Annals of Nuclear Energy 201 (2024) 110450

Contents lists available at ScienceDirect

Annals of Nuclear Energy

journal homepage: www.elsevier.com/locate/anucene

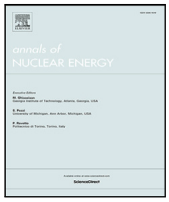

A memory-efficient neutron noise algorithm for reactor physics

Paul Cosgrove *, Maximilian Kraus, Valeria Raffuzzi

Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, United Kingdom

ARTICLE INFO

Keywords:
Neutron noise
The random ray method
Neutron transport
Method of characteristics

ABSTRACT

The neutron noise equation contains a fixed source term where the neutron angular flux is present. Deterministic noise solvers have tended to solve for and store this angular flux term in a preprocessing step. This can be a substantial memory burden as the angular flux is a function of space, energy, and angle. This can limit these solvers to smaller problems and can limit the obtainable angular resolution. This paper proposes solving the neutron noise equation and the standard eigenvalue neutron transport equation simultaneously, reducing the memory footprint of a noise solve while making frequency-domain method of characteristics solvers practical. This approach is demonstrated using the random ray method but is applicable to any sweep-based methods. Using the random ray method also has the advantage of exact geometry representation, no ray effects, and cheaper transport sweeps, all of which have been problematic. These advantages are illustrated on several noise problems.

1. Introduction

Recently, the analysis of neutron noise in nuclear reactors has become popular as an approach to reactor core diagnostics and monitoring (Pázsit and Demazière, 2010). While the field of neutronics is most often concerned with the time-independent, average neutron flux (estimated by the k-eigenvalue neutron transport equation), the flux in an operating reactor is not truly static due to, for example, flow-induced vibrations in fuel assemblies and control rods. Detectors in reactor cores produce time-varying signals which, assuming the flux is stationary, may be decomposed into a static average component and a time-dependent noise component. In principle, this noise component may be useful in identifying concerning behaviour during reactor operation in an unobtrusive manner. In general cases, this can be done provided a transfer function between possible noise sources and the detector reading can be obtained. While neutron diffusion-based methods have proven popular and successful (Demazière, 2004, 2011; Mylonakis et al., 2021), much recent work has focused on applying neutron transport methods to this task (Yamamoto, 2013; Rouchon, 2016; Rouchon et al., 2017; Gammicchia et al., 2020; Yi et al., 2021; Belanger et al., 2023; Kang and Joo, 2021; Carreño et al., 2022). However, some of these transport approaches to neutron noise (namely those based on deterministic methods) can entail a significant memory burden or be impractical to extend to standard algorithms, making them challenging to apply to reactor-scale problems which often require substantial discretisation to obtain acceptable results.

This paper proposes a means by which this memory burden can be eliminated at the cost of additional computational expense during a neutron noise solve. The approach taken is demonstrated using a stochastic version of the method of characteristics (MoC), the random ray method (TRRM), but the logic applies directly to the deterministic MoC and (with some minor modifications) to discrete ordinates methods.

The paper begins by introducing the relevant concepts in neutron transport and noise and highlighting where traditional deterministic approaches suffer a significant memory burden. Subsequently, in the framework of MoC, equations are derived which eliminate this memory burden at the cost of solving both the noise equation and the transport equation simultaneously. Finally, the computational features of the derived method are illustrated on several numerical problems.

2. Neutron transport and noise

The general neutron transport equation applied to nuclear reactors is:

$$\frac{1}{v}\frac{\partial \psi}{\partial t} + \boldsymbol{\Omega}\cdot\nabla\psi + \Sigma_{\mathrm{t}}\psi = \int_0^\infty \mathrm{d}E' \int_{4\pi} \mathrm{d}\Omega'\ \Sigma_{\mathrm{s}}(E'\to E, \boldsymbol{\Omega}'\to\boldsymbol{\Omega})\psi + \frac{\chi_{\mathrm{p}}}{4\pi}(1-\beta)\int_0^\infty \mathrm{d}E'\ \bar{\nu}\Sigma_{\mathrm{f}}\phi + \frac{1}{4\pi}\sum_j \chi_j \lambda_j C_j\ , \tag{1}$$

where ψ is the angular neutron flux (which varies with position in space, $\mathbf{r}$, neutron energy, E, flight direction, $\boldsymbol{\Omega}$, and time, t), v is the

* Corresponding author.
E-mail address: pmc55@cam.ac.uk (P. Cosgrove).

https://doi.org/10.1016/j.anucene.2024.110450
Received 9 January 2024; Received in revised form 19 February 2024; Accepted 21 February 2024
Available online 23 February 2024

neutron speed, Σ_t is the total macroscopic cross section (a function of position, neutron energy, and, generally, time), $\Sigma_s(E' \to E, \mathbf{\Omega}' \to \mathbf{\Omega})$ is the differential neutron scattering kernel which redistributes neutrons from their incoming momentum to an outgoing momentum, χ_p is the prompt fission neutron energy spectrum (a distribution over energy which varies with position), β is the delayed neutron fraction, $\bar{\nu}$ is the average number of neutrons produced per fission (a function of position and energy), Σ_f is the fission macroscopic cross section, and ϕ is the scalar neutron flux, defined as:

$$\phi = \int_{4\pi} \mathrm{d}\Omega \; \psi \; . \tag{2}$$

The final term on the right-hand side is the production of delayed neutrons due to precursor fission products. These precursors tend to be lumped into several groups where, here, j is the group index, χ_j is the delayed fission neutron energy spectrum of group j, λ_j is the group's decay constant, and C_j is the concentration of precursors, governed by:

$$\frac{\partial C_j}{\partial t} = \beta_j \int_0^\infty \mathrm{d}E' \; \bar{\nu}\Sigma_f \phi - \lambda_j C_j \; , \tag{3}$$

where β_j is the fraction of fission neutrons produced by the decay of precursor group j and $\beta = \sum_j \beta_j$.

The above dynamic equations are rarely solved due to their computational expense. Reactors also operate at very close to steady state, such that reactor analysts tend to zero the time derivative and enforce a steady state by introducing the k eigenvalue to scale the fission neutron production. The k-eigenvalue form of the neutron transport equation is:

$$\begin{aligned}\left(\mathbf{\Omega}\cdot\nabla + \Sigma_t\right)\psi = &\int_0^\infty \mathrm{d}E' \int_{4\pi} \mathrm{d}\Omega' \; \Sigma_s(E' \to E, \mathbf{\Omega}' \to \mathbf{\Omega})\psi \\ &+ \frac{\chi_{\mathrm{eff}}}{4\pi k} \int_0^\infty \mathrm{d}E' \; \bar{\nu}\Sigma_f \phi \; ,\end{aligned} \tag{4}$$

where χ_{eff} is the effective fission spectrum, defined as:

$$\chi_{\mathrm{eff}} = \chi_p(1-\beta) + \sum_j \chi_j \beta_j \; . \tag{5}$$

This form of the transport equation is most commonly used to model nuclear reactors, even when they are not truly steady state. However, provided the reactor is close to critical ($k \approx 1$), this approach is acceptable.

The neutron noise problem asserts that there is a time-varying neutron flux composed of a static flux governed by Eq. (4) and a small time-varying perturbation:

$$\psi(\mathbf{r}, E, \mathbf{\Omega}, t) = \psi_s(\mathbf{r}, E, \mathbf{\Omega}) + \delta\psi(\mathbf{r}, E, \mathbf{\Omega}, t) \; , \tag{6}$$

where ψ_s is the static component of the neutron flux (governed by Eq. (4)) and $\delta\psi$ is the noise component of the neutron flux. Integrating over angle gives an equivalent scalar flux noise, $\delta\phi(\mathbf{r}, E, t)$. This noise is induced by small, stationary fluctuations in the macroscopic cross sections. The equation governing the noise component is obtained by inserting the definition of flux given by Eq. (6) into Eqs. (1) and (3) (with the addition of a noise source term) and subtracting the static terms given by Eq. (4) while neglecting second-order noise terms (Pázsit, 1984). This results in a time-dependent equation for $\delta\psi$ which is similarly difficult to solve as Eqs. (1) and (3). However, given the noise is asserted to be stationary, this equation can be Fourier transformed to solve for different frequency components of the neutron noise, $\delta\psi(\mathbf{r}, E, \mathbf{\Omega}, \omega)$, which are associated with the noise source. Doing so results in the neutron noise equation:

