Correcting k_{eff}-Induced Bias in the Time-Dependent Random Ray Method

Maximilian Kraus [1,*], **Paul Cosgrove**[1], **John Tramm**[2], **Eugene Shwageraus**[1]

[1]Department of Engineering, University of Cambridge, Cambridge, UK;
[2]Argonne National Laboratory, IL, USA

[leave space for DOI, which will be inserted by ANS]

ABSTRACT

The Random Ray Method (TRRM) is a stochastic adaptation of the Method of Characteristics, calculating the neutron flux along randomly sampled rays rather than fixed tracks. This paper addresses a bias observed in transient simulations with previously introduced time-dependent TRRM, which stems from a mismatch between the ray layout used to calculate k_{eff} at $t = 0$ and the ray layout in later time steps. In this work, we discuss two methods to mitigate this bias: a previous approach that recalculates the steady-state k_{eff} at each time step and a new approach that repeats the same quadrature sequence in each time step. Tests on time-dependent C5G7 benchmark problems showed that both methods removed the bias. However, the new approach significantly improved efficiency, bringing the mean error down to 0.06-0.15 % with a shorter runtime and greatly reduced noise between time steps. This method represents a major advance in precision and speed for time-dependent TRRM.

Keywords: Transient Modelling, Random Ray Method, Method of Characteristics, Bias Correction

1. INTRODUCTION

1.1. The Significance of k_{eff} in Transient Simulations

The time-dependent neutron transport equation in its integro-differential form without external sources is the following:

$$\frac{1}{v}\frac{\partial}{\partial t}\psi + \mathbf{\Omega}\cdot\nabla\psi + \Sigma^{\text{t}}\psi = \int_{4\pi}\int_0^\infty \Sigma^{\text{s}}\psi dE d\Omega + \frac{\chi^{\text{p}}}{4\pi}(1-\beta)\int_0^\infty \nu\Sigma^{\text{f}}\phi dE + \sum_{j=1}^{N}\frac{\chi_j^{\text{d}}}{4\pi}\lambda_j C_j\ . \tag{1}$$

Here the dependencies on location $\mathbf{r}$, energy E, time t and direction $\mathbf{\Omega}$ are omitted for the sake of brevity. Variables ψ and ϕ represent the angular and scalar neutron fluxes. Quantities Σ^{t}, Σ^{s}, Σ^{f}, v and ν characterise the total, scattering and fission cross-sections, neutron velocity and average number of neutrons released per fission reaction, respectively. A small fraction β of the newly born neutrons emerge as delayed neutrons resulting from the decay of precursor nuclides (represented by $\lambda_j C_j$). Symbols χ^{p} and χ^{d} describe the energy spectrum of the released prompt and delayed neutrons.

Eq. (1) describes the time-dependent behaviour of the neutron flux distribution over time. When neutron productions (through prompt and delayed neutron releases) and losses (through absorption and leakage) even out, the reactor is critical and in a steady state. In numerical simulations, exact criticality is rare, so k_{eff} is used to scale the fission neutron production rate to $\frac{\nu\Sigma^{\text{f}}\phi}{k_{\text{eff}}}$ to enforce a balance between production and loss rates. Transient simulations typically begin from this steady state, and therefore scale the fission rate with the initial k_{eff} before any transient-triggering perturbations occur. This adjustment with the steady-state multiplication factor is applied across all time steps in a transient simulation (see e.g. Ref. [1]).

---

[*] mhgk4@cam.ac.uk

1.2. The Problem Observed in the Time-Dependent Random Ray Method

The Random Ray Method (TRRM) is a novel variation of the conventional, deterministic Method of Characteristics (MOC), which calculates the neutron flux along tracks (rays) that transect the geometry and accumulates the weighted angular contributions within each geometry cell towards the scalar flux together with the respective source term. Contrary to conventional MOC, TRRM uses a stochastically sampled quadrature, laying tracks with randomly sampled directions and starting points across the model geometry. However, it calculates all quantities via deterministic equations, which makes this method a hybrid of deterministic and probabilistic methods.

TRRM needs a number of inactive iterations (cycles) to allow the fission and scattering source distributions to become stationary, such that each iteration is statistically similar. Similar to Monte Carlo simulations, it also requires a number of active iterations, over which the results are eventually averaged to achieve sufficient statistical confidence. TRRM has several advantages over conventional MOC, e.g. higher accuracy by sampling angles from a continuous distribution, lower required ray density, reduced memory requirements and great geometric flexibility. For an in-depth description of this method, the interested reader is referred to Refs. [2, 3].

