Title
Implementation of weighted delta-tracking with scattering in the serpent 2 monte carlo code

Permalink
https://escholarship.org/uc/item/6bg1s71k

Authors
Rehak, JS
Kerby, LM
Dehart, MD
et al.

Publication Date
2017

Peer reviewed

Implementation of Weighted Delta-Tracking with Scattering in the Serpent 2 Monte Carlo Code

J. S. Rehak,[a,*] L. M. Kerby,[b,c] M. D. DeHart,[b] R. N. Slaybaugh,[a] J. Leppänen,[d]

[a] *University of California, Berkeley* [b] *Idaho National Lab* [c] *Idaho State University* [d] *VTT Technical Research Centre of Finland*
jsrehak@berkeley.edu

INTRODUCTION

Serpent 2, a continuous-energy Monte Carlo Reactor physics burnup code developed at the VTT Technical Research Centre of Finland [1], uses a common rejection sampling technique, known as Woodcock delta-tracking [2]. This method improves the efficiency of Monte Carlo in the complicated geometric layout of nuclear reactors and is widely used despite drawbacks that make it inefficient in specific cases, such as in the presence of strong absorbers. At present, Serpent actively uses delta-tracking, falling back to a surface tracking option in situations where delta-tracking is inefficient to partially offset these weaknesses. Morgan and Kotlyar [3] introduced a Weighted Delta Tracking (WDT) routine to address these issues, replacing the rejection sampling of absorption events with an implict event.

In this research, we examine the theoretical basis for rejection sampling and implict events, and how they relate to WDT. We then discuss implementation of a routine that extends the WDT routine to include scattering events. The routine is implemented and tested within Serpent. We examine various test cases including a Boiling Water Reactor (BWR) pin cell, and homogenous fuel media based on the content of the fuel of the Transient Test Reactor (TREAT) at Idaho National Lab (INL).

THEORY

The underlying algorithm for simulating neutron transport using Monte Carlo is ray tracing. Ray tracing explicitly tracks a particle between interaction points and different material boundaries. The value of Σ_t is a piece-wise discontinuous function that varies with position and the geometry of the problem. Monte Carlo particles must stop at boundaries to sample a new path length in a new material region. This simplifies the problem by ignoring the complicated nature of $\Sigma_t(\mathbf{r})$ by only sampling in regions where it is a constant value. However, in regions of complex geometries with many boundaries, ray tracing (or surface tracking) can be inefficient. To avoid this, Monte Carlo codes traditionally use a rejection sampling technique known as Woodcock delta-tracking [2].

Woodcock Delta-tracking

Woodcock delta-tracking introduces the concept of the majorant cross-section, chosen to be the maximum of all material total cross-sections in the region of interest:

$$\Sigma_{\text{maj}} \equiv \max_{\mathbf{r}\in\mathbb{V}}\{\Sigma_t(\mathbf{r})\}, \tag{1}$$

where $\mathbb{V}$ is the volume of interest. The majorant cross-section can also be represented as the summation of the total cross-section and a delta cross-section:

$$\Sigma_{\mathbf{maj}} = \Sigma_\delta(\mathbf{r}) + \Sigma_t(\mathbf{r}), \tag{2}$$

where at the position $\mathbf{r}$ where the maximum value of $\Sigma_t(\mathbf{r})$ occurs, the delta cross-section $\Sigma_\delta(\mathbf{r})$ is zero. Elsewhere, is it non-zero and positive. It therefore follows that the sampled path length s is

$$s = -\frac{1}{\Sigma_{\text{maj}}}\ln(\xi), \tag{3}$$

and these samples are accepted with probability

$$P = \frac{\Sigma_t(\mathbf{r})}{\Sigma_{\text{maj}}}. \tag{4}$$

Rejected collisions are considered "virtual" collisions, and a new subsequent path length is sampled. The advantage of delta tracking is that the path length can be sampled across multiple material regions of varying Σ_t without explicitly stopping the neutron at a given boundary. This method can become computationally inefficient in regions where the total cross-section is much less than the majorant cross-section, leading to oversampling of virtual collisions. This is seen in geometries that include localized absorbers, such as control rods. Another downside is that the track-length estimator (TLE) for flux cannot be used. The TLE requires calculating the track lengths within a particular material cell, and therefore does not work when the neutron path length can cross one or more material boundaries. The collision flux estimator (CFE) can be used in its place, but often results in inferior statistics as not every track length sampled ends in a collision [4].

