PUBLICATION III

Optimizing the implementation of the target motion sampling temperature treatment technique – How fast can it get?

In proc. M&C 2013, Sun Valley, ID, May 5–9 2013.

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)

OPTIMIZING THE IMPLEMENTATION OF THE TARGET MOTION SAMPLING TEMPERATURE TREATMENT TECHNIQUE — HOW FAST CAN IT GET?

Viitanen Tuomas and Leppänen Jaakko
VTT Technical Research Centre of Finland
P.O. Box 1000, FI-02044 VTT, Finland
tuomas.viitanen@vtt.fi; jaakko.leppanen@vtt.fi

ABSTRACT

This article discusses the optimization of the target motion sampling (TMS) temperature treatment method, previously implemented in the Monte Carlo reactor physics code Serpent 2. The TMS method was introduced in [1] and first practical results were presented at the PHYSOR 2012 conference [2]. The method is a stochastic method for taking the effect of thermal motion into account on-the-fly in a Monte Carlo neutron transport calculation. It is based on sampling the target velocities at collision sites and then utilizing the 0 K cross sections at target-at-rest frame for reaction sampling. The fact that the total cross section becomes a distributed quantity is handled using rejection sampling techniques.

The original implementation of the TMS requires 2.0 times more CPU time in a PWR pin-cell case than a conventional Monte Carlo calculation relying on pre-broadened effective cross sections. In a HTGR case examined in this paper the overhead factor is as high as 3.6. By first changing from a multi-group to a continuous-energy implementation and then fine-tuning a parameter affecting the conservativity of the majorant cross section, it is possible to decrease the overhead factors to 1.4 and 2.3, respectively. Preliminary calculations are also made using a new and yet incomplete optimization method in which the temperature of the basis cross section is increased above 0 K. It seems that with the new approach it may be possible to decrease the factors even as low as 1.06 and 1.33, respectively, but its functionality has not yet been proven. Therefore, these performance measures should be considered preliminary.

Key Words: on-the-fly, Doppler, temperature, Monte Carlo, neutron tracking

1. INTRODUCTION

Recently, there has been a lot of interest in combining Monte Carlo reactor physics with thermal hydraulics or other multi-physics applications [3–6]. Multi-physics is also a major topic in the future development of Serpent*, involving developent of a universal multi-physics interface that will provide for information exchange with thermal hydraulics and fuel performance codes [7]. The long-term target is to develop Serpent into a competitive neutronics solver, capable of both steady-state and transient analyses including thermal-hydraulics feedback [8].

When examining the temperature dependence of system parameters, as is the case in multi-physics applications, it is often necessary to model the temperature distributions within a system in high detail. However, when using the conventional approach relying on pre-broadened effective cross sections, the cross section data must be stored in computer memory separately for each nuclide and temperature appearing in the problem geometry. The memory capacity becomes a major limitation if the number of different temperatures in the system is large.

*A complete and up-to-date description of the Serpent code is found at the project website: http://montecarlo.vtt.fi

As a solution to this issue, on-the-fly Doppler-broadening techniques have been developed, meaning that the effective cross sections are calculated as they are needed. The very efficient technique implemented in MCNP6 is based on calculating the temperature-dependent cross sections using series expansions with pre-calculated coefficients [9]. This paper is about another promising on-the-fly temperature treatment technique which, in fact, does not involve Doppler-broadening in the sense that effective cross sections are not calculated at any point of the transport calculation.

The target motion sampling temperature treatment method, discussed in this paper, was first introduced in [1] with the name "explicit treatment of thermal motion", and the first implementation and first practical results were presented at PHYSOR 2012 [2]. The method is based on taking the thermal motion of target nuclei into account explicitly at each collision site and dealing with the collisions in target-at-rest frame using cross sections at 0 K temperature. The fact that the total cross section becomes a distributed quantity is handled using rejection sampling techniques. As a novel feature the method provides for rigorous modelling of continuous temperature distributions within homogeneous material zones.

The preliminary multi-group implementation of the TMS method, for which results were presented in [2], resulted in a significant increase in calculation time when compared to traditional transport methods utilizing effective cross sections. In a PWR pin-cell case with realistic temperature distribution the calculation time increased by a factor of 2 and in a HTGR case, which was slightly different from the one examined in this paper, the overhead factor was 4.2. To make the method practical for everyday usage, some optimization is required.

This paper describes the optimization effort that has so far been put in the implementation of the target motion sampling method in Serpent. Results of the evolvement of the calculational overhead are also provided. The tracking scheme of the TMS method is shortly introduced in Section 2, the means of optimization covered in this paper are discussed in Section 3, performance results are provided in Section 4 and Section 5 is left for conclusions.

