LA-UR-12-22089
Approved for public release; distribution is unlimited.

Title: Adjoint-Based k-Eigenvalue Sensitivity Coefficients to Nuclear Data Using Continuous-Energy Monte Carlo

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

Intended for: Nuclear Science & Engineering

Adjoint-Based k-Eigenvalue Sensitivity Coefficients to Nuclear Data Using Continuous-Energy Monte Carlo

Brian C. Kiedrowski and Forrest B. Brown

X-Computational Physics Division

Los Alamos National Laboratory

P.O. Box 1663, MS A143

Los Alamos, NM 87545 USA

Corresponding Author: Brian C. Kiedrowski (bckiedro@lanl.gov)

Page Count: 53

Table Count: 9

Figure Count: 9

Adjoint-Based k-Eigenvalue Sensitivity Coefficients to Nuclear Data Using Continuous-Energy Monte Carlo

Brian C. Kiedrowski, Forrest B. Brown

Abstract – A continuous-energy Monte Carlo method is developed to compute adjoint-based k-eigenvalue sensitivity coefficients with respect to nuclear data. The method is implemented into MCNP6 and is based upon similar methodologies used to compute other adjoint-weighted quantities. The Monte Carlo tallies employed are explained. Verification of the method is performed by comparing results to analytic solutions, direct density perturbations, and those from other software packages such as TSUNAMI-3D and MONK. Results of analytic solutions agree within a few tenths of a percent. Direct density perturbations and comparisons with other software generally agree within a few percent.

Keywords – MCNP; uncertainty analysis; perturbation; cross section

1 INTRODUCTION

Monte Carlo software packages have been used for computing sensitivity coefficients to the effective multiplication factor k for sensitivity and uncertainty analysis for several years [1, 2, 3, 4]. To date, most implementations, such as TSUNAMI-3D [5] that is part of the SCALE package developed by Oak Ridge National Laboratory (ORNL), have used multigroup cross-section data. This presents difficulties involving group collapse and self shielding, necessitating the calculation of implicit sensitivity coefficients, which is the impact of the multigroup approximation on the sensitivity coefficients. Because of these complications, and the desire for higher fidelity with a simpler work flow, there has been a push within the US Department of Energy National Nuclear Security Agency (DOE/NNSA) Nuclear Criticality Safety Program (NCSP) to develop continuous-energy sensitivity coefficient capabilities.

Such a capability, based on adjoint methodologies similar to those employed in TSUNAMI-3D, has been developed by Los Alamos National Laboratory (LANL) in MCNP6 [6] and is available in the initial production version. Also included in MCNP6 is a related capability, called KPERT, for adjoint-weighted perturbations from material substitutions, which have been employed for sensitivity coefficient calculations [7]. While material substitutions allow for more general perturbations, practical limitations lead to the introduction of an approximation in the handling of scattering laws that leads to large and unacceptable deviations in scattering sensitivities. Additionally, the interface is designed with a different intent, and some users may find it cumbersome to use for sensitivity coefficient calculations.

For these reasons, a similar capability, called KSEN, that is more accurate and efficient, and easier to use than KPERT for this purpose, has been developed. The details of the Monte Carlo implementation and the associated tallies are discussed. Verification of the capability is performed by comparing results against analytic solutions of simple problems, direct density perturbations, and those generated by other software. The benchmarks selected for inter-software comparisons are the Organisation for Economic Cooperation and Development/Nuclear Energy Agency, Working Party on Nuclear Criticality Safety Expert

Group on Uncertainty Analysis for Criticality Safety Assessment Phase III Benchmarks [8] (called OECD/NEA UACSA Phase III Benchmarks for short). Results of sensitivities to angular scattering distributions for ^{239}Pu in Jezebel are also given. Finally, performance and computational requirements of the method and implementation are discussed.

2 MCNP6 KSEN METHOD

The method for calculating sensitivity coefficients in MCNP6 employs the adjoint-based methodology that is used in TSUNAMI-3D. First, the theory and mathematical models governing the adjoint-based method are explained, and then the tallies employed in MCNP6 are discussed.

2.1 Sensitivity Coefficient Background

From linear perturbation theory, the following relation can find a differential change in k resulting from a differential change in cross section:

$$dk = -\frac{\left\langle \psi^{\dagger}, \left(d\Sigma_t - d\mathcal{S} - \lambda d\mathcal{F}\right)\psi \right\rangle}{\left\langle \psi^{\dagger}, \lambda^2 \mathcal{F}\psi \right\rangle}, \tag{1}$$

where ψ is the forward flux and $\psi^{\dagger}$ is its adjoint, $\lambda = 1/k$, Σ_t is the macroscopic total cross section, $\mathcal{S}$ is the integral scattering operator, $\mathcal{F}$ is the total integral fission operator, and the brackets denote integration over all space, angle, and energy variables.

For sensitivity and uncertainty analysis, the sensitivity coefficient for k, $S^j_{k,x}$, to some nuclear data x^j (x is a cross section, fission ν, etc. over some energy range, and j is a specific isotope) is desired. The sensitivity coefficient is defined as the ratio of the fractional change in k for a corresponding fractional change in x^j, or

$$S_{k,x} \equiv \frac{x^j}{k}\frac{dk}{dx^j}. \tag{2}$$

By manipulating Eq. (1),

$$S_{k,x}^{j} = \frac{\left\langle \psi^{\dagger}, \mathcal{P}_{x}^{j}\psi \right\rangle}{\left\langle \psi^{\dagger}, \lambda\mathcal{F}\psi \right\rangle}, \tag{3}$$

where $\mathcal{P}_{x}^{j}$ is the perturbation operator defined as

$$\mathcal{P}_{x}^{j} = \left(-N^{j}x^{j}\delta_{x,t} + \mathcal{S}_{x}^{j} + \lambda\mathcal{F}_{x}^{j}\right)\delta_{g}\delta_{z}, \tag{4}$$

$\mathcal{S}_{x}^{j}$ is the scattering operator for x^{j} exclusively,

$$\mathcal{S}_{x}^{j} = N^{j}\iint dE'\, d\mathbf{\Omega}' \sum_{l=1}^{L} \delta_{x,sl}\, f_{sl}^{j}(E, \hat{\mathbf{\Omega}}\,|E', \hat{\mathbf{\Omega}}')\, m_{l}\, \sigma_{sl}^{j}, \tag{5}$$

and $\mathcal{F}_{x}^{j}$ is the total fission operator for x^{j} exclusively,

$$\mathcal{F}_{x}^{j} = \frac{N^{j}}{4\pi}\iint dE'\, d\mathbf{\Omega}' \sum_{c=1}^{C}\left[\delta_{x,pc}\chi_{pc}^{j}(E|E')\nu_{p}^{j}\sigma_{fc}^{j}\right] + \sum_{i=1}^{I}\left[\delta_{x,di}\chi_{di}^{j}(E|E')\nu_{di}^{j}\sigma_{f}^{j}\right]. \tag{6}$$

The variables are as follows:

E' = incident energy variable,

$\hat{\mathbf{\Omega}}'$ = incident direction unit vector,

E = outgoing energy variable,

$\hat{\mathbf{\Omega}}$ = outgoing direction unit vector,

N^{j} = atomic density of isotope j,

$\delta_{x,t}$ = one if x is a cross section and zero otherwise,

δ_{g} = one if energy of neutron is within specified range g and zero otherwise,

δ_{z} = one if neutron is within spatial region z and zero otherwise,

l = index for the scattering reaction,

$\delta_{x,sl}$ = one if x is involved with scattering reaction l and zero otherwise,

f_{sl} = scattering energy-direction transfer function for reaction l,

m_{l} = multiplicity of the scattering event [e.g., (n,2n) has $m_{l} = 2$],

σ_{sl} = scattering cross section for reaction l,

c = index for the chance of prompt fission,

$\delta_{x,pc}$ = one if x is involved with prompt fission chance c and zero otherwise,

χ_{pc} = prompt fission energy transfer function for chance c,

ν_p = mean prompt neutron emission,

σ_{fc} = fission cross section for prompt chance c,

i = delayed neutron precursor index,

$\delta_{x,pi}$ = one if x is involved with delayed fission precursor group i and zero otherwise,

χ_{di} = delayed fission energy transfer function for precursor group i,

ν_{di} = mean delayed neutron emission for precursor group i,

σ_f = total fission cross section.

Another convention defined is that the energies used to determine range g are incident for cross sections and fission ν, and outgoing for scattering or fission-χ transfer functions. For the transfer functions, sensitivities may also be further restricted to incident energy E' (with bin index g'), and (for scattering) direction cosine change μ (with bin index n).

The energy-resolved sensitivities computed are actually bin-integrated. For the case of a cross-section or fission-ν sensitivity,

$$S^j_{k,x,g} = \int_{E_{g-1}}^{E_g} dE\; S^j_{k,x}(E), \tag{7}$$

where the integrand $S^j_{k,x}(E)$, with explicit dependence on E, is taken to the be the "sensitivity density" with units of per energy. The energy-integrated sensitivity is

$$S^j_{k,x} = \int_0^\infty dE\; S^j_{k,x}(E). \tag{8}$$

The sensitivities to the scattering law or fission-χ transfer functions are defined similarly, except they can have direction cosine change and incident-energy bins.

The transfer functions themselves are probability density functions conditional on the

incident energy (the standard notation for conditional probability densities is used where the variables to the left of the vertical bar are conditional on the variables to the right), and are normalized to unity. Because of this constraint, the total sensitivities over all outgoing energies and direction changes for a given incident-energy bin must sum to zero – increasing the transfer function somewhere must be offset by a corresponding decrease somewhere else. For this reason, MCNP6 computes the constrained sensitivity coefficient for transfer functions, $\hat{S}^j_{k,x}$. Note that the sensitivities for cross sections, such as fission or scattering, are not constrained, as there is no normalization condition to impose, and are therefore unaffected.

For scattering laws,

$$\hat{S}^j_{k,f}(E,\mu|E') = S^j_{k,f}(E,\mu|E') - f^j(E,\mu|E')S^j_{k,f}(E'), \tag{9}$$

and for fission χ,

$$\hat{S}^j_{k,\chi}(E|E') = S^j_{k,\chi}(E|E') - \chi^j(E|E')S^j_{k,\chi}(E'). \tag{10}$$

These relationships are continuous in energy (incident and outgoing) and angle. In practice, discrete forms of these are needed to do calculations. The respective discrete forms of these are

$$\hat{S}^j_{k,f,g,g',n} = S^j_{k,f,g,g',n} - f^j_{g,g',n}S^j_{k,f,g'}, \tag{11}$$

$$\hat{S}^j_{k,\chi,g,g'} = S^j_{k,\chi,g,g'} - \chi^j_{g,g'}S^j_{k,\chi,g'}. \tag{12}$$

To account for the incident energy dependence of f^j and χ^j within the chosen interval, a weighted-average transfer function is used, where the weighting function is the neutron secondary production from the reaction of interest:

$$f^j_{g,g',n} = \frac{\int_{\mu_{n-1}}^{\mu_n} d\mu \int_{E_{g-1}}^{E_g} dE \int_{E'_{g'-1}}^{E'_{g'}} dE'\; f^j m_l N^j \sigma^j_s}{\int_{-1}^{1} d\mu \int_{0}^{\infty} dE \int_{E'_{g'-1}}^{E'_{g'}} dE'\; f^j m_l N^j \sigma^j_s}, \tag{13}$$

$$\chi^j_{g,g'} = \frac{\int_{E_{g-1}}^{E_g} dE \int_{E'_{g'-1}}^{E'_{g'}} dE' \, \chi^j N^j \nu^j \sigma^j_f}{\int_0^{\infty} dE \int_{E'_{g'-1}}^{E'_{g'}} dE' \, \chi^j N^j \nu^j \sigma^j_f}. \tag{14}$$

There is a problem, however, with averaging and integrating over the incident energy dependence. The $f^j_{g,g',n}$ and $\chi^j_{g,g'}$ must be computed on a fine enough incident-energy resolution such that the transfer function shape does not change significantly within the incident-energy interval. This assumption works reasonably well for fission χ, which is a weak function of incident energy over the range of interest, but does not apply nearly as well for scattering laws, which can vary significantly as a function of incident energy.