$$\begin{aligned}\left(\mathbf{\Omega}\cdot\nabla + \Sigma_t + \frac{i\omega}{v}\right)\delta\psi = &\int_0^\infty \mathrm{d}E' \int_{4\pi} \mathrm{d}\Omega' \, \Sigma_s(E' \to E, \mathbf{\Omega}' \to \mathbf{\Omega})\delta\psi \\ &+ \frac{1}{4\pi k}\left[\chi_p(1-\beta) + \sum_j \chi_j \frac{\lambda_j\beta_j}{i\omega + \lambda_j}\right] \int_0^\infty \mathrm{d}E' \bar{\nu}\Sigma_f \delta\phi + S \; ,\end{aligned} \tag{7}$$

where S is the noise source, a function of space, energy, angle, and frequency. This is generally written as:

$$\begin{aligned}S = &-\delta\Sigma_t\psi + \int_0^\infty \mathrm{d}E' \int_{4\pi} \mathrm{d}\Omega' \delta\Sigma_s(E' \to E, \mathbf{\Omega}' \to \mathbf{\Omega})\psi \\ &+ \frac{1}{4\pi k}\left[\chi_p(1-\beta) + \sum_j \chi_j \frac{\lambda_j\beta_j}{i\omega + \lambda_j}\right] \int_0^\infty \mathrm{d}E' \delta\left(\bar{\nu}\Sigma_f\right)\phi \; ,\end{aligned} \tag{8}$$

where the cross section terms prefaced by δ are the frequency-dependent perturbations to each macroscopic cross section.

Eqs. (7) and (8) are solved by the same methods applied to solving the standard k-eigenvalue transport equation (Eq. (4)): Monte Carlo (Yamamoto, 2013; Rouchon et al., 2017; Belanger et al., 2023) and various deterministic methods (Rouchon, 2016; Yi et al., 2021; Carreño et al., 2022). This paper is concerned with deterministic approaches to solving the noise equation, i.e., approaches which discretise Eqs. (7) and (8). This discretisation is typically performed over space (creating a mesh), energy (using the multi-group approximation) and angle (using either the discrete ordinates approximation or an expansion in spherical harmonics).

In neutronics, the quantities of interest are angle-integrated reaction rates or ratios of reaction rates: reactor designers care about the rate at which reactions occur in a given volume but not whether they are caused by neutrons travelling left or right. As such, for static calculations, one very rarely needs to explicitly solve the transport equation to obtain ψ (a function of space, energy, and angle), instead solving for ϕ (a function only of space and energy) through what is known as a transport sweep.

The transport sweep is a solution algorithm used by both first-order discrete ordinates methods and the method of characteristics where the neutron angular flux is not stored but is used to accumulate estimates of the scalar flux sequentially in space along the neutron's flight direction. This is equivalent to inverting the operator on the left-hand side of Eq. (4) and integrating over all angles. One of the advantages of this approach is significant memory savings: one need only store the scalar flux vector and not the angular flux vector. This is because the right-hand side of Eq. (4) can be written as a function of only the scalar neutron flux (if isotropic scattering is assumed) or of the scalar flux and several of its higher spherical harmonic moments (if anisotropic scattering is allowed). Given that, in reactor calculations, one must discretise the angular variables into $\mathcal{O}(100)$ directions to obtain reasonably accurate results (Knott and Yamamoto, 2010), the scalar flux vector is much more compact than the angular flux vector.

Solving the neutron noise equations can be approached like the solution of a fixed-source problem where S is the source term and all static components are precomputed with a prior eigenvalue problem solution. Unfortunately, unlike the right-hand side of Eq. (4), the first term in Eq. (8) explicitly depends on the angular flux. This implies that one must store the angular flux vector from the static transport solution (at least in regions which are perturbed) in order to obtain a noise solution. Even if only small regions are perturbed, this can be a substantial memory burden which can limit the achievable angular resolution of both the static and noise solutions. This can also limit the achievable size of noise problems which are analysed. It is this problem which the present paper aims to address by demonstrating that it is possible to solve the transport and noise equations simultaneously.

2.1. The Random Ray Method of Characteristics (TRRM)

Before describing the new solution method, it is necessary to provide details of the neutron transport (and noise) algorithm that will be used for the remainder of the paper. This paper will use TRRM, a stochastic version of MoC. As mentioned previously, the approach to solve the noise equation can equally be applied to any sweep-based deterministic transport method (directly to deterministic MoC and with some small differences in algebra for discrete ordinates). It should be mentioned that, perhaps partly due to the memory limitations

highlighted above, MoC has not yet been used to directly solve the neutron noise equation, although it has been used in a time-dependent mode to deduce the neutron noise behaviour (Gammicchia et al., 2020; Kang and Joo, 2021). This is not to be confused with work using the Method of Short Characteristics (Rouchon, 2016; Vinai et al., 2021). In the standard method of characteristics featured presently (or the method of long characteristics), tracks are traced across the geometry and the angular flux balance is performed across these tracks – this algorithm is commonly implemented in lattice physics codes due to its geometric flexibility. The method of short characteristics is essentially a discrete ordinates method, but using a characteristic closure across a given mesh while assuming a flux shape on the surfaces of a given cell. This approach tends to have more limited geometric capability. The previous difficulty with the method of long characteristics in frequency-domain noise calculations is the necessity to store angular fluxes along each track (and there may be many tracks in each angle in each cell), while the method of short characteristics should only have to store the angular flux in each angle in each cell (which may still be a challenge).

TRRM has been well described across several publications already (Tramm et al., 2017), but the main equations will be briefly revisited here. Much of the discussion to follow is common to both MoC and TRRM, but where differences occur they will be highlighted.

First, we assume a multi-group discretisation of Eq. (4), i.e., we define:

$$\psi_g = \int_{E_g}^{E_{g-1}} \mathrm{d}E \psi \ , \tag{9}$$

and

$$\phi_g = \int_{E_g}^{E_{g-1}} \mathrm{d}E \phi \ , \tag{10}$$

where E_{g-1} and E_g are the upper and lower bounds, respectively, of the given energy group. As is common in reactor physics, we assert that in an energy group the cross sections and fission spectra are constant and chosen such that they preserve the energy-integrated reaction rates in Eq. (4). This gives:

$$\begin{aligned}\left(\boldsymbol{\Omega}\cdot\nabla + \Sigma_{\mathrm{t},g}\right)\psi_g = &\sum_{g'}\int_{4\pi}\mathrm{d}\Omega'\ \Sigma_{\mathrm{s},g'\to g}(\boldsymbol{\Omega}'\to\boldsymbol{\Omega})\psi_{g'} \\ &+ \frac{\chi_{\mathrm{eff},g}}{4\pi k}\sum_{g'}\bar{\nu}\Sigma_{\mathrm{f},g'}\phi_{g'} \ .\end{aligned} \tag{11}$$

Although not strictly necessary, for simplicity we assume isotropic scattering (along with the appropriate corrections to $\Sigma_{\mathrm{t},g}$ and $\Sigma_{\mathrm{s},g\to g}$ that result), such that the source term is wholly isotropic:

$$\left(\boldsymbol{\Omega}\cdot\nabla + \Sigma_{\mathrm{t},g}\right)\psi_g = \frac{1}{4\pi}\sum_{g'}\Sigma_{\mathrm{s},g'\to g}\phi_{g'} + \frac{\chi_{\mathrm{eff},g}}{4\pi k}\sum_{g'}\bar{\nu}\Sigma_{\mathrm{f},g'}\phi_{g'} = q_g \ , \tag{12}$$

where q_g is a compact form for the isotropic source. The method of characteristics transforms this equation by considering the distance, s, from a particular reference position in a particular neutron flight direction. In such a coordinate system, the characteristic form of the transport equation is:

$$\frac{\mathrm{d}}{\mathrm{d}s}\psi_g + \Sigma_{\mathrm{t},g}\psi_g = q_g \ . \tag{13}$$

This equation describes how ψ_g varies along the direction of the given characteristic of the transport equation. Following this one also tends to apply the 'Flat Source Approximation', where one asserts there are Flat Source Regions (FSRs) over which both q_g and all cross sections are spatially constant (these regions correspond to individual mesh elements). Given these approximations, the solution of Eq. (13) can be analytically obtained along a characteristic in an FSR as:

$$\psi_g(s) = \psi_g(0)\,\mathrm{e}^{-\Sigma_{\mathrm{t},g}s} + \frac{q_g}{\Sigma_{\mathrm{t},g}}\left(1-\mathrm{e}^{-\Sigma_{\mathrm{t},g}s}\right) \ . \tag{14}$$