Weighted Delta Tracking (WDT)

Morgan and Kotlyar [3] introduced the weighted delta tracking (WDT) method to improve the inefficiencies of Woodcock delta-tracking in the presence of large absorbers. The method replaces the rejection sampling of delta-tracking with a weight reduction.

The WDT method samples the particle path length in the same fashion as Woodcock delta-tracking. The WDT method then bypasses the rejection probability by accepting all collisions as real with a subsequent reduction in weight. Replacing this statistical rejection with a weight reduction requires calculation of the expected value of the weight following the collision. In this case, the two events are a real collision and a virtual collision.

$$E[w_f] = w_{f,\text{real}}P_{\text{real}} + w_{f,\text{virt}}P_{\text{virt}} \tag{5}$$

Where P_{real} and P_{virt} are the probabilities of a real and virtual collision, respectively, and $w_{f,\text{real}}$ and $w_{f,\text{virt}}$ are the corresponding post-collision weights. Prior to calculating this expected value, the type of interaction must be sampled, as this will determine the final weight of a real collision, $w_{f,\text{real}}$. The calculation of the expected value for absorption and scattering are described below.

Absorption

Morgan and Kotlyar examine a 1D test case with absorption. As an absorption event removes the particle, $w_{f,\text{real}} = 0$, and a virtual collision leaves the initial weight unchanged, $w_{f,\text{virt}} = w_i$. Inserting the appropriate values into Eq. (5) gives the expected value of the final weight for an absorption event.

$$\begin{aligned} E[w_f] &= w_{f,\text{real}} P_{\text{real}} + w_{f,\text{virt}} P_{\text{virt}} \\ &= 0 + w_i P_{\text{virt}} \\ &= w_i (1 - P_{\text{real}}) \\ &= w_i \left(1 - \frac{\Sigma_t}{\Sigma_{\text{maj}}} \right) \end{aligned} \tag{6}$$

The particle that is left following the collision continues propagating as if it underwent a virtual collision. In this case, the absorption is then scored using the expected value of the score.

$$\begin{aligned} S_{\text{absorption}} &= E[w_i - w_f] \\ &= E[w_i] - E[w_f] \\ &= w_i \left(\frac{\Sigma_t}{\Sigma_{\text{maj}}} \right) \end{aligned} \tag{7}$$

This is implemented by Kotlyar and Morgan in a 1D problem and the results are verified with the analytical solution. The authors point out that a rouletting routine should be implemented when this is used, to prevent the tracking of low-weight neutrons.

Scattering

In order to implement WDT in Serpent 2, the theory for scattering events must be developed. In a scattering event, the weight of the incident particle does not change, $w_{f,\text{real}} = w_i$. Therefore, application of the expectation value results in a final weight equal to the initial weight.

$$\begin{aligned} E[w_f] &= w_{f,\text{real}} P_{\text{real}} + w_{f,\text{virt}} P_{\text{virt}} \\ &= w_i P_{\text{real}} + w_i P_{\text{virt}} \\ &= w_i (P_{\text{real}} + 1 - P_{\text{real}}) \\ &= w_i \end{aligned} \tag{8}$$

The WDT method splits the weight of the colliding particle into a virtual portion and a real portion (dependent on the P_{real} from Eq. (4)). The real portion of the weight undergoes the collision; this is the portion that was used to score absorption. The virtual portion of the weight is left with the particle that continues propagating as if no collision had occured at all. In the event of absorption, the real portion of the weight is attributed to a particle that is then immediately killed, but this is not the case in a scattering event.

