LA-UR-12-25947
Approved for public release; distribution is unlimited.

Title: Methodology, Verification, and Performance of the Continuous-Energy Nuclear Data Sensitivity Capability in MCNP6

Author(s): Kiedrowski, Brian C.
Brown, Forrest B.

Intended for: International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, 2013-05-05/2013-05-09 (Sun Valley, Idaho, United States)

International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering (M&C 2013)
Sun Valley, Idaho, USA, May 5-9, 2013, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2013)

METHODOLOGY, VERIFICATION, AND PERFORMANCE OF THE CONTINUOUS-ENERGY NUCLEAR DATA SENSITIVITY CAPABILITY IN MCNP6

Brian C. Kiedrowski and Forrest B. Brown
Los Alamos National Laboratory
P.O. Box 1663 MS A143, Los Alamos, NM 87545
bckiedro@lanl.gov; fbrown@lanl.gov

ABSTRACT

A continuous-energy sensitivity coefficient capability has been introduced into MCNP6. The methods for generating energy-resolved and energy-integrated sensitivity profiles are discussed. Results from the verification exercises that were performed are given, and these show that MCNP6 compares favorably with analytic solutions, direct density perturbations, and comparisons to TSUNAMI-3D and MONK. Run-time and memory requirements are assessed for typical applications, and these are shown to be reasonable with modern computing resources.

Key Words: Monte Carlo, perturbation, uncertainty analysis

1. INTRODUCTION

For the previous decade, the U.S. DOE/NNSA Nuclear Criticality Safety Program has funded the development of sensitivity and uncertainty analysis tools for critical experiments and benchmark analysis. The primary analysis tool for this purpose is TSUNAMI-3D, which is part of the SCALE package [1,2]. In recent years, other similar capabilities throughout the world have been developed, and the first production release of MCNP6 includes such a capability [3].

The iterated fission probability (IFP) method [4] is used to compute the adjoint-weighted tallies needed for sensitivity coefficients. This method has been previously shown to generalize to continuous energy because only a forward simulation is required to perform the adjoint weighting. This same method is currently used in MCNP6 to estimate point-kinetics parameters or estimating reactivity changes from material substitutions. An overview of the IFP method is given, along with details of the adjoint-weighted tallies.

Before releasing this capability, verification was performed in three ways: (1) comparisons to analytic solutions, (2) direct density perturbations, and (3) comparisons with other software. Each of these verification methods is discussed, results are given, and they show that the routines in MCNP6 are capable of correctly estimating sensitivity coefficients. The run-time and memory requirements are measured for typical applications. Neither of these is prohibitive on modern computing platforms, demonstrating that the method as implemented is a practical tool.

Note that this paper specifically focuses on sensitivities for cross sections and fission ν. Little attention is given to fission-χ and scattering distribution sensitivities here, and this is covered more extensively in a companion paper [5].

2. METHODOLOGY

The sensitivity coefficient of k to nuclear data x of isotope j, x^j, is defined as the ratio of the resulting relative differential change in k to the relative differential change in x^j. This can be found from perturbation theory as a ratio of adjoint-weighted integrals:

$$S^j_{k,x} = \frac{x^j}{k}\frac{dk}{dx^j} = -\frac{\left\langle \psi^\dagger, (\Sigma^j_x - \mathcal{S}_x - \lambda\mathcal{F}_x)\psi \right\rangle}{\langle \psi^\dagger, \lambda F\psi\rangle}. \quad (1)$$

Here ψ is the angular (forward) flux and $\psi^\dagger$ is its adjoint function. Σ^j_x is the cross section corresponding to x^j if x^j is a cross section, and zero otherwise (e.g., fission χ). $\mathcal{S}_x$ is the integral scattering operator for x^j if x^j is a scattering cross section or law [includes elastic, inelastic, (n,2n), etc.], and zero otherwise. $\mathcal{F}_x$ is the integral fission operator for x^j if x^j is a fission cross section, fission ν, or fission χ and zero otherwise. The quantity $\lambda = 1/k$ and the brackets denote integration over all phase space.

2.1. The Adjoint Weighting

Sans the adjoint function, these terms can be estimated with continuous-energy Monte Carlo with standard techniques using reaction-rate estimators. Because this integral is adjoint-weighted, as a neutron traverses through space and undergoes reactions, it must not only be weighted by reaction cross sections, but the local value of the importance function as well, which is usually not known.

Physically, the adjoint function is proportional to its (and that of any progeny it may have) expected future performance toward contributing to a particular response [6]. In the eigenvalue problem, the adjoint function specifically is proportional to the expected number of neutrons that would be in a system infinitely many iterations into the future as a result of the current neutron [7] – Note that the consideration of sub- and super-criticality is irrelevant as the eigenvalue problem imposes a balance on the overall population. While it is impossible to directly simulate neutrons for infinitely many generations, it has been shown that for quantities estimated by system-wide or global integrals, 5-10 generations is effectively infinite [4].