This paper only covers the efficiency of the method in the neutron transport calculation, basically CPU time required per transported neutron. The possible effect of the method on the variance of reaction rate estimators is left as a future topic.

2. TARGET MOTION SAMPLING METHOD FOR TEMPERATURE TREATMENT

Neutron tracking scheme of the TMS method is shortly introduced in the following on a step-by-step basis. More complete description of the scheme can be found in References [1] and [2]. Since the method only governs the way the neutrons are transported within material zones, discussion on the handling of material boundaries is omitted as irrelevant. In any case, the TMS method is compatible with all kinds of geometry routines, including the surface tracking and the Woodcock delta tracking methods.

1. The neutron transport begins from point x_i by sampling a path length l using a majorant cross section $\Sigma_{\mathrm{maj}}(E)$ at laboratory-frame energy of the neutron E. This is done by normal inversion sampling with equation

$$l = -\frac{1}{\Sigma_{\mathrm{maj}}} \ln(\xi)\,, \tag{1}$$

where ξ is a uniformly distributed random variable on unit interval.

2. At the new collision point candidate, $x_{i+1} = x_i + l$, collision nuclide is sampled based on nuclide-wise majorants such that probability of colliding with nuclide n is simply

$$P_n = \frac{\Sigma_{\mathrm{maj},n}(E)}{\Sigma_{\mathrm{maj}}(E)} = \frac{\Sigma_{\mathrm{maj},n}(E)}{\sum_n \Sigma_{\mathrm{maj},n}(E)} \,. \tag{2}$$

3. Velocity and direction of the target nuclide V_t μ are sampled from distribution

$$f(V_\mathrm{t}, \mu) = \frac{v'}{2v} f_{\mathrm{MB}}(V_\mathrm{t}) \,, \tag{3}$$

where v is the velocity of the neutron and μ is the cosine of the angle between the directions of the neutron and target,

$$v' = \sqrt{v^2 + V_\mathrm{t}^2 - 2vV_\mathrm{t}\mu} \tag{4}$$

is the relative velocity of the neutron to the target nuclide and f_{MB} is the Maxwell-Boltzmann (MB) distribution for target speed

$$f_{\mathrm{MB}}(V_\mathrm{t}) = \frac{4}{\sqrt{\pi}} \gamma^3 V_\mathrm{t}^2 e^{-\gamma^2 V_\mathrm{t}^2} \,, \tag{5}$$

where

$$\gamma(T, A_n) = \sqrt{\frac{A_n M}{2kT}} \,. \tag{6}$$

Target-at-rest-frame energy E' corresponding to the relative velocity calculated with Equation (4) is calculated and used from here on.

4. The collision is accepted or rejected according to rejection sampling criterion

$$\xi < \frac{g_n\Big(E, T(x_{i+1}), A_n\Big)\Sigma^0_{\mathrm{tot},n}(E')}{\Sigma_{\mathrm{maj},n}(E)} \,, \tag{7}$$

where $T(x_{i+1})$ is the local temperature at x_{i+1}, normalization factor g is obtained from equation[†]

$$g_n(E, T, A_n) = \left(1 + \frac{1}{2\lambda_n(T)^2 E}\right) \mathrm{erf}\Big(\lambda_n(T)\sqrt{E}\Big) + \frac{e^{-\lambda_n(T)^2 E}}{\sqrt{\pi}\lambda_n(T)\sqrt{E}} \,, \tag{8}$$

$$\lambda_n(T) = \sqrt{\frac{A_n}{kT}} \,, \tag{9}$$

and $\Sigma^0_{\mathrm{tot},n}$ is the total cross section of nuclide n at 0 Kelvin temperature. If the collision point is rejected, the procedure restarts from 1. with $i = i + 1$.

5. If the collision point is accepted, the procedure continues by reaction sampling using the relative energy E' and zero-Kelvin cross sections. Reaction r is sampled with probability

$$P_r = \frac{\Sigma^0_{r,n}(E')}{\Sigma^0_{\mathrm{tot},n}(E')} \tag{10}$$

and the transport proceeds according to the sampled reaction.

[†]The sign of the exponential term in this equation was erroneously negative in previous publications [1] and [2]. The authors apologize any inconvenience this may have caused.