Extension of this methodology to scattering requires duplication of the particle at the point of collision. The virtual portion of the weight is carried away by a particle that propagates as if no collision has occured, and the real portion is carried away by a particle that undergoes scattering. In problems with scattering, this results in a rapid multiplication of neutrons. When implemented into Serpent 2, this multiplication very quickly filled any available neutron buffer in simulations of a BWR, ending the simulation.

To maintain proper statistics, scattering must take into account the possibility of a virtual collision while using delta-tracking. To achieve this goal, the delta-tracking rejection sampling that had been supplanted by WDT was moved into the scattering subroutine. Therefore, the new routine uses both rejection-sampling and implicit events to account for the possibility of real and virtual collisions. The algorithm is shown in Alg. 1. The algorithm used by Morgan and Kotlyar removes all rejection sampling, completely replacing it with a weight reduction. This carries the benefit of not requiring repeated path length sampling when virtual collisions take place. This benefit was shown in a pure-absorbing media. When we add scattering to the routine the rejection sampling must be reintroduced, as before. Unlike Woodcock delta-tracking, every absorption event is scored as a real collision; only scattering events have the chance of being virtual. Overall, this should reduce repeated sampling of virtual collisions, especially in the vicinity of heavy absorbers.

Algorithm 1 WDT with scattering

```
Sample path length
Sample collision type
if collision type == (capture or fission) then
    Score capture or fission ← $w_i P_{\text{real}}$
    Score collision ← $w_i P_{\text{real}}$
    $w_f \leftarrow w_i (1 - P_{\text{real}})$
    Execute virtual collision
else
    Sample random number $\xi \in [0, 1)$
    if $\xi < P_{\text{real}}$ then                ▷ Collision is real
        Score scattering ← $w_i$
        Score collision ← $w_i$
        Execute scattering collision
    else                                          ▷ Collision is virtual
        Execute virtual collision
    end if
end if
```

It is also important to note that the addition of WDT does not effect the decision tree that determines if surface tracking is used; surface tracking is still used when $\Sigma_t / \Sigma_{maj} < 0.1$. Future work plans to test replacing some surface tracking with WDT.

Rouletting Routine

Replacing a statistical process with a weight reduction often requires introduction of a Russian rouletting routine. This prevents the simulation from tracking low weight neutrons and near infinite loops that may result when there are few neutron removal mechanisms. A rouletting routine was added to Serpent 2 that is called automatically if WDT is used. Arbitrary values were chosen for the weight cutoff and rouletting probability. The tuning of these values to improve performance is a direction for future work.

RESULTS

Homogenized TREAT Fuel Element

Researchers at INL are using Serpent 2 to generate cross-sections for TREAT. These cross-sections are used by various deterministic codes that model the neutronics and coupled multi-physics of the reactor. Generation of low-error cross-sections with the current methods in a reasonable amount of time requires billions of simulated particles. Improved variance reduction methods, such as WDT, may provide a means to improve the ability of the code to quickly generate cross-sections with low error. The preliminary test case chosen to assess the ability of the WDT method to improve cross-section generation is a homogenized TREAT fuel element.

The energy spectrum is divided into eleven groups. The simulation was repeated with identical seeds both with and without using the WDT method. The runtime of both simulations are shown in Table I; the WDT method causes a slight increase in runtime. Most of the extra processing time is due to sampling the type of collision after every delta-tracking path length, instead of only sampling for real collisions. The Figure of Merit (FOM) used to asses this method is given by:

$$\text{FOM} = \frac{1}{\sigma(\hat{x})^2 T} \tag{9}$$

where T is the runtime of the simulation. Table II presents the criticality parameters for the simulation, and shows that the FOM when using the WDT method falls short or is comparable to running the simulation without WDT.

