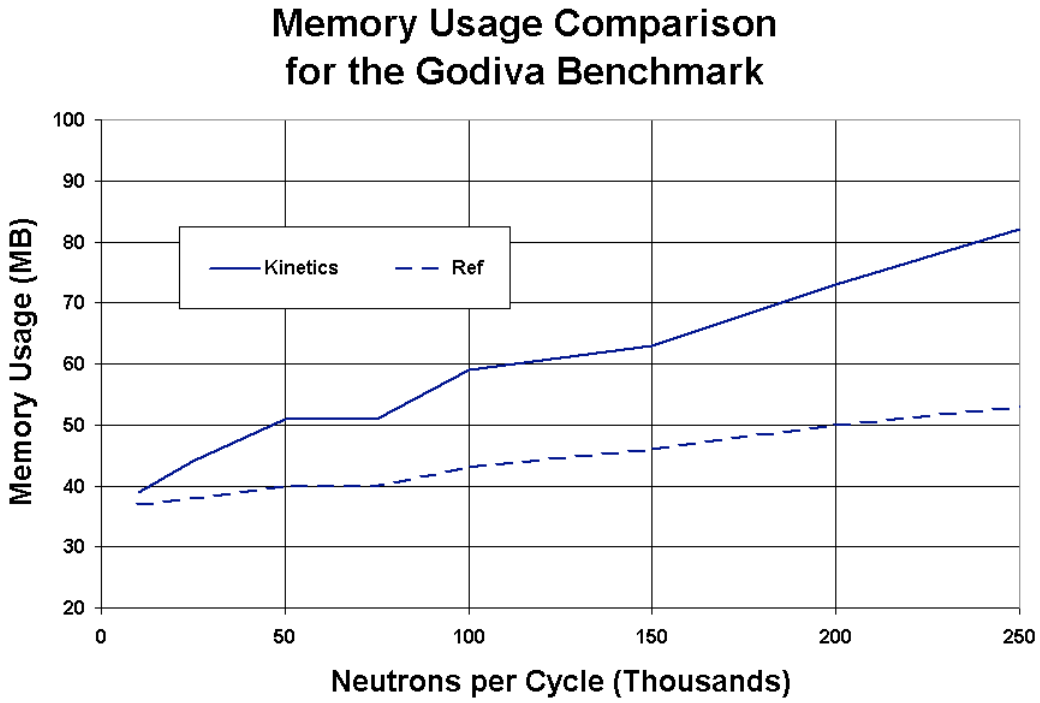

Figure 2: Comparison of memory usage for varying the batch size (neutrons per cycle) of the Godiva benchmark. The reference cases denote running an equivalent calculation without kinetics parameters.

6 Summary

MCNP v1.60 has the ability to compute the point-reactor kinetics parameters (Λ, the neutron generation time, β_{eff}, the effective delayed neutron fraction, and Rossi-α). Should the user desire, MCNP will also produce detailed precursor information consisting of the effective dealyed fraction of each precursor β_i and the average decay constant $\bar{\lambda}_i$. The default is to produce no kinetics parameter information and needs to be specified by way of the KINETICS and PRECURSOR keywords on the new KOPTS card.

The results of the calculations have been validated and verified using analytic, infinite-medium solutions, comparisons with discrete ordinates, and comparisons with experimental results. The comparisons indicate

that the MCNP calculations produce answers that match reference solutions or values, and does so better than similar, previously available methods. Testing indicates that the performance penalties tend to be fairly modest, on the order of around five percent. There is a fairly significant increase in memory usage that scales with the number of neutrons requested per cycle.

References

[1] X-5 Monte Carlo Team, "MCNP - A General N-Particle Transport Code, Version 5, Volume I: Overview and Theory," Los Alamos National Laboratory, LA-UR-03-1987 (2003).

[2] D.L. Hetrick, *Dynamics of Nuclear Reactors*, University of Chicago Press, Chicago, IL (1971).

[3] H. Hurwitz in *Naval Reactor Physics Handbook*, A. Radkowsky (ed.), vol. 1, pp. 864-849, Naval Reactors, U.S. Atomic Energy Commission (1964).

[4] R.E. Alcouffe, R.S. Baker, et. al., "Partisn Manual," Los Alamos National Laboratory, LA-UR-02-5633 (2002).

[5] B.C. Kiedrowski, "Analytic, Infinite-Medium Solutions for Point Reactor Kinetics Parameters and Reactivity Perturbations," Los Alamos National Laboratory, LA-UR-10-01803 (2010).

[6] J. Blair Briggs (ed.), *International Handbook of Evaluated Criticality Safety Benchmark Experiments*, Nuclear Energy Agency, NEA/NSC/DOC(95)03/I, Paris, France (2004).