The majorant cross section $\Sigma_{\mathrm{maj}}(E)$ is the maximum total cross section within the range of the target-at-rest energy E', which depends on the lab energy E and the thermal motion of the target nuclide. Since the MB distribution has an infinite tail, all possible nuclide velocities cannot be taken into account in the calculation of the majorant. As a quick and easy solution to this issue, a cut-off condition was adopted from the calculation of the Doppler-Broading Rejection Correction (DBRC) majorant [10]. This cut-off condition was first introduced in the SIGMA1 program [11] and is also used in the BROADR module of NJOY [12]. The majorant cross section is, thus, calculated according to equation

$$\Sigma_{\mathrm{maj},n}(E) = g(E, T_{\max}, A_n) \max_{E_\xi \in [(\sqrt{E}-\alpha)^2, (\sqrt{E}+\alpha)^2)} \Sigma^0_{\mathrm{tot},n}(E_\xi)\,, \tag{11}$$

where

$$\alpha = \frac{4}{\lambda_n(T_{\max})} \tag{12}$$

determines the proportion of nuclides that are omitted in the calculation of the majorant. It should be noted that in case a nuclide appears in the problem at different temperatures, the maximum of these temperatures $T_{\max}$ must be used in the calculation of the majorant to assure conservativity within the whole system.

A significant advantage of the TMS method is that the transport routine sees the temperature variable only through an arbitrary function $T(x)$, which can be continuous in both space and time.

3. DIFFERENT MEANS OF OPTIMIZATION

The performance of the TMS method is ultimately defined by the efficiency of the rejection sampling criterion (7). In this equation it can be seen that the probability of a successful sample depends on the ratio of the total 0 K cross section at target-at-rest energy E', multiplied by the g-factor that only depends on the L-energy of the neutron, to the majorant cross section. It should be noted that E' is a distributed quantity and, therefore, the true sampling efficiency for neutron energy E cannot be obtained directly from (7).

In practice the optimization of the TMS method is, thus, about decreasing the differences between the total cross section and the majorant cross section, which can be done in a few different ways. Two different means of decreasing the majorant are discussed in Sections 3.1–3.3. In addition, a new but not yet fully implemented optimization approach of increasing the temperature of the basis cross section above zero Kelvin is introduced in 3.4.

3.1. From Multi-group to Continous-energy Implementation

In the preliminary implementation the majorant cross section was constructed on a multi-group energy grid mainly because this way of implementation required the least changes in the source code. Since the multi-group energy grid is rather coarse, the majorant at an energy point corresponds to the maximum total cross section within quite a broad energy interval, thus making the majorant unnecessarily conservative.

Evidently, the efficiency of the majorant can be increased by adding grid points in its energy grid. The more grid points are used, the less conservative the majorant becomes. However, in practice there exists a limit after which adding new points does not anymore speed up the overall calculation. This is because after some point more CPU time will be lost in dealing with a very fine energy grid than what is saved due

to a less conservative majorant. Naturally, increasing the number of grid points also affects the memory requirement of the majorant cross sections.

As a comfortable compromise between efficiency and practicality, it was decided to generate the majorants on the unionized continuous-energy grid on which all of the cross sections are constructed in Serpent.

3.2. Energy Range Boundaries of the Majorant

Before proceeding to the next logical step of further decreasing the conservativity by manipulating parameter α, defined in Equation (12), it was noticed that the current "DBRC way" of choosing the energy boundaries for the majorant is not necessarily the best possible method for this purpose. The energy range boundaries, as determined by Equation (11), are basically the minimum and maximum target-at-rest energy corresponding to a collision with a target nuclide with velocity equal to the SIGMA1 cut-off condition. The condition, for its part, is based on the MB distribution such that the chosen value for the α parameter omits about a proportion of 5×10^{-7} of the most energetic nuclei.

However, the velocity distribution of a target nucleus interacting with a neutron with velocity v differs from the MB distribution by factor v'/v (see Equation (3)). This should also be taken into account in the generation of the majorant for the majorant to be optimal. Above thermal energies the deviance from the MB distribution is negligible, because the v'/v factor approaches unity, but at low energies with $v \sim V_\mathrm{t}$ the difference can be recognized. Particularly, the difference affects the high-energy tail of the distribution that specifies the energy boundaries of the majorant.

One way of choosing the energy range boundaries is to use the same confidence level $Q/2$ for each upper (E_max) and lower (E_min) majorant energy boundary associated with neutron L-frame energy E. In other words, a proportion of $Q/2$ of the least probable target velocities are omitted when determining both the upper and the lower boundary of the energy energy range in this approach.