Here $\psi_g(0)$ is the angular flux of the characteristic as it enters the given FSR. These characteristics are generated by ray tracing across the geometry of interest in many starting positions and directions. If this equation is solved along sufficiently many distinct characteristics (adequately covering the spatial and angular phase space of the problem) the solutions may be combined to give the scalar flux solution across the geometry. This is because the FSR volume-averaged scalar flux is defined as:

$$\phi_g = \frac{4\pi\sum_c w_c \int_0^{L_c}\mathrm{d}s\ \psi_g(s)}{V} \ , \tag{15}$$

where c is an index for a given characteristic track, w_c is the quadrature weight of a given track, L_c is the length of that track, and V is the volume of the given FSR. In deterministic MoC, quadrature weight is typically non-uniform, but in TRRM it is uniform and equal to the total volume of the problem divided by the total track length traversed during a transport sweep — given that this is an eigenvalue calculation, this can be set to 1, as it is uniform and normalisation will otherwise fix the amplitude of the flux. Combining Eqs. (14) and (15), one obtains:

$$\phi_g = \frac{4\pi q_g}{\Sigma_{\mathrm{t},g}} + \frac{4\pi}{\Sigma_{\mathrm{t},g}V}\sum_c \Delta\psi_g \ , \tag{16}$$

where $\Delta\psi_g$ is defined as:

$$\Delta\psi_g = \left(\psi_g(0) - \frac{q_g}{\Sigma_{\mathrm{t},g}}\right)\left(1-\mathrm{e}^{-\Sigma_{\mathrm{t},g}L_c}\right) = \psi_g(0) - \psi_g(L) \ . \tag{17}$$

These equations are extremely computationally convenient as they imply that, given a neutron source and an initial angular flux along a characteristic, one can proceed along each characteristic track, calculate $\Delta\psi_g$, increment the local scalar flux (ϕ_g += $4\pi\Delta\psi_g$), decrement the angular flux ($\psi_g(0)$ −= $\Delta\psi_g$) and proceed to the next FSR until all tracks have been traversed. The estimate of the scalar flux is then scaled by $1/\Sigma_{\mathrm{t},g}V$ and incremented by $4\pi q_g/\Sigma_{\mathrm{t},g}$ and a scalar flux solution will have been obtained.

In eigenvalue calculations, this must also be iterated upon to obtain the eigenvalue and eigenvector, typically by power iteration. Given an initial guess, the eigenvalue can be calculated for iteration $n+1$ as:

$$k^{(n+1)} = k^{(n)}\frac{\sum_r V_r\sum_g \bar{\nu}\Sigma_{\mathrm{f},g,r}\phi_{g,r}^{(n+1)}}{\sum_r V_r\sum_g \bar{\nu}\Sigma_{\mathrm{f},g,r}\phi_{g,r}^{(n)}} \ , \tag{18}$$

where r is the index of a given region. This is simply the ratio of volume-integrated fission rates in the present iteration to that in the previous iteration, multiplied by the previous estimate of the eigenvalue. This can then be used to update the source which, for the next iteration in an eigenvalue calculation, would be defined as:

$$q_g^{(n+1)} = \frac{1}{4\pi}\left(\sum_{g'}\Sigma_{\mathrm{s},g'\to g}\phi_{g'}^{(n+1)} + \frac{\chi_{\mathrm{eff},g}}{k^{(n+1)}}\sum_{g'}\bar{\nu}\Sigma_{\mathrm{f},g'}\phi_{g'}^{(n+1)}\right) \ . \tag{19}$$

There are several differences in TRRM compared to deterministic MoC which are worth briefly revisiting. Deterministic MoC tends to precompute the tracks distances which are followed and traverses the same tracks each iteration; these are started at the problem boundaries where boundary conditions provide initial angular fluxes (or one can iterate to obtain initial angular fluxes). TRRM performs ray tracing on-the-fly and does not traverse the same track twice. Rays are sampled uniformly in space and angle throughout the geometry and so the starting conditions are not known. However, they can be reasonably estimated using a 'dead length' where the angular flux of the ray is updated without accumulating to local estimates of the scalar flux. Once this dead length has been traversed, the scalar flux is accumulated towards as normal. In deterministic MoC, the fidelity of a calculation is determined by how finely to discretise angle and how fine is the spacing between different tracks. In TRRM, there is no angular discretisation error (given uniform sampling in angle), but the final error has a stochastic component which can be reduced by increasing the number of rays and the distance over which they are followed. The standard deterministic errors from both spatial and energy discretisation are also present.

A final difference which has some importance for the present work is determining when the solution has converged or, more correctly for TRRM, reached stationarity. For deterministic MoC, convergence can be measured in the standard manner, e.g., by checking that a relative L_1 norm on the difference between successive scalar flux vectors is less than some tolerance threshold. This cannot reliably be done with TRRM as stochastic noise between iterations tends to be significantly larger than conventional choices of this residual value. One successful approach to determining convergence has been using the Shannon entropy of the fission source. Tramm et al. (2018) used a moving window of entropy values across iterations, with a width of 200; the mean values of the entropy of the first and second 100 values are compared and, if they fall within one standard deviation, it is accepted that the solution has reached stationarity. However, once these inactive iterations have concluded in TRRM, active iterations must begin in order to accumulate flux estimates (as in Monte Carlo).

A more detailed discussion of TRRM can be found elsewhere (Tramm et al., 2017, 2018; Tramm, 2018; Tramm et al., 2020).

2.2. Neutron noise

The multi-group, isotropic version of Eq. (7) is:

$$\left(\boldsymbol{\Omega}\cdot\nabla+\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)\delta\psi_g=\frac{1}{4\pi}\sum_{g'}\Sigma_{\mathrm{s},g'\to g}\delta\phi_{g'}+\frac{1}{4\pi k}\left[\chi_{\mathrm{p},g}(1-\beta)+\sum_j\chi_{j,g}\frac{\lambda_j\beta_j}{i\omega+\lambda_j}\right]\sum_{g'}\bar{\nu}\Sigma_{\mathrm{f},g'}\delta\phi_{g'}+S_g\,, \tag{20}$$

where ω is the noise frequency, $\delta\psi_g$ is the angular flux noise component in group g, $\delta\phi_g$ is the scalar flux noise component in group g, and S_g is the isotropic neutron noise source in group g. The noise source is given by:

$$S_g=-\delta\Sigma_{\mathrm{t},g}\psi_g+\frac{1}{4\pi}\sum_{g'}\delta\Sigma_{\mathrm{s},g'\to g}\phi_{g'}+\frac{1}{4\pi k}\left[\chi_{\mathrm{p},g}(1-\beta)+\sum_j\chi_{j,g}\frac{\lambda_j\beta_j}{i\omega+\lambda_j}\right]\sum_{g'}\delta\left(\bar{\nu}\Sigma_{\mathrm{f},g'}\right)\phi_{g'}\,, \tag{21}$$

where δ terms denote perturbations in the cross section. The isotropic assumption is not necessary as flux and flux noise spherical harmonic moments could also be calculated, though isotropy significantly simplifies the computation.

As discussed previously, the angular flux remains in the total cross section perturbation term in Eq. (21); this could result in an impractically large memory footprint even for relatively small problems. This will be handled by developing an analogous MoC algorithm for the noise problem and demonstrating how this can be solved simultaneously with the static problem.

The MoC version of Eq. (20) is:

$$\left(\frac{\mathrm{d}}{\mathrm{d}s}+\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)\delta\psi_g=\delta Q_g\,, \tag{22}$$

where s is the distance along the characteristic and δQ_g is the total noise source in a given direction, equal to the right-hand side of Eq. (20). The novelty in this work is in how the usual characteristic FSR approximation is applied. Here it is applied to all terms except for the angular flux-dependent term, splitting the source into a flat component (a function of the volume-averaged scalar flux and scalar flux noise) and a term which varies in space and angle:

$$\delta Q_g=\delta q_g-\delta\Sigma_{\mathrm{t},g}\psi_g\,, \tag{23}$$

where δq_g contains all flat source terms, i.e.,

$$\begin{aligned}\delta q_g=&\frac{1}{4\pi}\sum_{g'}\Sigma_{\mathrm{s},g'\to g}\delta\phi_{g'}+\frac{1}{4\pi k}\left[\chi_{\mathrm{p},g}(1-\beta)+\sum_j\chi_{j,g}\frac{\lambda_j\beta_j}{i\omega+\lambda_j}\right]\\&\times\sum_{g'}\bar{\nu}\Sigma_{\mathrm{f},g'}\delta\phi_{g'}\\&+\frac{1}{4\pi}\sum_{g'}\delta\Sigma_{\mathrm{s},g'\to g}\phi_{g'}+\frac{1}{4\pi k}\left[\chi_{\mathrm{p},g}(1-\beta)+\sum_j\chi_{j,g}\frac{\lambda_j\beta_j}{i\omega+\lambda_j}\right]\\&\times\sum_{g'}\delta\left(\bar{\nu}\Sigma_{\mathrm{f},g'}\right)\phi_{g'}\,.\end{aligned} \tag{24}$$