Refs. [4, 5] recently introduced two versions of time-dependent TRRM for transient simulations. One of them, the so-called time-implicit approach, uses a backwards Euler time discretisation with an isotropic approximation for the time derivative (see also Sec. 2). In that approach, an eigenvalue calculation for the steady state is performed prior to the transient calculation. The fission cross-sections for all subsequent time steps are then divided by the value for k_{eff} before introducing the perturbations that start the transient (see Fig. 1). In Ref. [4], a bias of several percent in some cases was discovered in the transient TRRM results for the C5G7-TD3-4 benchmark exercise, which is depicted in Fig. 2.

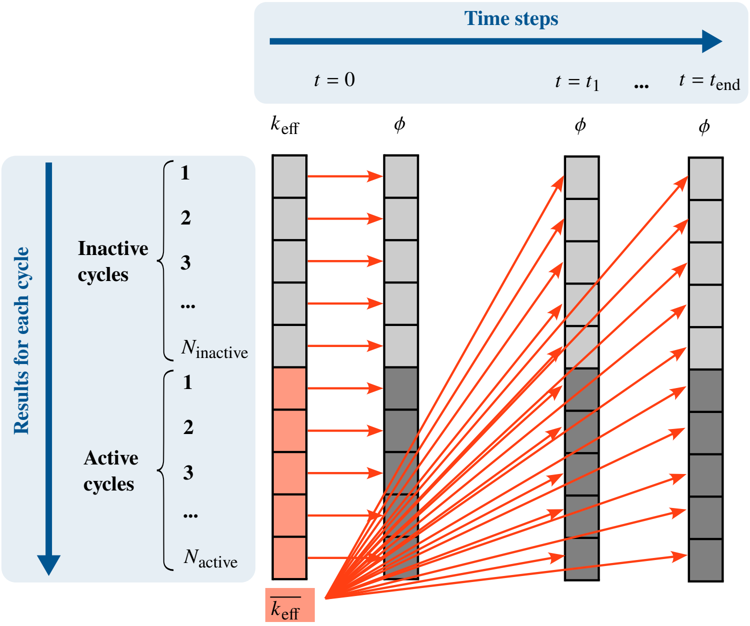

Figure 1. Visual representation of the time-implicit TRRM calculation scheme. k_{eff} is calculated for $t = 0$, averaged over the active cycles and then used to scale the fission rate in all subsequent time steps.

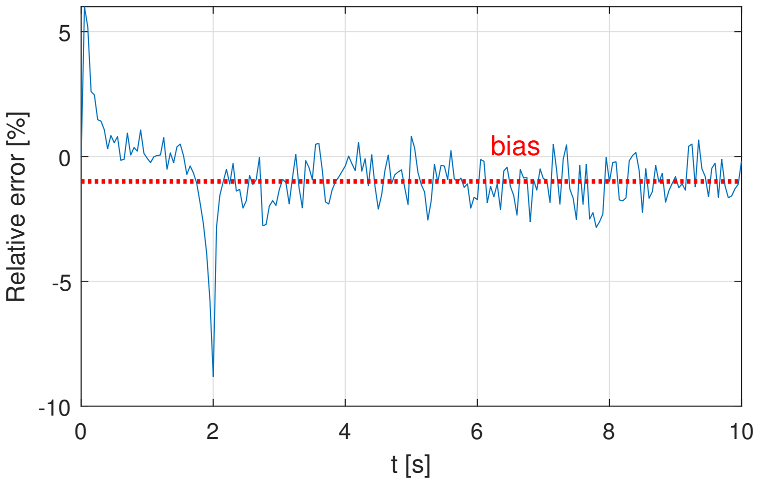

Figure 2. Error between results obtained with time-dependent (time-implicit) TRRM and the reference solution for the C5G7-TD3-4 benchmark exercise.

The origin of this bias was found to be a mismatch between k_{eff} that is obtained as an average over the active cycles in the steady state on a certain quadrature and the reactivity associated with the specific, changing ray layout in each subsequent time step. Note that this mismatch only concerns the steady-state reactivity and is unrelated to the external reactivity insertion that triggers the transient. As a consequence, this discrepancy leads to the fission rate not being adjusted properly, resulting in small additional upward or downward shifts of the neutron flux on top of the leading transient perturbation.

Ref. [5] presented a work-around for this problem by recalculating the steady-state k_{eff} for the specific ray layout in each iteration of each time step. The work at hand introduces a new approach, where the pseudo-random number seed of the first time step is reused for all subsequent time steps of the transient. Thereby the same quadrature sequence is generated in each time step and the k_{eff} values from $t = 0$ can be applied without giving a bias.