This leads to the IFP method. First, the eigenvalue iterations are separated into "blocks" of fixed size (typcially 5-10). In the first generation of the block, the simulated neutrons are tagged and non-adjoint weighted tally contributions are recorded and linked to each tag. The neutron tags are passed onto any progeny and kept until the last cycle in the block. During this last cycle, estimates of neutron production are made (MCNP6 uses a track-length estimator of fission neutron production), which, when summed over all neutrons sharing the same tag is a random sample of its adjoint function. This estimate is multiplied by the original, non-adjoint weighted estimate, to get an estimate of the adjoint-weighted tally.

The tagging scheme must be set up to account for branches in the neutron history within the individual iteration within the block, which may arise from (n,2n) reactions, particle splitting, etc. Many of the details and theory behind the IFP have been omitted here for space, but these are discussed elsewhere in the literature [4].

Note that the IFP method is one of many methods that have been developed in recent years. For instance, the differential operator method for perturbation theory may be used with a fission source correction to produce the same results [8,9]. Other proposals involve using discrete approximations to the adjoint

functions using a Green's function approach [10] , or precomputing approximate adjoint functions with deterministic codes.

2.2. The Forward Tallies

Correctly tallying the quantities on the right side of the integrals has, in itself, some subtleties. First, to be more explicit, there are three terms that must be estimated. The first is the loss of neutrons from interactions because of data x. This can be done with a standard estimators (e.g., track length, collision), and MCNP6 uses the (unweighted) track-length estimator:

$$\Sigma_x^j \ell \delta_g. \tag{2}$$

Here Σ_x is the macroscopic cross section for data x^j (this is zero in the case of fission ν, for example), ℓ is the length of track, and δ_g is one if the particle energy is within energy band g and zero otherwise. The tally contribution is not multiplied by the statistical weight of the particle because it is built into the adjoint weighting factor; this is because the fission neutron production in the current (first) generation is multiplied by the particle weight at the time of collision.

Estimating the gains from scattering and fission sources because of x^j requires the adjoint function be taken for the exiting energy and direction weighted by the relative probability of the interaction from data x^j. The options for this are analog or expected-value estimators. For the scattering source, all but a few of the nonelastic scattering reactions involve Dirac-delta functions in their probability density functions (e.g., for elastic scattering, knowing the exiting direction uniquely determines the exiting energy). For this reason, the analog and expected-value estimators are often identical, and therefore the analog estimator is used in MCNP6:

$$\delta_{s,x} \delta_g. \tag{3}$$

Namely, one is added to the tally if scattering reaction x^j occurs within energy band g.

The fission source term does not have this type of behavior, having a continuum of energies independent of direction, and therefore an expected-value estimate is most appropriate. This is

$$\frac{\chi^j \nu^j \Sigma_f^j}{\chi \nu \Sigma_f} \delta_{f,x} \delta_g. \tag{4}$$

Here the numerator is for fission reaction x^j and the denominator is for the sum of all fissionable isotopes. Like with the others, the δ terms ensure the tally is only scored if it is an appropriate reaction (i.e., elastic scattering is not scored) and within the energy band g.

The forward tally in the denominator is straightforward as it is just the total number of fission neutrons produced each iteration. This serves only as a normalization, and is independent of which data x^j sensitivities are desired.

3. VERIFICATION

Verification is performed using analytic solutions, direct density perturbations, and comparisons to other codes. For the analytic solution, a multigroup, infinite medium problem is used. Direct density

Table I. Nuclear Data for the Analytic Test Problem

g	σ_t	σ_c	σ_f	ν	χ	σ_{sg1}	σ_{sg2}	σ_{sg3}
1	2	1/2	0	–	5/8	1	1/2	0
2	4	1	0	–	1/4	0	1	2
3	4	1/2	3/2	8/3	1/8	0	0	2

perturbations are obtained in continuous-energy calculations with spheres containing homogeneous mixtures of uranium-tetrafluoride and polyethylene. Code comparisons are performed with published results [11] of TSUNAMI-3D and MONK [12]. Note that the MOX lattice and polyethylene spheres are part of a code comparison exercise by OECD/NEA Expert Group on Uncertainty Analysis for Criticality Safety Assessment (EG UACSA). This exercise is Phase III of a series of work by the Expert Group. The MOX lattice is the first case, and the polyethylene spheres are the third case of Phase III.

3.1. Analytic Test Problem

The analytic test problem used is an infinite medium case with three energy groups. Because of the simple nature of this problem, it is straightforward to obtained a solution for k algebraically. Reference solutions may be obtained by differentiating the solution with respect to each of the data.