From standard MoC, we know how ψ_g varies across a FSR from Eq. (14):

$$\psi_g(s)=\psi_g(0)\mathrm{e}^{-\Sigma_{\mathrm{t},g}s}+\frac{q_g}{\Sigma_{\mathrm{t},g}}\left(1-\mathrm{e}^{-\Sigma_{\mathrm{t},g}s}\right)\,.$$

This equation can be inserted into Eq. (22) and integrated to give:

$$\begin{aligned}\delta\psi_g(s)=&\delta\psi_g(0)\mathrm{e}^{-\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)s}+\frac{\delta q_g-\frac{\delta\Sigma_{\mathrm{t},g}}{\Sigma_{\mathrm{t},g}}q_g}{\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}}\left(1-\mathrm{e}^{-\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)s}\right)\\&+\frac{iv_g}{\omega}\delta\Sigma_{\mathrm{t},g}\left(\psi_g(0)-\frac{q_g}{\Sigma_{\mathrm{t},g}}\right)\left(\left(1-\mathrm{e}^{-\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)s}\right)-\left(1-\mathrm{e}^{-\Sigma_{\mathrm{t},g}s}\right)\right)\,.\end{aligned} \tag{25}$$

The volume-average scalar flux noise in a mesh region is defined as a sum over all tracks that cross that region:

$$\delta\phi_g=\frac{4\pi\sum_c\int_0^{L_c}\mathrm{d}s\;\delta\psi_g(s)}{V}\,. \tag{26}$$

The integral of $\delta\psi_g$ over a given track is:

$$\begin{aligned}\int_0^{L_c}\mathrm{d}s\;\delta\psi_g(s)=&L_c\frac{\delta q_g-\frac{\delta\Sigma_{\mathrm{t},g}}{\Sigma_{\mathrm{t},g}}q_g}{\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}}\\&+\left(\delta\psi_g(0)-\frac{\delta q_g-\frac{\delta\Sigma_{\mathrm{t},g}}{\Sigma_{\mathrm{t},g}}q_g}{\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}}\right)\left(\frac{1-\mathrm{e}^{-\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)L_c}}{\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)}\right)\\&+\frac{iv_g}{\omega}\delta\Sigma_{\mathrm{t},g}\left(\psi_g(0)-\frac{q_g}{\Sigma_{\mathrm{t},g}}\right)\\&\times\left[\left(\frac{1-\mathrm{e}^{-\Sigma_{\mathrm{t},g}L_c}}{\Sigma_{\mathrm{t},g}}\right)-\left(\frac{1-\mathrm{e}^{-\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)L_c}}{\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)}\right)\right]\,.\end{aligned} \tag{27}$$

This equation can be made more compact by defining an additional delta term (analogously with that in Eq. (17)):

$$\begin{aligned}\Delta\delta\psi_g=&\left(\delta\psi_g(0)-\frac{\delta q_g-\frac{\delta\Sigma_{\mathrm{t},g}}{\Sigma_{\mathrm{t},g}}q_g}{\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}}-\frac{iv_g}{\omega}\delta\Sigma_{\mathrm{t},g}\left(\psi_g(0)-\frac{q_g}{\Sigma_{\mathrm{t},g}}\right)\right)\\&\times\left(1-\mathrm{e}^{-\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)L_c}\right)\,.\end{aligned} \tag{28}$$

Hence:

$$\int_0^{L_c}\mathrm{d}s\;\delta\psi_g(s)=\frac{\delta q_g-\frac{\delta\Sigma_{\mathrm{t},g}}{\Sigma_{\mathrm{t},g}}q_g}{\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}}L_c+\frac{\Delta\delta\psi_g}{\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)}+\frac{\frac{iv_g}{\omega}\delta\Sigma_{\mathrm{t}}\Delta\psi_g}{\Sigma_{\mathrm{t},g}}\,, \tag{29}$$

and so:

$$\delta\phi_g=4\pi\frac{\delta q_g-\frac{\delta\Sigma_{\mathrm{t},g}}{\Sigma_{\mathrm{t},g}}q_g}{\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}}+4\pi\sum_c\left(\frac{\Delta\delta\psi_g}{\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)L_c}+\frac{\frac{iv_g}{\omega}\delta\Sigma_{\mathrm{t}}\Delta\psi_g}{\Sigma_{\mathrm{t},g}L_c}\right)\,. \tag{30}$$

These equations are strikingly similar to those used for the standard static transport solve and imply that one can track both ψ_g and $\delta\psi_g$ simultaneously, i.e., one can solve both the static transport equation and the noise equation at the same time. While a given simulation becomes more computationally expensive, this may be cheaper than doing two separate simulations and results in a memory reduction as there is no longer a need to store and access the angular flux in all groups and angles in all perturbed regions. The most involved differences in these equations from the standard MoC equations are requiring complex arithmetic and the evaluation of a second exponential term, $1-\mathrm{e}^{-\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)L_c}$.

As an aside, for TRRM, storing the angular flux along tracks is essentially inconceivable as it would require following the exact same track laydown for the noise solve as for the transport solve, which differs on every iteration. This approach would be slightly less onerous for deterministic MoC (where the track laydown does not vary between iterations) but would still incur a dramatic memory penalty.

To be clear, a combined eigenvalue/noise algorithm would begin with assumed sources (e.g., resulting from the initial guesses $\phi_g = 1$ and $\delta\phi_g = 0$ in all meshes). Both scalar flux terms would then be set to zero and the transport sweep would commence. As normal, $\Delta\psi_g$ would be incremented upon ϕ_g:

$$\phi_g \mathrel{+}= 4\pi\Delta\psi_g \,, \tag{31}$$

while ψ_g would be decremented during the sweep:

$$\psi_g \mathrel{-}= \Delta\psi_g \,. \tag{32}$$

In addition to this, $\delta\phi_g$ would be incremented:

$$\delta\phi_g \mathrel{+}= 4\pi\Delta\delta\psi_g \,, \tag{33}$$

while $\delta\psi_g$ would be decremented as:

$$\delta\psi_g \mathrel{-}= \left(\Delta\delta\psi_g + \frac{iv_g}{\omega}\delta\Sigma_{\mathrm{t},g}\Delta\psi_g\right) \,. \tag{34}$$

On the conclusion of the sweep these would be normalised as:

$$\phi_g = \frac{\phi_g}{\Sigma_{\mathrm{t},g}L} \,, \tag{35}$$

and

$$\delta\phi_g = \frac{\delta\phi_g}{\left(\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}\right)L} \,, \tag{36}$$

where L is the total length traversed by all tracks during the transport sweep. The noise term would be further incremented by:

$$\delta\phi_g \mathrel{+}= \left(4\pi\frac{\delta q_g - \frac{\delta\Sigma_{\mathrm{t},g}}{\Sigma_{\mathrm{t},g}}q_g}{\Sigma_{\mathrm{t},g}+\frac{i\omega}{v_g}} + \frac{iv_g}{\omega}\delta\Sigma_{\mathrm{t},g}\phi_g\right) \,, \tag{37}$$

after which the scalar flux would be incremented as:

$$\phi_g \mathrel{+}= \frac{4\pi q_g}{\Sigma_{\mathrm{t},g}} \,. \tag{38}$$

Both ϕ_g and $\delta\phi_g$ can then be used to calculate the corresponding source terms (and k eigenvalue) and the process can be repeated until some measure of convergence is reached.

In this work it was found that the Shannon entropy metric (Tramm et al., 2018) is not a robust indicator of convergence in stochastic neutron noise problems. This metric was applied to the noise amplitude and indicated convergence much earlier than it should, resulting in noticeable disagreement from other noise solutions. We believe this is because the noise amplitude converges more rapidly than the phase. The solution we found is instead to apply a similar metric to the noise phase, which does evolve visibly even after many iterations. Hence, to determine convergence, each iteration we calculated the quantity:

$$\theta = \sum_r V_r \sum_g \arctan\left(\frac{\mathfrak{I}(\delta\phi_{r,g})}{\mathfrak{R}(\delta\phi_{r,g})}\right) \,, \tag{39}$$

where the sum over r is a sum over mesh cells, V_r is the volume of a given cell, the sum over g is over energy groups, and $\mathfrak{I}$ represents the imaginary part of a complex number and $\mathfrak{R}$ represents the real part. This does not have any intrinsic meaning other than being a volume weighted sum of all estimates of the noise phase in all regions and energy groups. A window size of 400 iterations was used. The average and standard deviation of θ was computed for both of the 200 iteration intervals. If the averages of each interval fell within one standard deviation of each other it was decided that convergence was reached. This metric tended to reproduce a number of inactive iterations comparable to those reported in literature. It is probable that even more reliable metrics for convergence may exist, but this seemed successful in producing agreement with previous noise results.