The previous and new procedures are described in Sec. 2. Both approaches are tested on the two-dimensional C5G7 benchmark geometry [1], first simulating a null transient (critical system without perturbation) and then the time-dependent cases TD1-2 (control rod insertion) and TD3-4 (moderator density change). The obtained results are compared to the biased results of the original time-dependent TRRM in terms of accuracy and runtime.

2. ADJUSTMENTS TO THE TIME-DEPENDENT RANDOM RAY METHOD

The time-dependent random ray approach used in this work corresponds to the "time-implicit" approach in Ref. [4]. This approach uses an isotropic approximation for the time derivative in Eq. (1) where $\frac{\partial \psi}{\partial t} = \frac{1}{4\pi}\frac{\partial \phi}{\partial t}$. After discretising the problem into spatial cells and time steps, a first-order backwards differentiation formula is used to represent the time derivative. The contribution of the delayed neutrons from the decay of the precursor nuclei is obtained by solving the delayed neutron precursor differential equation analytically, making use of a second-order approximation for the fission source as suggested in Ref. [6]. The time-implicit algorithm calculates the spatial scalar flux solution time step by time step, always completing the solution of the current time step before moving on to the subsequent time step, where it then uses the previous result in the approximation of the time derivative. In the following, two approaches are presented to remove the bias discussed under Sec. 1.2.

2.1. Option 1: Recalculation of k_{eff}

The procedure proposed in Ref. [5] recalculates the steady-state k_{eff} in each cycle at each time step on the same quadrature for which the transient flux results are calculated (see Fig. 3). Thus, for each cycle, the fission source is adjusted by the corresponding recalculated steady-state multiplication factor of that cycle. This procedure removes the bias that originated from the mismatch between k_{eff} and the ray layout in the transient simulations. However, this approach constitutes an additional computational burden since the transport sweeps have to be carried out for both the steady-state and the time-dependent flux at each time step. As the transport sweeps for both are very similar, they can make efficient use of many shared functionalities, which alleviates the runtime increase slightly.

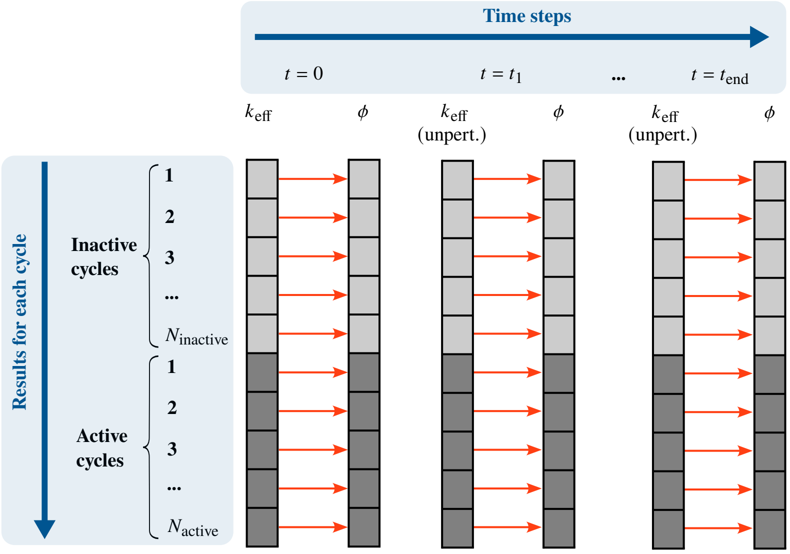

Figure 3. Visual representation of the time-dependent TRRM scheme with recalculated k_{eff}. The recalculation is performed for the unperturbed problem (unpert.) with the same ray layout that is used for the transient computation of the neutron flux ϕ at each respective time step.

2.2. Option 2: Using a Time-Consistent Seed

A new option to remove the bias reuses the same pseudo-random number to sample the starting locations of the rays in each time step, such that each time step traverses the same ray sequence during its inactive and active cycles. Note that the quadrature between individual cycles within a given time step still changes as the "re-seeding" is only done once at the start of each time step. The random number generator then produces the same sequence of pseudo-random numbers to generate subsequent rays, replicating the ones of the steady-state k-eigenvalue calculation. In this procedure, it is important that all values of k_{eff} from the initial steady state calculation are retained and stored in an array rather than just the averaged value. Thus, in subsequent time steps, the specific k_{eff} for each corresponding cycle (and hence ray layout) can be loaded and used to adjust the fission source. Fig. 4 shows a schematic representation of this approach.

