Deterministically Estimated Fission Source Distributions for Monte Carlo k-Eigenvalue Problems[☆]

Elliott D. Biondo[a,1], Gregory G. Davidson[a,1], Tara M. Pandya[a,1], Steven P. Hamilton[a,1], Thomas M. Evans[a,1]

[a] *Oak Ridge National Laboratory, 1 Bethel Valley Rd., Oak Ridge, TN 37831 U.S.A.*

Abstract

The standard Monte Carlo (MC) k-eigenvalue algorithm involves iteratively converging the fission source distribution using a series of potentially time-consuming *inactive cycles* before quantities of interest can be tallied. One strategy for reducing the computational time requirements of these inactive cycles is the Sourcerer method, in which a deterministic eigenvalue calculation is performed to obtain an improved initial guess for the fission source distribution. This method has been implemented in the Exnihilo software suite within SCALE using the SP_N or S_N solvers in Denovo and the Shift MC code. The efficacy of this method is assessed with different Denovo solution parameters for a series of typical k-eigenvalue problems including small criticality benchmarks, full-core reactors, and a fuel cask. It is found that, in most cases, when a large number of histories per cycle are required to obtain a detailed flux distribution, the Sourcerer method can be used to

[*]Corresponding Author. Tel: +1 (865) 241-4431

Email addresses: biondoed@ornl.gov (Elliott D. Biondo), davidsongg@ornl.gov (Gregory G. Davidson), pandyatm@ornl.gov (Tara M. Pandya), hamiltonsp@ornl.gov (Steven P. Hamilton), evanstm@ornl.gov (Thomas M. Evans)

[1]Radiation Transport Group, Reactor and Nuclear Systems Division

reduce the computational time requirements of the inactive cycles.

Keywords: `fission source convergence, Monte Carlo transport, hybrid methods`

---

1. Introduction

The design and operation of nuclear fission reactors requires a detailed solution to the neutron transport k-eigenvalue problem to obtain the neutron multiplication factor (k_{eff}) for criticality safety analyses and pin- and subpin-resolved flux distributions for depletion, thermohydraulic-coupling, and radiation damage analyses. In some cases, the accuracy required for these analyses necessitates the use of Monte Carlo (MC) methods for neutron transport. The traditional MC power iteration algorithm for solving the k-eigenvalue problem involves carrying out a series of *cycles*, in which a collection of neutron histories are simulated and the induced fission sites are recorded. Since the fission source distribution is not known *a priori*, an initial guess must be used for the first cycle. For subsequent cycles, the fission source can be formed from the fission sites from the previous cycle. *Inactive* cycles—in which no events are tallied—are first carried out to obtain a converged fission source. *Active* cycles are then performed, and an estimate of k_{eff} and the flux distribution are recorded for each cycle. The final estimate for k_{eff} and the flux distribution can then be ascertained by averaging together the estimates from the active cycles.

Full-core reactor analysis generally requires a large number of inactive cycles that are computationally expensive. For eigenvalue-only calculations inactive cycles may account for approximately half of the total runtime. For calculations that require more detailed tallies the relative expense of inactive cycles is reduced. Numerous strategies have been proposed to reduce the computational time spent on inactive cycles.

Acceleration techniques such as the fission matrix method [1, 2] attempt to improve how a new fission source is generated from the previous iteration by using additional information garnered from solving an approximate problem. Preliminary analysis using a deterministic-based fission matrix acceleration method yielded promising results, but further investigation is required before adopting this method for production-level use [3, 4]. With the particle ramp-up method [5], the number of histories per cycle is gradually increased during inactive cycles so that less computational time is wasted simulating particle

histories when the fission source is far from converged. However, it is possible that this method would exacerbate issues with neutron clustering [6]. With the MC version of Wielandt's method the fission source is formed from a combination of the induced fission sites from the previous cycle and also the current cycle, in order to increase the interaction between decoupled regions of a problem [7].

Another strategy is to use an improved initial guess for the fission source distribution such that fewer inactive cycles are required to obtain a converged fission source. As opposed to acceleration methods, this technique does not modify the rate of convergence of power iteration. In principle this could be done in concert with the other two strategies, but in this work this third strategy is studied in isolation. The default initial guess for the fission source for k-eigenvalue problems is usually generated from a flux distribution that is uniform throughout all space. The conversion between the flux distribution and the fission source is later described in Eq. (4). For reactor problems, flux distributions that are uniform radially and have a cosine or flattened cosine[2] shape axially can be used. The Watt fission spectrum is typically used for the energy distribution, and isotropic emission is used for the angular distribution. The complexity of the spatial distribution of the fission source is not captured by these typical initial guesses, especially in the presence of inhomogeneities such as control and burnable absorber rods, mixed fuels, and variable burnup states.

The Sourcerer method [8] can be used to obtain an improved initial guess for arbitrary configurations, thereby providing an automatic means of reducing the number of inactive cycles. With this method, a deterministic neutron transport calculation is first carried out as a preprocessing step. The resulting approximate mesh-based flux distribution is used to calculate a mesh-based fission source distribution that is used as an initial guess. This method has previously been implemented in the Sourcerer module within the SCALE modeling and simulation suite [9] using Denovo [10] as a discrete ordinates (S_N) solver and the KENO MC transport code [9]. The method was demonstrated to be effective for a spent fuel canister problem with both transport codes run in serial [8]. Another method similar to the Sourcerer method has previously been proposed, but does not appear to have ever been implemented [1].

[2]The flattened cosine shape is defined as $f(x) = 1 - (1 - cos(x))^2$.

The Sourcerer method is effective only if a deterministic estimate of the flux distribution can be obtained quickly relative to the time required for the MC inactive cycles. As a practical consideration, the computer memory requirements for the deterministic solution should not be so large that an analyst must use additional computer resources solely to employ the Sourcerer method. Though most deterministic techniques are likely to satisfy these criteria for small reactor problems, full-core analysis poses significant challenges. Full-core modeling with the S_N method requires leadership-class computer resources [11], making the application of the Sourcerer method to full-core analysis impractical in most cases. The Simplified P_N (SP_N) method is a low-order approximation of the transport equation that can provide estimates of the flux distribution for full-core problems in significantly less processor time and with smaller memory requirements than S_N methods [12]. When a high degree of accuracy is not required, the SP_N method is preferable to the S_N method provided that the problem does not contain regions with extremely long mean free paths (e.g., air or void regions) that cannot be handled by the former method.

In this work, the performance of the Sourcerer method is assessed using deterministic estimates of the neutron flux with varying degrees of accuracy for a collection of typical k-eigenvalue problems. The Sourcerer method has been implemented using the Denovo SP_N [12] and S_N solvers, so the Sourcerer method can be applied to a wide range of problems, including full-core reactor analysis. This implementation is done in the Exnihilo software suite [13] within SCALE using the Shift MC transport code [14]. Shift, like Denovo, was designed to be massively parallel and has demonstrated excellent scaling behavior from laptops to leadership-class supercomputers [10, 14]. Both codes support a variety of geometry and physics packages. Shift and Denovo are easily coupled through the Omnibus front-end [15], which is common to both codes. This implementation requires the user to choose the solution parameters for the Denovo solver, including the SP_N order or S_N quadrature set, resonance self-shielding treatment for cross section generation, number of energy groups, and solution mesh. The accuracy of the Denovo solution—which is governed by this choice of parameters—could affect the efficiency of the Sourcerer method.

The rest of this paper proceeds as follows. Section 2 provides the theoretical background for the Sourcerer method and the use of Shannon entropy [16] to qualitatively assess fission source convergence. Section 3 first details how the Sourcerer method is implemented within Exnihilo and how cross sections are

generated with different resonance self-shielding treatments. Criteria are then proposed for quantitatively assessing the cycle in which an MC simulation can be deemed converged, based on Shannon entropy. Using results from these convergence criteria, a quantitative metric $N_{\text{break-even}}$ is then described that can be used to determine when the use of the Sourcerer method is justified. This parameter considers the number of inactive cycles saved using the Sourcerer method and the computational time requirements of the two transport steps.

Section 4 describes the models used for experimentation, which include two small criticality benchmark problems, three full-core reactor problems, and a fuel cask problem:

1. A modified version of the 2003 C5G7 mixed-oxide (MOX) fuel benchmark problem [17].
2. Core 17 of the Babcock & Wilcox® (B&W) 1810 benchmark [18].
3. Watts Bar Nuclear Power Station, Unit 1 (WBN1) [19].
4. Westinghouse AP1000® reactor [20].
5. Westinghouse AP1000® reactor, unrodded configuration [20].
6. NAC International Inc. Universal Multi-Purpose Canister System (NAC UMS®) fuel cask [21].

For each of these problems, the number of inactive cycles saved using the Sourcerer method, as well as $N_{\text{break-even}}$, are determined. The results of these experiments are discussed in Section 5. Section 6 provides concluding remarks.

This work demonstrates that the Sourcerer method might provide performance enhancements in a subset of problems—especially those where a large number of histories per cycle are required—but further advancements in the speed and accuracy of deterministic methods are necessary for Sourcerer to be advantageous for arbitrary use cases.

2. Theory

MC neutron transport can be used to estimate k_{eff} via the power iteration method, which involves successively solving the neutron transport equation to obtain convergent estimates of the fission source. The neutron transport equation can be formulated in operator notation as

$$(\hat{\mathrm{T}} - \hat{\mathrm{S}})\psi = \frac{1}{k_{\text{eff}}}\hat{\chi}\hat{\mathrm{F}}\psi \,, \tag{1}$$

where $\hat{\mathrm{T}}$ is the transport operator, $\hat{\mathrm{S}}$ is the scattering operator, $\hat{\chi}$ is the fission spectrum, $\hat{\mathrm{F}}$ is the fission operator, and ψ is the neutron flux. Assuming steady-state operation, these operators are independent of flux and depend only on the geometry, material compositions, and temperature. By collecting these operators in a single term,

$$\hat{\mathrm{A}} \equiv \hat{\mathrm{F}}(\hat{\mathrm{T}} - \hat{\mathrm{S}})^{-1}\hat{\chi}\,. \tag{2}$$

Eq. (1) can be formulated as a standard eigenvalue problem,

$$k_{\mathrm{eff}}\mathrm{f} = \hat{\mathrm{A}}\mathrm{f}, \tag{3}$$

where k_{eff} is the dominant eigenvalue of $\hat{\mathrm{A}}$ and f, the eigenvector, is the fission source given by

$$\mathrm{f} = \hat{\mathrm{F}}\psi. \tag{4}$$

The power iteration method estimates the fission source at iteration $n+1$ as

$$\mathrm{f}^{(n+1)} = \frac{1}{k_{\mathrm{eff}}^{(n)}}\hat{\mathrm{A}}\mathrm{f}^{(n)}. \tag{5}$$

Subsequently, k_{eff} for iteration $n+1$ can be formulated in terms of $\mathrm{f}^{(n+1)}$, $\mathrm{f}^{(n)}$, and $k_{\mathrm{eff}}^{(n)}$ as

$$k_{\mathrm{eff}}^{(n+1)} = \frac{\langle \hat{\mathrm{A}}\mathrm{f}^{(n)} \rangle}{\langle \mathrm{f}^{(n)} \rangle} = k_{\mathrm{eff}}^{(n)} \frac{\langle \mathrm{f}^{(n+1)} \rangle}{\langle \mathrm{f}^{(n)} \rangle}, \tag{6}$$

where $\langle \cdot \rangle$ denotes integration over all phase space. With MC neutron transport, Eq. (5) is implemented by recording fission sites during each cycle and using them to form the fission source for the next cycle.

The convergence rate of power iteration is equivalent to the dominance ratio, ρ, which is defined as the ratio of the second and first largest eigenvalues of $\hat{\mathrm{A}}$:

$$\rho = \frac{|k_1|}{k_0}. \tag{7}$$

In other words, the error is reduced by a factor of $1-\rho$ each iteration. The dominance ratio is a physical property of a system that describes the extent to which fissionable regions are decoupled. A reactor configuration with a large

dominance ratio (i.e., near 1) indicates that neutrons born in one fissionable region are unlikely to induce fission in some other fissionable region. The dominance ratio scales with the physical size of the reactor and increases with the presence of strong absorbers such as control rods. High burnup can create configurations with disparate regions with high concentrations of fissionable material (i.e., the top and bottom of fuel rods) that increase the dominance ratio. This decoupling results in slow power iteration convergence.

With power iteration, the number of cycles required for fission source convergence depends not only on the dominance ratio but also on the quality of the initial guess for the flux distribution used to form the initial fission source in Eq. (5). With the Sourcerer method, a deterministic transport calculation is carried out to obtain an approximate flux distribution that is then used as an initial guess. Provided this distribution is closer to the converged distribution than typical initial guesses (i.e., uniform, axial cosine, or axial flattened cosine flux distributions), this method will reduce the number of cycles required for fission source convergence compared with these typical cases. It is emphasized that even if the deterministically estimated flux distribution is inaccurate, since it is only used as an initial guess it will not affect the converged MC result for k_{eff} or the MC flux distribution.