The boundaries can be resolved from the velocity distribution of the target nuclide (3). The minimum target-at-rest energy is obtained for a parallel collision with $\mu = 1$ for which the unnormalized cumulative distribution is

$$
\begin{aligned}
f_1(x) \quad &= \int_x^\infty \frac{\sqrt{v^2 + V_\mathrm{t} - 2vV_\mathrm{t}}}{2v} f_\mathrm{MB}(V_\mathrm{t}) dV_\mathrm{t} \\
&= \begin{cases}
-\frac{1}{2\gamma\sqrt{\pi}v}\Big(-4e^{-\gamma^2 v^2} + \gamma\sqrt{\pi}v + e^{-\gamma^2 x^2}(2 + 2\gamma^2 x(x-v)) \\
\qquad\qquad - 2\gamma\sqrt{\pi}v\,\mathrm{erf}(\gamma v) + \gamma\sqrt{\pi}v\,\mathrm{erf}(\gamma x)\Big) & \text{if } \; v > x \\
\frac{1}{2}\Big(\frac{e^{-\gamma^2 x^2}(2 + 2\gamma^2 x(x-v))}{\gamma v\sqrt{\pi}} + \mathrm{erf}(\gamma x) - 1\Big) & \text{if } \; v \leq x
\end{cases}
\end{aligned}
$$

The corresponding distribution for the maximum energy resulting from a head-on collision with $\mu = -1$ is

$$f_{-1}(x) = \int_x^\infty \frac{\sqrt{v^2 + V_\mathrm{t} + 2vV_\mathrm{t}}}{2v} f_\mathrm{MB}(V_\mathrm{t}) dV_\mathrm{t} = \frac{e^{-\gamma^2 x^2}(1 + \gamma^2 x(v + x))}{\gamma v \sqrt{\pi}} + \frac{1}{2} \operatorname{erfc}(\gamma x) \,.$$

Using conditions

$$\frac{f(x)}{f(0)} = \frac{Q}{2} \tag{13}$$

it is possible to numerically resolve the boundaries x for the target velocity and, furthermore, the corresponding boundary energies E_min and E_max. Determining the boundaries this way increases the complexity and computational cost of the majorant generation compared to the much simpler DBRC method, but advantageously the majorant is based on the same distribution from which the target velocities are actually sampled, and a uniform and easily adjustable cut-off condition can be used throughout the energy spectrum. The advantages of these properties are emphasized in the next phase, described in Section 3.3, in which the conservativity of the majorant will be decreased by manipulating the Q parameter.

3.3. Energy Range Width of the Majorant

The proportion Q, as introduced in the previous section, is by definition the expected proportion of target-at-rest energy samples outside the energy region used in the generation of majorant cross sections for the two extreme cases of a head-on and a parallel collision. Hence, it also gives a very conservative upper limit for the proportion of unphysical or erroneous samples P_err at which the sampled total cross section exceeds the majorant. In practice $P_\mathrm{err} \ll Q$, because only a small fraction of all collisions are close to head-on or parallel collisions and, additionally, the total cross section at an "erroneously sampled" target-at-rest energy E' quite rarely exceeds the majorant cross section.

The Q parameter has a large effect on the efficiency of the majorant. The larger the Q parameter, the higher the sampling efficiency, but on the other hand the laws of physics are broken more often due to increasing P_err. Thus, the Q parameter should be given as high a value as possible such that the proportion of erroneous samples P_err still remains small enough not to, in practice, distract the solution of neutronics. Since playing with the Q parameter is basically cheating and too large values result in serious errors, the final value for the parameter should be chosen conservatively.

3.4. Increasing the Temperature of the Basis Library

The range of temperatures appearing in reactor conditions is not very wide. Usually, the smallest temperature appearing in reactor physical calculations is that of the coolant material, which is almost always larger than room temperature, about 300 K. On the other hand, the fuel temperature may reach a few thousand degrees, but in many cases the range of different fuel temperatures appearing in a system is quite narrow. Considering this, it is quite obvious that the current approach of using the basis cross sections at zero Kelvin and at each collision making the on-the-fly correction to temperatures above 300 K might

not be an optimal one. A much better idea would be to apply the on-the-fly Doppler-broadening to a cross section at the minimum temperature of a nuclide. For example, in PWR conditions this would spare about 700–800 K of on-the-fly Doppler-broadening "effort" (as seen through a lowered rejection sampling efficiency) at each collision in the fuel region.