There are certain subtleties that require further explanation (which are not shown in Fig. 4 for the sake of simplicity): The steady-state calculation is typically initialised with $k_{\mathrm{eff}} = 1$ and a uniform flux distribution, and needs a number of inactive cycles to converge. Thus, the cycle-wise k_{eff} values from this initial run will not scale the fission sources correctly in later time steps. Therefore, the steady-state calculation is repeated a second time. The second run is initialised with the averaged k_{eff} and flux scores from the first. The k_{eff} values from this second run are then stored for reuse in subsequent time steps.

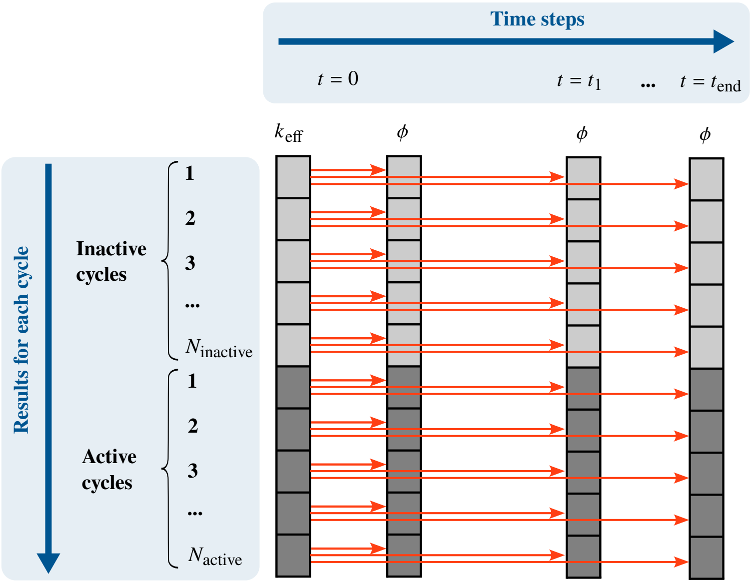

Figure 4. Visual representation of the time-dependent TRRM scheme with a time-consistent random number seed. The ray sequence traversed over the course of the inactive and active iterations of each time step is the same as at $t = 0$. Each cycle uses the corresponding k_{eff} value from the same cycle of the steady-state calculation.

Since a dynamic convergence criterion is chosen for the analyses below (see Sec. 3.2), each time step requires a different number of iterations. In order to generate a large enough array of k_{eff} values, the steady-state calculations are performed for a large, fixed number of inactive and active cycles. In this work, the steady-state number of cycles is set to twice the usual amount needed to converge and score results for the analysed C5G7 problems. The additional runtime for the repeated steady-state k-eigenvalue calculations and the extra memory for storing the cycle-wise results of k_{eff} are negligible compared to the overall computational requirements of a transient simulation.

3. COMPARISON ON THE C5G7-TD BENCHMARK

3.1. Benchmark Description

The unaltered time-dependent Random Ray Method, along with the modified versions introduced in the previous section are tested and compared using the C5G7-TD geometry [1]. This benchmark comprises a small light-water reactor with 16 fuel assemblies, one half containing uranium dioxide fuel and the other half mixed oxide fuel in three different degrees of enrichment (4.3, 7.0 and 8.7 %). Each fuel assembly accommodates a 17 × 17 grid filled with fuel pins and guide tubes. A water reflector surrounds the core region, with the model boundaries being set to vacuum. Due to its symmetric layout, often only a quarter-core geometry is analysed. The benchmark is specified for 7 neutron energy groups and 8 delayed neutron groups. It describes five transient problem exercises, each with several variations. In this study, two of the two-dimensional exercises are simulated, both starting with all control rods fully extracted.

The first transient problem (TD1-2) characterises a control rod movement in one rod bank. Those rods are inserted at constant speed to a depth of 1 % of the total core height over one second and subsequently extracted again over another second. Since this is a two-dimensional benchmark exercise, the insertion is represented by a proportional mixture of the guide tube and control rod cross-sections. In the second exercise (TD3-4), the transient is introduced by a gradual change of the moderator density (except for the reflector) to 80 % of its original value over one second, followed by a gradual return to the original value over one more second. Since the moderator cross-sections scale linearly with the density change, the perturbation is imposed via a corresponding variation of the respective cross-sections.