Though the ultimate goal of k-eigenvalue problems is to obtain a converged estimate of k_{eff}, it is not sufficient to judge the convergence of a calculation by considering k_{eff} alone. When the dominance ratio is near unity, k_{eff} will converge faster than the fission source [22]. As a result, k_{eff} might falsely appear to be converged before the fission source is fully resolved. To avoid false convergence, the convergence of the fission source can be assessed with the Shannon entropy,

$$H = -\sum_i f_i \log_2(f_i), \tag{8}$$

where f_i is the fraction of source neutrons born in cell i [16]. As the cycle-wise estimates of f_i improve, the Shannon entropy converges to a constant, problem-dependent value.

Other metrics for assessing convergence have been proposed more recently including the center of mass of the neutron population [23] and higher-order moments of the Shannon entropy [6]. These metrics have x, y, and z components and may be more effective in detecting false convergence due to neutron clustering [6]. The Shannon entropy is used in this work because it is currently the most widely-used metric, and it is not directionally-dependent,

simplifying the method that will be described in Section 3.3 for assessing convergence.

Plotting the Shannon entropy versus cycle was originally proposed as a visual metric for assessing fission source convergence [16]. Both k_{eff} and H (as well as the other aforementioned metrics) are autocorrelated, meaning cycle-wise estimates of these quantities depend on results from previous cycles and are therefore not statistically independent. This poses challenges for quantitatively determining whether a simulation is converged [24]. For this reason, both visual and quantitative analysis of Shannon entropy convergence are used in this work.

3. Methodology

To assess how the fidelity of the deterministic solution affects the performance of the Sourcerer method, a series of experiments were conducted wherein the deterministic solution parameters were varied and the Shannon entropy was analyzed. In Section 3.1, implementation of the Sourcerer method within Exnihilo is discussed. Section 3.2 describes how broad-group cross sections were generated for deterministic transport, with different resonance self-shielding treatments. Section 3.3 describes how the Shannon entropy convergence of MC power iteration is quantified for the purposes of this work. Section 3.4 introduces a quantitative method for determining when the use of the Sourcerer method is justified.

3.1. Implementation of the Sourcerer Method

The Sourcerer method has been implemented in the Exnihilo software suite within SCALE, using Denovo as an SP_N or S_N solver and Shift for the MC k-eigenvalue calculation. This implementation is facilitated by the Omnibus front-end that allows Denovo and Shift to be run interchangeably from a common input format. When the Sourcerer method is initiated, Omnibus ray traces the geometry onto a Cartesian mesh for use with Denovo. Denovo is run, and the converged estimate of the flux distribution is automatically sent to Shift. Shift generates the initial fission source, $q(\vec{r}, E)$, based on one of two methods. The fission spectrum can be generated with a volume-weighted mixed fission spectrum, calculated automatically from the XSProc module

[9] within SCALE. Defining this spectrum as $\chi(\vec{r}, E)$, the fission source is

$$q(\vec{r}, E) = \chi(\vec{r}, E) \int_0^\infty \nu(E')\, \Sigma_f(\vec{r}, E')\, \phi(\vec{r}, E')\, \mathrm{d}E'. \tag{9}$$

Alternatively, the fission spectrum can be assumed to be a Watt fission spectrum in energy, $\chi_{\mathrm{Watt}}(E)$, regardless of the material. This yields the following fission source:

$$q_{\mathrm{Watt}}(\vec{r}, E) = \underbrace{\chi_{\mathrm{Watt}}(E)}_{q(E)} \underbrace{\int_0^\infty \nu(E')\, \Sigma_f(\vec{r}, E')\, \phi(\vec{r}, E')\, \mathrm{d}E'}_{q(\vec{r})}. \tag{10}$$

The Watt fission spectrum is nuclide specific. As a result, this assumption is valid only if all fissions come from the single nuclide used to form the Watt fission spectrum. A ^{235}U Watt fission spectrum might provide a reasonable approximation for a variety of typical reactor cases. As denoted in Eq. (10), this version of the fission source is fully separable in space and energy. In other words, it has a component that depends only on space, $q(\vec{r})$, and a component that depends only on energy, $q(E)$. Instead of storing a unique fission spectrum for each mesh cell based on the material composition, the Watt fission spectrum is stored only once. This is advantageous because it saves a significant amount of computer memory when forming the initial MC source probability distribution.

The user must specify the cross section library and parameters to be used by Denovo, including the $\mathrm{SP_N}$ order or $\mathrm{S_N}$ quadrature set and solution mesh. The impact that these choices have on performance is addressed in this work. In addition, the user must select the fission spectrum (i.e., Eq. (9) or (10)). Since the memory savings associated with Eq. (10) are significant and ^{235}U fission is expected to dominate for the problems considered, this latter method will be used exclusively in this work.

3.2. Cross Section Generation

The implementation of the Sourcerer method within Exnihilo requires users to specify the cross section library used by Denovo. The choice of energy group structure and cross section treatment could affect performance, so experimentation is done using two different strategies. Cross sections were

generated for broad group structures with 2 and 24 groups using two resonance self-shielding treatments: an infinite homogeneous medium (IHM) assumption and Wigner-Seitz (WS) pin cells [25]. The 24-group structure was created by starting with a 23-group structure from the LANCR code [26]. An additional energy group bound was added at 5 eV—the purpose of which is explained below—to form a 24-group structure used for experimentation, as shown in Appendix A. A 2-group structure was created with bounds at 20 MeV, 5 eV and 10^{-5} eV. Typical k-eigenvalue problems generally require finer group structures to capture important physics. However, with the Sourcerer method, the goal is to quickly obtain an improved initial guess for the flux distribution, so coarser group structures are considered in this work.

For both sets of groups, a Doppler-broadened library was created from an ENDF/B-VII 999-group neutron library [9] using MALOCS2 [9]. This was done using "MT=1099" flux data [9]. This data is material specific but generally uses a Maxwellian distribution in the thermal region, a $1/E$ distribution in the epithermal region, and a Watt spectrum in the fast region. One limitation of the MALOCS2 module is that a broad-group energy boundary must be present at 5 eV, which is used as a thermal cutoff energy [27], hence the addition of an extra group in the LANCR structure.

Using these Doppler-broadened libraries, IHM resonance self-shielded cross sections are generated inline within Exnihilo using Omnibus, which relies on XSProc for this operation. For each material in the problem, XSProc performs a 0-D S_N calculation, assuming an infinite homogeneous medium of the material. This is done using an energy group structure that uses every point in the continuous energy data (typically between 10^2 and 10^5, depending on the nuclide). The resulting neutron flux energy distributions are used to calculate the IHM cross sections for each material. For each mesh cell in the problem, the cross section within the cell is calculated using the "pin-resolved" method [28], in which the cross sections of the materials in each zone within the mesh cell are mixed by volume fraction.

The Wigner-Seitz approximation is used to model a pin within a lattice as a set of concentric cylinders. With this approximation the outermost annular region, representing the moderator, is sized such that the moderator volume is preserved. By generating these pin cells for each unique pin in a problem, deterministic neutron transport calculations can be performed to obtain zone-specific neutron flux distributions to create self-shielded cross sections. In this work, all WS cross sections are generated with a "pin-homogenized" method [28], in which the cross section within a pin cell (σ) is calculated from the

zone-specific fluxes (ϕ_z), Doppler-broadened cross sections (σ_z), and volume fractions within the pin cell (f_z) as given by

$$\sigma = \frac{\sum_z f_z \phi_z \sigma_z}{\sum_z f_z \phi_z}. \tag{11}$$

If there are multiple mesh cells per pin cell, all mesh cells are assigned this same cross section. For the small critical benchmark experiments in this work, unique pin cells were manually input into XSProc.

For the three full-core reactor problems considered in this work, models were available in the reactor-aware geometry format used by the Virtual Environment for Reactor Applications (VERA) software suite [29] produced by the Consortium for Advanced Simulation of Light Water Reactors (CASL) [30]. Using the VERA geometry format, WS cross sections were generated automatically using the `vera2omn` utility within Exnihilo. This utility automatically builds multizone pin cells for each unique fuel pin configuration and executes XSProc [28]. A 252-group master library is used, and the group collapse and self-shielding calculation are done in the same operation.

3.3. Convergence Criteria

As mentioned in Section 2, Shannon entropy was originally proposed as a visual metric for assessing fission source convergence. Visualization is a useful technique for qualitatively evaluating whether a calculation is converged; therefore, Shannon entropy plots are presented as a component of the results in Section 5. A quantitative method, proposed here, is also used to determine the cycle in which a simulation can be declared converged.

Criteria for determining whether a calculation has converged using the Shannon entropy have previously been proposed by Brown et al. [24]. These criteria consist of six statistical checks to be applied to an interval of q cycles, where values for q were suggested to be 10, 20, and 30. The objective was to use these criteria to automatically assess whether the Shannon entropy is converged during a simulation. These criteria did not use *a priori* knowledge of the converged Shannon entropy. In this work, MC simulations were run with a significant number of cycles such that the converged Shannon entropy is known conservatively. Based on this converged Shannon entropy, a simpler set of metrics were used to evaluate convergence. For the purpose of this work, power iteration is declared to be converged if the following Shannon entropy criteria are met:

1. The difference between the average of the previous q cycles and the average over the active cycles is within one standard deviation of zero.
2. The slope of the previous q cycles is less than 10^{-4}.

These criteria, as well as those proposed by Brown et al., are *ad hoc*. Since the Shannon entropy is autocorrelated, statistical methods that treat cycle-wise estimates of the Shannon entropy as independent samples are not valid. More research in this area is necessary but is beyond the scope of this work.

3.4. Computational Time Requirements

The Sourcerer method should reduce the number of cycles necessary for the fission source to converge. However, the use of this method is justified only if the deterministic neutron transport calculation can be performed in less time than that required to perform the extra cycles with MC. The computational time requirements for these extra cycles is directly proportional to the number of histories run per cycle. Different applications will require different numbers of histories per cycle. For example, a calculation where k_{eff} is the only desired result will require fewer histories per cycle than if detailed subpin-resolution results are desired. Likewise, the use of Sourcerer might be justified in the latter case but not the former. For this reason, the performance of the Sourcerer method is assessed in terms of the number of histories per cycle required to "break even" such that the time requirements to obtain a converged fission source with and without Sourcerer are equivalent. To calculate this quantity, the number of cycles saved using Sourcerer (C_{saved}) is defined as

$$C_{\text{saved}} = C_{\text{standard}} - C_{\text{Sourcerer}}, \tag{12}$$

where $C_{\text{Sourcerer}}$ and C_{standard} are the number of cycles required to obtain a converged fission source as defined by the criteria in Section 3.3 with and without the Sourcerer method, respectively. For the purpose of this work, C_{standard} will refer to the number of cycles required for convergence using the most appropriate typical initial guess for the flux, whether it be a uniform, axial cosine, or axial flattened cosine distribution. The number of particles required to break-even, $N_{\text{break-even}}$, is then defined by the relation

$$t_{\text{det}} = N_{\text{break-even}}\, C_{\text{saved}}\, t'_{\text{MC}}, \tag{13}$$

where t_{det} is the computational time required for deterministic transport and t'_{MC} is the computational time required for MC transport per particle. Sourcerer is run with a moderate, but arbitrarily chosen, number of particles per cycle, and $N_{\text{break-even}}$ is then calculated by rearranging Eq. (13):

$$N_{\text{break-even}} = \frac{t_{\text{det}}}{C_{\text{saved}}\, t'_{\text{MC}}}. \tag{14}$$

This quantity can then be compared with numbers of histories per cycle used for different analyses for each tested model, as found in the literature. For a given model, if the number of histories per cycle typically used is significantly larger than $N_{\text{break-even}}$, the use of the Sourcerer method is justified. This assumes that the number of histories per cycle is a fixed value, which is the case for most MC codes. This is not the case if the particle ramp-up method is being used. The $N_{\text{break-even}}$ is specific to a particular MC and deterministic transport code combination. Thus, results presented here are applicable only to the implementation of the Sourcerer method with Denovo and Shift but should provide insight into the applicability of the Sourcerer method in general.

4. Models

This section describes the models used for experimentation. Two small criticality benchmark experiments were selected: a modified version of the 2003 C5G7 benchmark and core 17 of the B&W 1810 benchmark. Two full-core models were also chosen: the WBN1 reactor and the rodded and unrodded versions of the Westinghouse AP1000®. Shift has previously been validated using all four of these problems [14, 31]. Finally, a NAC International Inc. (NAC) UMS® cask was selected, similar to one used for the initial analysis of the Sourcerer method conducted by Ibrahim et al. [8].