Picking an inappropriately large incident-energy bin will yield a different constrained sensitivity than the sum of constrained sensitivities computed from finely resolved incident-energy bins. This is different than unconstrained sensitivities (or constrained sensitivities over outgoing energy bins or direction changes), which are always additive. For now, it is up to the user to pick an appropriately fine incident-energy grid to perform the calculations. In the future, research will be done to automate the process so that appropriate values of the transfer functions will be used and the sums will come out correct regardless of the incident-energy binning chosen.

One final comment on the sensitivities of transfer functions is that MCNP6 uses outgoing energies and directions that are on the table. For all but one inelastic scattering law, correlated energy-angle scattering (law 67 in MCNP), this is in the center-of-mass frame. Note that some of the results in this paper explicitly are in lab-frame coordinates to test the method, in contradiction to the usual MCNP6 convention. Which reference frame used for each scattering law sensitivity calculation is explicitly stated throughout this paper.

2.2 Tallies and Scoring

As seen from Eqs. (3) and (4), there are adjoint-weighted tallies involving total interactions, scattering, and fission for x^j and one for the adjoint-weighted fission source in the denominator – a total of four, but, in practice, only one is needed because the three tallies

in the numerator can be combined and the denominator does not require any extra storage. MCNP6 uses the Iterated Fission Probability methodology [9] to compute adjoint-weighted tallies.

The Iterated Fission Probability methodology relies on the interpretation of the adjoint function as the expected number of neutrons in a system (or, more generally, any response), after infinitely many generations, caused by a neutron at a location in phase space [10, 11] (e.g., arising from a reaction of some sort in this case). To facilitate computing the adjoint-weighted tallies, the active cycles (batches, generations) are grouped into B contiguous blocks with index b of fixed size. In the first cycle within the block, original (non-adjoint weighted) contributions T_q of tallies are evaluated and stored (the variables on the right side of $< \cdot , \cdot >$) and the neutrons are tagged with index q (the total number of these indices within a block is Q) so that later which neutron is associated with which contribution is known. The index q is distinct from the history index, because whenever a neutron history branches, such as in an (n,2n) reaction or when an implicit capture event creates a fission neutron and the original neutron continues, a new index is needed. These tags are inherited by subsequent progeny within the block of generations. In the final cycle within the block, fission neutron production estimates (asymptotic populations) R_q are made and summed for all neutrons with the same tag and multiplied by the associated contributions T_q to obtain a score for the adjoint-weighted tally.

With regards to the original contributions, either track-length or collision estimators may be used for the total interaction term, and either analog or expected-value estimators may be used for the scattering or fission terms. In MCNP6, the track-length estimator is used for the total interaction term, the analog estimator is used for the scattering term, and the expected-value estimator is used for the fission term. The reasons for this are solely based upon convenience. Regardless of which types MCNP6 actually uses, all the mentioned estimators are discussed. The first three are grouped in the (adjoint-weighted) numerator $A^{(b)}$ for block b, and the corresponding (adjoint-weighted) denominator is denoted by $D^{(b)}$.

2.2.1 Total Interaction Term

The total interaction term is computed in MCNP6 by a track-length estimator. For each simulated Monte Carlo track,

$$\delta_{x,t} N^j x^j \ell, \tag{15}$$

(ℓ is the length of the track) is subtracted from the original contribution in the numerator T_q, because this represents a net loss in neutrons.

The collision estimator is

$$\delta_{x,t} \frac{N^j x^j}{\Sigma_t}. \tag{16}$$

The collision estimator is not used by MCNP6 for this purpose but is included for completeness.

2.2.2 Scattering Term

At each collision where a non-fission secondary emerges, the scattering reaction z occurs with some isotope h. The expected-value estimator is the ratio of the double-differential scattering cross section for all x^j for the energy and direction change that occurred to the corresponding total double-differential scattering cross section.

$$\delta_{x,sl} \frac{f_{sl}^j m_l \Sigma_{sl}^j}{f_{sl} m_l \Sigma_{sl}} \tag{17}$$

is added each collision to the numerator original contribution T_q. To clarify, the variables in the numerator with the j superscript are for specific isotopes, and the terms in the denominator sans superscript are for all possible colliding isotopes whether sensitivity coefficients to those isotopes are desired for them or not. Furthermore, all l and j are accumulated independent of the actual collision that occurred; hence, it is called an expected-value estimator.

For most scattering reactions, however, such as elastic scattering with a stationary target or inelastic-level scattering, the double-differential scattering cross section involves a Dirac

delta function, and the ratio is one if $l = z$ and $j = h$ and zero otherwise. While other scattering reactions may have ratios between zero and one, often times these are difficult and time consuming to compute in practice because of the various representations of the Evaluated Nuclear Data File (ENDF) scattering laws used by MCNP6, and therefore an analog estimator may be more appropriate, and that is what is actually used.

The analog estimator added to T_q at each collision is

$$\delta_{x,z}\delta_{j,h}. \tag{18}$$

In other words, one is added only if $x = z$ and $j = h$. Of course, in the limit of infinitely many collisions, both the analog and expected-value estimators are the same.

2.2.3 Fission Term

Like with scattering, neutrons emerge from fission from reaction z (e.g., prompt or delayed and considering various chances of fission) with isotope h. The fission term is calculated for each source emission at the beginning of the block of generations.

Much like for the scattering term, the expected-value estimator for the fission term is

$$\delta_{x,f}\frac{\chi^j\nu^j\Sigma_f^j}{\chi\nu\Sigma_f}, \tag{19}$$

and is added to T_q at each simulated source emission event. Note that since this is performed at source emission at MCNP6, the incident energy causing fission must be stored and transferred between cycles in addition to the outgoing energy that is normally stored.

The analog estimator for the fission term is identical to that of scattering term in Eq. (18) except here z and h are for fission reactions and isotopes.

MCNP6 uses the expected-value estimator for the fission term because fission χ is given in MCNP6 as one of three fairly simple scattering laws where the outgoing energies are specified as tabular data (most common in the newer evaluations), or are computed by evaluating a

function (e.g., Watt fission) where the fitting parameters are specified by tabular data. Furthermore, fission χ is both continuous and univariate (in outgoing energy), being far easier to calculate than the bivariate probability density for scattering, which is dependent on both the outgoing energy and direction cosine change.

2.2.4 Denominator Term

The denominator term does not require storing any information that is not already being computed. Simply finding the average of total asymptotic populations of all neutrons gives λ times the adjoint-weighted fission source.

2.2.5 Scoring and Computing Means and Variances

Once the asymptotic populations R_q are known from the forward simulation, the adjoint-weighted estimates for $A^{(b)}$ and $D^{(b)}$ can be found by

$$A^{(b)} = \frac{1}{Q} \sum_{q=1}^{Q} R_q T_q, \tag{20}$$

$$D^{(b)} = \frac{1}{Q} \sum_{q=1}^{Q} R_q. \tag{21}$$

Taking the ratio gives an estimate for each $S_{k,x}$, which are the tally scores for block b,

$$S_{k,x}^{j,(b)} = \frac{A^{(b)}}{D^{(b)}}. \tag{22}$$

The average of these over numerous blocks should converge to the true sensitivity coefficient for a given set of inputs. The mean value reported by MCNP6 is the simple arithmetic mean over all the blocks B,

$$S_{k,x}^{j} = \frac{1}{B} \sum_{b=1}^{B} S_{k,x}^{j,(b)}. \tag{23}$$

The variance or standard deviation is computed using basic sample statistics like any other tally. Note that correlation between the numerator and denominator is implicitly included because the averaging occurs after the division; however, the correlation that arises because a block impacts subsequent blocks is not included. Neglecting inter-block correlation is a reasonable assumption, since blocks are sequences of typically five to ten neutron cycles and correlation between cycles that separated is usually small. Therefore, the expected bias of the estimated statistical uncertainties is small.

3 VERIFICATION & RESULTS

Verification is performed in three ways: (1) using simple two-group, infinite-medium problems with readily obtainable analytic solutions, (2) comparisons of sensitivities calculated from a Δk obtained from direct perturbations of isotopic densities, and (3) by comparisons to results of other software packages such as TSUNAMI-3D and MONK [12]. The reference problems used for software comparison are those of the OECD/NEA UACSA Phase III Sensitivity benchmarks, which has three cases, each having subcases. Results are also given for sensitivities to angular scattering distributions for the Jezebel benchmark.

3.1 Comparisons to Analytic Solutions

To test that MCNP6 can compute sensitivities to nuclear data correctly, three simple, multigroup[1], infinite-medium test problems are constructed. The goal of the first is to test the ability of MCNP6 to compute sensitivities to cross sections. The second, in addition to cross-section sensitivities, tests the computation of both the unconstrained and constrained fission-χ sensitivities. The third analytic test specifically verifies the calculation of unconstrained and constrained scattering transfer function sensitivities.

Note that doing verification with multigroup calculations does not stress all the portions

[1]While MCNP6 (and earlier versions of MCNP) are capable of multigroup calculations, this feature exists for purposes of code comparisons or method validation; there are no self shielding or implicit sensitivity capabilities as the software is focused on being continuous energy.

of relevant MCNP6 coding that would be used for continuous-energy calculations. There is, however, a significant overlap in the internal mechanics between them, even if some of the details for handling different reaction numbers are somewhat different. So while the verification is not complete when applied to continuous energy, it still covers significant sections of code.

3.1.1 Analytic Problem 1

A simple test problem is constructed, an infinite-medium problem with two energy groups. The test problem has all fission neutrons born in the group 1, the top energy group, and there is no scattering from group 2 to 1. The nuclear data are given in Table I and are chosen so that $k = 1$.

The analytic solution for k is

$$k = \frac{\nu_1 \sigma_{f1}}{\sigma_{R1}} + \frac{\nu_2 \sigma_{f2}}{\sigma_{R2}} \frac{\sigma_{s12}}{\sigma_{R1}}, \tag{24}$$

where σ_{Rg} is the removal cross section for energy group g, which is $\sigma_{Rg} = \sigma_{tg} - \sigma_{sgg}$.