In fact, increasing of the temperature of the basis library above zero Kelvin is quite straightforward with the TMS method. This is due to the fact that the well-known definition of the effective cross section [13]

$$\sigma_{\mathrm{eff}}(v) = \frac{1}{v} \int |\mathbf{v} - \mathbf{V}_{\mathrm{t}}| \sigma_0(|\mathbf{v} - \mathbf{V}_{\mathrm{t}}|) f_{MB}(\mathbf{V}_{\mathrm{t}}, T) \mathbf{d}\mathbf{V}_{\mathrm{t}} \tag{14}$$

on which also the TMS method is based, allows the increasing of the cross section temperature in two or more steps, i.e. the cross section temperature can first be risen from 0 K to T' and then from T' to T utilizing basically the same methodology. This can be shown using either Fourier transform techniques [13] or with straightforward calculation [14]. There is, however, one drawback in the splitting of the Doppler-broadening in two parts: the sampled target velocities can no longer be recycled in the solution of scattering kinetics, somewhat decreasing the elegance of the TMS method. The theoretical considerations behind this optimization approach are under examination and to be published in the near future.

A quick-and-dirty implementation purely for performance studying purposes was implemented in Serpent by making the following modifications:

1. Every nuclide appearing in the problem is associated with a minimum and maximum temperature such that $T_{\max} = T_{\min} + \Delta T$.
2. The cross sections of each nuclide are Doppler-broadened to $T_{\min}$ in the pre-processing phase using the Doppler-preprocessor of Serpent [15].
3. The majorant cross sections are calculated similar to Equation (11), but using effective cross sections at $T_{\min}$ instead of zero Kelvin and ΔT instead of $T_{\max}$ in the generation.
4. The tracking routine is modified such that the target velocities are sampled from a distribution corresponding to temperature $T - T_{\min}$ instead of T.

The solving of scattering kinetics is based on the standard free gas treatment without Doppler-Broadening Rejection Correction in this very first implementation. For this reason, the results are not even expected to match perfectly with the benchmark.

4. EFFECTS ON PERFORMANCE

The performance of the TMS method with different levels of optimization is measured with the following key parameters:

- REA_SAMPLING_EFF, a standard output parameter of Serpent 2 that describes the proportion of accepted reactions, i.e. the average rejection sampling efficiency of (7). Denoted with η_{rea} in the following.

- The measure t_{rel}, defined as the ratio of the transport cycle calculation time to that of a Serpent 2 calculation and NJOY-broadened effective cross sections, free gas thermal treatment at all energies and Doppler-Broadening Rejection Correction used between 0.4 eV and 210 eV for $^{238}\mathrm{U}$.

- Proportion of erroneous samples, P_{err}, for which the majorant is exceeded at least by a factor of 1E-3, i.e. $\Sigma^0(E') > 1.001\ \Sigma_{\mathrm{maj}}(E)$.

Results are provided for two test cases: a PWR pin cell and a HTGR system. The PWR pin cell test case is identical to that used in the PHYSOR2012 paper [2]. The temperature of the coolant water and the cladding in the PWR pin-cell model is 579 K and the temperature distribution within the fuel pin is modelled with a step function ranging from 772 to 1275 K. The HTGR system was changed from the PHYSOR test case to an infinite cubic lattice of TRISO-particles, because, with the current version of Serpent, the TRISO system was found to reflect much better the effect of the tracking method on performance[‡]. The exact composition of the TRISO particles is described in Table I and the lattice pitch is 0.16341 cm.

Table I. Composition of the TRISO particle used in the HTGR test case.

Outer radius cm	Density g/cm^3	Temperature K	Material
0.0125	10.4	1800 K	UO_2 enriched to 8.2 w-% ^{235}U with 1 ppm boron
0.0250	10.4	1200 K	UO_2 enriched to 8.2 w-% ^{235}U with 1 ppm boron
0.0340	1.05	1200 K	carbon
0.0380	1.90	1200 K	pyrolytic carbon
0.0415	3.18	1200 K	silicon carbide
0.0455	1.90	1200 K	pyrolytic carbon
surroundings	1.75	1200 K	graphite

All of the calculations were made using a JEFF-3.1.1 based cross section library with 0.001 reconstruction tolerance. Results for performance comparison are presented in Tables II and III for a PWR pin cell and a HTGR system, respectively. Ten million active neutron histories were calculated in these test cases.

The statistics of these calculations are poor, leading to high statistical deviations in k_{eff}. DBRC was not used when calculating the results with the elevated basis cross section temperature, which causes significant bias in the corresponding eigenvalues. It should also be noted that the η_{rea} parameter is scored at all collisions also in the EBT case. Thus, also the collisions with nuclides appearing in one temperature only, for which the reaction sampling efficiency is 100 %, affect the η_{rea} parameter.