4.1. 2003 C5G7 Benchmark

The C5G7 computational benchmark is an eighth-core MOX fuel problem originally created to assess how effectively deterministic neutron transport can model heterogeneous reactor models without the use of spatial homogenization. This work used the 2003 version [17] of this problem rather than the more commonly used 2005 version, which is half the size in the axial dimension. Since the 2003 version is larger, it has a higher dominance ratio (approximately 0.978 [32]) and is therefore a more challenging fission source convergence

problem. An image of the model appears in Figure 1. The model consists of four 17×17 square pitch lattice pin assemblies. Each assembly is 21.42 cm × 21.42 cm × 192.78 cm. Including the moderator, the total size of the geometry is 64.26 cm × 64.26 cm × 214.20 cm. Reflecting boundaries are used at $x = 0$ cm, $y = 64.26$ cm, and $z = 0$ cm. Vacuum boundary conditions are used elsewhere.

The model contains two assemblies with UO_2 fuel and two assemblies fueled with three different enrichments of MOX: 4.3 wt.%, 7.0 wt.%, and 8.7 wt.%. Each assembly has a fission chamber in the center pin location along with 24 guide tubes arranged radially around the center pin. In the original benchmark, each fuel pin contained a homogenized fuel-clad mixture, for which seven-group cross sections were provided. This benchmark was modified for use with continuous energy MC by creating detailed fuel pins that explicitly model the fuel, clad, and gap regions, based on dimensions and compositions provided by preliminary documents pertaining to the benchmark [33].

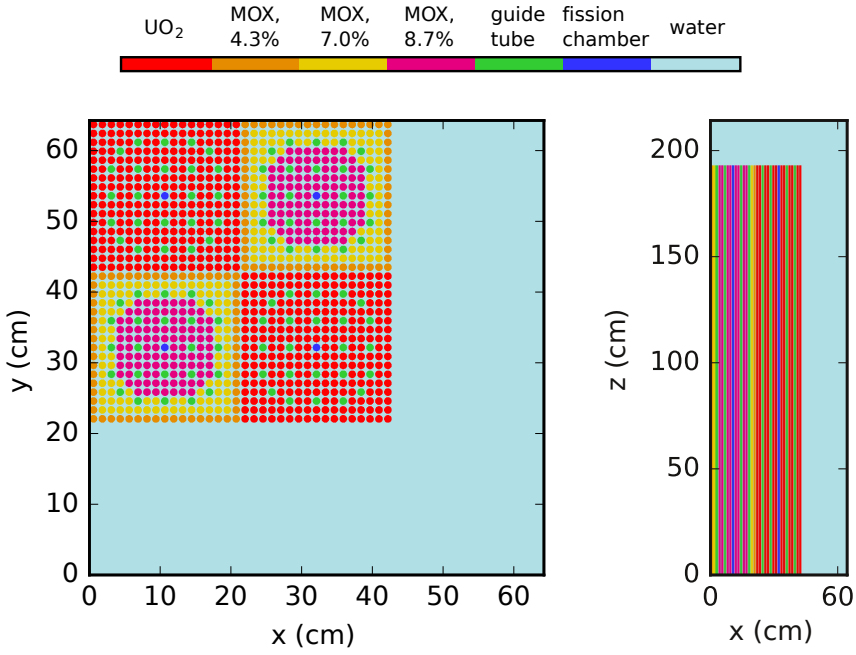

(a) Fuel loading

(b) x-z slice at $y = 32$ cm

Figure 1: Eighth-core model of the 2003 version of the C5G7 geometry with reflecting boundaries at x = 0 cm, y = 64.26 cm, and z = 0 cm. Cladding and gaps within pins have been omitted for visualization.

4.2. Babcock and Wilcox 1810 Benchmark

The B&W 1810 benchmark [18] consists of a series of critical experiments performed on a small light water reactor. The benchmark was originally created to compare computational results from a two-group diffusion model to experimental results in the presence of a UO_2-GdO_3 burnable absorber fuel mixture. A total of 23 different core configurations were studied. In this work, core XVII was chosen for experimentation. A diagram of this core configuration appears in Figure 2. This core was chosen because of the presence of additional absorbers and multiple enrichment zones. These heterogeneities increase the complexity of the fission source distribution.

The computational model used for the benchmark is a quarter core with reflecting boundaries on the x and y centerlines and vacuum boundaries elsewhere. The total size of the model is approximately 107.14 cm in x and y and 145 cm in z. Core XVII has an inner zone with fuel enriched to 4.0 wt.%, an outer zone with an enrichment of 2.46 wt.%, and 10 UO_2-Gd_2O_3

rods that contain 4.00 wt.% Gd_2O_3 and 96.00 wt.% UO_2 with a 1.94 wt.% enrichment. The model also contains four B_4C control rods and aluminum grid plates at the bottom of the geometry (minimum z), which provide axial heterogeneity. The water moderator in this model contains 1,432 ppm boron. Center and top grid plates, as well as the core tank, are not modeled.

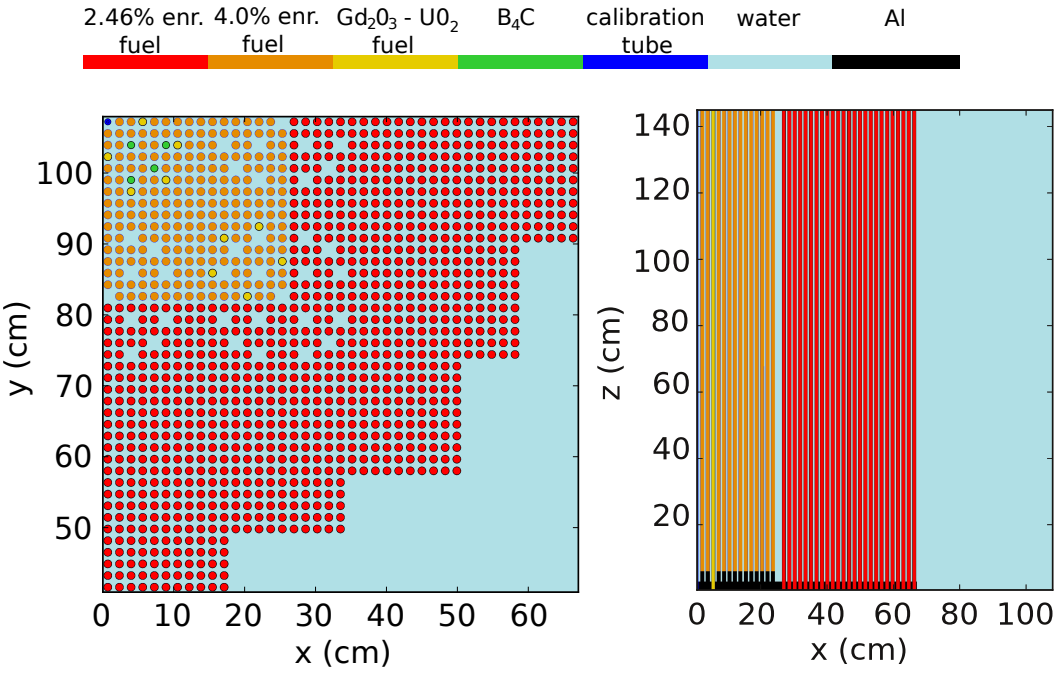

(a) Fuel loading

(b) x-z slice on the y centerline

Figure 2: Fuel loading of core XVII of the B&W 1810 benchmark. Cladding and gaps within pins have been omitted for visualization. Aluminum regions have several different compositions, but are collapsed into a single color for simplicity.

4.3. Watts Bar Nuclear Power Station, Unit 1

WBN1 is a Westinghouse-designed pressurized water reactor (PWR) in Rhea County, Tennessee, that began operating in 1996. A full-core model of this reactor (shown in Figure 3) was created to simulate initial startup physics as a component of the validation of the VERA software suite [19]. There are three varieties of 17×17 fuel assemblies containing fuel with ^{235}U enrichments of 2.110 wt.%, 2.619 wt.%, and 3.100 wt.%. Each assembly contains 24 control rod guide tube locations and a central instrumentation

tube. Both Pyrex® burnable absorber rods and rod cluster control assembly banks of AIC/B_4C rods are used, as well as moderator with a 1,285.2 ppm soluble boron chemical shim.

The reactor has a significant degree of axial heterogeneity with eight spacer grids as well as top and bottom assembly nozzles and core plates. The pressure vessel and containment are not modeled. This reactor is considerably larger than the 2003 C5G7 and B&W 1810 models, with the quarter core measuring approximately 365.5 cm × 365.5 cm × 418.94 cm. WBN1 has a high dominance ratio (approximately 0.987 [32]), making this a challenging fission source convergence problem. Additionally, the axial variation in this problem perturbs the axial flux distribution. A simplified approach, such as assuming that the flux distribution is a cosine axially, is not expected to be effective for this problem, which makes this problem a good candidate for application of the Sourcerer method.

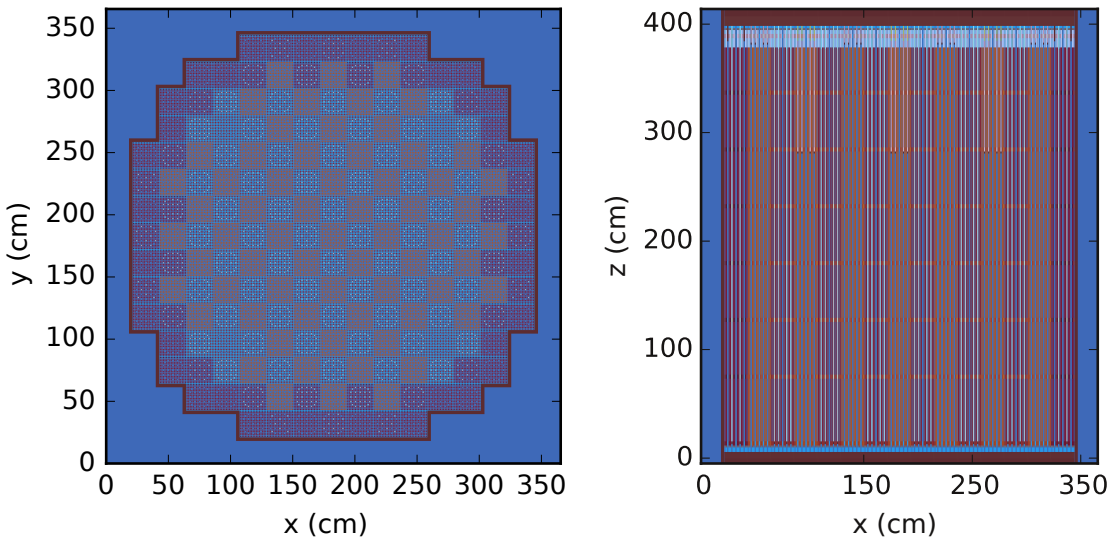

(a) Fuel loading

(b) x-z slice on the y centerline

Figure 3: WBN1 model. Red, orange, and purple represent fuel, and blue represents moderator.

4.4. Westinghouse AP1000®

The Westinghouse AP1000® is a Generation III+ PWR reactor design currently being constructed at six sites in China and the United States. A

full-core model of the AP1000® was created as a test stand for the VERA software suite [20], as shown in Figure 4. The reactor has five different 17 × 17 assembly types with ^{235}U enrichments ranging from 0.7 wt.% (natural) to 4.8 wt.%. The AP1000® uses as many as 12 Al_2O_3-B_4C wet annular burnable absorber rods per assembly. In addition, as many as 124 rods per assembly consist of fuel coated with a ZrB_2 burnable absorber known as Westinghouse integral fuel burnable absorber (IFBA). The mechanical shim operational strategy is also employed using tungsten control banks, Ag-In-Cd control banks, and boron chemical shim. Simulations were performed with these control rods inserted and also removed. The latter case is referred to as the "unrodded" configuration. The unrodded configuration was also chosen because it is known to have a top-peaked flux distribution from axial variation in the burnable absorbers. Since uniform, axial cosine, and axial flattened cosine initial guesses do not capture this shape, it might be a good candidate for the Sourcerer method.

Axially, the reactor has significant heterogeneities including 12 ZIRLO® grids, 8 intermediate mixing vane grids, 4 intermediate flow mixing grids, Inconel grids above and below the fuel, and an Inconel grid above the bottom nozzle [20]. The pressure vessel and containment are not modeled. This model is similar in size to the WBN1 model, measuring approximately 365.56 cm × 365.56 cm × 494.48 cm, which means this problem is also challenging for fission source convergence.