Another change made from previous implementations of TRRM is the use of double precision arithmetic during transport, rather than single precision. This was done because some of the test problems to follow were sensitive to numeric precision when, e.g., the imaginary component of noise was very small. Although the noise amplitude would be well predicted, the phase would be poorly predicted, sometimes resulting in premature termination of the inactive iterations due to the use of the phase metric described above. For the problems tested, the effects of this are an increase in runtime and memory, but these are relatively minor due to the small number of energy groups and relatively low arithmetic intensity of the test problems. For problems with $\mathcal{O}(10)$s of groups, the resulting slowdown and memory impact would be more noticeable.

3. Numerical investigations

The described algorithm has been implemented in SCONE (Kowalski et al., 2021). SCONE is a Monte Carlo code, but recent work has also developed an eigenvalue TRRM solver capability (Cosgrove and Tramm, 2023). This was straightforward to modify to also solve for the neutron noise equation.

The algorithm was tested on two popular neutron noise benchmarks and compared against previously reported results in terms of runtime, accuracy, and memory consumption. All calculations were performed on 128 threads of an AMD Threadripper PRO 5995WX (2.7 GHz) node.

3.1. C5G7

C5G7 is a popular deterministic transport benchmark (Lewis et al., 2001) with a recent extension to noise problems (Vinai et al., 2019; Yi et al., 2021), where the kinetics data is provided by Hou et al. (2017). The most common version is 2D, featuring four assemblies (two UO_2, two MOX), surrounded on two sides by a water reflector and vacuum boundaries and with reflective boundaries on the other two sides (shown in Fig. 1(a)). There are several variants which introduce noise sources. The variant which will be shown here is a 'variable strength absorber-type' problem, described in Yi et al. (2021), where there is a capture cross section oscillation in all groups in one of the MOX pins (highlighted in Fig. 1). In Yi et al. (2021), the time domain oscillation takes the form:

$$\delta\Sigma_{\mathrm{c},g}(t) = \frac{0.05\Sigma^0_{\mathrm{c},g}}{\pi}\cos\omega_0 t \,, \tag{40}$$

where $\Sigma^0_{\mathrm{c},g}$ is the nominal capture cross section and ω_0 is the angular frequency equal to 2π rad/s. Fourier transforming this oscillation gives:

$$\delta\Sigma_{\mathrm{c},g}(\omega) = 0.05\Sigma^0_{\mathrm{c},g}\left(\delta(\omega-\omega_0)+\delta(\omega+\omega_0)\right) \,, \tag{41}$$

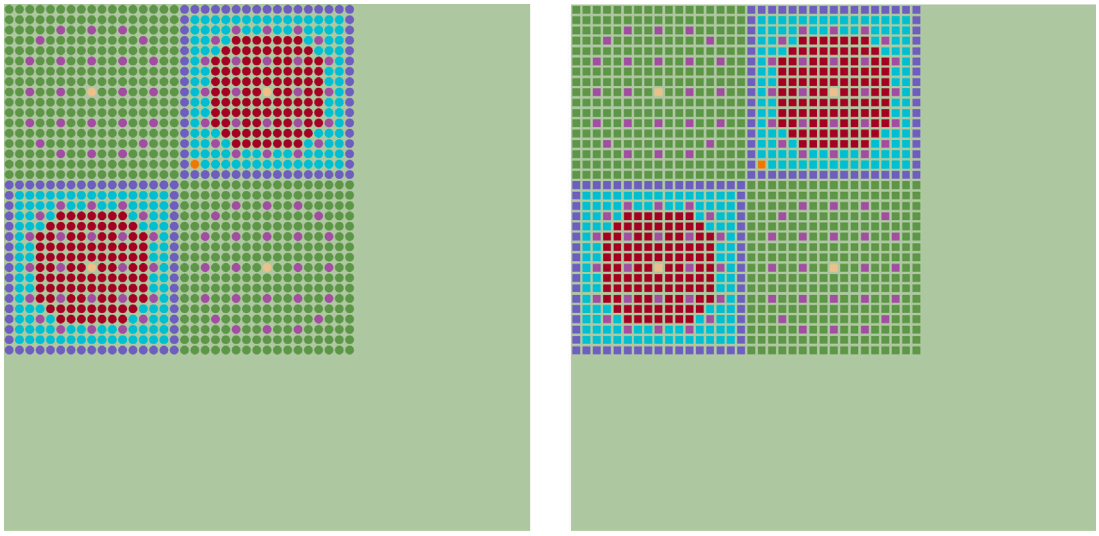

(a) Circular pin geometry.

(b) Square pin geometry.

Fig. 1. C5G7 geometry. The layout is conventional, but the orange pin in the top right MOX assembly is subject to a cross section oscillation.

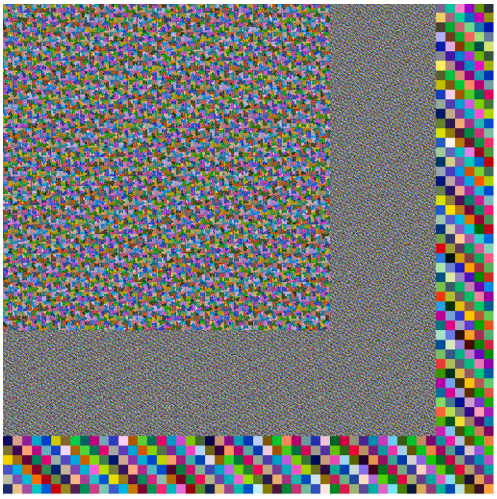

(a) Mesh elements in the circular C5G7 geometry.

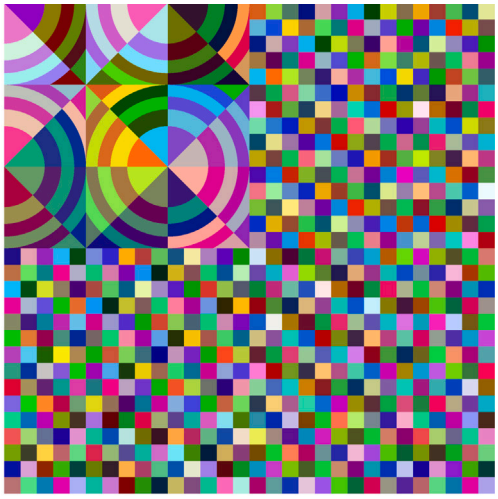

(b) Mesh elements in the circular C5G7 geometry in the vicinity of the corner fuel pin.

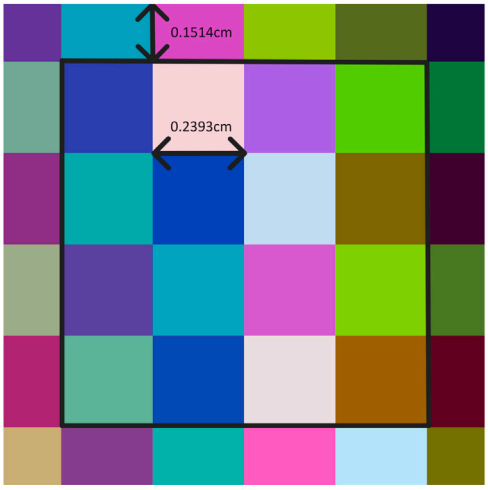

(c) Mesh elements in a single pin of the square C5G7 geometry. Inside the thick boundary is the fuelled region.

Fig. 2. Spatial discretisation of the 2D C5G7 problem.