The features of the problem physics are as follows: fission may only occur in group 3 (the lowest energy group), fission neutrons may appear, however, in all energy groups, downscattering is restricted to subsequent groups (e.g., groups 1 to 2), and there is no upscattering. Given this, the analytic solution for k is

$$k = \frac{\nu_3 \sigma_{f3} \sigma_{s23}}{\sigma_{R2} \sigma_{R3}} \left[\frac{\sigma_{s12}}{\sigma_{R1}} \chi_1 + \chi_2 + \frac{\sigma_{R2}}{\sigma_{s23}} \chi_3 \right]. \tag{5}$$

Note that σ_{Rg} is the microscopic removal cross section, which is the total cross section minus the within-group scattering cross section. The nuclear data are chosen such that $k = 1$, and are given in Table I. The sensitivity coefficients to the cross sections and fission ν are as follows:

$$S_{k,\sigma_{c1}} = -\frac{\sigma_{c1}}{k} \frac{\nu_3 \sigma_{f3} \sigma_{s23}}{\sigma_{R2} \sigma_{R3}} \frac{\sigma_{s12}}{\sigma_{R1}} \chi_1 = -\frac{5}{24}, \tag{6}$$

$$S_{k,\sigma_{c2}} = -\frac{\sigma_{c2}}{\sigma_{R2}} \left[1 - \frac{\chi_3}{k} \frac{\nu_3 \sigma_{f3}}{\sigma_{R3}} \right] = -\frac{1}{4}, \tag{7}$$

$$S_{k,\sigma_{c3}} = -\frac{\sigma_{c3}}{\sigma_{R3}} = -\frac{1}{4}, \tag{8}$$

$$S_{k,\sigma_{s12}} = \frac{\sigma_{s12}}{k} \frac{\nu_3 \sigma_{f3} \sigma_{s23}}{\sigma_{R1} \sigma_{R2} \sigma_{R3}} \left[1 - \frac{\sigma_{s12}}{\sigma_{R1}} \right] \chi_1 = \frac{5}{24}, \tag{9}$$

$$S_{k,\sigma_{s23}} = \frac{\sigma_{s23}}{k} \frac{\nu \sigma_{f3}}{\sigma_{R2} \sigma_{R3}} \left[1 - \frac{\sigma_{s23}}{\sigma_{R2}} \right] \left[\frac{\sigma_{s12}}{\sigma_{R1}} \chi_1 + \chi_2 \right] = \frac{1}{4}, \tag{10}$$

$$S_{k,\sigma_{f3}} = 1 - \frac{\sigma_{f3}}{\sigma_{R3}} = \frac{1}{4}, \tag{11}$$

$$S_{k,\sigma_\nu} = 1. \tag{12}$$

Table II. Sensitivity Results for the Analytic Problem

x	Exact $S_{k,x}$	MCNP6 $S_{k,x}$	C/E
σ_{c1}	$-5/24$	$-0.20868 \pm 0.10\%$	1.002
σ_{c2}	$-1/4$	$-0.24993 \pm 0.07\%$	0.999
σ_{c3}	$-1/4$	$-0.24985 \pm 0.05\%$	0.999
σ_{s12}	$+5/24$	$+0.20810 \pm 0.16\%$	0.999
σ_{s23}	$+1/4$	$+0.25083 \pm 0.15\%$	1.003
σ_{f3}	$+1/4$	$+0.25045 \pm 0.16\%$	1.002
ν_3	$+1$	$+1.00000 \pm 0.00\%$	1.000

MCNP6 calculations of the sensitivity coefficients are made and the results are compared to the analytic solutions in Table II. The C/E value is the calculated to the expected value, and measures the discrepancy between the calculated and reference answer. The MCNP6 calculations all agree with their corresponding analytic solutions within a few tenths of a percent.

Note that it is also possible to take derivatives with respect to the fission χ as well. Because the distribution must be renormalized, there is an additional step involved. This is done specifically in the companion paper.

3.2. Direct Density Perturbations

The sensitivity to the total cross section can be approximated by increasing the density of a particular nuclide by some small fractional amount f, finding the perturbed eigenvalue (call it k' such that the unperturbed eigenvalue is simply k) from a direct calculation (in this case with MCNP6), and taking

$$S_{k,x} \approx \frac{1}{f}\frac{\Delta k}{k}, \tag{13}$$

where $\Delta k = k' - k$. So long as f is sufficiently small such that k is an almost linear function of σ_t^j, the approximation should be valid. Unfortunately, at times it can be difficult to directly resolve $S_{k,x}$ with a Monte Carlo calculation if that sensitivity is quite small. This is especially true when the reason for the small sensitivity is because of large, but equally offsetting positive and negative effects. In order to resolve the difference, a fairly large f must be chosen such that the linear assumption becomes questionable, as the offsetting effects do not necessarily vary at the same rate.