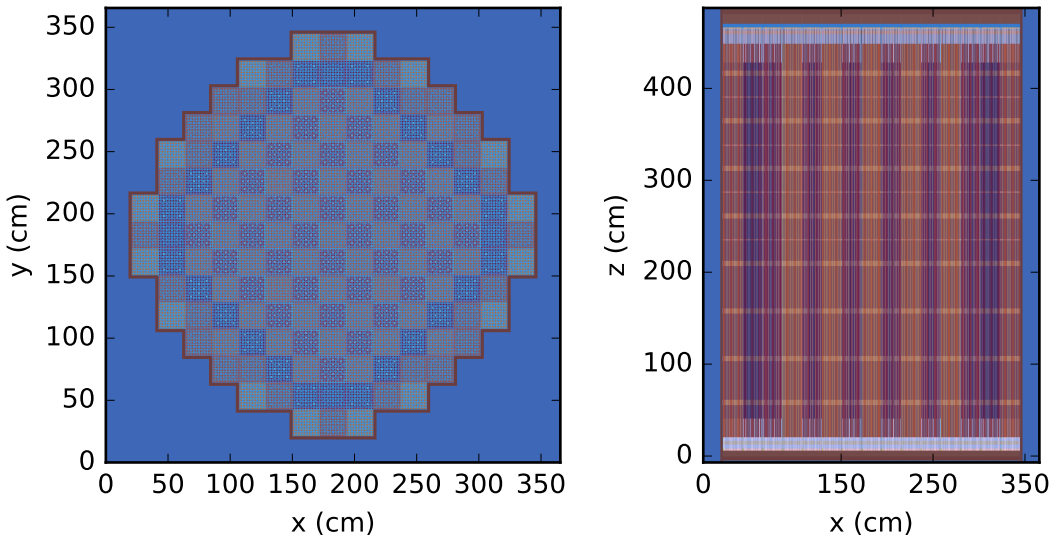

(a) Fuel loading

(b) x-z slice on the y centerline

Figure 4: Westinghouse AP1000® model. Red, orange, and purple represent fuel, and blue represents moderator.

4.5. NAC UMS® Cask

This model (shown in Figure 5) consists of a NAC International Inc. Universal Multi-Purpose Canister System (UMS®) Model TSC-24 fuel cask loaded with 24 Combustion Engineering 14×14 spent fuel assemblies from the Maine Yankee PWR [21]. The Maine Yankee power station was located in Wiscasset, Maine, and was decommissioned in 1996. This cask is serial number TSC-24-TSC-11 and contains assemblies that were discharged over the course of 15 years with initial ^{235}U enrichments between 3.03 and 3.91% and discharge burnups between 25.9 GWd/MtHM and 49.2 GWd/MtHM. The fuel compositions have been adjusted to simulate the expected compositions in the year 2100 (i.e., a decay time of 104 years after the last discharge date). Each assembly is surrounded by Boral plates, which are effective at absorbing thermal neutrons. The water regions between assemblies form a "flux trap" where neutrons are thermalized and thereby prevented from entering adjacent assemblies. The total size of the model is approximately 270.33 cm × 270.33 cm × 544.57 cm, which includes a 60 cm annular water reflector region. Vacuum boundaries are used on all sides of the model.

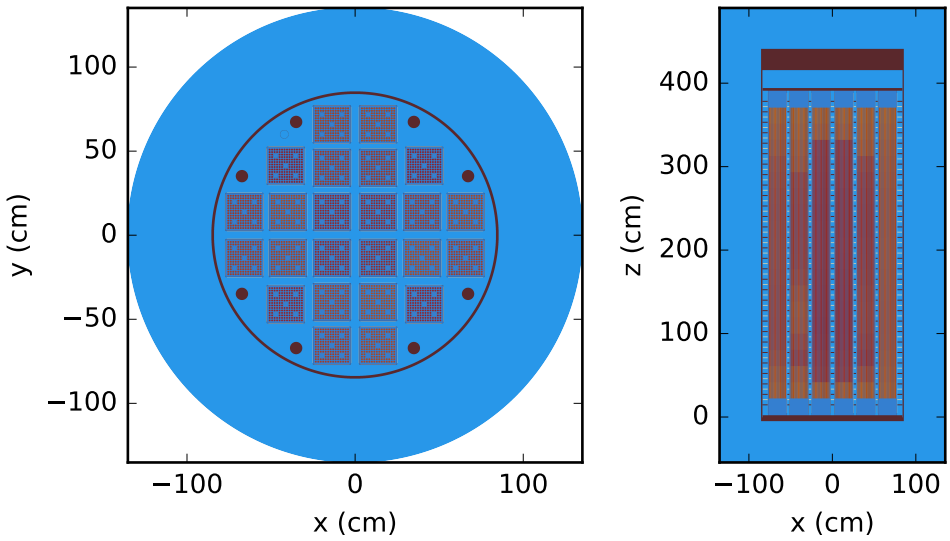

(a) Fuel loading

(b) x-z slice on the y

Figure 5: INF model. Red, orange, and purple represent fuel, and blue represents water.

5. Results

For each model described in Section 4, a Shift k-eigenvalue calculation was performed using a uniform, axial cosine, or axial flattened cosine initial flux distribution, as well as flux distributions from Denovo solutions with different resolutions. These starting sources are compared with the MC solution from Shift. Plots of k_{eff} and Shannon entropy convergence, in combination with convergence and timing results, are used to assess the efficacy of the Sourcerer method. All trials were run on the CADES compute cluster at ORNL, which is composed of Intel® Xeon® CPU E5-2698 v3 compute nodes with 32 cores per node and 128 GB RAM per node. All Denovo SP_N calculations were performed with an L_2 convergence tolerance of 1×10^{-6}, which is moderately tight for the SP_N solution method [12].

5.1. C5G7

For the C5G7 problem, the low-resolution Sourcerer trial used an SP_1 angular expansion, a P_0 scattering expansion solver order, 2-group IHM cross sections, and a solution mesh with 2 pins per mesh element radially and

5 cm mesh divisions axially ($26 \times 26 \times 43 = 29{,}068$ mesh cells). The high-resolution trial used an SP_3 angular expansion, a P_1 scattering expansion, 24-group WS cross sections, and a solution mesh with 1 pin per mesh element radially and 5 cm mesh divisions axially ($51 \times 51 \times 43 = 111{,}843$ mesh cells). These trials are compared against uniform and axial cosine starting sources. All trials were run with 500 inactive cycles and 500 active cycles with 1×10^7 histories per cycle on 2 compute nodes with 32 MPI tasks per node.

The radial and axial flux distributions for the axial cosine distribution and two Sourcerer trials are compared against the MC solution in Figure 6. The two Sourcerer trials match the MC solution nearly exactly axially, but the low-resolution Sourcerer trial fails to capture the radial shape with the same fidelity as the high-resolution trial.

The k_{eff} and Shannon entropy are plotted as a function of cycle for each trial in Figure 7. As expected, k_{eff} appears to converge significantly faster than the Shannon entropy, an effect that is most pronounced in the uniform and axial cosine trials. Qualitatively it is apparent that both Sourcerer trials converge in significantly fewer cycles than the uniform and axial cosine trials. This is corroborated by the results from the convergence criteria described in Section 3.3, which are marked on the Shannon entropy plot and also tabulated in Table 1. For these convergence tests, a step size (q) of 20 was used. According to this metric, approximately 100 cycles are saved by using either Sourcerer trial when compared with using an axial cosine. This suggests that the reduction of inactive cycles is insensitive to the choice of Denovo solution parameters for this problem.

As expected, Table 1 indicates that the Shift runtime per particle history is not significantly affected by the use of a more accurate starting fission source; all of the potential runtime savings result from the ability to run fewer inactive cycles to converge the fission source. The Denovo run times vary by more than two orders of magnitude, and likewise the $N_{\text{break-even}}$ values are 1.68×10^3 and 2.97×10^5 for the low- and high-resolution Sourcerer trials, respectively. It is likely that typical analysis using this problem would require significantly more than 10^3 – 10^5 histories per cycle. For example, Shift was previously benchmarked with this problem by running 1×10^6 histories for 250 inactive cycles and 1,000 active cycles to compare k_{eff}, integrated assembly powers, and maximum pin powers to published results [31].

Because the convergence metric used in this study is affected by the statistical noise in the Shannon entropy, additional analysis was performed to estimate the uncertainty in the converged cycle results. This was done by

repeating each trial 9 times with different random number seeds and applying the convergence criteria, resulting in a total of 10 statistically independent samples of the converged cycle for each trial. The average and standard deviation of these 10 samples are compared to the original sample in Table 2. From these results, it appears that there is considerable statistical variation in converged cycle results, but the general trend is the same: both Sourcerer trials result in significantly faster convergence. Obtaining 10 samples for each trial for the remaining problems in this work was not feasible from the standpoint of computational resource requirements. For the remaining problems, only a single sample was obtained, with the understanding that there is some degree of uncertainty with these results.

Table 1: Convergence results (using the criteria described in Section 3.3 with $q = 20$) and $N_{\text{break-even}}$ values (as described in Section 3.4) for the C5G7 problem. Cycles saved values are relative to the cosine distribution.

Method	Conv. cycle	Cycles saved	Denovo wall time (s)	Shift wall time per hist. (s)	$N_{\text{break-even}}$
Uniform	220	–	–	3.12×10^{-6}	–
Cosine	138	–	–	3.12×10^{-6}	–
Sourcerer low	39	99	0.53	3.18×10^{-6}	1.68×10^{3}
Sourcerer high	36	102	95.6	3.15×10^{-6}	2.97×10^{5}

Table 2: Comparison of the single-sample converged cycles results in Table 1 to the average and standard deviation of the converged cycles of 10 samples.

Method	Conv. cycle, initial sample	Conv. cycle, average of 10 samples
Uniform	220	260 ± 40
Cosine	138	150 ± 20
Sourcerer low	39	80 ± 60
Sourcerer high	36	50 ± 20

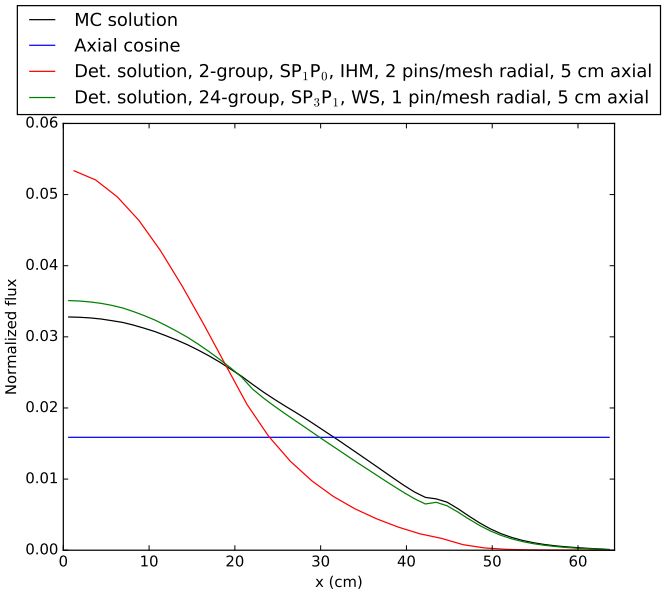

(a) Radial, y and z fuel midplanes

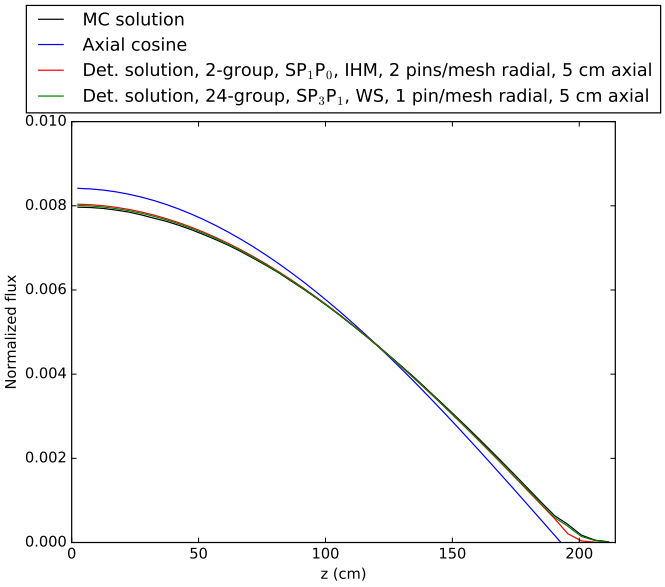

(b) Axial, centerline

Figure 6: Flux distributions used as initial guesses for MC power iteration, compared with the converged MC solution, for the C5G7 problem.