Sensitivity coefficients to the nuclear data are obtained by differentiating Eq. (24) with respect to the corresponding data. These reference sensitivity coefficients and the results computed by MCNP6 are given in Table II, and all agree within a few tenths of a percent as determined by the C/E ratio (note that the E in C/E in this paper means "expected" and not "experimental" result, as it usually does, and the C/E is only a valid measure when that reference value is non-zero). The sensitivities to total cross sections, within group scattering, and fission χ are not shown, because they are zero, and MCNP predicts a value that approaches zero for them.

3.1.2 Analytic Problem 2

A second infinite-medium problem with three energy groups is constructed with the following features: fission only occurs in group 3, fission neutrons may be produced in all groups, downscattering is restricted to the subsequent group, and there is no upscattering. The nuclear data are given in Table III, and are chosen again to make $k = 1$.

The analytic solution for k is

$$k = \frac{\nu_3 \sigma_{f3} \sigma_{s23}}{\sigma_{R2} \sigma_{R3}} \left[\frac{\sigma_{s12}}{\sigma_{R1}} \chi_1 + \chi_2 + \frac{\sigma_{R2}}{\sigma_{s23}} \chi_3 \right]. \tag{25}$$

Sensitivity coefficients to the nuclear data are obtained by differentiating Eq. (25) with respect to the corresponding data, and for constained fission χ, applying Eq. (12). Those sensitivity coefficients that are non-zero are displayed in Table IV along with the corresponding MCNP6 calculated results. These sensitivities include cross sections, fission ν, and both unconstrained and constrained fission χ. Both cross-section and fission-χ sensitivities (constrained and unconstrained), like with the cross-section sensitivities in Analytic Problem 1, agree within a few tenths of a percent as indicated by the C/E ratios.

The unconstrained fission-χ sensitivities are an effective increase in the multiplicity, but at the neutron emission energy as opposed to the incident neutron energy for fission ν. This is an increase in fission χ in some energy region without renormalizating the overall distribution, which is neither mathematically nor physically meaningful. Since an increase in the number of neutrons per fission always leads to an increase in k by definition, the unconstrained fission-χ sensitivity is always positive. The constrained fission-χ sensitivity is that of the renormalized distribution. Increases somewhere must be offset by decreases somewhere else (in this case, everywhere else), and therefore the constrained fission-χ sensitivity is both positive and negative.

For this problem, increasing the fission χ in group 1 leads to a decrease in reactivity. This is because the average importance of neutrons emitted from fission decreases, as the number

of neutrons emitted in group 3 decreases while the number emitted in group 1 increases. Since fission only occurs from neutrons in group 3 and are therefore the most important, having fewer of them and more less important neutrons leads to a decrease in the overall importance. Conversely, increasing the fission χ in group 3 while decreasing fission χ in the other groups, leads to an increase in reactivity, and therefore the constrained fission-χ sensitivity there is positive. The constrained fission-χ sensitivity for group 2 happens to come out slightly positive, and is a consequence of the combination of both the positive effect of having fewer neutrons having to scatter from group 1 to 3 to cause fission, and the negative effect of having fewer neutrons emitted directly in group 3 and causing fission.

3.1.3 Analytic Problem 3

The third infinite-medium problem tests constrained sensitivities for scattering laws. The problem contains four groups and the data are given in Table V. This problem is somewhat unphysical in that neutrons in group 1 can downscatter into any of the three other groups, but downscattering in groups 2 and 3 can only occur into the subsequent group. There is no upscattering, fission neutrons are only produced in group 1, and fission can only occur in group 4.

The analytic solution for k, which equals unity when the data in Table V are used, is

$$k = \left(\frac{\nu_4 \sigma_{f4} \sigma_{s1}}{\sigma_{R1} \sigma_{R2} \sigma_{R3} \sigma_{R4}} \right) \left[f_{14} \sigma_{R2} \sigma_{R3} + f_{34} \sigma_{s3} \left(f_{13} \sigma_{R2} + f_{12} f_{23} \sigma_{s2} \right) \right]. \tag{26}$$

Here f_{ij} is the probability of scattering from group i to j, and $\sigma_{sij} = f_{ij} \sigma_{si}$.

Solutions for the sensitivity to the scattering laws from group 1, the f_{1j}'s, are desired. These are obtained analytically and, to test the constraining methodology, from MCNP6 calculations with direct perturbations to the multigroup scattering data. Direct perturbations are done by perturbing the isotopic number density by some small amount (in this case 1%), calculating a new k to find a Δk, and then using a discrete approximation of Eq.

(2) to estimate the sensitivity. Both the analytic and MCNP6 calculated unconstrained and constrained (both from direct calculations and with the adjoint-based method) scattering law sensitivities are given in Table VI. The results of the direct MCNP6 calculations agree within 10% of the analytic results, confirming the constraining is done consistently. The adjoint-based sensitivity coefficients agree within the 2-σ statistical uncertainties, which are all under 2%.

Note that the sensitivity for the probability of within-group scattering in group 1 is non-zero in both the unconstrained and constrained cases, whereas the sensitivity to the within-group scattering cross section σ_{s11} is zero. This may be a counterintuitive result. Increasing a within-group scattering cross section in an infinite-medium problem may change the relative number of collisions in that group, but it does not have any effect on the reactivity because there is no change in the group's steady neutron population nor is there a change in the rate that neutrons transfer to groups with a different importance – this is not true for finite medium problems because leakage is present. Mathematically, from Eq. (1), the positive effect from increasing scattering is directly offset by the negative effect of additional interactions by increasing the total cross section.

Increasing the scattering probability, however, is an effective increase in the multiplicity of scattering events, as unphysical as that may be [in multigroup, this could be thought of as increasing the (n,2n) rate, which is often grouped with scattering, but this explanation does not apply as well for continuous-energy where all reactions are treated explicitly]. From Eq. (1), the positive effect from scattering source is present, but the negative effect from increasing the total cross section is absent, meaning that the unconstrained sensitivities for distributions (both scattering and fission χ) are always positive. Like with fission χ, since the probability density must be normalized, an increase of scattering into one group must be offset by decreases of scattering into other groups. This causes portions of the constrained sensitivity to be negative.

For this problem, increasing the rate of within-group scattering for group 1 has a relatively

strong negative effect because fewer neutrons are scattering directly into group 4 where they may cause fission, or into the other groups where their chances are better. Increasing scattering into group 2 is negative for similar reasons, but is less so because neutrons are more likely to reach group 4 from group 2 and cause fission, having escaped capture in group 1. Increasing scattering into group 3 is slightly positive, because neutrons scattering to group 3 have a better chance of causing fission in group 4 and they are doing less scattering into the less favorable groups 1 and 2 causing a positive effect, but this is mostly offset by the decrease of scattering directly into group 4. Increasing the scattering directly into group 4 has the strongest positive effect because it both brings neutrons into the group where they can cause fission, and decreases the amount of neutrons scattering to less favorable groups.

3.2 OECD/NEA UACSA Phase III Benchmark Comparisons

Comparisons to either other software, direct density perturbations, or both are made using the OECD/NEA UACSA Phase III Sensitivity benchmark. The purpose of the benchmark is to provide a series of common models by which various Monte Carlo software packages can be compared. The benchmark has three cases or phases: Phase III.1 is a finite square lattice of MOX fuel pins reflected by water, Phase III.2 is a series of infinite pin cell arrays with varying pitches, and Phase III.3 consists of two spheres of UF_4 mixed with polyethylene, one with low-enriched uranium (LEU) and the other with intermediate-enriched uranium (IEU). Details of the benchmarks and reference results may be found in Ref. [8].

All MCNP6 calculations use ENDF/B-VII.0 nuclear data [13]. A selection of the results generated for the benchmark are presented, with particular attention given to those with reported results from other Monte Carlo software packages found in Ref. [8]. Specifically, the co-authors of Ref. [8] who generated the results for ORNL TSUNAMI-3D and MONK were kind enough to share numerical results, and those are used as the basis for comparison. The Monte Carlo software packages that are used in Ref. [8] for comparison for the benchmarks are TSUNAMI-3D, MONK, DRAGON-SUSD3D [14], APOLLO2-MORET 5 [15], and

MMKKENO [2].

3.2.1 OECD/NEA UACSA Benchmark Phase III.1

Phase III.1 is an array of MOX fuel pins immersed in light water, simulating a criticality accident where a lattice of MOX fuel in a storage or shipping cask is flooded. This is based on the benchmark within the International Handbook of Evaluated Criticality Safety Benchmark Experiments (ICSBEP) having the identifier MIX-COMP-THERM-001-001 [16]. Within III.1, there are also four subcases: a 3-D detailed model, a homogenized Cartesian model, a homogenized cylindrical model, and an infinite array of fuel pins. MCNP6 results are presented only for the detailed model. The homogenized cases involve requiring special treatment generating something called implicit sensitivities from the homogenization process; MCNP6 does not support homogenization, and is specifically designed for detailed models.

The detailed 3-D model is a square lattice with a pitch of 0.9525 cm. 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. A light-water scattering kernel is used for thermal scattering with hydrogen in the water.

Energy-integrated sensitivity coefficients are computed for each isotope in the problem for the following nuclear data: total, (n,γ) capture, elastic scattering, inelastic scattering (all 40 discrete levels plus continuum), fission, (n,2n), and fission ν. The ten most sensitive nuclear data (except for the totals) are given in Table VII. For comparison, C/E values are given for TSUNAMI-3D results generated by ORNL (238-group ENDF/B-VII.0 cross sections are used). The C/E values agree within about five percent.

Energy-resolved sensitivities to the ^{238}U total cross section computed by MCNP6, TSUNAMI-3D, and MONK (using JEF-2.2 nuclear data) are shown in Fig. 1. All results mostly agree, with some small discrepancies between TSUNAMI-3D and the results from either MONK or MCNP6.

Note that the original ORNL TSUNAMI-3D values published in Ref. [8] show significant disagreement of about 50% with MONK and MCNP6. The differences arise from the ^{238}U capture resonances. The ORNL participants were contacted by the authors and new, more accurate results were provided. The original discrepancies were from not including the implicit sensitivity coefficients, which account for the self-shielding effects. Note that, MONK, like MCNP6, uses continuous-energy data, except that it uses the differential operator or Taylor series technique, which should converge to the same result [17]. Since MCNP6 agrees with both MONK and TSUNAMI-3D (with implicit sensitivities included), this supports the conclusion that the software is calculating energy-resolved sensitivity coefficients correctly.

This affirms the importance of doing multiple code comparisons with established benchmarks, as is being done with these OECD/NEA UACSA Phase III sensitivity benchmarks. Note that the originally reported ORNL TSUNAMI-3D results were not the only ones in Ref. [8] that disagree with with MCNP6, MONK, and the revised ORNL TSUNAMI-3D results. Many of the other participants reported similar discrepant results, while others reported ones that agree with those shown here. There does not seem to be any trend between which set of results were calculated and code, version number, or nuclear data between the multigroup methods. This illustrates the advantages of a continuous-energy implementation, as there is no need for the generation of implicit sensitivity coefficients and the possible error of not including them. The work flow is significantly simpler, and the chance of undetected user error is reduced.

Fig. 2 shows energy-resolved elastic scattering sensitivities for ^{1}H computed by MCNP6, TSUNAMI-3D, and MONK. Agreement between the software packages observed, with a few minor differences for some of the ^{239}Pu resonances for MONK, which may be attributable to using JEF-2.2 nuclear data rather than ENDF/B-VII.0.