	T (min)
No WDT	118.883
WDT	121.212

TABLE I. Runtime for the treat homogenous fuel element.

Fig. 1 shows a plot of the convergence of FOM for Analog k_{eff} versus simulation cycles and Fig. 2 shows the FOM convergence for the highest energy group in the homogenous fuel cell. These are two cases where the FOM when using WDT converges to a higher value than when WDT is not used.

Note that a homogenized fuel element will not leverage the expected advantages of WDT, as there are no complex geometries or regions of semi-voids and strong absorbers. Future work includes testing the WDT implementation on a more realistic TREAT simulation. This simulation will include air cooling channels and a large air channel used with a hodoscope, which is for real-time observation of fuel dynamics.

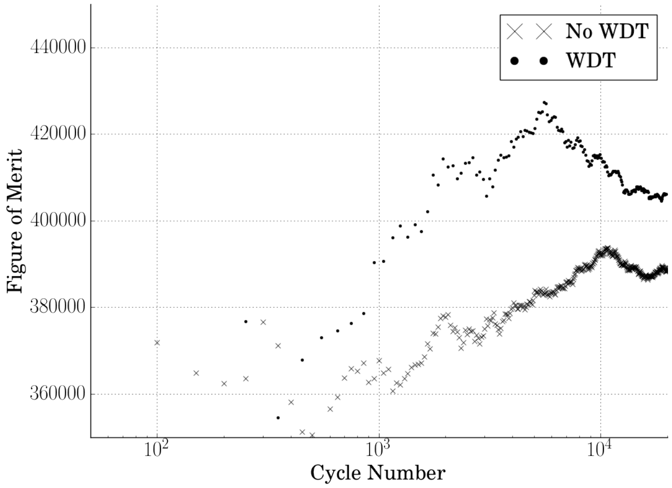

Fig. 1. Convergence of FOM for Analog k_{eff} for the homogenous fuel element.

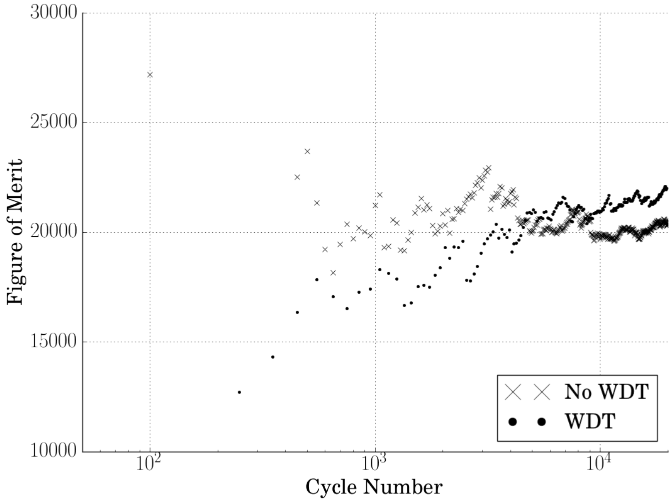

Fig. 2. Convergence of the FOM for infinite flux for the highest energy group in the homogenous fuel element.

Parameters	No WDT			WDT			Ratios	
	$\hat{x}$	$\sigma^2(\hat{x})$	FOM	$\hat{x}$	$\sigma^2(\hat{x})$	FOM	$\sigma^2(\hat{x})$	FOM
Implicit k_{eff}	1.77209	0.000008	1.382582×10^8	1.77208	0.000009	1.018520×10^8	1.15	0.74
Analog k_{eff}	1.77272	0.000240	1.460353×10^5	1.77254	0.000230	1.559548×10^5	0.96	1.07

TABLE II. Criticality parameters for the TREAT homogenous fuel pin. The ratios column is expressed as the value with WDT divided by the value without WDT.