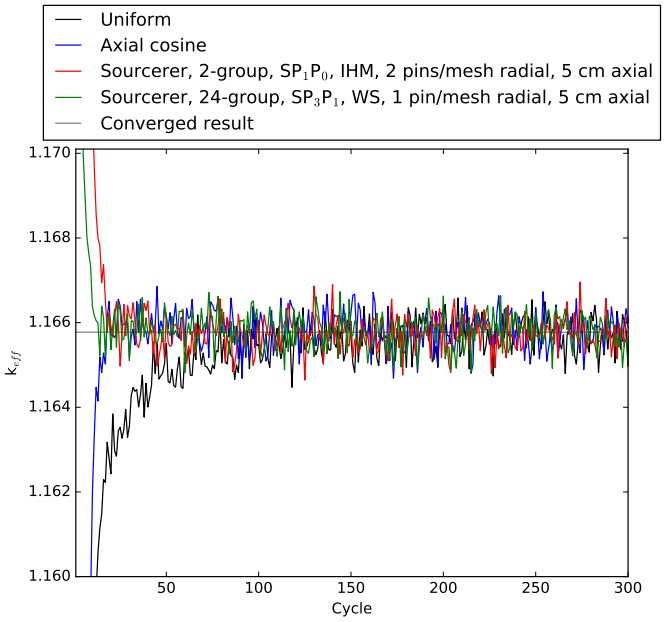

(a) k_{eff}.

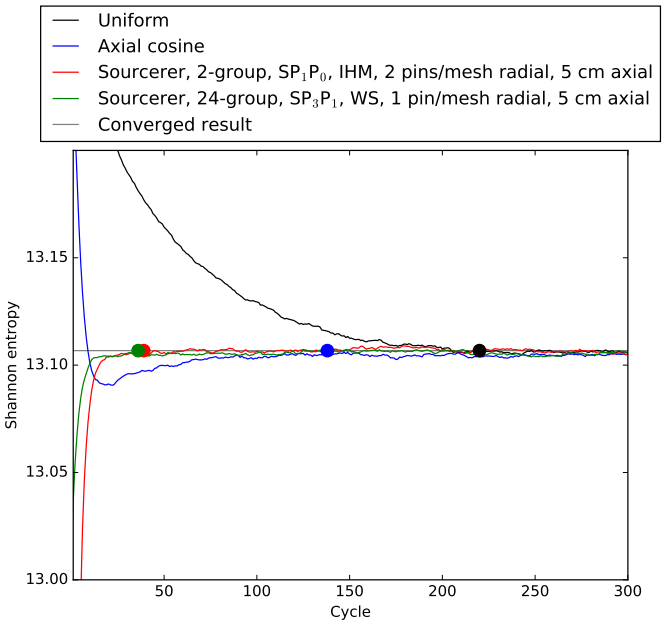

(b) Shannon entropy.

Figure 7: Convergence plots for the C5G7 problem. Markers represent the converged cycles tabulated in Table 1.

5.2. B&W 1810 Core 17

For the B&W 1810 core 17 problem, the low-resolution Sourcerer trial used an SP_1 angular expansion, a P_0 scattering expansion, 2-group IHM cross sections, and a solution mesh with 2 pins per mesh element radially and 5 cm mesh divisions axially ($30 \times 30 \times 30 = 27{,}000$ mesh cells). The high-resolution trial used an SP_3 angular expansion, a P_1 scattering expansion, 24-group WS cross sections, and a solution mesh 1 pin per mesh element radially and 5 cm mesh divisions axially ($58 \times 58 \times 30 = 100{,}920$ mesh cells). These trials are compared against uniform and axial cosine starting sources. All trials were run with 200 inactive cycles and 200 active cycles with 1×10^7 histories per cycle on two compute nodes with 32 MPI tasks per node.

The radial and axial flux distributions for the axial cosine distribution and two Sourcerer trials are compared against the MC solution in Figure 8. As with the C5G7 trial, the high-resolution solution matches the MC solution nearly exactly, and the low-resolution solution captures the correct axial shape but not the correct radial shape.

Figure 9 shows the Shannon entropy as a function of cycle. Visually, it is apparent that the high-resolution Sourcerer trial converges in significantly fewer cycles than the other trials. Quantitative convergence analysis was performed with $q = 10$ for this problem (instead of $q = 20$, as was done for all other problems) because the high-resolution Sourcerer trial appeared to converge in under 20 cycles. Visually it appears that the axial cosine trial converges in the second fewest number of cycles, followed by the low-resolution Sourcerer trial. However, using the quantitative metrics, this order is reversed. This highlights the shortcomings of the convergence criteria described in Section 3.3 and suggests there is a considerable degree of uncertainty in the number of cycles saved and $N_{\text{break-even}}$ as tabulated in Table 3.

Previous Shift benchmarking with this problem involved running 250 inactive cycles and 1,000 active cycles with 1×10^6 histories per cycle to compare k_{eff} and pin powers to experimental results [31]. Though the $N_{\text{break-even}}$ for the low-resolution trial is considerably less than 1×10^6, it is not clear that this result is reliable, as previously mentioned. The $N_{\text{break-even}}$ for the high-resolution Sourcerer trial is 6.50×10^5. Since this value is close to 1×10^6, it appears that even if the value for the number of cycles saved is accurate for the high-resolution trial, Sourcerer will provide very little benefit. If more detailed analysis were to be required for this problem—e.g., subpin flux distributions for depletion analysis—perhaps 1×10^7 histories per cycle would be run and the use of the Sourcerer method would be well justified.

Table 3: Convergence results (using the criteria described in Section 3.3 with $q = 10$) and $N_{\text{break-even}}$ values (as described in Section 3.4) for the B&W 1810 core 17 problem. Cycles saved values are relative to the cosine distribution.

Method	Conv. cycle	Cycles saved	Denovo wall time (s)	Shift wall time per hist. (s)	$N_{\text{break-even}}$
Uniform	51	–	–	2.71×10^{-6}	–
Cosine	42	–	–	2.73×10^{-6}	–
Sourcerer low	29	13	0.38	2.75×10^{-6}	6.43×10^{3}
Sourcerer high	14	28	50.1	2.75×10^{-6}	6.50×10^{5}

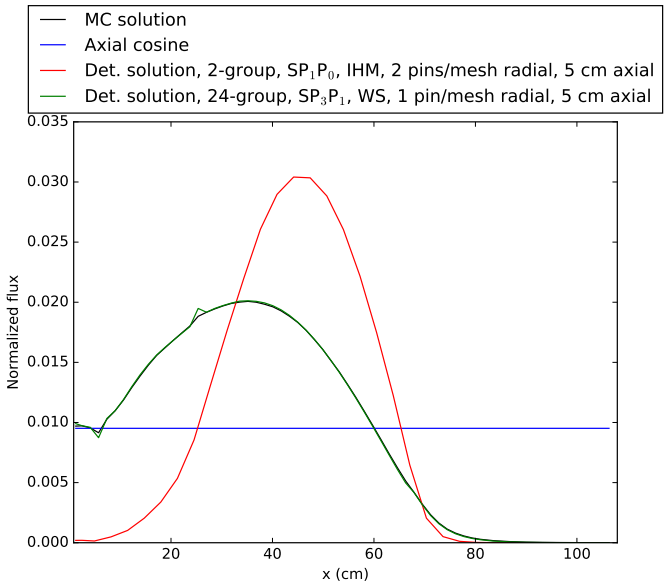

(a) Radial, y and z fuel midplanes

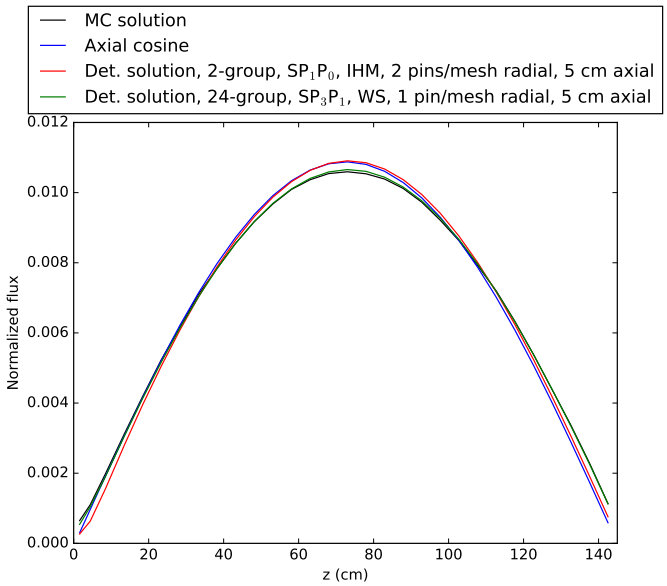

(b) Axial, centerline

Figure 8: Flux distributions used as initial guesses for MC power iteration, compared with the converged MC solution, for the B&W 1810 core 17 problem.

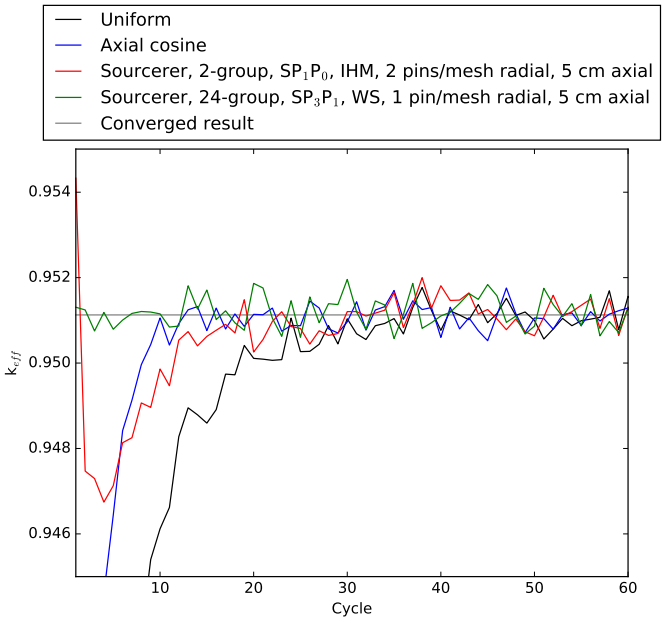

(a) k_{eff}

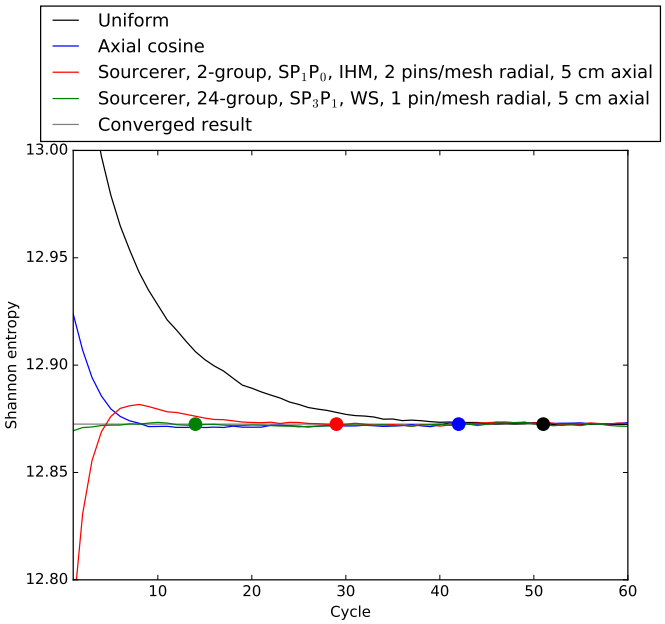

(b) Shannon entropy

Figure 9: Convergence plots for the B&W 1810 core 17 problem. Markers represent the converged cycles tabulated in Table 3.

5.3. Watts Bar Nuclear Unit 1

Early experimentation with the WBN1, AP1000®, and unrodded AP1000® problems revealed that low-resolution Denovo solutions were ineffective in improving fission source convergence. Therefore, the trials presented for these problems all used the 24-group energy structure with an SP_3 angular expansion and a P_1 scattering expansion. For this problem, three Sourcerer trials are compared with uniform and axial cosine initial flux distributions. The low-resolution trial used IHM cross sections and a solution mesh with 8 mesh cells per assembly radially (i.e., 2.125 pins per mesh cell) and 5 cm mesh divisions axially ($128 \times 128 \times 88 = 1{,}441{,}792$ mesh cells). The medium-resolution trial used IHM cross sections and a solution mesh with 1 pin per mesh radially and 5 cm mesh divisions axially ($263 \times 263 \times 88 = 6{,}086{,}872$ mesh cells). The high-resolution trial used WS cross sections generated automatically from a VERA model as described in Section 3.2. The solution mesh was generated automatically by the `vera2omn` utility. The mesh resolution was 1 pin per mesh radially, with additional mesh divisions within assembly gaps. Axially, the maximum mesh division size was 5 cm. The total mesh size was $305 \times 305 \times 115 = 10{,}697{,}875$ mesh cells.