Constrained fission-χ sensitivities of ^{239}Pu for MCNP6 and TSUNAMI-3D are displayed in Fig. 3. The agreement is very good; both software packages exhibit the same general features and are largely within the 2-σ statistical uncertainty of the MCNP6 results. This

agreement suggests MCNP6 is computing constrained fission-χ sensitivities correctly.

Results of the constrained sensitivities (per lethargy) to the total elastic scatter transfer distribution for ^{1}H are given in Figs. 4-6 (fast, epithermal, and thermal energy ranges) as a function of both incident and outgoing energy. The sensitivities given here are for exiting neutron energies in the lab frame. For the fast region, most of the neutrons with incident energies between 1-4 MeV have sizable negative constrained sensitivities for scattering within that same range, and small and more widely distributed positive sensitivities for downscattering into the 100's of keV energy range. The trend tends to continue downward along the diagonal in the fast region. The epithermal region shows relatively high negative sensitivities for downscattering into low-lying ^{238}U resonances, with the one at about 6.67 eV being most dominant. For the thermal region, scattering within an energy range or slightly higher appears to have a sizable positive sensitivity, as does upscattering in the 0.01-0.03 eV range. Downscattering at thermal energies has the largest negative sensitivity in this region.

The constrained total scattering sensitivity [including elastic, inelastic, (n,2n), etc.] to ^{16}O is displayed for the fast region in Fig. 7. This shows a relatively strong sensitivity (negative for small energy losses and positive for larger ones) for fission neutron energies. The scattering distributions of ^{16}O for the rest of the energy range are not all that significant and, as such, are not displayed.

To the authors' knowledge, the ability to compute sensitivities to scattering transfer distributions is unique to MCNP6 as of the writing of this paper, and therefore no reference results are available to use for comparison.

3.2.2 OECD/NEA UACSA Benchmark Phase III.2

Phase III.2 is an infinite array of the same MOX fuel pin in Phase III.1. For this phase, there are seven different subcases (pitch configurations) with lattice pitches of 0.586 cm, 0.60 cm, 0.66 cm, 0.73 cm, 0.9525 cm, 1.05 cm, and 1.15 cm. Sensitivity coefficients are computed for each subcase. The primary purpose of this benchmark is to test the implicit sensitivity

coefficient calculation, which is the differential change in k from the energy group collapse. Since this is a requirement of multigroup codes, and MCNP's purpose is continuous-energy physics, there is no implicit sensitivity coefficient. In either case, MCNP6 should be able to match final results.

Decent agreement is observed with the other benchmark participants in Ref. [8], especially those by the MONK code, which uses continuous-energy physics. Unfortunately, there is a fairly large spread in the ^{238}U total cross-section sensitivities ranging from -0.03 to less than -0.01 for pitch configuration 1. Like with those of Phase III.1, it appears the results can be grouped into two sets, suggesting an issue as to whether or not implicit sensitivities are included. This is complicated, however, by different data sets and the lack of quantified statistical uncertainties in many cases.

Despite the efforts of the authors, the reasons for these differences could not be identified. Therefore, rather than reporting code-to-code comparisons, comparisons with direct perturbations are performed. Perturbations of ^{1}H , ^{16}O , ^{238}U , ^{239}Pu are done for each pitch configuration. The isotopic number densities are increased as follows: ^{1}H by 1%, ^{16}O by 2%, ^{238}U by 10%, and ^{239}Pu by 1%. The relative magnitudes of the perturbations are chosen to ensure Δk can be estimated with enough statistical significance.

The results of the MCNP6 total cross-section sensitivities computed both by KSEN and directly are displayed in Fig. 8. 26 of the 28 results agree within the 2-σ statistical uncertainty bands. The two that disagree are the ^{238}U sensitivity for configuration 2 and the ^{239}Pu sensitivity of configuration 5. The former overpredicts (in magnitude) by about 2.7σ and the latter by about 2.2σ. The 2-σ band forms a 95% confidence interval, so it is expected about 1 out of 20 cases should fall outside of that interval. Therefore, the outliers here can likely be attributed to random chance, and not some systematic bias in the results.

3.2.3 OECD/NEA UACSA Benchmark Phase III.3

Phase III.3 has two subcases: one using LEU (2 atom percent ^{235}U) and the other using IEU (50 atom percent ^{235}U). This is based on the test problem in Ref. [1]. The radii of the two subcases are 36 cm and 18.2 cm respectively. Neutron thermal scattering with hydrogen is treated with a polyethylene scattering kernel.

Comparisons are made with MCNP6 KSEN generated total cross-section sensitivities to those generated by direct perturbations in MCNP6, adjoint-based multigroup methods in TSUNAMI-3D, and continuous-energy differential operator methods in MONK. For the direct perturbations, the isotopic number densities are increased as follows: 1% for ^{1}H , 5% for C, 2% for ^{19}F , 1% for ^{235}U , and 1% for ^{238}U in the LEU subcase and 20% in the IEU subcase.

The results of these are given in Table VIII. They largely agree, except for the ^{238}U total cross-section sensitivity in the IEU subcase, which is difficult to calculate because there are positive and negative effects on the reactivity that roughly cancel. The difference of MCNP6 KSEN result is about a factor of three higher than the MCNP6 direct perturbation result. Additionally, TSUNAMI-3D predicts a positive impact, whereas all the other calculations predict a negative one. This case aside, the MCNP6 KSEN and direct perturbation results all agree within about 5%. Note that the ^{1}H result for the LEU subcase is slightly outside the 2-σ confidence interval, but tends to agree well with the other codes. Likewise, (excluding the ^{238}U IEU result) MCNP6 agrees with both TSUNAMI-3D and MONK within 5%. Unfortunately, no statistical uncertainties are provided for the results of the other codes, so no quantification of statistical significance can be made.

As for the outlier ^{238}U IEU result and perhaps the ^{1}H LEU result, remember there is a particular caveat to comparing with direct perturbations. The implicit assumption is that the Δk is a linear function of $\Delta\sigma$. So long as the perturbation is sufficiently small, this should be satisfied. Here, even the 1% increase in hydrogen density may have been enough to create a small, non-linear effect in Δk because of spectral changes from additional

downscattering (the results still agree within 3%, so this effect would be quite small if it exists at all). The 20% increase in the ^{238}U density for the IEU subcase may also be stressing the linear assumption. While the overall effect of ^{238}U on k is small, this again is the result of competing effects, which may not change at the same rate with increasing the ^{238}U density, leading to a non-linear effect on Δk.

3.3 Jezebel Scattering Distribution Sensitivities

To provide a further assessment of the ability of MCNP6 to produce sensitivities scattering laws, analysis with the Jezebel benchmark (ICSBEP identifier PU-MET-FAST-001 [16]) is performed. Jezebel is a nearly-spherical mass of plutonium (mostly ^{239}Pu), which is often approximated as a simple sphere. A detailed benchmark specification of the experiment was obtained [18], and sensitivities to the elastic and inelastic scattering laws of ^{239}Pu were computed with this model.

In contrast to the OECD/NEA UACSA Phase III benchmarks, which have thermal or epithermal spectra, Jezebel has a fast spectrum and is much more leakage dominated. This should magnify the effect of angular distributions upon the reactivity, and therefore have higher sensitivity coefficients. The angle-resolved ^{239}Pu elastic and inelastic scattering distribution sensitivities are given in Fig. 9 – the direction cosines used are those of the data table, which for this case are in the center-of-mass frame. This resultant profile is integrated over all incident energies. Individual sensitivity profiles are obtained for each 0.1 MeV incident energy bin ranging from 0 to 20 MeV, and the profile in Fig. 9 is obtained by summing them together. This removes the bias produced by integrating over too large of an incident energy bin and then introducing the constraint.

First, the angular distribution of ^{239}Pu elastic scattering has a much greater effect than that of inelastic scattering. This makes sense as the energy-integrated cross-section sensitivities for elastic and inelastic scattering are 0.062 and 0.038 respectively. From Fig. 9, there appears to be a relatively strong negative effect on reactivity for increasing the

probability of forward-peaked elastic scattering of ^{239}Pu . This is an expected result as a majority of neutron losses occur from neutrons near the edge traveling toward the boundary. Forward-peaked scattering does little to help these neutrons survive and cause fission, and increasing the forward-peaked scattering must correspondingly decrease the less forward-peaked or backward scattering, leading to a net increase in the leakage of neutrons and therefore lower reactivity. Similar, but smaller trends are observed for inelastic scattering.

Unfortunately, to the knowledge of the authors, there are currently no published results available for quantitative comparison for this benchmark. Nonetheless, the results do match physical intuition, and offer some confidence that the MCNP6 results for this kind of calculation are reasonable. If the angle-resolved elastic sensitivity is integrated for scattering angles from 0° to 30°, the net sensitivity for that range of -0.0167. Depending on the relative uncertainties of the elastic scattering distribution (which tend not be known very well), it may have a significant impact on the overall uncertainty of this benchmark. Future analysis using these techniques will be used to quantify this uncertainty.

4 USER INTERFACE & METHOD PERFORMANCE

For a method to be of practical use, it must produce reliable results and be able to do so with reasonable computational requirements. The interface should be easy to use and offer flexibility, such that a user can generate needed results without too much effort. The input format is a simple line (or "card" in the MCNP parlance) in the MCNP6 input file called "ksen" where the user specifies a list of isotopes, reactions, an energy grid, and possibly some other things (e.g., cosine bins) depending on what the user wants. Multiple instances of the "ksen" card are possible. Specific details of the input file format are provided with the MCNP6 documentation.

It should also be difficult to unintentionally generate incorrect results that are difficult to detect. For example, the need for implicit sensitivity coefficients with multigroup methods can lead to confusion, as appears to be the case for some of the results for the OECD/NEA

UACSA Phase III benchmarks in Ref. [8]. Using a continuous-energy method, as MCNP6 does, removes this source of potential error.

Of course, even the most reliable method with the best user interface is not useful if it takes too long to produce statistically significant results or requires too much memory. This depends on both the method itself and the implementation within a software package. For this reason, quantitative assessments of method performance are given. The results that follow demonstrate is that MCNP6 can, for typical applications, produce sensitivities in a reasonable time without demanding too much memory.

4.1 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. To measure performance of the method, the OECD/NEA UACSA Phase III.3 benchmark, the spheres of uranium fluoride in polyethylene, are used. A 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 are performed on Intel Xeon E5-2670 processors at 2.6 GHz with 32 GB RAM. KSEN works both with Message Passing Interface (MPI) and Open Multi-Processing (OMP) threading forms of parallelism in MCNP, 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 subcases with and without KSEN, 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 IX.

The increased overhead for employing KSEN for this problem leads to increases in wall-clock time by factor of 4.2 for the LEU subcase and 3.4 for the IEU subcase. Computing a relatively large number sensitivities is not a trivial task for MCNP6, and calculations may take significantly longer than a traditional eigenvalue calculation. Note that obtaining just the 15 sensitivities displayed in Table IX 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. This explains the difference in the factors between the LEU and IEU subcases, with the LEU subcase being a more thermal, near-critical system, and therefore requiring more downscattering than the IEU one. While these increases may seem large, it is worth considering that acquiring a large number of normal flux results or differential operator perturbations in MCNP also incurs run time penalties as well. Also, the alternative of perturbing the nuclear data files to obtain these results from direct perturbation calculations is a very time consuming process on the part of the user, and may even take a greater amount of simulation time to resolve small differences in k.