The test case consists of two polyethylene spheres surrounded by vacuum boundaries. The first sphere contains low-enriched uranium (LEU) at 2 atom percent ^{235}U . The second sphere contains intermediate-enriched uranium (IEU) at 50 atom percent ^{235}U . The radii of the spheres are chosen such that they are approximately critical.

Isotopic densities are increased by f as shown in Table III. The magnitudes of f are chosen such that Δk can be resolved statistically with a reasonable amount of computational effort.

Table III. Comparison of Adjoint and Direct Methods for Computing Total Cross-Section Sensitivities with MCNP6

Sphere	Isotope	f	Adjoint	Direct	C/E
LEU	^{1}H	0.01	$+2.40 \times 10^{-1} \pm 0.4\%$	$+2.34 \times 10^{-1} \pm 1.3\%$	1.03
	C	0.05	$+2.74 \times 10^{-2} \pm 0.7\%$	$+2.74 \times 10^{-2} \pm 2.2\%$	1.00
	^{19}F	0.02	$+4.38 \times 10^{-2} \pm 0.5\%$	$+4.17 \times 10^{-2} \pm 3.7\%$	1.05
	^{235}U	0.01	$+2.53 \times 10^{-1} \pm 0.1\%$	$+2.55 \times 10^{-1} \pm 1.2\%$	0.99
	^{238}U	0.01	$-2.01 \times 10^{-1} \pm 0.2\%$	$-2.03 \times 10^{-1} \pm 1.5\%$	0.99
IEU	^{1}H	0.01	$+4.54 \times 10^{-1} \pm 0.2\%$	$+4.57 \times 10^{-1} \pm 0.7\%$	0.99
	C	0.05	$+6.54 \times 10^{-2} \pm 0.3\%$	$+6.46 \times 10^{-2} \pm 1.0\%$	1.01
	^{19}F	0.02	$+1.18 \times 10^{-1} \pm 0.2\%$	$+1.17 \times 10^{-1} \pm 1.3\%$	1.01
	^{235}U	0.01	$+1.30 \times 10^{-1} \pm 0.2\%$	$+1.25 \times 10^{-1} \pm 2.5\%$	1.04
	^{238}U	0.20	$-1.57 \times 10^{-3} \pm 5.4\%$	$-5.22 \times 10^{-4} \pm 20.3\%$	3.01

MCNP6 calculations of the sensitivities are made for each sphere using ENDF/B-VII.0 data. These are given in Table III as well, and listed as "Adjoint"; the direct results are (intuitively) listed under "Direct". Except for the outlier ^{238}U sensitivity in the IEU sphere, all of the adjoint-based sensitivities agree within 5% of the direct density perturbation results. The ^{1}H sensitivity for the LEU sphere is slightly outside of the 2-σ confidence band, but all the others (except the ^{238}U outlier) are within.

The sensitivity of ^{238}U in the IEU sphere is quite small relative to the those of the other isotopes, and it required a 20% increase in the ^{238}U density to get a statistically significant Δk. The reason ^{238}U is so small is not because of a lack of abundance, but because the effect from elastic scattering and capture almost exactly cancels. There is no expectation that these effects should both increase linearly at the same rate for a 20% perturbation, and the results are expected to be anomalous.

3.3. Software Comparisons

Other software packages exist for calculating sensitivity coefficients, and some, like TSUNAMI-3D, have been used regularly for about a decade using multigroup physics. The software package MONK has recently released a continuous-energy coefficient capability based on the differential operator method with a fission source correction. Both of TSUNAMI-3D and MONK participated in the EG UACSA Phase III benchmark exercise for comparing codes. It therefore seems logical to see how MCNP6 results compare with other software packages with similar capabilities. The results used for comparison are those generated by Oak Ridge National Laboratory (ORNL) for TSUNAMI-3D with ENDF/B-VII.0 with and the Atomic Weapons Establishment (AWE) for MONK with JEF-2.2.

The polyethylene spheres discussed in the last section are part of the Phase III benchmark exercise (third problem, to be precise). The results of the total cross section sensitivities for MCNP6 versus TSUNAMI-3D and MONK are given in Table IV. The statistical uncertainties of the MCNP6 results are identical to those in Table III. Unfortunately, no statistical uncertainties for either TSUNAMI-3D or MONK are provided. The C/E values in Table IV are MCNP6 with respect to the code in the column immediately to the left.