To show the effect of parameter Q on the calculation, differences in flux spectra compared to a NJOY-based benchmark are plotted in Figure 1 for both of the cases and three different values of Q. The temperatures of the cross sections were adjusted using the Doppler-preprocessor routine in the PWR case, while the HTGR case was based purely on NJOY-broadened cross sections. The number of active neutron histories used in the spectrum comparisons was 2.5 billion in the PWR case and 1 billion in the HTGR case.

[‡]The computational cost of the boundary condition handling routine of Serpent was slightly increased in version 2.1.8. Due to high number of boundary crossings in the original HTGR test case the effect of the tracking routine was blurred.

Table II. Performance measures for a PWR pin-cell calculation are provided in this table. Results are calculated with Serpent 2 and NJOY cross sections, multi-group (MG) on-the-fly method introduced in [2] and continuous-energy (CE) on-the-fly with different means of optimization as described in this paper, including the optimization method involving an elevated basis cross section temperature (EBT). The statistical deviations Δk_{eff} correspond to one sigma.

Case	Q	η_{rea}	t_{rel}	P_{err}	k_{eff}	Δk_{eff} (pcm)
NJOY XS, opti 2	-	0.99[§]	1.00	0	1.34604	40
MG on-the-fly	-[¶]	0.26	2.00	0.0E+0	1.34686	40
CE on-the-fly	1E-6	0.27	1.82	5.7E-9	1.34714	40
CE on-the-fly	1E-5	0.29	1.78	2.3E-8	1.34693	40
CE on-the-fly	1E-4	0.32	1.68	7.4E-8	1.34724	40
CE on-the-fly	1E-3	0.36	1.57	2.5E-7	1.34704	39
CE on-the-fly	1E-2	0.43	1.43	3.3E-6	1.34724	40
CE on-the-fly	1E-1	0.53	1.29	7.7E-5	1.34785	42
CE on-the-fly + EBT	1E-6	0.72	1.06	0.0E+0	1.35013	41
CE on-the-fly + EBT	1E-5	0.75	1.05	2.7E-9	1.35029	39
CE on-the-fly + EBT	1E-4	0.77	1.01	5.5E-9	1.34957	39
CE on-the-fly + EBT	1E-3	0.79	1.00	2.2E-7	1.34894	40
CE on-the-fly + EBT	1E-2	0.82	0.98	3.0E-6	1.34909	42

The benchmark calculation of the PWR case resulted in an analog k_{eff} estimate of $1.34651 \pm 2\ \mathrm{pcm}$, while the TMS method with $Q = 10^{-6}$ resulted in a 51 pcm larger value, $1.34702 \pm 2\ \mathrm{pcm}$. In the HTGR case the corresponding estimates were $1.20271 \pm 5\ \mathrm{pcm}$ and $1.20298 \pm 4\ \mathrm{pcm}$, the difference being 27 pcm. The multiplication factors of the TMS with Q-values 10^{-6} and 10^{-2} were equal within statistical accuracy.

In Figure 1 it can be seen that deviance curves for Q values 10^{-6} and 10^{-2} are practically the same over the whole energy spectrum, but when Q is increased to 10^{-1}, the spectrum starts to differ from the two other cases. This indicates that a value of $Q = 10^{-2}$ leading to P_{err} of magnitude 10^{-7}–10^{-6} seems tolerable, because with this proportion of erroneus samples the results of the transport calculation are still unaffected. On the other hand, 10^{-1} is evidently too high a value for the Q parameter and the transport calculation is notably disturbed if P_{err} exceeds 10^{-5}.

The PWR case with elevated basis temperature and the largest acceptable Q-value 10^{-2} performed, surprisingly, even better than normal Serpent in optimization mode 2. This is due to the fact that DBRC was not used in the calculations with EBT. Since the slow-down of DBRC in this PWR case is about 8 % [2], t_{rel} for the on-the-fly + EBT case with $Q = 10^{-2}$ is expected to become about 1.06 after introduction of DBRC. In the HTGR case the slow-down of DBRC is about 14 % and the expected t_{rel} for the fastest $Q = 10^{-2}$ case is 1.33.

[§] Since optimization mode 2 involves a combined multi-group/continuous-energy tracking scheme with rejection sampling, the reaction sampling efficiency differs from 1.00.

[¶] Multi-group on-the-fly uses the DBRC cut-off condition as defined in Equation (11).

Table III. Performance measures for a HTGR system.