BWR

A threshold value for $P_{\text{real}} = \Sigma_t/\Sigma_{maj}$, above which regular delta-tracking is used instead of WDT may improve performance. A simulation was run on a BWR and various values of this WDT threshold were tested. The surface tracking routine was not modified, so when $\Sigma_t/\Sigma_{maj} \leq 0.10$ normal surface tracking is used. The total FOM for all cross-sections of interest was calculated and results are shown in Fig. 3.

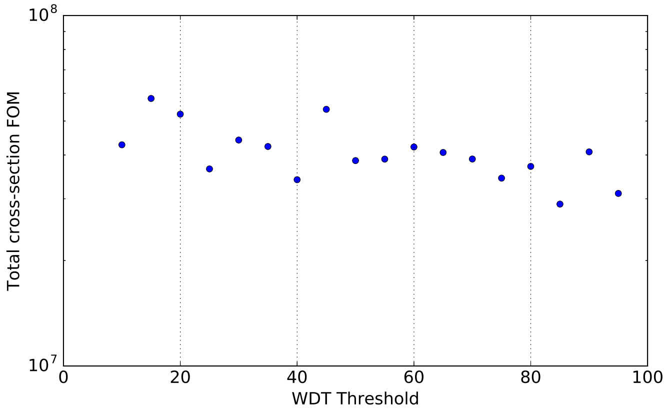

Fig. 3. FOM for the BWR for different WDT threshold values (percentages).

The total FOM is scattered but shows a decreasing trend as the threshold value rises. Below 25%, there is increasing value in limiting WDT use to only the more extreme cases of cross section disparities. This trend abruptly ends at the 10% value. Note, however, that the cutoff value corresponding to 10% means that no WDT is used–the simulation is the original Serpent with surface tracking and delta tracking. Future work exploring the use of WDT in place of surface tracking is planned.

CONCLUSION

Weighted delta tracking (WDT) has been implemented within Serpent 2 and preliminarily tested. WDT has the potential to make up for a major disadvantage of delta-tracking: that virtual collisions offer no statistical benefit. Results indeed show promise, but further exploration is necessary. Implementation of WDT with scattering requires the introduction of a rouletting routine and a threshold value, both of which need to be tuned.

Preliminary results show that the WDT method is as good as normal delta-tracking, returning a higher FOM in some cases. Therefore, with tuning and adjustment, the method may return better statistical results in future studies.

FUTURE WORK

Many of the FOM values and results presented need to be evaluated through multiple batches of runs to determine if they are consistently observed. Importantly, different energy groups need to be used and evaluated, as the top of the eleven energy groups used for the TREAT reactor consists of the entire resonance region. We anticipate that WDT will provide better results in this region, where there can be many virtual collisions that do not contribute to the statistics. This may explain the improvement seen in the highest energy group FOM shown, but this needs to be verified. Smaller energy bins will give more insight into the possible benefits of the method.

The rouletting routine and WDT threshold values require tuning. Although some testing was done by changing the WDT threshold value, this needs to be examined closer. In addition, WDT may offer advantages over surface tracking and this should also be explored. Lastly, the rouletting routine uses an arbitrary weight cutoff, which could be tuned to provide better results.

REFERENCES

1. J. Leppänen, M. Pusa, T. Viitanen, V. Valtavirta, and T. Kaltiaisenaho, “The Serpent Monte Carlo code: Status, development and applications in 2013”, Ann. Nucl. Energy **82**: pp 142–150 (2015).
2. E. Woodcock, et al., “Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry”, ANL-7050, Argonne National Laboratory (1965).
3. L. Morgan and D. Kotlyar, “Weighted-delta-tracking for Monte Carlo particle transport”, Ann. Nucl. Energy **85**: pp 1184–1188 (2015).
4. J. Leppänen, “Modeling of Nonuniform Density Distributions in the Serpent 2 Monte Carlo Code”, Nuclear Science and Engineering **174**: pp 318–325 (2013).