Table IV. Software Comparisons for Total Cross-Section Sensitivities of the Polyethylene Spheres

Sphere	Isotope	MCNP6	TSUNAMI-3D	C/E	MONK	C/E
LEU	^{1}H	$+2.40 \times 10^{-1}$	$+2.41 \times 10^{-1}$	1.00	$+2.45 \times 10^{-1}$	0.98
	C	$+2.74 \times 10^{-2}$	$+2.77 \times 10^{-2}$	0.99	$+2.77 \times 10^{-2}$	0.99
	^{19}F	$+4.38 \times 10^{-2}$	$+4.56 \times 10^{-2}$	0.96	$+4.20 \times 10^{-2}$	1.04
	^{235}U	$+2.53 \times 10^{-1}$	$+2.53 \times 10^{-1}$	1.00	$+2.58 \times 10^{-1}$	0.98
	^{238}U	-2.01×10^{-1}	-1.95×10^{-1}	1.03	-2.01×10^{-1}	1.00
IEU	^{1}H	$+4.54 \times 10^{-1}$	$+4.55 \times 10^{-1}$	1.00	$+4.52 \times 10^{-1}$	1.00
	C	$+6.54 \times 10^{-2}$	$+6.60 \times 10^{-2}$	0.99	$+6.61 \times 10^{-2}$	0.99
	^{19}F	$+1.18 \times 10^{-1}$	$+1.19 \times 10^{-1}$	0.99	$+1.15 \times 10^{-1}$	1.03
	^{235}U	$+1.30 \times 10^{-1}$	$+1.26 \times 10^{-1}$	1.03	$+1.36 \times 10^{-1}$	0.96
	^{238}U	-1.57×10^{-3}	$+1.35 \times 10^{-3}$	-1.16	-1.30×10^{-3}	1.21

Aside from the outlier ^{238}U IEU sensitivity, all codes agree within 5%. There does not appear to be any trend in terms of one code being consistently high or consistently low, although this sample size is too small to make a statistically significant determination on such.

The continuous-energy results shown so far are energy-integrated and only for total cross sections. MCNP6 can, however, generate energy-resolved sensitivities for any reaction on an arbitrary, user-defined energy grid.

Energy-resolved results for comparison are available for both TSUNAMI-3D and MONK for the MOX lattice benchmark – problem 1 of the EG UACSA Phase III benchmark exercise. The problem is a lattice of MOX fuel pins, with a pitch of 0.9525 cm, submerged in water. The lattice is 28 $\times$ 22, except for the top row, which has only 17 pins, for a total of 605 pins. The fuel pins have a diameter of 0.5842 cm and a height of 91.44 cm. The areas axially above and below are buffers that are treated as homogenized mixtures of water and pin materials.

The energy-resolved sensitivities are computed by MCNP6 and MONK on a 238-group energy grid that aligns with the standard SCALE multigroup library. MCNP6, TSUNAMI-3D, and MONK results for the ^{238}U total cross-section sensitivities are given in Fig. 1, and the results for ^{1}H elastic scattering cross-section sensitivities are given in Fig. 2. The results mostly agree in shape and magnitude, indicating all three codes are calculating energy-resolved sensitivities consistently.

4. CODE PERFORMANCE

Even the most accurate method is worthless to an end user if it cannot obtain results in a reasonable amount of time or its memory requirements are prohibitively large. An assessment of the wall-clock time requirements is made using the polyethylene spheres. Then, a problem with a very large number of sensitivity coefficients (about 1.6 million) is obtained in the MOX lattice benchmark and the memory requirements are given. Both the run-time and memory requirements are shown to be reasonable by modern standards.

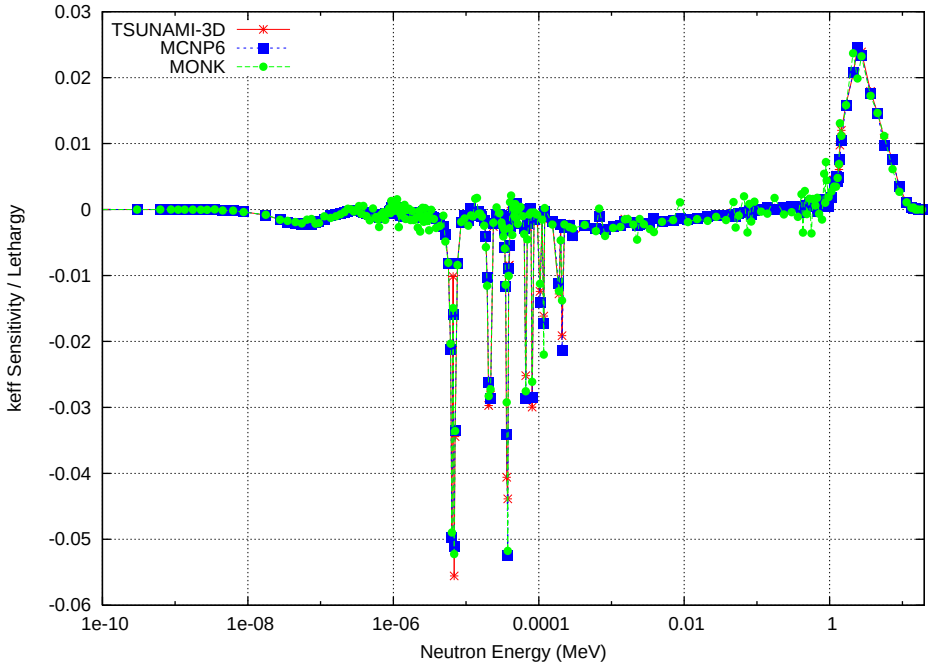

Figure 1. Code comparison of energy-resolved ^{238}U total cross-section sensitivities for the MOX lattice.