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. Future research may focus on variance reduction techniques that would improve the sampling efficiency of particular isotopes and reactions.

4.2 Memory Usage

It has been suggested [19] that the memory requirements for this technique can be very high when a very large number of sensitivity coefficients (such as those involving numerous isotopes and reactions, a fine energy group structure, and many regions) are desired. Additionally, large memory requirements harm parallel scaling for MPI because passing large amounts of data across the network is expensive. Observing the fact that 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. For many problems, this reduces memory requirements by a factor of 10-100 (of course, individual problems may vary significantly) and improves parallel scaling.

To illustrate memory savings, consider the OECD/NEA UACSA benchmark problem III.1 (the MOX lattice) 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 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, not including any additional overhead information that must be stored to carry out the standard eigenvalue calculation or for constraining the fission χ or scattering law sensitivities.

Because much of this memory is either not needed at all, or will, most likely, not be required each history, significant savings can be made. 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. If region-specific sensitivities are desired, the memory requirements would surely increase; however, the increase would likely not be prohibitively large because within an individual history, neutrons are likely to only visit a small portion of the cells in a complicated geometry and only have a few energy states while visiting each of those. 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

A continuous-energy k-eigenvalue sensitivity coefficient capability called KSEN is implemented MCNP6 and helps meet the need of sensitivity and uncertainty analysis of critical experiments within the US DOE/NNSA NCSP. The method is derived from adjoint-based perturbation theory, as is TSUNAMI-3D, an ORNL software package for similar purposes. The capability in MCNP6 allows users to compute energy-resolved nuclear-data sensitivity coefficients to cross sections, fission ν, fission χ, and scattering energy-angle transfer distributions – the last of which is, at the time of writing this and to the knowledge of the authors, the only such capability that exists in a production Monte Carlo software package.

The methods are verified against analytic solutions of three multigroup, infinite-medium problems, direct density perturbations, and against results of other Monte Carlo software packages using the three cases of the OECD/NEA UACSA Phase III Benchmarks. The results of the analytic calculations agree within a few tenths of a percent, verifying the underlying routines in MCNP6. The continuous-energy calculations are verified using direct density perturbations and TSUNAMI-3D and MONK calculations of the OECD/NEA UACSA Phase III benchmarks, and the results generally agree within a few percent. Comparisons of calculated results from different codes show that while multigroup methods can certainly find answers consistent with continuous-energy methods, the requirement of implicit sensitivity coefficients to correct for the multigroup approximation introduces a source of potential user error that continuous-energy methods avoid. Sensitivities to ^{239}Pu elastic and inelastic angular distributions in Jezebel are also given, and they follow physical intuition.

Performance tests show that the computational resources needed for MCNP6 k-eigenvalue sensitivity coefficient calculations are reasonable. As expected, MCNP6 calculations needing sensitivity results takes longer and requires more memory than standard eigenvalue calculations. A test on OECD/NEA UACSA Benchmark Phase III.3 shows that MCNP6 requires about 5-10% longer to compute 15 sensitivities, and about a factor of 4 increase in run

time to compute about 63,000 sensitivities. A test of the memory requirements using the OECD/NEA UACSA Benchmark Phase III.1 with over 1.6 million sensitivities and a batch size of 50,000 neutrons shows that MCNP6 requires about 10 GB of RAM, a large, but not prohibitive amount on modern computers.

Future research and development will focus on uncertainty propagation with ENDF covariance data. This will allow criticality safety practitioners to use MCNP6 in addition to TSUNAMI-3D to perform uncertainty analysis. Another possible application is to apply these techniques to the calculation of temperature coefficients for reactor analysis. Recent research allows fast and efficient Doppler broadening to be performed in-line with the transport simulation [20]. It should be possible, at least in theory, to adapt these approaches to compute temperature coefficients as well. The sensitivity routines for nuclear data can also be extended to handle the case of sensitivities to interface locations or boundaries, to address the problem of propagating uncertainties arising from manufacturing tolerances or geometric measurement uncertainties. A prototype of this is already done [21], and results mostly agree with other methods; however, more research is needed to ensure that the techniques are correct and, if so, to make it a production-level capability within MCNP6. Another question that remains open are the related issues of intra-block convergence of the importance weight and inter-block correlations and their impact on resulting values and uncertainty estimates. Such a study is important to understand theoretical issues of adjoint-weighted tallies needed for sensitivity coefficients as well as other results such as point-kinetics parameters.

ACKNOWLEDGMENTS

Funding for this work was provided by the US DOE/NNSA Nuclear Criticality Safety Program. Special thanks to Tatiana Ivanova at IRSN for her support in providing the OECD/NEA UACSA Phase III benchmark specifications as well as helping facilitate the collection of results from other participants. The authors also wish to acknowledge Keith Bledsoe of ORNL and James Dyrda of AWE for providing numerical results of OECD/NEA

UACSA Phase III benchmark results for TSUNAMI-3D and MONK respectively, and all other benchmark participants who agreed to share numerical results of their data.

References

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

[2] 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).

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

[4] H.J. PARK, H.J. SHIM, C.H. KIM, "Uncertainty Quantification in Monte Carlo Criticality Calculations with ENDF/B-VII.1 Covariance Data," *Trans. Am. Nucl. Soc.*, **106**, 796-797 (2012).

[5] 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).

[6] 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.*, Accepted for Publication, also LA-UR-11-07082, Los Alamos National Laboratory (2011).

[7] B.C. KIEDROWSKI, J.A. FAVORITE, F.B. BROWN, "Verification of k-Eigenvalue Sensitivity Coefficient Calculations Using Adjoint-Weighted Perturbation Theory in MCNP," *Trans. Am. Nucl. Soc.*, **103**, 409-411 (2010).

[8] 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. BLEDSOE, B. 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).

[9] 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**, 226-241 (2011).

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

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

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

[13] M.B. CHADWICK, et. al., "ENDF/B-VII.0: Next Generation Evaluated Nuclear Data Library for Nuclear Science and Technology," *Nucl. Data Sheets*, **107** (12), 2931-3060 (2006).

[14] G. MARLEAU, A. HEBERT, R. ROY, "DRAGON Programmers Manual," IGE-251 (2002).

[15] L. HEULERS, et al, "MORET 5 Overview of the New Capabilities Implemented in the Multigroup/Continuous-Energy Version," *Proc. Intl. Conf. on Nuclear Criticality (ICNC 2011)*, Edinburgh, Scotland, Sep. 19-22 (2011).

[16] J. B. BRIGGS (Ed.), "International Handbook of Evaluated Criticality Safety Benchmark Experiments", Nuclear Energy Agency, NEA/NSC/DOC(95)03/I, Paris, France (2004).

[17] H.J. SHIM, C.H. KIM, "Equivalence of the Adjoint-Weighted Monte Carlo Perturbation Method and the First Order Differential Operator Sampling Method with Fission Source Perturbation," *Trans. Am. Nucl. Soc.*, **104**, 789-790 (2011).

[18] R.W. BREWER, J.A. FAVORITE, *Personal Correspondence* (2012).

[19] 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).

[20] G. YESILYURT, W.R. MARTIN, F.B. BROWN, "On-the-Fly Doppler Broadening for Monte Carlo Codes," *Nucl. Sci. Eng.*, **171**, 226-241 (2012).

[21] B.C. KIEDROWSKI, J.A. FAVORITE, F.B. BROWN, "Monte Carlo Calculations of Eigenvalue Sensitivities to System Dimensions," *Proc. Intl. Conf. on Nuclear Criticality (ICNC 2011)*, Edinburgh, Scotland, Sep. 19-22 (2011).

Table I: Nuclear Data for Analytic Problem 1

g	σ_t	σ_c	σ_f	ν	χ	σ_{sg1}	σ_{sg2}
1	2	1/2	1/2	3/4	1	1/2	1/2
2	3	1	1	9/2	0	0	1

Table II: Sensitivity Results for Analytic Problem 1

x	Exact $S_{k,x}$	MCNP6 $S_{k,x}$	C/E
σ_{c1}	$-1/3$	$-0.33343 \pm 0.05\%$	1.000
σ_{c2}	$-3/8$	$-0.37480 \pm 0.06\%$	0.999
σ_{f1}	$-1/12$	$-0.08347 \pm 0.34\%$	1.002
σ_{f2}	$+3/8$	$+0.37524 \pm 0.08\%$	1.001
ν_1	$+1/4$	$+0.24995 \pm 0.08\%$	1.000
ν_2	$+3/4$	$+0.75005 \pm 0.03\%$	1.000
σ_{s12}	$+5/12$	$+0.41707 \pm 0.09\%$	1.001

Table III: Nuclear Data for Analytic Problem 2

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

Table IV: Sensitivity Results for Analytic Problem 2

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_{f3}	$+1/4$	$+0.25045 \pm 0.16\%$	1.002
ν_3	$+1$	$+1.00000 \pm 0.00\%$	1.000
σ_{s12}	$+5/24$	$+0.20810 \pm 0.16\%$	0.999
σ_{s23}	$+1/4$	$+0.25083 \pm 0.15\%$	1.003
χ_1	$+5/12$	$+0.41688 \pm 0.09\%$	1.001
χ_2	$+1/3$	$+0.33345 \pm 0.07\%$	1.000
χ_3	$+1/4$	$+0.24967 \pm 0.06\%$	0.999
	Exact $\hat{S}_{k,x}$	MCNP6 $\hat{S}_{k,\chi_g}$	C/E
χ_1	$-5/24$	$-0.20805 \pm 0.12\%$	0.999
χ_2	$+1/12$	$+0.08339 \pm 0.28\%$	1.001
χ_3	$+1/8$	$+0.12465 \pm 0.17\%$	0.997

Table V: Nuclear Data for Analytic Problem 3

g	σ_t	σ_c	σ_f	ν	χ	σ_{sg1}	σ_{sg2}	σ_{sg3}	σ_{sg4}
1	3	1	0	–	1	1	1/2	1/4	1/4
2	4	1	0	–	0	0	2	1	0
3	4	2	0	–	0	0	0	1	1
4	6	3	2	12	0	0	0	0	1

Table VI: Group-1 Scattering Law Sensitivity Results for Analytic Problem 3

	$S_{k,x}$		$\hat{S}_{k,x}$			
		MCNP6		MCNP6		
x	Exact	KSEN	Exact	Direct	KSEN	C/E
f_{11}	$+1/2$	$+0.504 \pm 1.5\%$	$-1/4$	-0.257	$-0.2503 \pm 0.1\%$	1.001
f_{12}	$+1/5$	$+0.199 \pm 1.1\%$	$-7/40$	-0.180	$-0.1748 \pm 0.1\%$	0.999
f_{13}	$+1/5$	$+0.203 \pm 1.1\%$	$+1/80$	$+0.013$	$+0.0124 \pm 1.4\%$	0.994
f_{14}	$+3/5$	$+0.598 \pm 0.4\%$	$+33/80$	$+0.405$	$+0.4126 \pm 0.1\%$	1.000