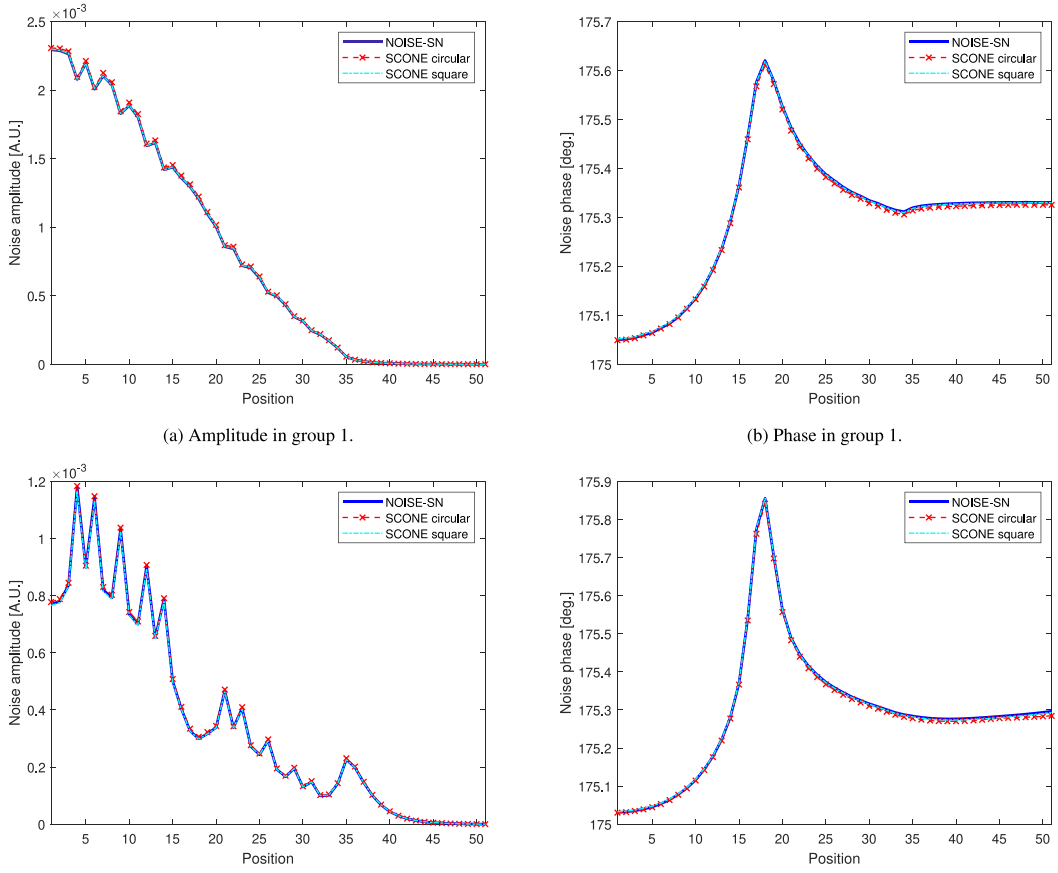

(a) Amplitude in group 1.

(b) Phase in group 1.

(c) Amplitude in group 7.

(d) Phase in group 7.

Fig. 3. Amplitudes and phases produced by SCONE and NOISE-SN for C5G7 with an oscillating capture cross section.

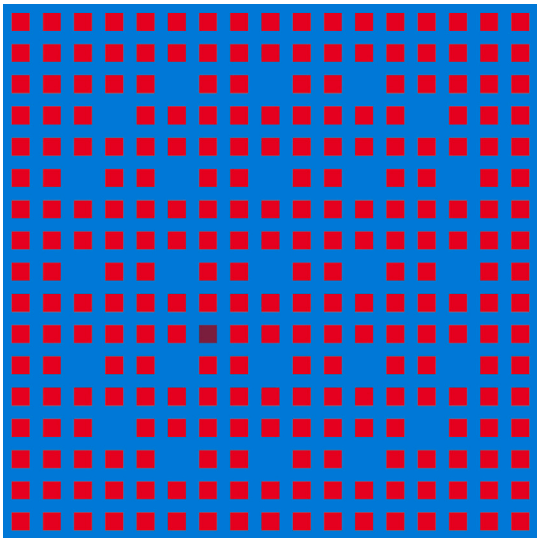

Fig. 4. UOX assembly geometry where blue is coolant, red is fuel, and bordeaux is the perturbed fuel pin.

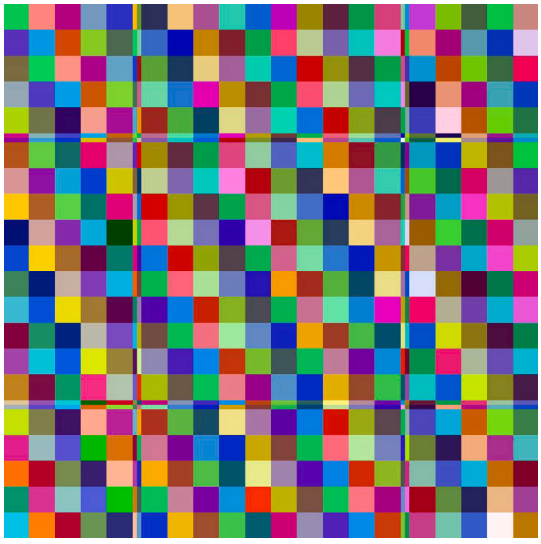

Fig. 5. UOX assembly pin cell discretisation showing two pin pitches from the centre of the central guide tube.

where $\delta(...)$ is a Dirac delta distribution. The noise solution of a $\delta(\omega+\omega_0)$ cross section source problem is simply the complex conjugate of the solution to a problem with a $\delta(\omega - \omega_0)$ cross section source. Hence, it is customary to only present the solution subject to a single δ source. For comparison, pin-integrated amplitudes and phases from Yi et al. (2021) were kindly provided along the top-left-to-bottom-right diagonal, including in the reflector. In the study by Yi et al. (2021), the fuel pins were square due to the constraints of using a discrete ordinates

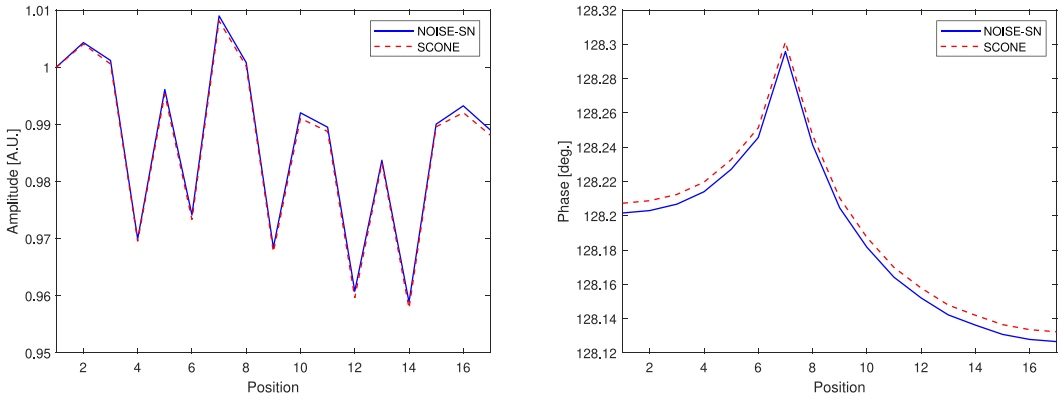

(a) Amplitude in group 1.

(b) Phase in group 1.

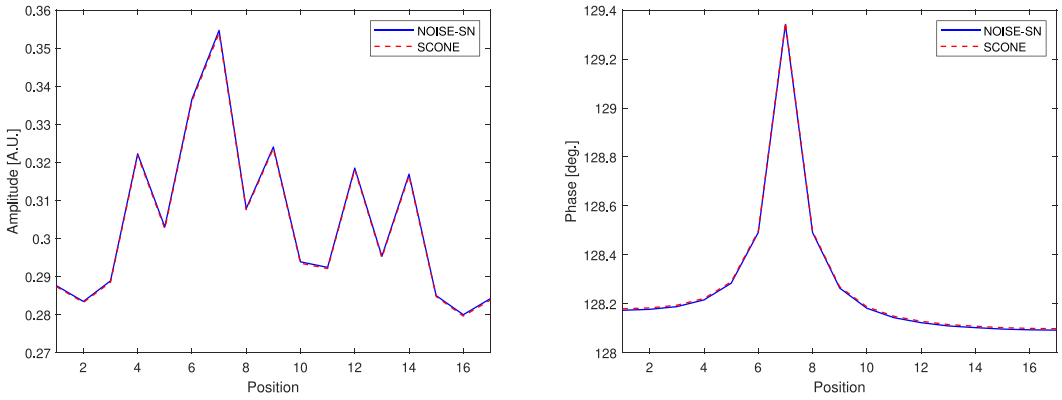

(c) Amplitude in group 2.

(d) Phase in group 2.