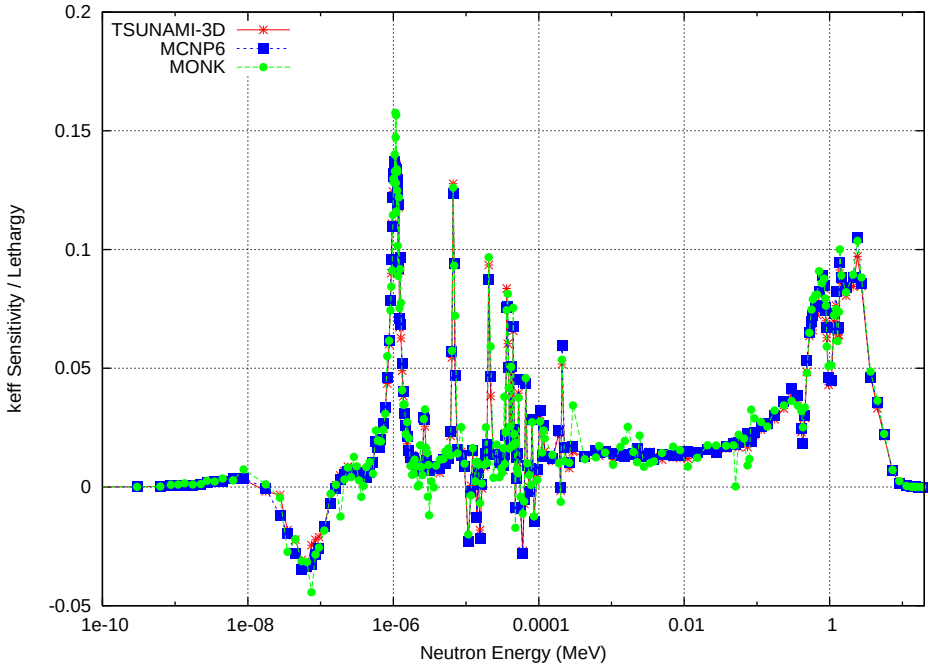

Figure 2. Code comparison of energy-resolved ^{1}H elastic scattering cross-section sensitivities for the MOX lattice.

4.1. Run-Time Performance and Computational Efficiency

The goal of a Monte Carlo calculation is to achieve the lowest statistical uncertainties for the least amount of computational time – depending on the situation this can be either wall-clock or computer (clock-cycle) time. The polyethylene spheres are used with batch size of 50,000 neutrons, 50 skip cycles, 1000 active cycles, and a block size of 10 generations for the adjoint-weighted sensitivity tallies are used. The calculations were performed on Intel Xeon E5-2670 processors at 2.6 GHz with 32 GB RAM. The adjoint methods work both with Message Passing Interface (MPI) and Open Multi-Processing (OMP) threading forms of parallelism in MCNP6, and all calculations are run with 16 OMP threads. The energy-resolved and energy-integrated sensitivities for 53 different reactions – total, total elastic, total inelastic, all 40 individual inelastic levels in ENDF, inelastic continuum, (n,2n), fission, fission ν, total capture, (n,γ), (n,d), (n,t), (n,^{3}He), and (n,α) – are calculated for each of the five isotopes on the standard 238-energy group structure in TSUNAMI-3D (total of 63,335 sensitivities).

First a calculation is performed for both the LEU and IEU spheres with and without the adjoint-weighted tallies, and the increase in wall-clock time from the additional overhead is measured (using the Linux/UNIX "time" utility). Second, the percent uncertainties of the energy-integrated total, elastic, and capture cross-section sensitivities for each isotope are given along with the figure of merit (FOM) – the FOM is the inverse of the product of the computation time and the square of the relative uncertainty, where the computational time used here is the wall-clock time multiplied by the number of threads (in this case, 16). These results are shown in Table V.

The increased overhead for employing the adjoint methods for this problem leads to increases in wall-clock time by factor of 4.2 for the LEU sphere and 3.4 for the IEU sphere. Computing a relatively large number sensitivities is not a trivial task for MCNP6, and calculations may take significantly longer than a traditional eigenvalue calculation. Obtaining just the 15 sensitivities displayed in Table V increases computational run times by 5-10%. Exactly how much extra time is needed depends on the number and type of sensitivities required and the physics of the problem – it tends to scale with the expected number of collisions per history.