3.2. Input Specification

Time-dependent TRRM is implemented within the SCONE Monte Carlo code [7], written in Fortran 2008. A reference solution for the C5G7-TD transients is provided by Ref. [8], obtained with the code MPACT. MPACT employs the Transient Multilevel Method, which combines point kinetics, coarse mesh finite differences diffusion and MOC solvers to simulate time-dependent problems.

The transients are run with a time step size of 0.01 s, matching the value in the MPACT study. The ray population (650), ray length (628 cm) and dead length (13 cm) are adopted from Ref. [3], where these values were given as optimum values for the steady-state C5G7 case. The spatial discretisation is chosen such that it corresponds to the one used for the reference solution. In the pin cells, 5 equal-area rings are used for the fuel region and 3 for the moderator region, each with 8 azimuthal divisions. The reflector region is divided up into 5 x 5 square cells per pin cell equivalent.

To ensure an efficient calculation scheme, the inactive cycles are stopped once a specified convergence criterion is met. The chosen convergence method calculates the fission source distribution F_i over mesh-cells i in each cycle N_{it} and stores the resulting values over a certain number of inactive iterations N_{w} in a "moving window". The criterion then checks if the root-mean-square (RMS) error between the cell-wise averages taken over the first and second halves of the moving window becomes sufficiently small such that it falls below the requested threshold value RMS_{prec}:

$$\text{RMS} = \sqrt{\frac{\sum_i^{N_{\text{cell}}} \left(\frac{\overline{F}_i^{\text{new}} - \overline{F}_i^{\text{old}}}{\overline{F}_i^{\text{new}}} \right)^2}{N_{\text{cell,fissile}}}} \overset{!}{<} \text{RMS}_{\text{prec}} \tag{2}$$

$$\text{with} \quad \overline{F}_i^{\text{new}} = \frac{\sum_{it=N_{\text{it}}-0.5N_{\text{w}}+1}^{N_{\text{it}}} \left(F_i^{it} \right)}{0.5N_{\text{w}}} \quad \text{and} \quad \overline{F}_i^{\text{old}} = \frac{\sum_{it=N_{\text{it}}-N_{\text{w}}+1}^{N_{\text{it}}-0.5N_{\text{w}}} \left(F_i^{it} \right)}{0.5N_{\text{w}}} \, .$$

The threshold is set to $\text{RMS}_{\text{prec}} = 1.1\text{E} - 2$ because it is the lowest that can be achieved in the analysed problems before numerical and statistical noise begin to interfere.

For the active cycles, the source average relative error (SARE) is monitored to determine when enough samples have been scored to obtain a sufficiently low variance for the solution [3]. The SARE criterion averages the relative standard deviation of the accumulated fission source term over all cells containing fissile material during the active iterations, updating with each new iteration until the uncertainty is smaller than the specified threshold. This value is set to $\text{SARE}_{\text{prec}} = 2.2\text{E} - 3$, representing a balanced trade-off between runtime and statistical variance that falls within the range recommended in Ref. [3].

The criteria presented here evaluate only the fission source values, following the approach in Ref. [3]. In theory, these procedures can be easily modified to track the evolution of other quantities, such as the scalar flux. However, for the simulated C5G7 problems, monitoring the fission rate was found to offer a more stable convergence method, while calculating the RMS error over scalar flux values introduced significant noise, making it harder to determine if the calculations had converged.

All calculations are run with 112 threads on a dual-socket Intel Xeon Platinum 8480+ node. In each simulated case, the same pseudo-random number seed is used to initialise runs with the original and adjusted time-dependent TRRM versions.

3.3. Results on C5G7 Null Transient

First, the three versions of time-dependent TRRM, i.e., the original version where k_{eff} is only calculated during the first (steady-state) time step, the version that recalculates k_{eff} and the version that uses a time-consistent random number seed, are tested on the C5G7 benchmark, however without any perturbation over time. Fig. 5 shows the corresponding results and the error for each approach over time. In this unperturbed transient case, the fission rate, normalised to its initial value at $t = 0$, is expected to remain at a constant value of 1 over time. This behaviour is accurately reproduced by the time-consistent seed version, with almost no error at all. The approach that recalculates k_{eff} gives results that are centred around the correct value of 1, however they are superimposed with noise. The original, uncorrected TRRM version, shows an average bias of 1.5 to 2 %, along with some noise.

These results again underline the problem of the original version described in the introduction and show that both presented adjustments remove the bias. As an additional benefit that enhances precision, the noise between time steps is removed from the results when fluxes are calculated on the same quadrature sequence each time step. Noise in the spatial domain, however, remains.