Table VII: Most Sensitive Nuclear Data for OECD/NEA UACSA Benchmark Phase III.1

x^j		$S^j_{k,x}$	C/E (TSUNAMI-3D)
^{239}Pu	ν	$+0.9248 \pm 0.00\%$	1.000
^{1}H	σ_{el}	$+0.4145 \pm 0.27\%$	1.029
^{239}Pu	σ_f	$+0.3777 \pm 0.02\%$	0.992
^{239}Pu	σ_c	$-0.2610 \pm 0.02\%$	0.999
^{16}O	σ_{el}	$+0.0860 \pm 0.29\%$	1.024
^{1}H	σ_c	$-0.0799 \pm 0.05\%$	1.015
^{240}Pu	σ_c	$-0.0590 \pm 0.04\%$	1.005
^{238}U	σ_c	$-0.0502 \pm 0.04\%$	0.948
^{241}Pu	ν	$+0.0283 \pm 0.03\%$	1.004

Table VIII: Comparison of OECD/NEA UACSA Benchmark Phase III.3 Total Cross-Section Sensitivities

		MCNP6			
		KSEN	Direct	TSUNAMI	MONK
LEU	^{1}H	$+2.40 \times 10^{-1} \pm 0.4\%$	$+2.34 \times 10^{-1} \pm 1.3\%$	$+2.41 \times 10^{-1}$	$+2.45 \times 10^{-1}$
	C	$+2.74 \times 10^{-2} \pm 0.7\%$	$+2.74 \times 10^{-2} \pm 2.2\%$	$+2.77 \times 10^{-2}$	$+2.77 \times 10^{-2}$
	^{19}F	$+4.38 \times 10^{-2} \pm 0.5\%$	$+4.17 \times 10^{-2} \pm 3.7\%$	$+4.56 \times 10^{-2}$	$+4.20 \times 10^{-2}$
	^{235}U	$+2.53 \times 10^{-1} \pm 0.1\%$	$+2.55 \times 10^{-1} \pm 1.2\%$	$+2.53 \times 10^{-1}$	$+2.58 \times 10^{-1}$
	^{238}U	$-2.01 \times 10^{-1} \pm 0.2\%$	$-2.03 \times 10^{-1} \pm 1.5\%$	-1.95×10^{-1}	-2.01×10^{-1}
IEU	^{1}H	$+4.54 \times 10^{-1} \pm 0.2\%$	$+4.57 \times 10^{-1} \pm 0.7\%$	$+4.55 \times 10^{-1}$	$+4.52 \times 10^{-1}$
	C	$+6.54 \times 10^{-2} \pm 0.3\%$	$+6.46 \times 10^{-2} \pm 1.0\%$	$+6.60 \times 10^{-2}$	$+6.61 \times 10^{-2}$
	^{19}F	$+1.18 \times 10^{-1} \pm 0.2\%$	$+1.17 \times 10^{-1} \pm 1.3\%$	$+1.19 \times 10^{-1}$	$+1.15 \times 10^{-1}$
	^{235}U	$+1.30 \times 10^{-1} \pm 0.2\%$	$+1.25 \times 10^{-1} \pm 2.5\%$	$+1.26 \times 10^{-1}$	$+1.36 \times 10^{-1}$
	^{238}U	$-1.57 \times 10^{-3} \pm 5.4\%$	$-5.22 \times 10^{-4} \pm 20.3\%$	$+1.35 \times 10^{-3}$	-1.30×10^{-3}

Table IX: Performance Data Using the OECD/NEA UACSA Benchmark Phase III.3

	Time (min)			Total		Elastic		Capture	
	Base	KSEN		% 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}

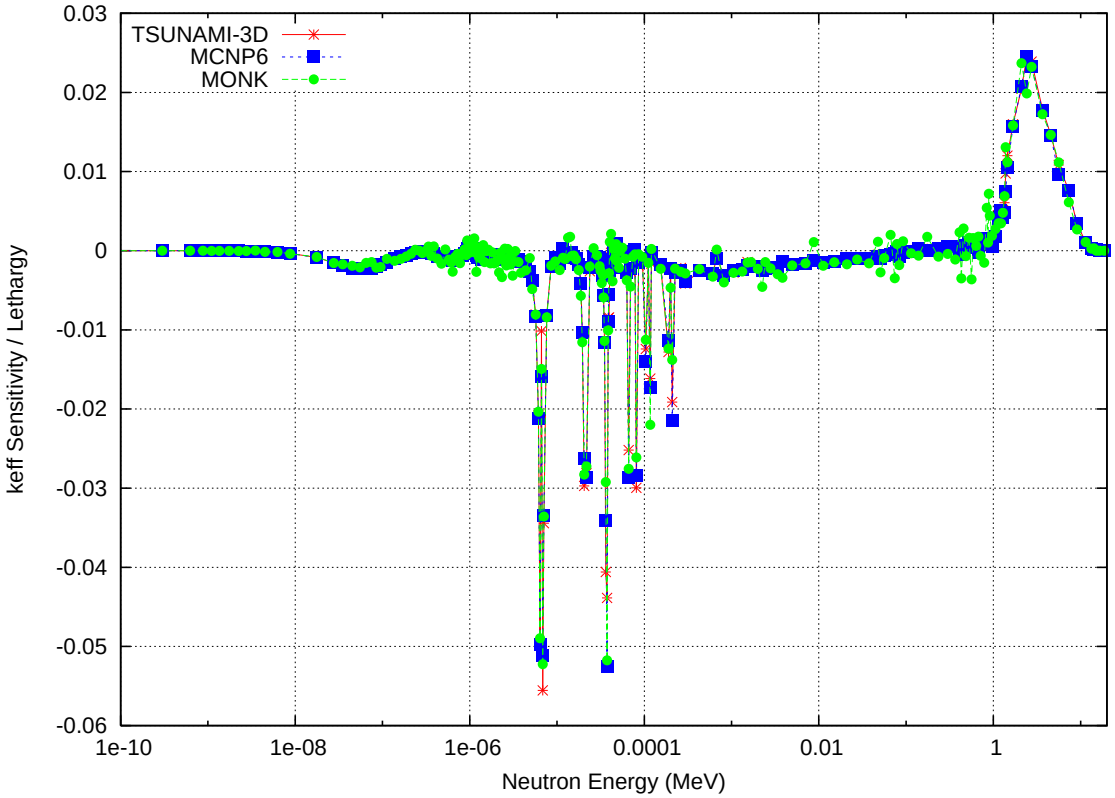

Figure 1: Comparison of ^{238}U total cross-section sensitivities for OECD/NEA UACSA Benchmark Phase III.1

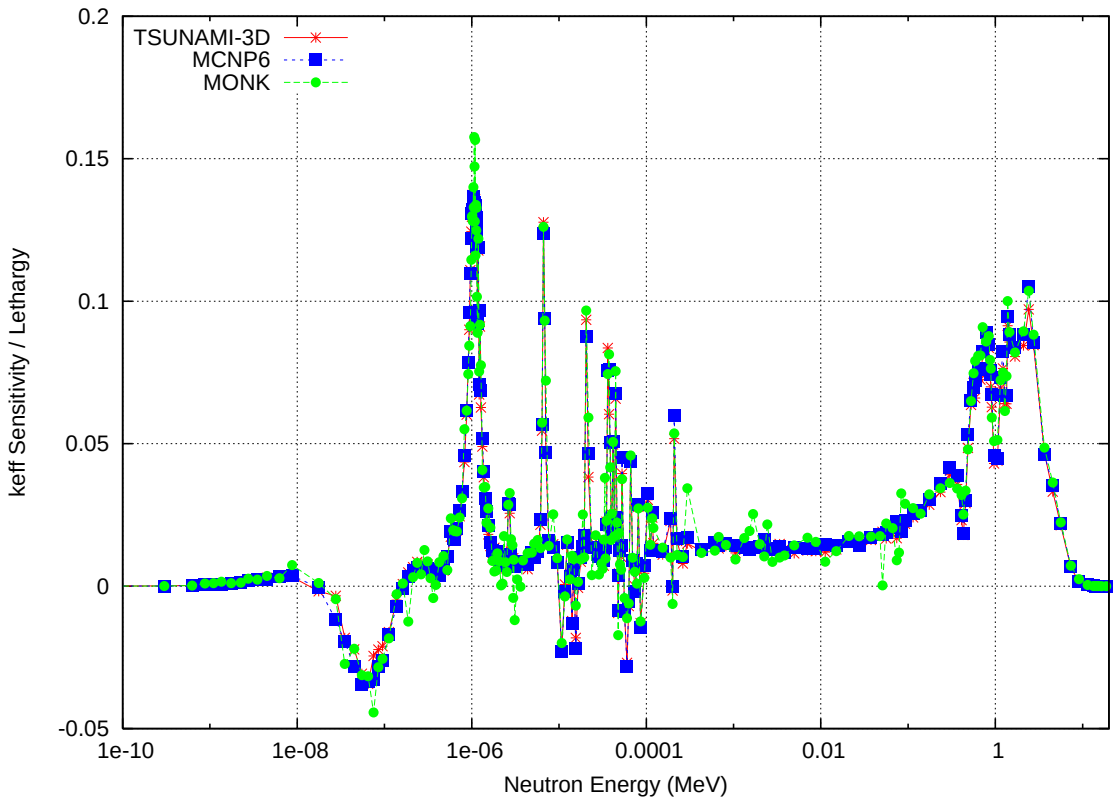

Figure 2: Comparison of ^{1}H elastic scattering cross-section sensitivities for OECD/NEA UACSA Benchmark Phase III.1

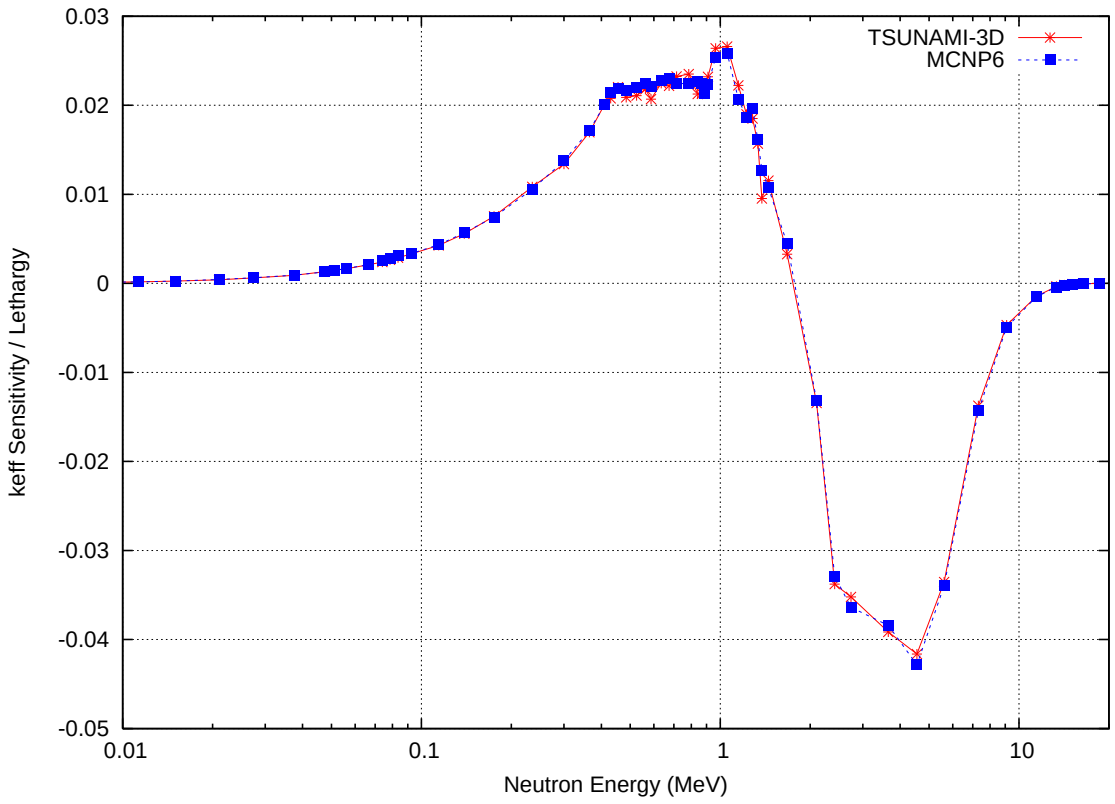

Figure 3: Comparison of constrained ^{239}Pu fission-χ sensitivities for OECD/NEA UACSA Benchmark Phase III.1