Table V. MCNP6 Run-Time Performance Data

	Time (min)			Total		Elastic		Capture	
	Base	Adjoint		% Unc	FOM	% Unc	FOM	% Unc	FOM
LEU	9.9	41.5	^{1}H	5.7	4.7×10^{-1}	2.8	1.9×10^{0}	0.2	3.1×10^{2}
			C	10.9	1.3×10^{-1}	10.8	1.3×10^{-1}	0.9	1.9×10^{1}
			^{19}F	9.2	1.8×10^{-1}	12.4	9.9×10^{-2}	0.5	6.5×10^{1}
			^{235}U	0.7	3.6×10^{1}	75.4	2.6×10^{-3}	0.2	3.8×10^{2}
			^{238}U	1.9	4.1×10^{0}	12.7	9.3×10^{-2}	0.2	2.8×10^{2}
IEU	4.3	14.7	^{1}H	1.3	2.5×10^{1}	1.3	2.4×10^{1}	0.4	3.1×10^{2}
			C	4.0	2.7×10^{0}	4.0	2.6×10^{0}	2.7	5.9×10^{0}
			^{19}F	2.8	5.6×10^{0}	3.7	3.1×10^{0}	0.8	7.2×10^{1}
			^{235}U	2.3	8.3×10^{0}	11.7	3.1×10^{-1}	0.3	5.8×10^{2}
			^{238}U	111.4	3.4×10^{-3}	8.8	5.5×10^{-1}	0.6	1.1×10^{2}

As for the figure of merit values, the absolute magnitudes are not all that relevant since they are system dependent. The trend, however, is for capture cross-section sensitivities to be easier to calculate than the others, as indicated by their higher figures of merit. This is because the term is purely negative, and such quantities are far easier to estimate than those that may be either positive or negative. While not given here, the same trend applies to fission-ν sensitivities, which is a purely positive quantity. Scattering cross-section sensitivities tend to be more difficult to compute, and usually have significantly lower figures of merit. Fission cross-section sensitivities (not shown here) tend to be easier than those for scattering cross sections, but more difficult than capture or fission-ν sensitivities. Total cross-section sensitivities are completely problem dependent, as evidenced by the large spread in figures of merit, with the pathological ^{238}U total cross-section sensitivity in the IEU subcase being particularly notable. Whenever positive and negative effects almost exactly balance, that corresponding sensitivity tends to be very difficult to estimate. As a general trend, isotopes with greater abundance tend to have higher figures of merit, as expected.

4.2. Memory Requirements

To illustrate memory requirements, consider the MOX lattice problem where a large number of sensitivity coefficients are desired. The problem specification includes 36 isotopes, 9 of which are fissionable. Suppose for each of these 36 isotopes, 53 different reactions (same as the run-time performance calculations) are desired using the standard 238-energy group structure in TSUNAMI-3D. Additionally, fission-χ sensitivities are desired for the 9 fissionable isotopes as a function of both incident and outgoing energy on a 238×238 energy grid (same group structure). Finally, total, elastic, and inelastic scattering law energy transfer distributions are desired for ^{1}H , ^{16}O , ^{56}Fe , and ^{238}U on the same 2-D energy grid. This leads to a total of about 1.64 million sensitivity coefficients. Note that many of these will not be used because many of the reactions have thresholds or, for the case of the scattering energy transfer distributions, are kinematically impossible.

Just multiplying the number of sensitivity coefficients by a reasonable batch size of 50,000 neutrons would lead to a base memory requirement of about 615 GB. This presumes one 8-byte real number for each sensitivity and source neutron, neglecting overhead. Because the arrays for each original neutron tend to be very sparse, special handling for such sparse data is employed that only stores which data are needed. A test of the present sparse data handling scheme currently implemented in MCNP6 performed on the Turing cluster at LANL (Quad-Core AMD Opteron model 8354 at 2.2 GHz or model 8356 at 2.3 GHz) using one node having 16 CPUs and using all 16 OMP threads shows that the memory requirements, as determined by the UNIX/Linux “top” utility, are only about 10 GB (including all of the normal overhead, of which storing the original contributions for the sensitivity coefficients are the vast majority), or a savings of over 98% relative to the theoretically required 615 GB. This shows that even for a over a million sensitivities, the implementation in MCNP6 is such that while the memory requirements can be large, they are likely not prohibitive.

5. CONCLUSIONS & FUTURE WORK

MCNP6 has a new capability for computing sensitivity coefficients to nuclear data using continuous-energy Monte Carlo methods. Specifically, the IFP method is used to account for the adjoint weighting. The method works by recording contributions to tallies and later weighting them by their respective importance functions. Verification on the method is performed by comparisons to analytic solutions, direct density perturbations, and to other software. Results of comparisons to analytic solutions

show agreement within 2-σ confidence bands, which are within a few tenths of a percent. Comparisons with direct density perturbations and other software usually agree within 5%, except for cases that are known to be problematic. Code performance of run-time and memory requirements are assessed and found to be reasonable.