Case	Q	η_{rea}	t_{rel}	P_{err}	k_{eff}	Δk_{eff} (pcm)
NJOY XS, opti 2	-	0.99	1.00	0	1.20281	38
MG on-the-fly	-	0.67	3.62	0.0E+0	1.20288	40
CE on-the-fly	1E-6	0.68	3.24	0.0E+0	1.20323	40
CE on-the-fly	1E-5	0.69	3.11	0.0E+0	1.20279	38
CE on-the-fly	1E-4	0.72	2.92	0.0E+0	1.20223	38
CE on-the-fly	1E-3	0.75	2.77	3.0E-8	1.20348	39
CE on-the-fly	1E-2	0.78	2.50	4.3E-7	1.20345	39
CE on-the-fly	1E-1	0.83	2.26	1.8E-5	1.20419	40
CE on-the-fly + EBT	1E-6	0.96	1.32	0.0E+0	1.21097	39
CE on-the-fly + EBT	1E-5	0.96	1.29	0.0E+0	1.21043	39
CE on-the-fly + EBT	1E-4	0.96	1.28	3.4E-9	1.21065	38
CE on-the-fly + EBT	1E-3	0.97	1.23	6.6E-8	1.21093	38
CE on-the-fly + EBT	1E-2	0.97	1.17	6.0E-7	1.21079	40

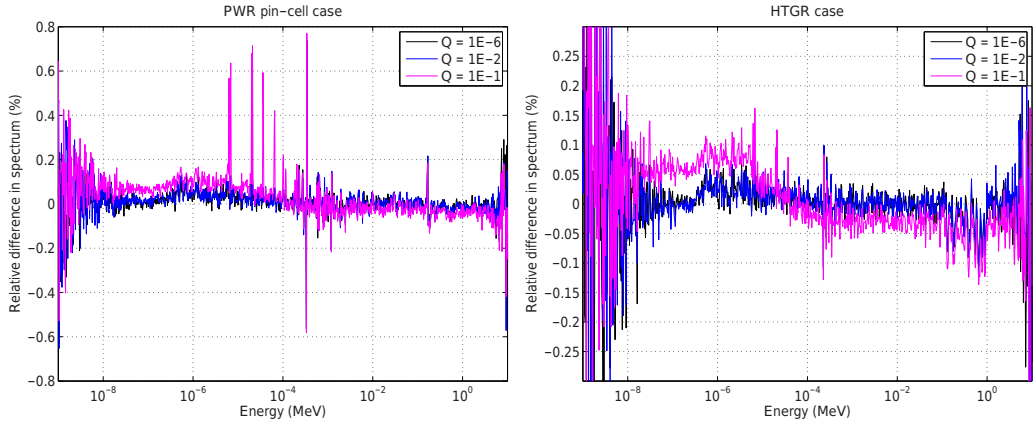

Figure 1. Spectra corresponding to 0 Kelvin basis temperature and three different Q-values are compared to a benchmark solution calculated using Serpent 2.1.9 with NJOY-broadened cross sections. Results are provided for two cases: a PWR pin-cell and a HTGR system.

5. CONCLUSIONS

In Figure 1 it can be seen that even with the lowest values of the Q parameter there is some systematical difference between the spectrum calculated with the on-the-fly treatment and the benchmark. The comparison of the results at this accuracy is unreasonable in the PWR case, since the accuracy of the cross section libraries is unknown due to the usage of the Doppler-preprocessor of Serpent. The preprocessor does not modify the energy grid of the cross section during Doppler-broadening and, hence, the reconstruction tolerance of the resulting cross section is higher than in the original one. In the HTGR case, small deviances can be recognized around energy intervals 2E-7–1E-5 MeV and 1E-1–1E0 MeV, which was found to result in a significant 27 pcm difference in the multiplication factor. These deviances most likely originate from small methodological differences combined with the limited accuracy of the cross section representation, but more results are needed for confirmation. The differences existed already in the previous multi-group implementation, but were not paid attention to before now.

From the performance measures it can be seen that the effect of changing from multi-group to continuous-energy majorant reduces the required CPU time about 5–10 %, depending on the case. By increasing the value of the Q parameter, which controls the conservativity of the majorant cross section, to 10^{-2} the CPU time savings increases to 28–30 %. With these means of optimization the relative calculational effort of the TMS method to an ordinary transport calculation can be decreased from 2.0–3.6 to 1.4–2.5 in the specific test cases.

The most significant enhancement in performance is, however, achieved with the still incomplete method of raising the temperature of the basis library above 0 K. Since the results obtained with the elevated basis library temperature do not match with the benchmark yet, it is too early to make any final conclusions about the efficiency of the method. Nevertheless it seems that, after introduction of DBRC, this means of optimization may decrease the relative calculational effort even as low as 1.06–1.33 in the specific test cases. This is still somewhat worse than the incredible efficiency of the on-the-fly methodology used in MCNP [9], but this level of slow-down can in many cases be considered acceptable and sometimes even negligible.