The principle difference between the medium- and high-resolution trials is the use of WS cross sections. With the WBN1 and AP1000® models there are too many unique pin cells for it to be practical to generate WS cross sections by hand. Since automatic generation of WS cross sections can be done only on a per-pin basis using VERA, the effect of WS cross sections was investigated specifically with these trials to understand the expected performance for cases when a VERA model is not available. All trials were run with 800 inactive cycles and 200 active cycles with 4×10^7 histories per cycle on 32 compute nodes with 32 MPI tasks per node. Note that in production simulations, considerably more active cycles would typically be performed. However, since the aim of this work is to study the convergence of the fission source, it is the inactive cycles that are of primary importance.

Radial and axial flux distributions for the initial guesses are compared with the MC solutions in Figure 10. Unlike the C5G7 and B&W 1810 problems, the radial flux distribution is captured best by the low-resolution distribution, followed by the medium- and high-resolution distributions. It is not apparent why this is the case, but the low-resolution distribution could be coincidentally more accurate due to a cancellation of error effect. Axially, the medium- and high-resolution distributions closely match the MC solution, whereas the low-resolution solution is considerably less accurate.

For the low-resolution Sourcerer trial, the Shannon entropy does not visually appear to converge any faster than the axial cosine trial, as seen in Figure 11. As a result, the $N_{\text{break-even}}$ value for this trial in Table 4 is probably not reliable. Both the medium- and high-resolution trials visually appear to converge in significantly fewer cycles than the axial cosine as corroborated by the quantitative convergence results. This suggests that capturing this axial shape is more important for Sourcerer performance. The $N_{\text{break-even}}$ values for the medium- and high-resolution trials are both larger than the number of histories used for previous analyses of this problem. A Shift benchmark study that compared k_{eff}, control rod/bank worth, and boron worth to experimental results used 300 inactive cycles and 700 active cycles with 1×10^7 histories per cycle [14]. In other words, similar to the B&W 1810 problem, the use of the Sourcerer method is justified for this problem only if considerably more histories per cycle are required for highly resolved flux distributions.

Table 4: Convergence results (using the criteria described in Section 3.3) and $N_{\text{break-even}}$ values (as described in Section 3.4) for the WBN1 problem

Method	Conv. cycle	Cycles saved	Denovo wall time (s)	Shift wall time per hist. (s)	$N_{\text{break-even}}$
Uniform	223	–	–	2.84×10^{-7}	–
Cosine	292	–	–	2.72×10^{-7}	–
Sourcerer low	249	43	117	2.72×10^{-7}	1.00×10^7
Sourcerer medium	81	211	773	2.66×10^{-7}	1.38×10^7
Sourcerer high	63	230	1770	2.64×10^{-7}	2.91×10^7

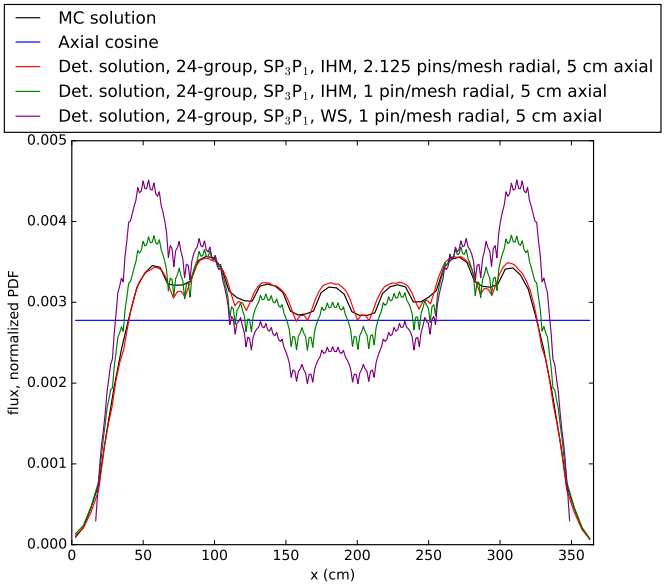

(a) Radial, y and z fuel midplanes

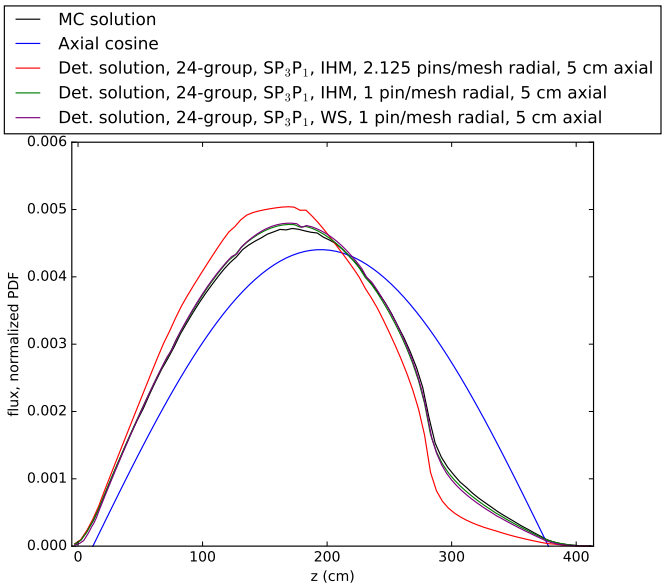

(b) Axial, centerline

Figure 10: Flux distributions used as initial guesses for MC power iteration, compared with the converged MC solution, for the WBN1 problem.

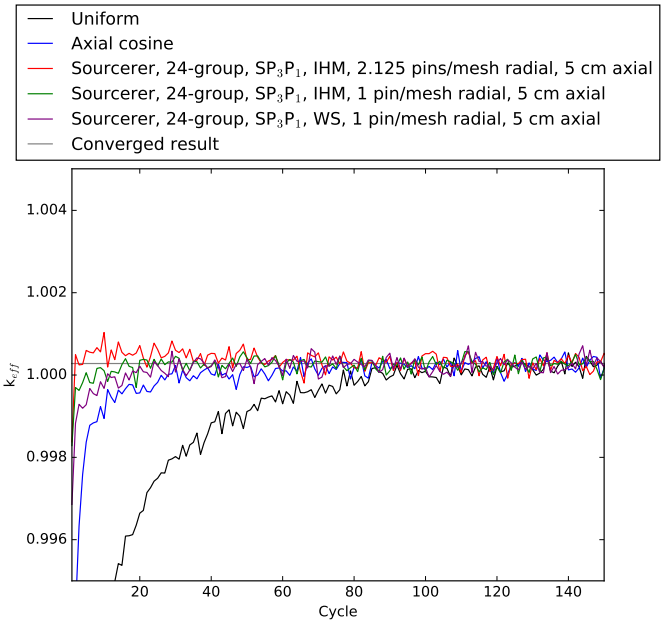

(a) k_{eff}

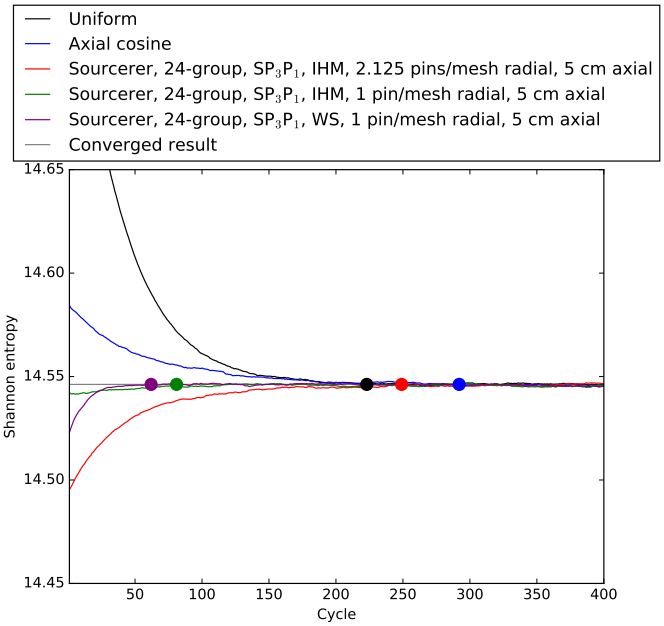

(b) Shannon entropy

Figure 11: Convergence plots for the WBN1 problem. Markers represent the converged cycles tabulated in Table 4. Note that the k_{eff} plot shows only the first 150 cycles.

5.4. Westinghouse AP1000®

For the AP1000® problem, a similar set of Sourcerer trials were used as with the WBN1 problem. All trials used a 24-group energy structure with an SP_3 angular expansion and a P_1 scattering expansion. The low-resolution trial used IHM cross sections with 8 mesh cells per assembly radially (2.125 pins per mesh radially) and 5 cm mesh divisions axially ($128 \times 128 \times 99 =$ 1,622,016 mesh cells). The medium-resolution trial used IHM cross sections and a solution mesh with 1 pin per mesh radially and 5 cm mesh divisions axially ($263 \times 263 \times 99 =$ 6,847,731 mesh cells). The high-resolution trial used WS cross sections generated from a VERA model. As was the case with the WBN1 problem, the generated mesh had an radial resolution of 1 pin per mesh cell with additional divisions in the assembly gaps and a maximum axial division size of 5 cm ($285 \times 285 \times 145 =$ 11,777,625 mesh cells). All trials were run with 1000 inactive cycles and 500 active cycles with 4×10^7 histories per cycle on 32 compute nodes with 32 MPI tasks per node. These trials are compared against uniform and axial flattened cosine initial guesses.

Radial and axial flux distributions for the initial guesses are compared against the MC solution in Figure 12. As was the case with the WBN1 problem, the higher-resolution trials do not improve the radial flux distribution. Unlike the WBN1 problem, none of the Sourcerer trials were able to closely match the axial flux distribution. As a result, none of the Sourcerer trials provided faster convergence compared with an axial flattened cosine or uniform initial guess as seen in Figure 13, with converged cycle values in Table 5.

The inability of the deterministic solution to capture the correct flux distribution might be because of the presence of IFBA pins in the AP1000®, as discussed in Section 4.4. The presence of a strong absorber directly surrounding the fuel pins significantly perturbs the local flux distribution. It is possible that the spatial and/or energy discretization are not fine enough to capture this effect, which in turn affects the flux distribution throughout the whole problem. The spatial/energy resolution could be further increased; however, the high-resolution solution for this problem required significantly more processor time than the high-resolution WBN1 trial, which had a high $N_{\text{break-even}}$ relative to the number of histories per cycle previously used for WBN1 analysis. In other words, a more detailed Denovo solution could potentially improve the performance of Sourcerer for this problem, but it is unlikely that this would be practical unless an extremely highly resolved flux distribution is required.

Table 5: Convergence results (using the criteria described in Section 3.3) and $N_{\text{break-even}}$ values (as described in Section 3.4) for the AP1000® problem

Method	Conv. cycle	Cycles saved	Denovo wall time (s)	Shift wall time per hist. (s)	$N_{\text{break-even}}$
Uniform	336	–	–	1.95×10^{-7}	–
Flattened cosine	422	–	–	2.02×10^{-7}	–
Sourcerer low	620	–	124	1.89×10^{-7}	–
Sourcerer medium	788	–	846	1.88×10^{-7}	–
Sourcerer high	616	–	2670	1.92×10^{-7}	–

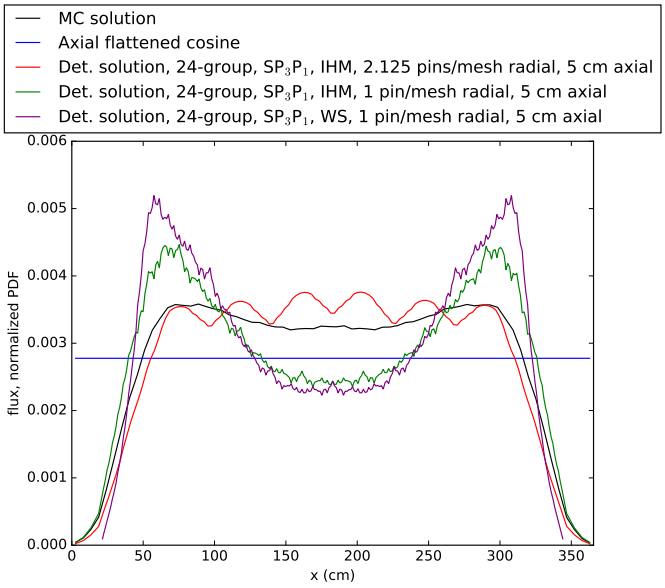

(a) Radial, y and z fuel midplanes

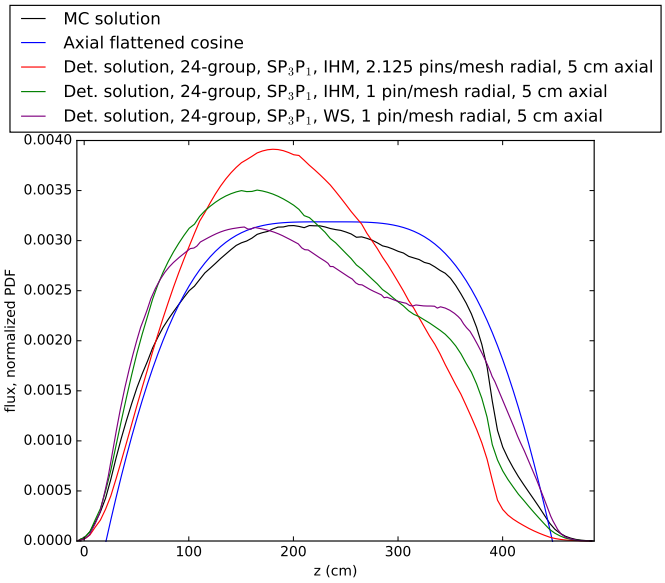

(b) Axial, centerline

Figure 12: Flux distributions used as initial guesses for MC power iteration, compared with the converged MC solution, for the AP1000® problem.