In this paper, specific focus is given to only cross-section and fission-ν sensitivities. MCNP6, however, has the capability to compute sensitivities to fission χ and scattering distributions, and this is covered in a companion paper.

Future efforts will be devoted to developing a capability within MCNP6 to use sensitivity profiles in conjunction with covariance data to perform estimates on the uncertainties of k. Further research will also go into improving the efficiency of these methods, and extending their capability.

ACKNOWLEDGMENTS

Funding was provided by the U.S. DOE/NNSA Nuclear Criticality Safety Program. Special thanks to Tatiana Ivanova at IRSN, Keith Bledsoe and ORNL, and James Dyrda at AWE for providing reference results for code comparisons.

REFERENCES

[1] B.T. Rearden, "Perturbation Theory Eigenvalue Sensitivity Analysis with Monte Carlo Techniques," *Nucl. Sci. Eng.*, **146**, pp. 367-382 (2004).

[2] B.T. Rearden, "TSUNAMI-3D: Control Module for Three-Dimensional Cross-Section Sensitivity and Uncertainty Analysis for Criticality," ORNL/TM-2005/39 Version 6, Oak Ridge National Laboratory (2009).

[3] MCNP6 Development Team: J.T. Goorley, M.R James, T.E. Booth, F.B. Brown, J.S. Bull, L.J. Cox, J.W. Durkee, J.S. Elson, M.L. Fensin, R.A. Forster, J.S. Hendricks, H.G. Hughes, R.D. Johns, B.C. Kiedrowski, R.L. Martz, S.G. Mashnik, G.W. McKinney, D.B. Pelowitz, R.E. Prael, J.E. Sweezy, L.S. Waters, T. Wilcox, T.J. Zukaitis, "Initial MCNP6 Release Overview", *J. Nucl. Technol.*, **180**, pp. 298-315 (2012).

[4] B.C. Kiedrowski, F.B. Brown, P.P.H. Wilson, "Adjoint-Weighted Tallies for k-Eigenvalue Calculations with Continuous-Energy Monte Carlo," *Nucl. Sci. Eng.*, **168**, pp. 226-241 (2011).

[5] B.C. Kiedrowski, F.B. Brown, "K-Eigenvalue Sensitivities of Secondary Distributions of Continuous-Energy Data," *Proc. ANS Mathematics & Computation Topical Mtg.*, Sun Valley, ID, USA, May 5-9 (2013).

[6] J.D. Lewins, "The Time-Dependent Importance of Neutrons and Precursors," *Nucl. Sci. Eng.*, **7**, pp. 268-274 (1960).

[7] H. Hurwitz, *Naval Reactor Physics Handbook*, A. Radkowsky (Ed.), Vol I, 864-849, Naval Reactors, U.S. Atomic Energy Commission (1964).

[8] K.F. Raskach, A.A. Blyskavka, "An Experience of Applying Iterated Fission Probability Method to Calculation of Effective Kinetics Parameters and keff Sensitivities with Monte Carlo," *Proc. PHYSOR 2010 – Advances in Reactor Physics to Power the Nuclear Renaissance*, Pittsburgh, Pennsylvania, USA, May 9-14 (2010).

[9] H.J. Shim, C.H. Kim, "Adjoint Sensitivity and Uncertainty Analyses in Monte Carlo Forward Calculations," *J. Nucl. Sci. Technol.*, **48**, pp. 1453-1461 (2011).

[10] C.M. Perfetti, W.R. Martin, B.T. Rearden, M.L. Williams, "Development of Continuous-Energy Eigenvalue Sensitivity Coefficient Calculation Methods in the Shift Mote Carlo Code," *Proc. PHYSOR 2012 – Advances in Reactor Physics*, Knoxville, TN, USA, Apr. 15-20 (2012).

[11] T. Ivanova, C. Laville, J. Dyrda, D. Mennerdahl, Y. Golovko, K. Raskach, A. Tsiboulia, G. Soo Lee, S.W. Woo, A. Bidaud, P. Sabouri, A. Patel, K.C. Bledsoe, B.T. Rearden, J. Gulliford, F. Michel-Sendis, "OECD/NEA Expert Group on Uncertainty Analysis for Criticality Safety Assessment: Results of Benchmark on Sensitivity Calculation (Phase III)," *Proc. PHYSOR 2012 – Advances in Reactor Physics*, Knoxville, TN, USA, Apr. 15-20 (2012).

[12] "MONK9: A Monte Carlo Program for Nuclear Criticality Safety and Reactor Physics Analyses," ANSWERS (2008).