Fig. 6. Amplitudes and phases produced by SCONE and NOISE-SN for the UOX benchmark with a fixed noise source.

solver (shown in Fig. 1(b)). Here we analyse both a square case for consistency with Yi et al. (2021), as well as a circular case to produce results which are more consistent with the benchmark definition and exercise the advantages of MoC.

Based on previous investigations on the C5G7 benchmark (Cosgrove and Tramm, 2023), TRRM used 1750 rays per iteration, an active length of 200 cm, and 2000 active iterations. The number of inactive iterations was decided by the stochastic convergence method outlined previously. On experimenting, it was found that the noise solution is significantly more sensitive to the choice of dead length than pure neutron transport: for C5G7, a dead length of 20 cm has been successfully used previously, but resulted in noticeable errors for this problem. Hence, the dead length was increased to 40 cm, giving a total ray length of 240 cm. The need for a larger dead length in noise problems implies a stronger influence of distant regions upon each other than in transport.

The mesh discretisation for the circular pin case is conventional for C5G7 and shown in Figs. 2(a) and 2(b), with five radial and eight azimuthal regions in the pins, and a refined moderator discretisation adjacent to the assemblies, coarsening towards the periphery of the problem, and totalling about 142k mesh cells; for the eigenvalue benchmark this discretisation previously produced excellent agreement when used with TRRM (Cosgrove and Tramm, 2023). Notably, MoC allows for exact treatment of the C5G7 geometry, as opposed to needing to approximate curved elements as squares and rectangles, which should allow for an improvement in fidelity compared to previous work (Yi et al., 2021). For the square pin case, a similar discretisation to Yi et al. (2021) was used preserving fuel volume, but with an additional refinement. This is necessary because flat source MoC has only first-order spatial convergence, whereas the diamond difference scheme applied by Yi et al. (2021) has second-order spatial convergence. The discretisation of a single pin is shown in Fig. 2(c). This discretisation was also used at all points in the moderator and resulted in about 96k mesh elements.

Running the circular pin case with the given settings took 9 min, required 25,469 inactive iterations to reach convergence, and occupied 146 MB of memory. The more coarsely discretised square pin case took 7 min 15 s, required 26,885 inactive iterations, and occupied 118 MB of memory. The noise amplitude and phase in groups 1 and 7 of each simulation are compared with results produced by NOISE-SN (Yi et al., 2021) using a refined quadrature setting. The results from Yi et al. (2021) are pin-integrated along the diagonal line extending from the top-left of the C5G7 geometry to the bottom-right. This comparison is shown in Fig. 3. Despite the differences in geometry, there is excellent agreement between NOISE-SN and the circular pin TRRM solution, with relative differences smaller than 1% in all amplitudes and phases at all points. As one would expect, there is even better agreement between NOISE-SN and the more consistent square pin TRRM solution.

3.2. UOX assembly benchmark

The UOX assembly benchmark is a problem which has been used in comparing various neutron noise solvers recently (Vinai et al., 2021, 2023). The geometry is a 2D 17 × 17 assembly with given two-group cross sections/kinetic parameters and square fuel pins, with a cell consisting of fuel and water only. Each pin cell has a pitch of 1.26 cm and the square fuel pins have sides of length 0.7314 cm. Guide tube pins are represented as a homogeneous water region. There is also a thin water blade region surrounding the assembly with a thickness of 0.08 cm. The geometry is shown in Fig. 4. The spatial discretisation of the pins is shown in Fig. 5. This was chosen for convenience of input: each mesh element of each pin had a uniform width of 0.1219 cm, resulting in a truncation at the boundaries of pin cells, where the thin

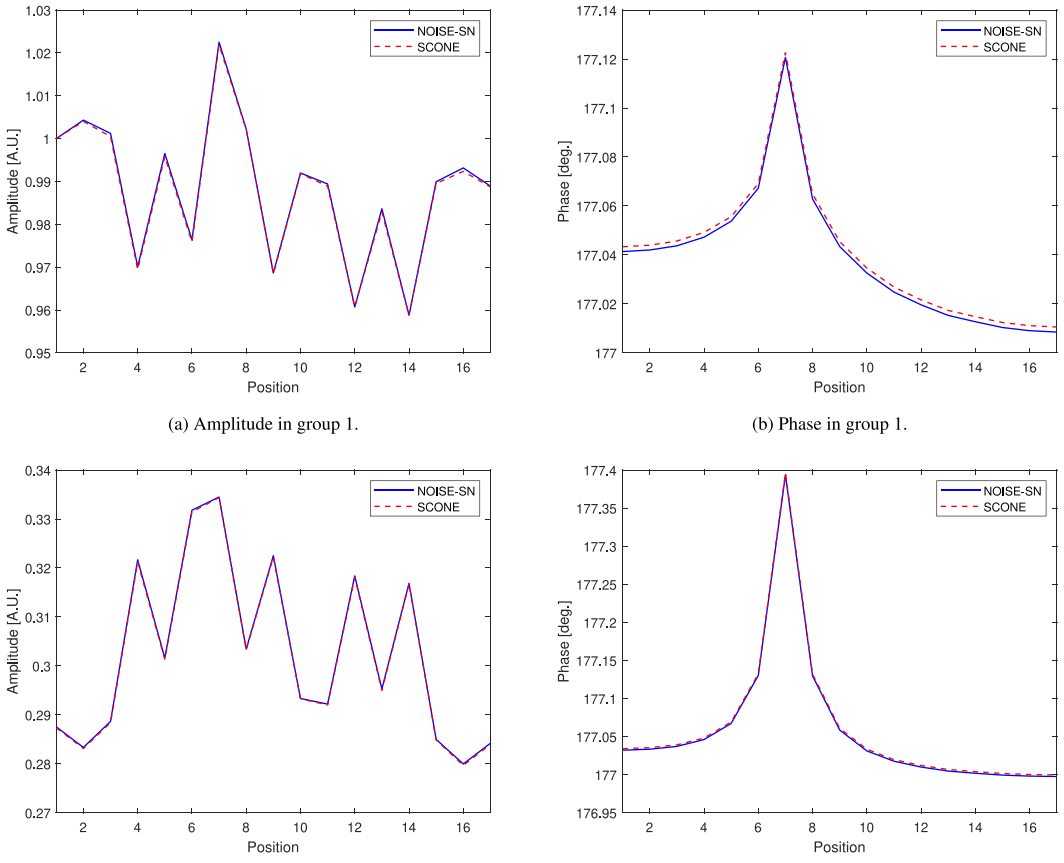

(a) Amplitude in group 1.

(b) Phase in group 1.

(c) Amplitude in group 2.

(d) Phase in group 2.

Fig. 7. Amplitudes and phases produced by SCONE and NOISE-SN for the UOX benchmark with oscillating cross sections.

elements are present. The discretisation of the water blade is similar, but with a width limited to 0.08 cm in the direction perpendicular to the assembly boundary. For all problems to follow, TRRM is run with 1250 rays per iteration, a dead length of 40 cm, a termination length of 240 cm, and 2000 active iterations.

Performing a standard eigenvalue calculation of the geometry took 18 s, requiring 415 inactive iterations and 67 MB of memory, and produced a k_{eff} of 0.99931 ± 2 pcm, which is in excellent agreement with the Monte Carlo reference results (0.99912 ± 8) (Vinai et al., 2021).

There are several variants of the problem and two are presented here, both are types of variable strength absorber problem: one with a 'fixed' noise source (an inhomogeneous source in the noise equation) and one with a cross section oscillation (an oscillation of a cross section, which produces a noise source by multiplying the scalar flux). Results for each problem are presented along the bottom-left-to-top-right diagonal of the geometry, normalised to the fast flux noise amplitude in the bottom left pin cell. The presentation style differs slightly from other publications as the quantities are averaged over the entirety of each pin cell along the diagonal, rather than being the values in mesh elements along the diagonal — this was done simply as it is more convenient for the SCONE solver and due to the differences in spatial discretisation.

For the fixed noise source, there is no cross section variation and so no dependence of the neutron noise on the scalar flux. To be precise, the S_g in Eq. (20) is replaced by a constant term in the source region and 0 elsewhere. This constant source is only in the thermal group and is given as:

$$S_2 = -1 + i \ .$$

The noise source is designated as having a frequency of 3 Hz or 6π rad/s. When solving with SCONE, this took 7 min and 43 s, requiring 38,762 inactive iterations and 89 MB of memory. Both runtime and memory could have been reduced given that the eigenvalue component of the calculation is entirely redundant here. The phase and amplitude results are plotted against results produced by NOISE-SN (which had excellent agreement with other noise solvers) and shown in Fig. 6. Likewise, there is an excellent agreement between SCONE and NOISE-SN, with any differences produced being much lower than 1% in all phase and amplitude results.