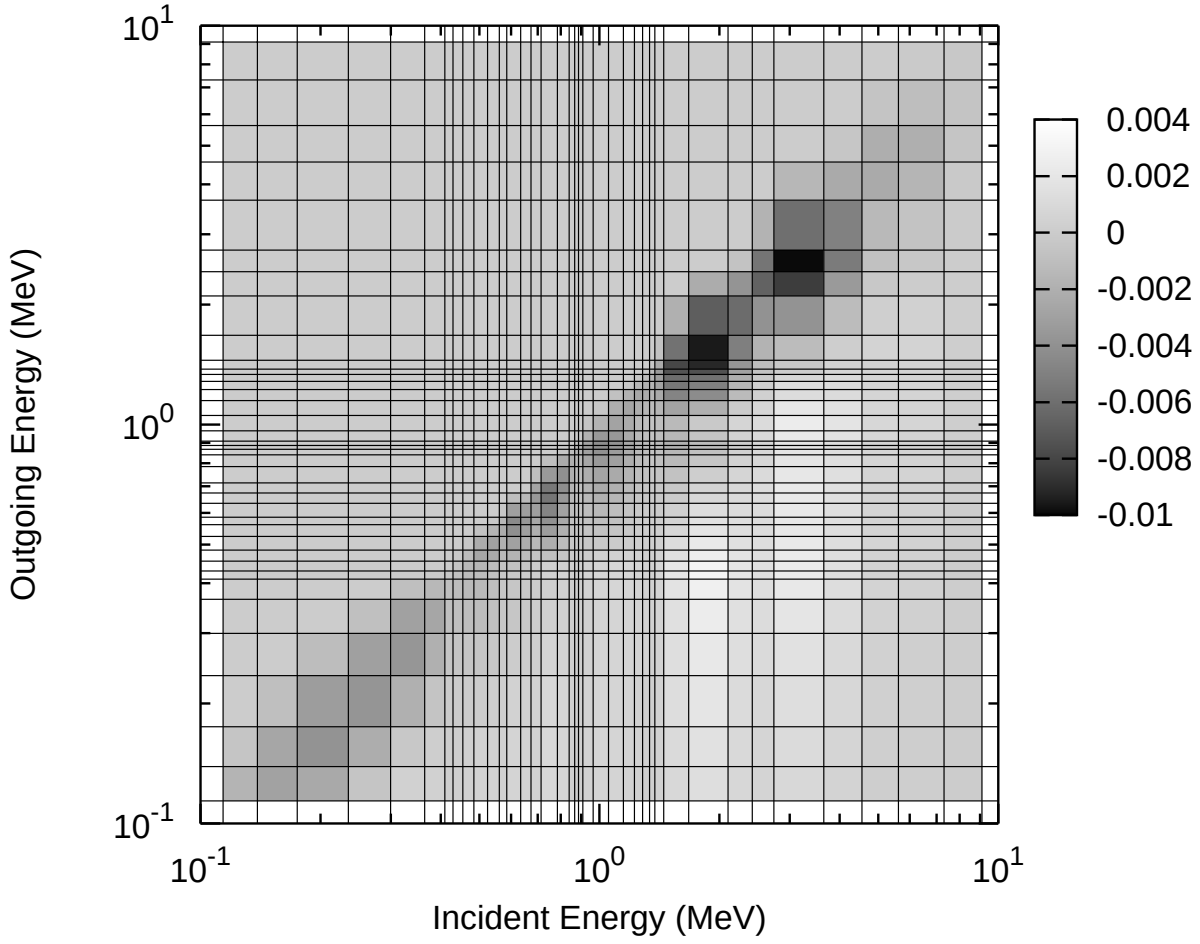

Figure 4: Energy-resolved constrained elastic scattering law sensitivity per lethargy of ^{1}H in the fast energy range for OECD/NEA UACSA Benchmark Phase III.1.

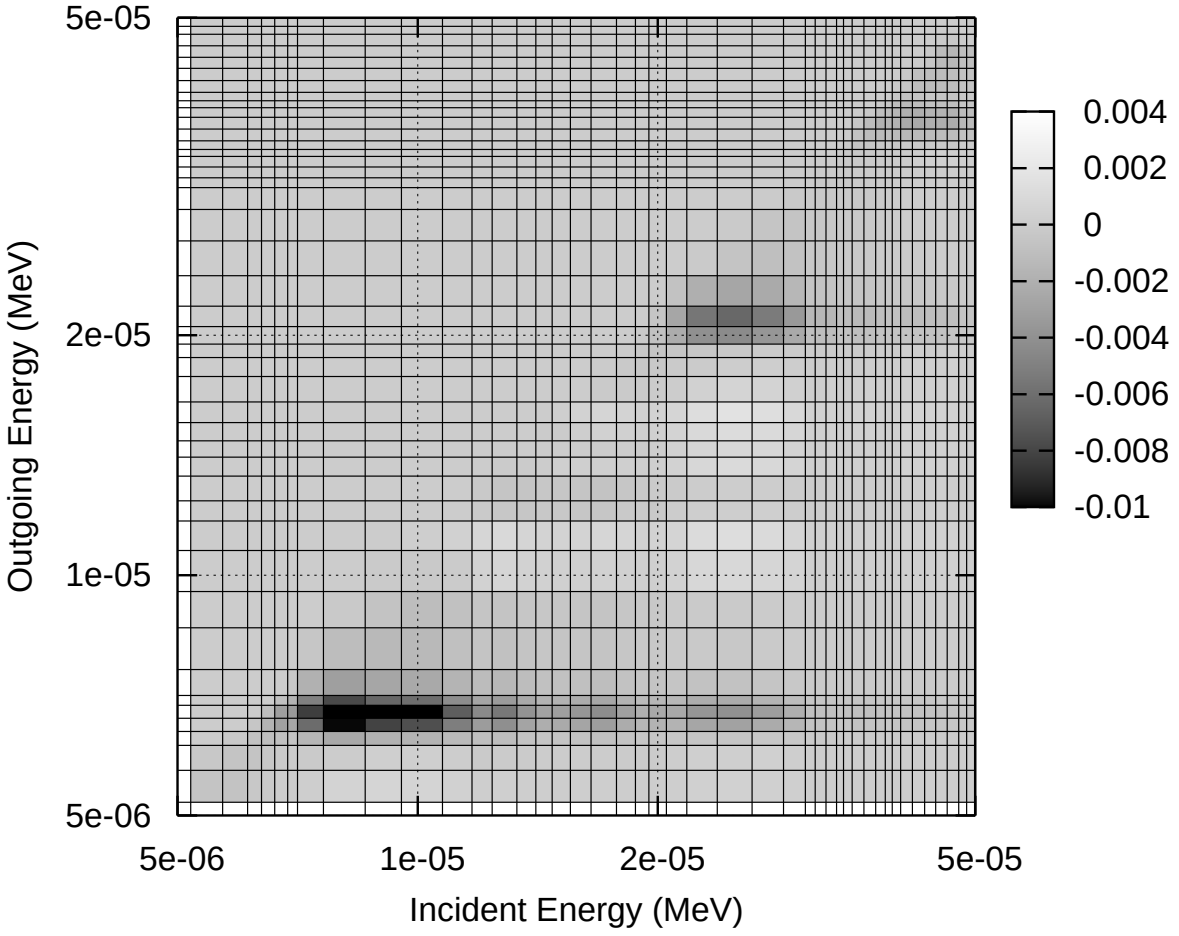

Figure 5: Energy-resolved constrained elastic scattering law sensitivity per lethargy of ^{1}H in the epithermal energy range for OECD/NEA UACSA Benchmark Phase III.1.

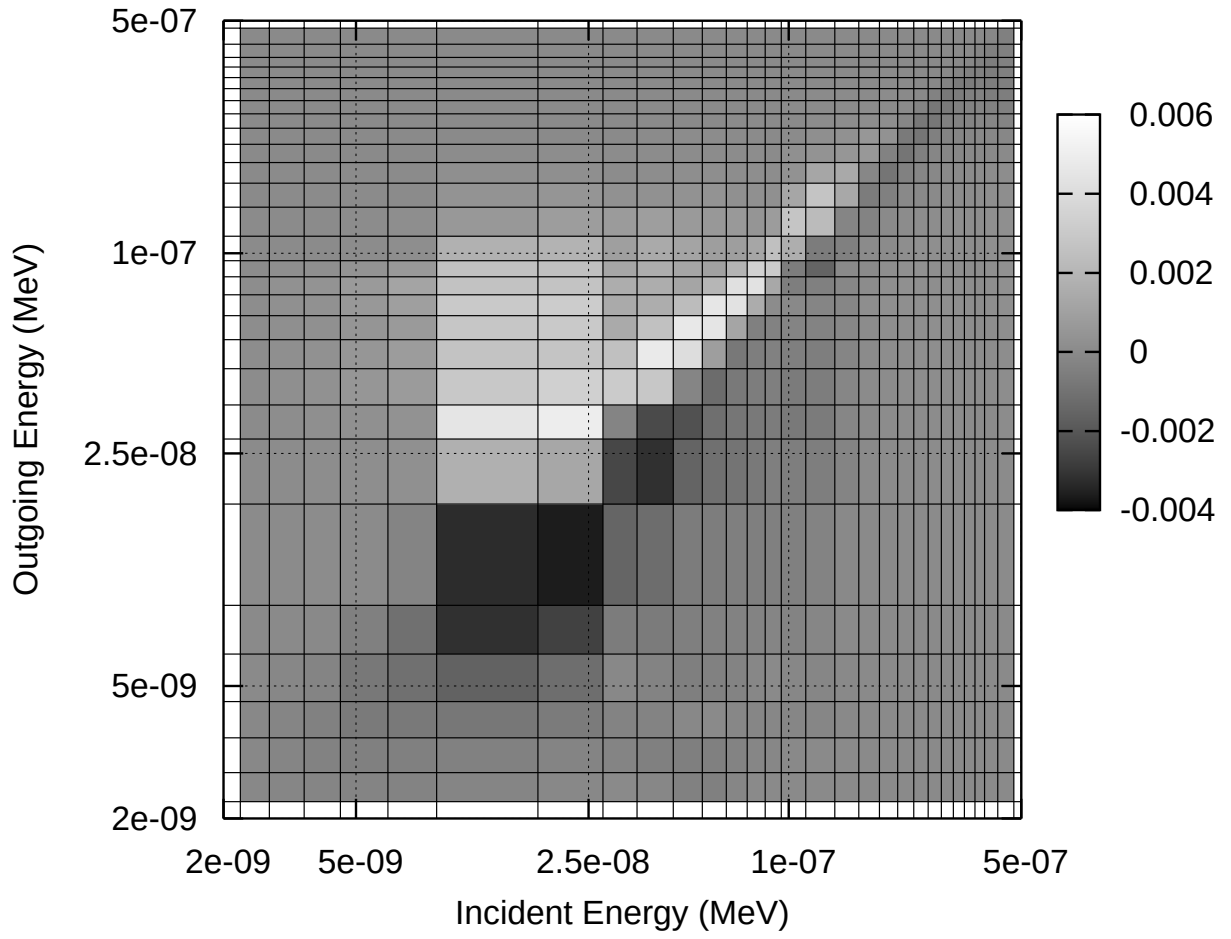

Figure 6: Energy-resolved constrained elastic scattering law sensitivity per lethargy of ^{1}H in the thermal energy range for OECD/NEA UACSA Benchmark Phase III.1.