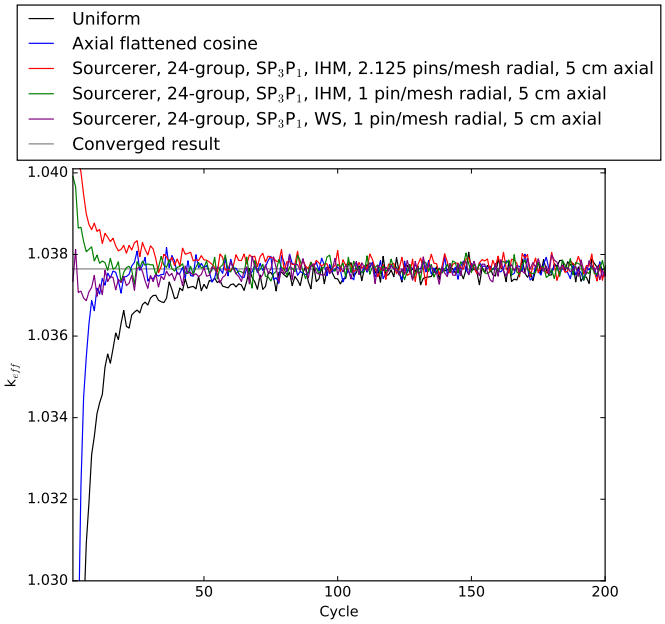

(a) k_{eff}

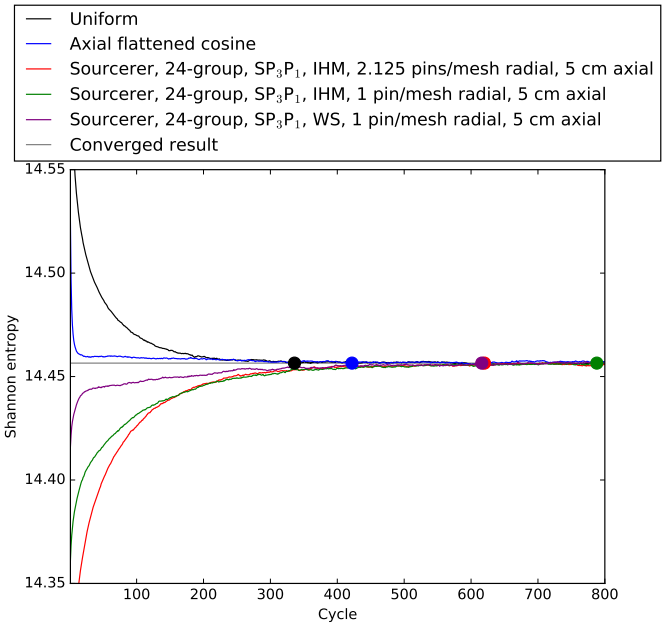

(b) Shannon entropy

Figure 13: Convergence plots for the AP1000® problem. Markers represent the converged cycles tabulated in Table 5. Note that the k_{eff} plot shows only the first 200 cycles.

5.5. Westinghouse AP1000®, Unrodded Configuration

As discussed in Section 4.4, this problem is identical to the standard AP1000® problem except that tungsten and Ag-In-Cd control rods have been removed. For this problem, the same three trials were run as the standard AP1000® problem. The only exception is that the mesh generated by `vera2omn` for the unrodded high-resolution trial ($285 \times 285 \times 143 = 11{,}615{,}175$ mesh cells) has two fewer axial divisions than the high-resolution trial in the standard AP1000® problem.

Radial and axial flux distributions for the initial guesses are compared against the MC solution in Figure 14. The radial flux distribution plot shows the same performance trend as the standard AP1000® problem. Axially, the low-resolution trial produces a flux distribution that resembles a cosine. By increasing the mesh resolution with the medium-resolution trial, the radial flux distribution becomes less accurate but the axial flux distribution improves. However, WS cross sections do not yield any further improvement, but rather the asymmetry of the axial flux distribution is overestimated.

As seen in Figure 15, qualitatively it appears that all trials converge in a similar number of cycles and that Sourcerer does not appear to provide any benefit. The order in which the cycles converge, tabulated in Table 6, is likely an artifact of the convergence criteria rather than a legitimate effect. As a result, the uncertainty in C_{saved} as given by Eq. 12 is high, meaning the $N_{\text{break-even}}$ value for the high-resolution trial is not reliable. As was the case with the standard AP1000® problem, a more detailed mesh or energy group energy structure might allow the Sourcerer method to outperform a uniform distribution; but it is unclear whether the Denovo solve time requirements would be justified for anything but Shift simulations which require an extremely large number of histories per cycle.

Table 6: Convergence results (using the criteria described in Section 3.3) and $N_{\text{break-even}}$ values (as described in Section 3.4) for the unrodded AP1000® problem

Method	Conv. cycle	Cycles saved	Denovo wall time (s)	Shift wall time per hist. (s)	$N_{\text{break-even}}$
Uniform	672	–	–	2.02×10^{-7}	–
Cosine	815	–	–	2.03×10^{-7}	–
Sourcerer low	763	–	177	1.91×10^{-7}	–
Sourcerer medium	688	–	1130	1.89×10^{-7}	–
Sourcerer high	642	30	2660	1.93×10^{-7}	4.59×10^{8}

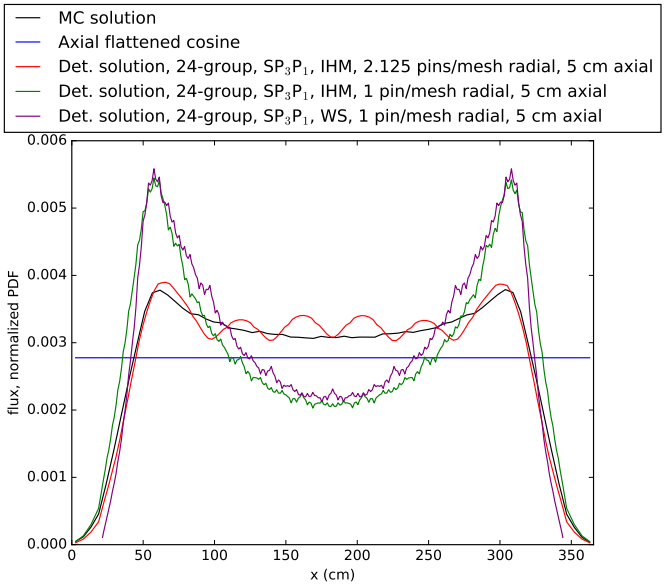

(a) Radial, y and z fuel midplanes

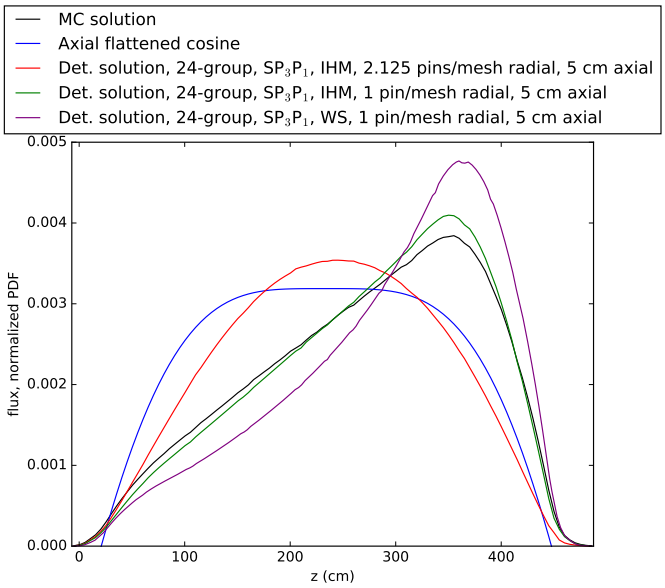

(b) Axial, centerline

Figure 14: Flux distributions used as initial guesses for MC power iteration, compared with the converged MC solution, for the unrodded AP1000® problem.

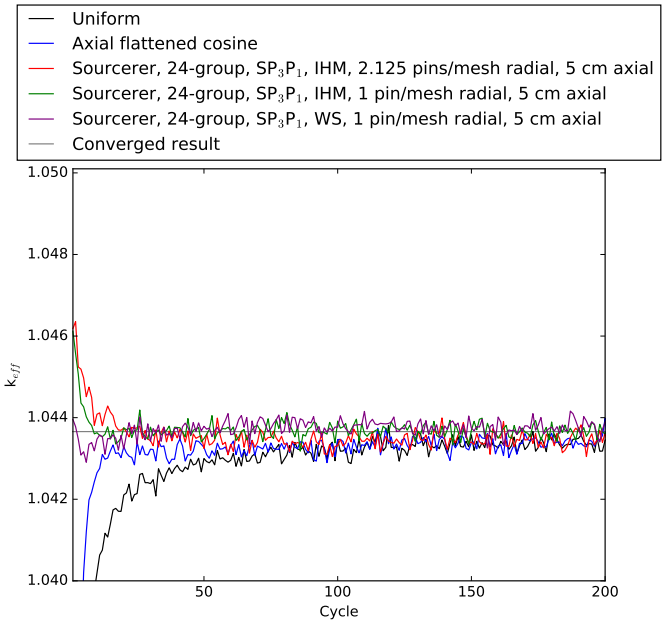

(a) k_{eff}

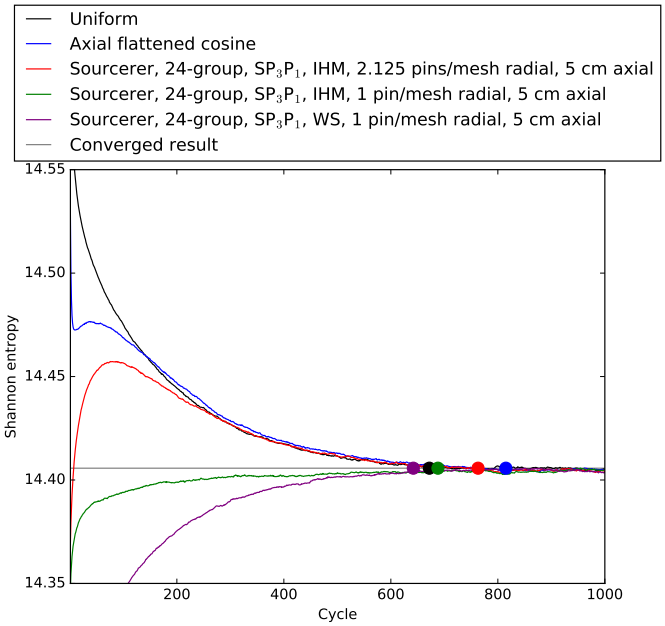

(b) Shannon entropy

Figure 15: Convergence plots for the unrodded AP1000® problem. Markers represent the converged cycles tabulated in Table 6. Note that the k_{eff} plot shows only the first 200 cycles.