3.4. Results on C5G7-TD Exercises

Next, the two adjusted time-dependent TRRM versions are tested on transient exercises TD1-2 and TD3-4, and compared to the MPACT reference solution, along with the original, uncorrected version. Table I shows the performance data for these different approaches on both benchmark exercises, listing the maximum and mean error from the reference solution, runtime and standard deviation. The transient results of each variant are plotted along with the respective error over time for TD1-2 and TD3-4 in Figs. 6 and 7, respectively.

In both figures, one can observe that the uncorrected approach shows a clear bias, which is removed by the two approaches described in Sec. 2. Again, the time-consistent seed approach shows no noise. The apparent

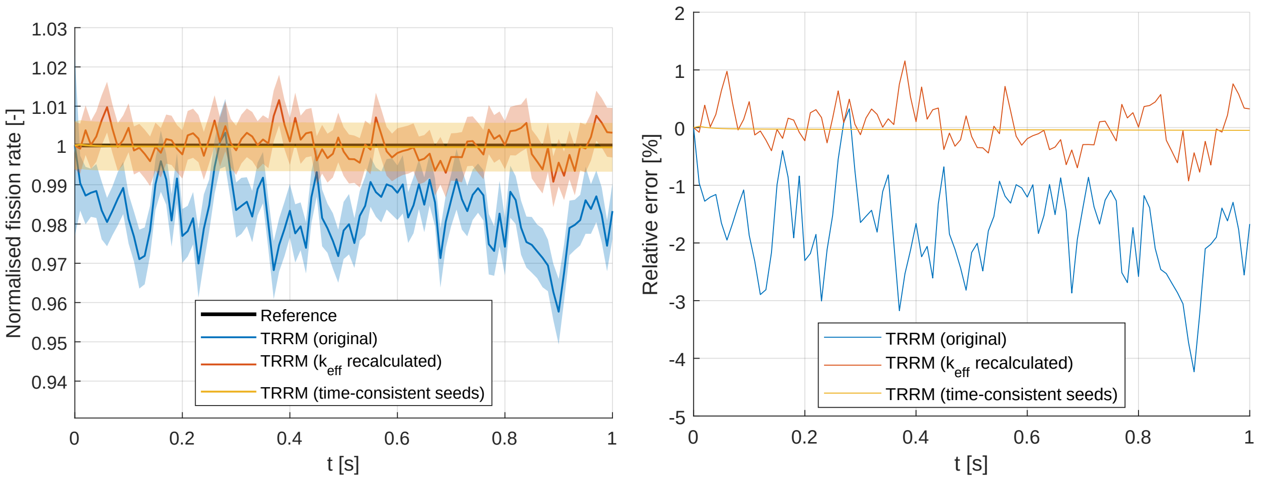

(a) Normalised fission rate results. **(b) Relative error from reference solution.**

Figure 5. Results for a transient simulation without perturbation on C5G7 geometry. The shaded area in subfigure (a) marks the respective confidence interval of one standard deviation.

saw-tooth pattern in the error plots is a result of the resolution of the available MPACT reference data (in steps of $1\mathrm{E}-3 = 0.1\,\%$). Both mean and maximum error in TD1-2 are drastically reduced (mean error 98 % lower than the original and 82 % lower than the recalculation version). The results of TD3-4 show a strong improvement of the mean error as well, with a 91 % lower error than the original approach. However, the maximum error has grown slightly compared to the method that recalculates k_{eff}. It is important to note that the changes in accuracy shown here are only examples. In other cases, the bias in the original version may be less pronounced as a favourable ray layout can obscure the issue.

The maximum error occurs just before $t = 2$ s, when the moderator density has almost recovered to its initial value. At this point, the difference between the fluxes from the previous and current time steps, used to approximate the time derivative, is largest, as indicated by the steep slope. Time-dependent TRRM can take a long time to converge fully when large flux changes occur between time steps and the inactive iterations might end prematurely. Choosing a stricter target precision value that the convergence criterion has to meet would reduce the error in these regions. However, imposing a higher precision for the convergence criterion

Table I. Performance indicators for the original and adjusted time-dependent TRRM methods for C5G7-TD1-2 and TD3-4 benchmark exercises.