The second problem features a cross section oscillation in the same location as the first. In particular, the oscillation occurs in all energy groups for all cross sections. In the time domain it takes the following form:

$$\delta\Sigma_{\text{f},g} = 0.021\,\Sigma^0_{\text{f},g}\cos(\omega_0 t) \ ,$$

$$\delta\Sigma_{\text{s},g\to g'} = 0.034\,\Sigma^0_{\text{s},g\to g'}\cos(\omega_0 t) \ ,$$

$$\delta\Sigma_{\text{t},g} = 0.041\,\Sigma^0_{\text{t},g}\cos(\omega_0 t) \ ,$$

where $\omega_0 = 2\pi$ rad/s and the 0 superscript represents the nominal values of the cross sections. Solving this problem took 8 min and 1 s, requiring 41,274 inactive iterations and 97 MB. The comparison against NOISE-SN is presented in Fig. 7 and once again the results are in strong agreement at all points.

4. Summary and conclusions

A new approach to solving neutron noise problems was introduced where the transport equation and noise equation are solved simultaneously. This means that angular fluxes do not need to be stored to

produce noise solutions. This also makes it straightforward to implement a frequency-domain neutron noise solver using the method of characteristics, allowing improved geometric fidelity in deterministic noise solvers. A random ray characteristics solver using this algorithm produced results rapidly and accurately on several popular neutron noise problems featuring absorbers of variable strength.

The most significant remaining work to attempt is solving problems featuring spatial vibrations, e.g., the third exercise of the assembly benchmark (Vinai et al., 2023). These problems are notoriously challenging due to featuring noise sources of similar magnitude but opposite sign, particularly for Monte Carlo solvers (Rouchon et al., 2017; Vinai et al., 2023; Belanger et al., 2023). This problem has been attempted with the noise solver, albeit with less success than the other problems presented. One might expect that this is due to stochastic noise introduced by the random ray solver although the effects of numeric precision also appear to be prominent — it is hoped that these challenges can be resolved with further investigation.

Although the solver was performant, it required many thousands of inactive iterations — if these iterations became more expensive (say, when analysing a larger geometry) then this approach would not be feasible. Further acceleration could be expected were an acceleration algorithm to be implemented for the noise solution, such as was done by Yi et al. (2021), resulting in only $\mathcal{O}(10)$ inactive iterations needing to be performed for C5G7. This should be pursued in future to improve the efficiency of these calculations.

CRediT authorship contribution statement

Paul Cosgrove: Writing – original draft, Software, Methodology, Investigation, Conceptualization. **Maximilian Kraus:** Writing – review & editing, Software. **Valeria Raffuzzi:** Writing – review & editing, Software.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

We are very grateful to P. Vinai for providing NOISE-SN solutions to the C5G7 and UOX assembly problems and for helpful correspondence. P. Benie assisted through maintaining the Cambridge Nuclear Energy Group workstations. This work was supported by the EPSRC, United Kingdom grant MCSIMus (EP/T022159/1).

References

Belanger, H., Mancusi, D., Rouchon, A., Zoia, A., 2023. Variance reduction and noise source sampling techniques for Monte Carlo simulations of neutron noise induced by mechanical vibrations. Nucl. Sci. Eng. 197 (4), 534–557.
Carreño, A., Vidal-Ferràndiz, A., Ginestar, D., Verdú, G., 2022. Frequency-domain models in the SPN approximation for neutron noise calculations. Prog. Nucl. Energy 148, 104233.
Cosgrove, P., Tramm, J.R., 2023. The Random Ray Method versus Multigroup Monte Carlo: The Method of Characteristics in OpenMC and SCONE. Nucl. Sci. Eng.
Demazière, C., 2004. Development of a 2-D 2-group neutron noise simulator. Ann. Nucl. Energy 31 (6), 647–680.
Demazière, C., 2011. CORE SIM: A multi-purpose neutronic tool for research and education. Ann. Nucl. Energy 38 (12), 2698–2718.
Gammicchia, A., Santandrea, S., Zmijarevic, I., Sanchez, R., Stankovski, Z., Dulla, S., Mosca, P., 2020. A MOC-based neutron kinetics model for noise analysis. Ann. Nucl. Energy 137, 107070.
Hou, J., Ivanov, K.N., Boyarinov, V.F., Fomichenko, P.A., 2017. OECD/NEA benchmark for time-dependent neutron transport calculations without spatial homogenization. Nucl. Eng. Des. 317, 177–189.
Kang, J., Joo, H.G., 2021. NTRACER solutions to C5G7-TD neutron noise simulation benchmark. Trans. Amer. Nucl. Soc. 124, 584–587.
Knott, D., Yamamoto, A., 2010. Lattice physics computations. In: Cacuci, D.G. (Ed.), Handbook of Nuclear Engineering. Springer, pp. 918–1239.
Kowalski, M.A., Cosgrove, P., Broman, J., Shwageraus, E., 2021. SCONE: A student-oriented modifiable Monte Carlo particle transport framework. J. Nucl. Eng. 2, 57–64.
Lewis, E.E., Smith, M.A., Tsoulfanidis, N., Palmiotti, G., Taiwo, T.A., Blomquist, R.N., 2001. Benchmark Specification for Deterministic 2-D/3-D MOX Fuel Assembly Transport Calculations Without Spatial Homogenisation (C5G7 MOX). OECD-NEA.
Mylonakis, A., Vinai, P., Demazière, C., 2021. CORE SIM+: A flexible diffusion-based solver for neutron noise simulations. Ann. Nucl. Energy 155, 108149.
Pázsit, I., 1984. The linearization of vibration-induced noise. Ann. Nucl. Energy 11 (9), 441–454.
Pázsit, I., Demazière, C., 2010. Noise techniques in nuclear systems. In: Cacuci, D.G. (Ed.), Handbook of Nuclear Engineering. Springer US, Boston, MA, pp. 1629–1737.
Rouchon, A., 2016. Analyse et Développement D'outils Numériques Déterministes Et Stochastiques Résolvant les Équations du Bruit Neutronique et Applications aux Réacteurs Thermiques et Rapides (Ph.D. thesis). Université Paris-Saclay.
Rouchon, A., Zoia, A., Sanchez, R., 2017. A new Monte Carlo method for neutron noise calculations in the frequency domain. Ann. Nucl. Energy 102, 465–475.
Tramm, J.R., 2018. Development of the Random Ray Method of Neutral Particle Transport for High-Fidelity Nuclear Reactor Simulation (Ph.D. thesis). Massachusetts Institute of Technology.
Tramm, J.R., Siegel, A.R., Lund, A.L., Romano, P.K., 2020. A comparison of stochastic mesh cell volume computation strategies for the random ray method of neutral particle transport. In: Proc. PHYSOR 2020. Cambridge, UK.
Tramm, J.R., Smith, K.S., Forget, B., Siegel, A.R., 2017. The Random Ray Method for neutral particle transport. J. Comput. Phys. 342, 229–252.
Tramm, J.R., Smith, K.S., Forget, B., Siegel, A.R., 2018. ARRC: A random ray neutron transport code for nuclear reactor simulation. Ann. Nucl. Energy 112, 693–714.
Vinai, P., Demazière, C., Hou, J., Ivanov, K., 2019. Deterministic Time-Dependent Neutron Transport Benchmark without Spatial Homogenization (C5G7-TD) Volume I: Kinetics Phase Part B: Neutron Noise Simulation. OECD-NEA.
Vinai, P., Yi, H., Demazière, C., Rouchon, A., Zoia, A., Vidal-Ferràndiz, A., Carreño, A., Ginestar, D., Verdú, G., 2023. On the simulation of neutron noise induced by vibrations of fuel pins in a fuel assembly. Ann. Nucl. Energy 181, 109521.
Vinai, P., Yi, H., Mylonakis, A., Demazière, C., Gasse, B., Rouchon, A., Zoia, A., Vidal-Ferràndiz, A., Ginestar, D., Verdú, G., Yamamoto, T., 2021. Comparison of neutron noise solvers based on numerical benchmarks in a 2-D simplified UOX fuel assembly. In: Proc. M&C 2021. Raleigh, NC.
Yamamoto, T., 2013. Monte Carlo method with complex-valued weights for frequency domain analyses of neutron noise. Ann. Nucl. Energy 58, 72–79.
Yi, H., Vinai, P., Demazière, C., 2021. On the simulation of neutron noise using a discrete ordinates method. Ann. Nucl. Energy 164, 108570.