In the results it seemed that the value of Q may be chosen surprisingly large: even as high value as 10^{-2} was small enough not to affect the transport calculation significantly. Before hard-coding any values in the implementation, more systems should be investigated. It is possible that in some systems the Q parameter must be chosen smaller than 10^{-2}, deteriorating the performance of the method. The efficiency of the method should also be investigated with a system containing burned fuel. The performance of the method may drop dramatically if the number of nuclides in a material increases to several hundreds.

The optimization method based on elevated basis library temperature will be completed and implemented in the official version of Serpent 2 in the near future. Thereafter, the main topics for future research considering the TMS method involve the implementation of reaction rate estimators, the resonance treatment at the unresolved region and the handling of bound-atom scattering.

ACKNOWLEDGEMENTS

This work was partly funded through the NUMPS project of the Finnish Academy and KÄÄRME project of the Finnish Research Programme on Nuclear Power Plant Safety, SAFIR2014.

REFERENCES

[1] T. Viitanen and J. Leppänen, "Explicit Treatment of Thermal Motion in Continuous-energy Monte Carlo Tracking Routines," *Nucl. Sci. Eng.*, **171**, pp. 165–173 (2012).

[2] T. Viitanen and J. Leppänen, "Explicit Temperature Treatment in Monte Carlo Neutron Tracking Routines – First Results," *Proc. PHYSOR-2012*, Knoxville, TN, 15–20 April (2012).

[3] D. Kotlyar, et al., "Coupled neutronic thermo-hydraulic analysis of full PWR core with Monte-Carlo based BGCore system," *Nucl. Eng. Design*, **241**, pp. 3777–3786 (2011).

[4] M. Vazquez, et al., "Coupled neutronics thermal-hydraulics analysis using Monte Carlo and sub-channel codes," *Nucl. Eng. Design*, **250**, pp. 403–411 (2012).

[5] A. Ivanov, V. Sanchez and J. E. Hoogenboom, "Single Pin BWR Benchmark problem for Couled Monte Carlo – Thermal Hydraulics Analysis," *Proc. PHYSOR-2012*, Knoxville, TN, 15–20 April (2012).

[6] F. P. Espel, et al., " New developments of the MCNP/CTF/NEM/NJOY code system Monte Carlo based coupled code for high accuracy modeling," *Ann. Nucl. Energy*, **51**, pp. 18–26 (2013).

[7] J. Leppänen, T. Viitanen and V. Valtavirta, "Multi-physics Coupling Scheme in the Serpent 2 monte Carlo Code", *Trans. Am. Nucl. Soc.*, **107**, pp. 1165–1168 (2012).

[8] J. Leppänen, "Development of a Dynamic Simulation Mode in Serpent 2 Monte Carlo Code", *Proc. M&C 2013*, Sun Valley, ID, 5–9 May, (2013).

[9] G. Yesilyurt, W. R. Martin and F. B. Brown, "On-the-Fly Doppler Broadening for Monte Carlo Codes", *Nucl. Sci. Eng.*, **171**, pp. 239–257 (2012).

[10] B. Becker, R. Dagan and G. Lohnert, "Proof and implementation of the stochastic formula for ideal gas, energy dependent scattering kernel," *Ann. Nucl. Energy*, **36**, pp. 470–474 (2009).

[11] D. E. Cullen, "Program SIGMA1 (version 79-1): Doppler broaden evaluated cross sections in the evaluated nuclear data file/version B (ENDF/B) format," UCRL-50400 Part B., Lawrence Livermore National Laboratory (1979).

[12] R. E. MacFarlane and D. W. Muir, "NJOY99.0 Code System for Producing Pointwise and Multigroup Neutron and Photon Cross Sections from ENDF/B Data," PSR-480, Los Alamos National Laboratory (2000).

[13] R. V. Meghreblian and D. K. Holmes, *Reactor analysis*, pp. 137140, McGraw-Hill Book Company, New York (1960).

[14] T. Viitanen, "Implementing a Doppler-Preprocessor of Cross Section Libraries in Reactor Physics Code Serpent", Masters thesis, Helsinki University of Technology (2009).

[15] T. Viitanen and J. Leppänen, "New Data Processing Features in the Serpent Monte Carlo Code," *J. Korean Phys. Soc.*, **59**, pp. 1365–1368 (2011).