	uncorrected	**k_{eff} recalculated**	**time-consistent seed**
	TD1-2		
Maximum error [%]	5.24	1.38	0.27
Mean error [%]	2.51	0.34	0.06
Runtime [h]	8.35	11.40	7.28
Mean standard deviation [%]	0.70	0.62	0.62
	TD3-4		
Maximum error [%]	4.60	2.01	2.31
Mean error [%]	1.60	0.38	0.15
Runtime [h]	8.12	11.76	7.77
Mean standard deviation [%]	0.70	0.62	0.62

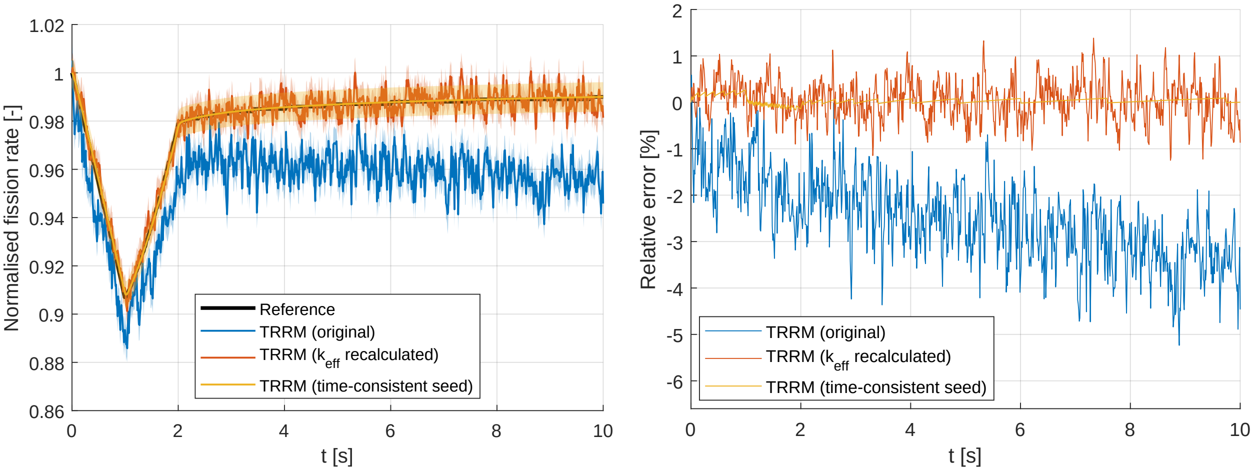

(a) Normalised fission rate results. **(b) Relative error from MPACT reference solution.**

Figure 6. Results for the transient simulation of benchmark exercise C5G7-TD1-2. The shaded area in subfigure (a) marks the respective confidence interval of one standard deviation.

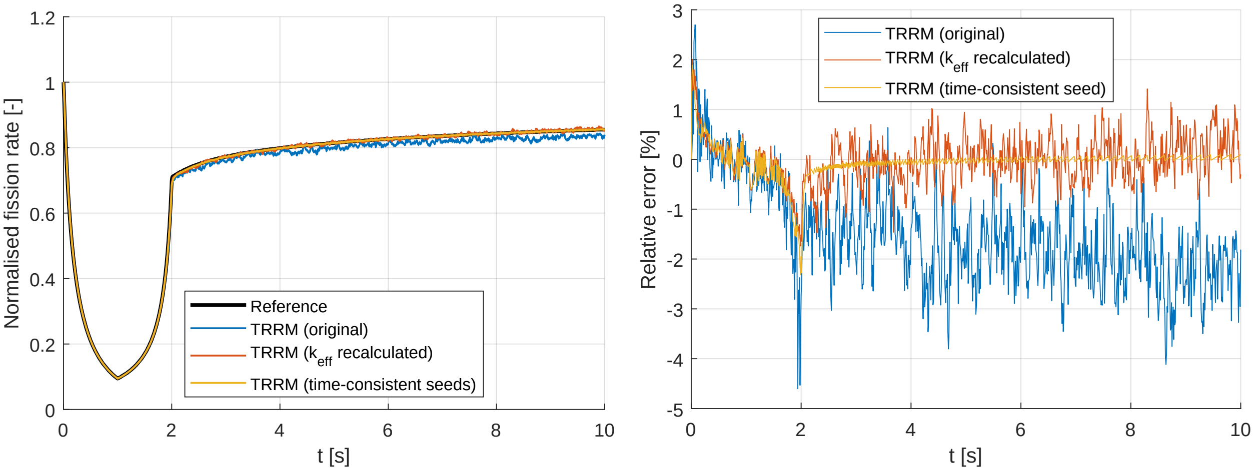

(a) Normalised fission rate results. **(b) Relative error from MPACT reference solution.**

Figure 7. Results for the transient simulation of benchmark exercise C5G7-TD3-4. The confidence interval of one standard deviation (shaded) is barely visible in subfigure (a).

in Eq. (2) will also lead to more iterations at times that are already sufficiently converged with the current setting, thereby increasing the overall runtime unnecessarily.