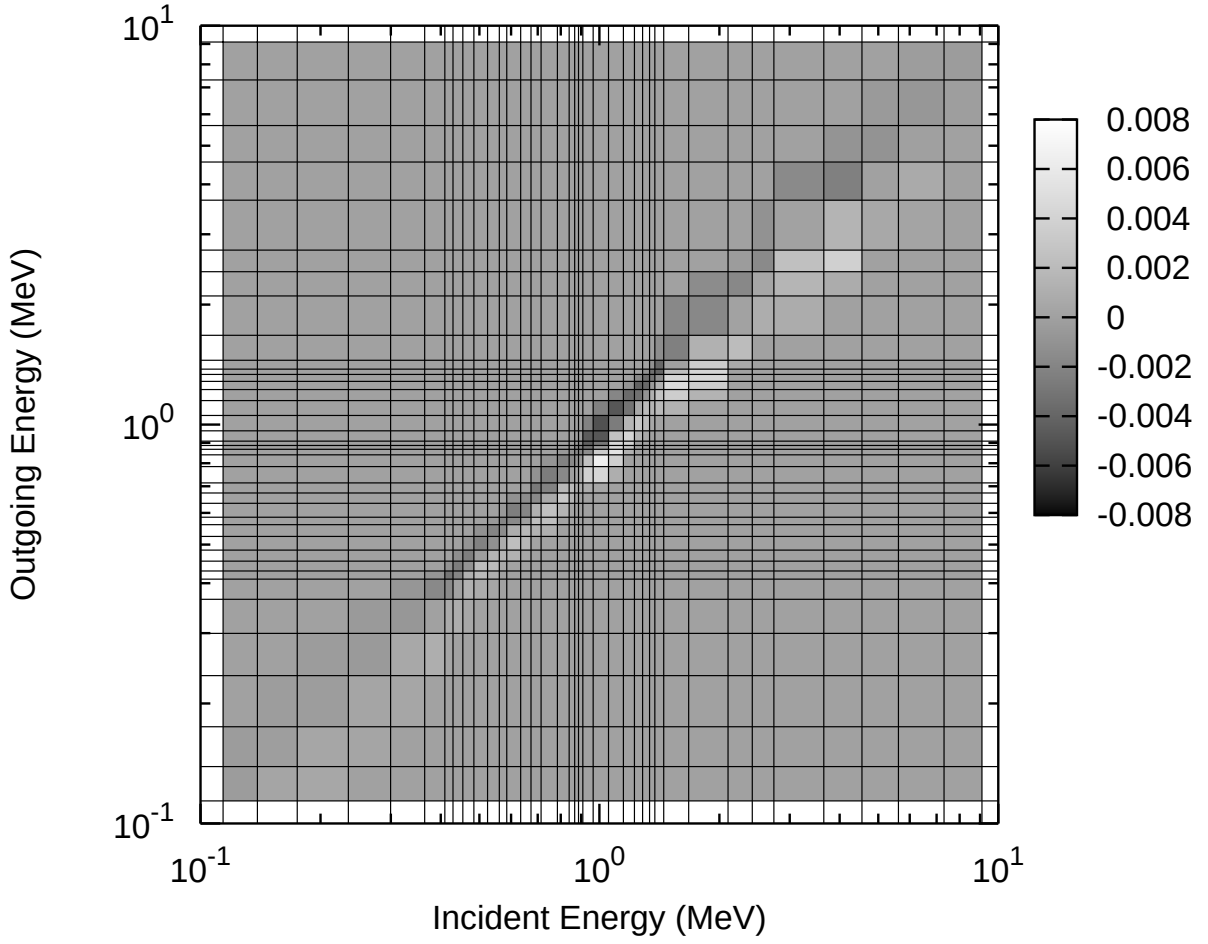

Figure 7: Energy-resolved constrained total scattering law sensitivity per lethargy of ^{16}O in the fast energy range for OECD/NEA UACSA Benchmark Phase III.1.

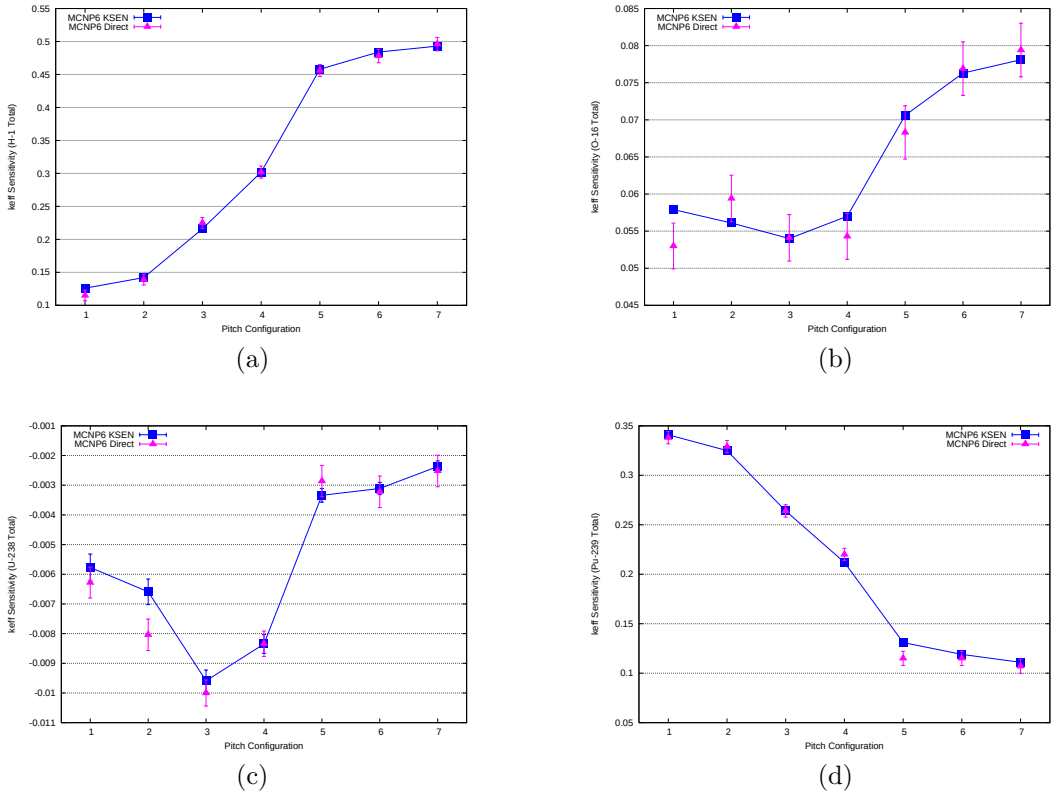

Figure 8: Comparison of MCNP6 KSEN estimates with those obtained with direct density perturbations for OECD/NEA UACSA Benchmark Phase III.2.

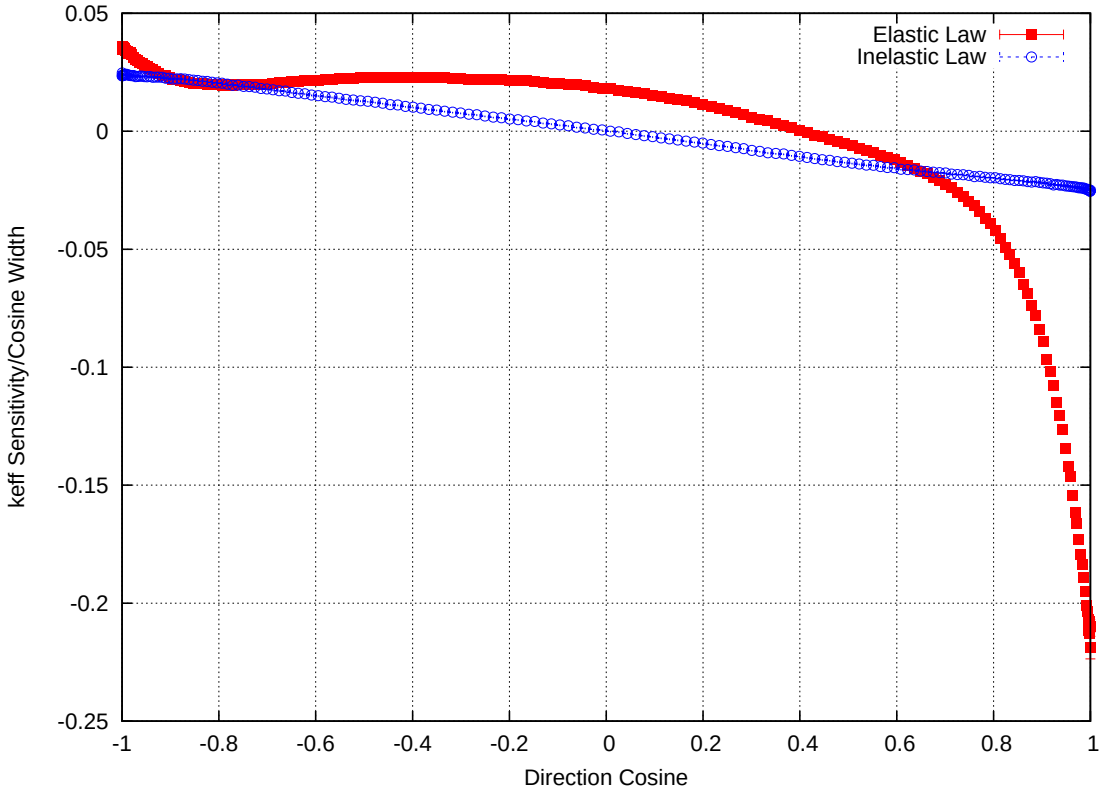

Figure 9: Constrained ^{239}Pu scattering law sensitivities for the Jezebel benchmark experiment.