Regarding the runtimes, it is evident that recalculating k_{eff} induces a penalty in computation time (36 and 45 % higher than the original for TD1-2 and 3-4, respectively), because it involves not only a time-dependent transport sweep but also a transport sweep for the unperturbed problem for each ray. On the other hand, the time-consistent seed approach reduces the runtime compared to the original version slightly by 5 to 10 %.

4. CONCLUSIONS

This work discussed the problem of a bias that was found in the initial (time-implicit) implementation of the time-dependent Random Ray Method as a result of a mismatch between k_{eff}, obtained at $t = 0$, and the steady-state reactivity associated with the varying ray layout in later time steps. A first approach to remove this bias was introduced in a previous publication [5]. This approach requires the recalculation of the steady-state k_{eff} at each time step of the transient, using the same quadrature that is used for calculating the transient neutron flux. The work at hand presented a new method: The pseudo-random number seed used to generate the ray layout sequence in the first (steady-state) time step is reused in all subsequent time steps to ensure the same ray sequence is repeated. The k_{eff} values obtained at $t = 0$ are stored for each cycle and then applied to adjust the fission sources in subsequent time steps as well, such that the quadrature used for their calculation matches the one of the current iteration.

The original, uncorrected approach and the adjusted versions were applied to simulate various transients on the C5G7 benchmark geometry. While the approach that recalculates k_{eff} was able to remove the bias, it also led to an increased runtime because transport sweep and scoring subroutines had to be carried out in each iteration for steady-state quantities as well as transient quantities. The time-consistent seed approach, however, proved highly efficient: Not only did it remove the bias, it also eliminated the noise between time steps that the other variants suffered from. Altogether, this approach reduced the mean error to 0.06 % in TD1-2 and 0.15 % in TD3-4. At the same time, it shortened the runtime by 5 to 10 % compared to the original approach.

This study has shown that the time-consistent seed adjustment for time-dependent TRRM results in a significant improvement in computational efficiency. This approach should therefore be implemented as a default procedure in TRRM for future simulations and further development of this method.

ACKNOWLEDGEMENTS

This work was supported by the University of Cambridge Harding Distinguished Postgraduate Scholars Programme and the EPSRC grants MaThRad (EP/W026899/1) and MCSIMus (EP/W037165/1). This work was performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk).

REFERENCES

[1] J. J. Hou, K. N. Ivanov, V. F. Boyarinov, and P. A. Fomichenko. "OECD/NEA benchmark for time-dependent neutron transport calculations without spatial homogenization." *Nuclear Engineering and Design*, **volume 317**, pp. 177–189 (2017).

[2] J. R. Tramm, K. S. Smith, B. Forget, and A. R. Siegel. "The Random Ray Method for neutral particle transport." *Journal of Computational Physics*, **volume 342**, pp. 229–252 (2017).

[3] J. R. Tramm. *Development of The Random Ray Method of Neutral Particle Transport for High-Fidelity Nuclear Reactor Simulation*. PhD thesis, Massachusetts Institute of Technology (2018).

[4] M. Kraus, P. Cosgrove, and E. Shwageraus. "Development and Benchmarking of Two Distinct Time-Dependent Random Ray Methods." In *International Conference on Physics of Reactors (PHYSOR 2024)*. San Francisco, CA (2024).

[5] M. Kraus. "Time-Dependent Variants of the Random Ray Method for Transient Simulations." *Nuclear Science and Engineering* (2024, submitted, not yet published).

[6] T. Downar, Y. Xu, V. Seker, and N. Hudson. "Theory Manual - PARCS v3.0 U.S. NRC Core Neutronics Simulator." Technical report, Department of Nuclear Engineering and Radiological Sciences, Ann Arbor, MI (2009).

[7] M. A. Kowalski, P. Cosgrove, J. Broman, and E. Shwageraus. "SCONE: A Student-Oriented Modifiable Monte Carlo Particle Transport Framework." *Journal of Nuclear Engineering*, **volume 2**, pp. 57–64 (2021).

[8] Q. Shen, Y. Wang, D. Jabaay, B. Kochunas, and T. Downar. "Transient analysis of C5G7-TD benchmark with MPACT." *Annals of Nuclear Energy*, **volume 125**, pp. 107–120 (2019).