5.6. NAC UMS® Cask

One limitation of the SP_N method is that it cannot be used in the presence of material voids. The NAC UMS® cask model has a void region surrounding the cylindrical reflector, as seen in Figure 5. As a result, the S_N method was used as the deterministic solver for this problem. In addition, a VERA model was not available, so only IHM cross sections were tested. Low- and high-resolution trials both used a step characteristics spatial differencing scheme, a quadruple range quadrature set with four polar and four azimuthal angles per level per octant (S_8), a P_1 scattering expansion, and an L_2 convergence tolerance of 1×10^{-3}. The low-resolution trial used a solution mesh with 5 cm mesh divisions both radially and axially ($54 \times 54 \times 109 = 317{,}844$ mesh cells). The high-resolution trial used a solution mesh with 1 pin per mesh element radially and 5 cm mesh divisions axially ($106 \times 106 \times 109 = 1{,}224{,}724$ mesh cells). All trials were run with 1,000 inactive cycles and 500 active cycles with 4×10^7 histories per cycle on 2 compute nodes with 32 MPI tasks per node. These trials are compared against a uniform initial guess.

As seen in Figure 16, the radial flux distribution has an asymmetric shape. This is expected because the loaded assemblies have different initial enrichments and discharge burnups. Axially, the flux is top-peaked. During reactor operation there is a higher moderator density at the bottom of the core due to the axial temperature distribution, which leads to lower burnup in the axial top portion of the fuel assemblies. The lower burnup fuel at the top of the assembly causes this asymmetry in the flux distribution.

The high-resolution trial captures the radial shape better than the low-resolution trial, but the opposite is true axially. As seen in Figure 17 and Table 7, the high-resolution trial converges significantly faster than the low-resolution and uniform trials. Unlike all of the previous problems, the trial with the most accurate axial shape did not perform the best. This is likely because of the Boral plates that form the flux traps between assemblies, as described in Section 4.5. These flux traps cause the assemblies to be significantly more neutronically decoupled. In other words, if the initial guess for the fission source is inaccurate radially, the power iteration process cannot correct this inaccuracy at the same rate as in a reactor problem.

The fact that both Sourcerer trials quantitatively and qualitatively offer improvements over a uniform initial guess demonstrates that sufficiently accurate Denovo solutions can be obtained over a range of solution parameters for this problem. However the $N_{\text{break-even}}$ values of 5.71×10^5 and 1.68×10^7 for the low- and high-resolution trials are greater than the number of histories

used for previous analysis. A production-level determination of k_{eff} for this problem used 150 inactive cycles and 400 active cycles with $5{\times}10^5$ histories per cycle [34]. In other words, the use of Sourcerer for this problem is justified only if more highly resolved flux distributions are required.

These results can also be compared with previous Sourcerer results from another NAC UMS® problem as reported by Ibrahim et al. [8]. Based off of published values from Ibrahim et al., $N_{\text{break-even}}$ was found to be 2.53×10^4, which is more than an order of magnitude less than the $N_{\text{break-even}}$ for the low-resolution case in this work. This discrepancy is from the sensitivity of $N_{\text{break-even}}$ to both the problem and combination of deterministic and MC codes used with the Sourcerer method. Though Ibrahim et al. used the same cask model as the one used in this work, it was loaded with different fuel and a different decay time was considered. In addition, Ibrahim et al. used Denovo and KENO, both run in serial. The Shift computational time per particle was found to be 2.74 times less in this work than the KENO computational time per particle calculated from results from Ibrahim et al. The Sourcerer method provides a greater benefit when the MC code requires more computational time per particle because the deterministic solve is comparatively less expensive. Finally, Ibrahim et al. used different convergence metrics and different Denovo solution parameters, which might contribute to the discrepancy in $N_{\text{break-even}}$.

Table 7: Convergence results (using the criteria described in Section 3.3) and $N_{\text{break-even}}$ values (as described in Section 3.4) for the NAC UMS® problem

Method	Conv. cycle	Cycles saved	Denovo Wall Time (s)	Shift wall time per hist. (s)	$N_{\text{break-even}}$
Uniform	159	–	–	4.96×10^{-7}	–
Low	114	45	12.7	4.94×10^{-7}	5.71×10^5
High	87	72	608	5.02×10^{-7}	1.68×10^7

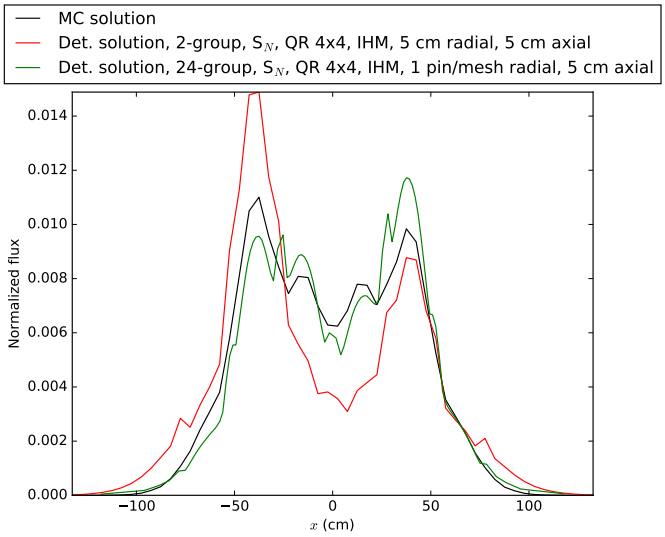

(a) Radial, y and z fuel midplanes

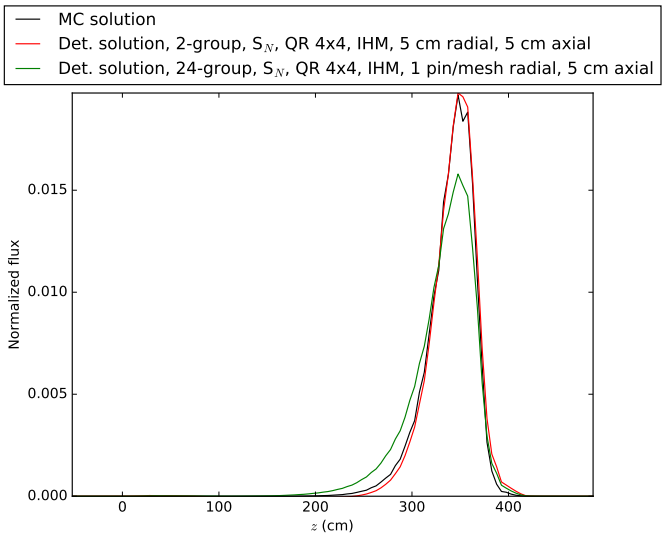

(b) Axial, centerline

Figure 16: Flux distributions used as initial guesses for MC power iteration, compared with the converged MC solution, for the NAC UMS® cask problem.

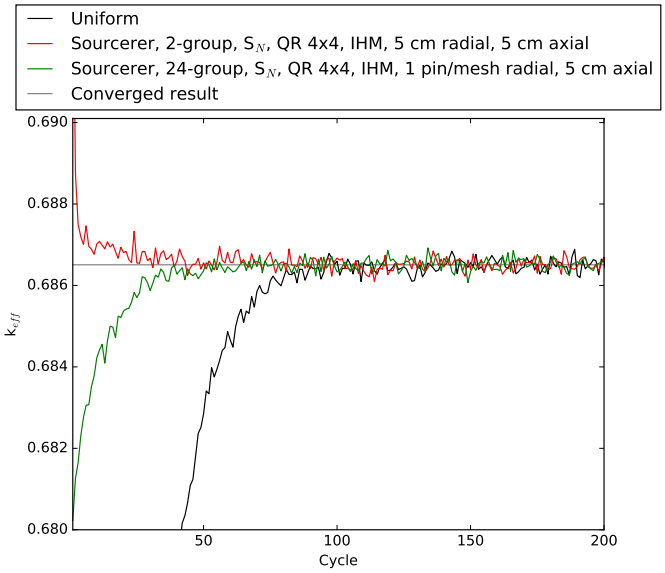

(a) k_{eff}

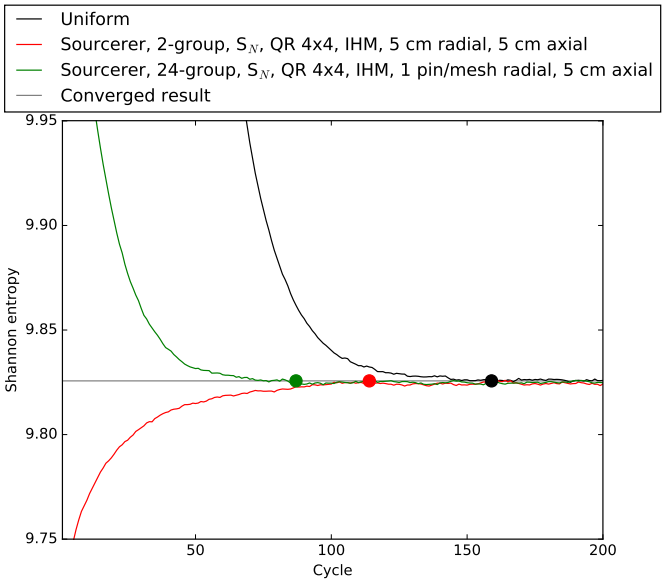

(b) Shannon entropy

Figure 17: Convergence plots for the NAC UMS® cask problem. Markers represent the converged cycles tabulated in Table 7.

6. Conclusion

In this work the Sourcerer method has been implemented in the Exnihilo transport suite within SCALE using Denovo's SP_N and S_N solvers and the Shift MC transport code. An *ad hoc* convergence metric was used to determine the number of cycles required for convergence with and without Sourcerer. Provided that the Sourcerer method is successful in reducing the number of cycles required for convergence, Denovo and Shift timing results can be used to calculate $N_{\text{break-even}}$ to estimate the number of particles per cycle for which the use of the Sourcerer method is justified.

Convergence and $N_{\text{break-even}}$ results were presented for a variety of typical k-eigenvalue problems for Sourcerer trials with low- and high-resolution Denovo solutions, as well as uniform and axial cosine or axial flattened cosine initial flux distributions. These Denovo solutions varied the SP_N or S_N order, resonance self-shielding treatment for cross section generation, number of energy groups, and solution mesh. For the C5G7 problem, both low- and high-resolution Denovo solutions resulted in convergence in significantly fewer cycles, with $N_{\text{break-even}}$ values that were small compared with the number of histories used in previous analysis. This indicates the Sourcerer method can reliably reduce the number of cycles required for convergence for this problem, insensitive to the choice of Denovo solution parameters.

For the B&W 1810 core 17, WBN1, and NAC UMS® cask problems, the higher-resolution Denovo solutions yielded the best performance, though all of these trials had $N_{\text{break-even}}$ values close to or exceeding the number of histories per cycle used for analyses appearing in the literature. In these cases, the use of the Sourcerer method is justified only if a large number of histories per cycle are required for more highly resolved flux distributions. For the standard and unrodded AP1000® problems, none of the Sourcerer trials reduced the number of cycles required for convergence relative to an axial flattened cosine initial flux distribution. Unlike the other reactor cases considered in this work, even the high-resolution Denovo solution did not capture the correct axial flux distribution, which appears to be necessary for the Sourcerer method to be effective for reactor problems. This could be caused by IFBA pins in the AP1000®, but further analysis is needed to confirm this possibility.

In several cases, the convergence results using the convergence criteria presented in Section 3.3 were inconsistent with visual expectations. Future work might involve developing improved convergence criteria that take into

account the autocorrelation of the Shannon entropy. The Sourcerer method as implemented in Exnihilo requires a marginal amount of additional effort from the user compared with a standard Shift k-eigenvalue simulation. Assuming that the number of histories per cycle is fixed, the use of the Sourcerer method appears to be justified in cases where highly resolved flux distributions are desired and therefore a large number of histories per cycle are used. Further advancements in the speed and accuracy of deterministic methods are necessary for Sourcerer to be advantageous for arbitrary use cases.

Acknowledgments

This work was funded in part by the US Nuclear Regulatory Commission (NRC) Office of Nuclear Regulatory Research under Task Order NRCHQ6015T0026 "Shift – Integration of SCALE Nuclear Stochastic Methods," under NRC Agreement No. NRCHQ6014D0005 "SCALE Support for Reactor and Spent Fuel Analyses" (DOE Interagency Agreement No. 1886-V779-14). Work for this paper was supported by Oak Ridge National Laboratory, which is managed and operated by UT-Battelle, LLC, for the US Department of Energy under Contract No. DEAC05-00OR22725. This research used resources of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the US Department of Energy under Contract No. DE-AC05-00OR22725.

References

[1] T. Urbatsch, Iterative acceleration methods for Monte Carlo and deterministic criticality calculations, Ph.D. thesis, University of Michigan, Ann Arbor, MI (1995).

[2] S. Carney, F. Brown, B. Kiedrowski, W. Martin, Theory and applications of the fission matrix method for continuous-energy Monte Carlo, Annals of Nuclear Energy 73 (2014) 423–431.

[3] S. Hamilton, T. Evans, Deterministic fission matrix acceleration of Monte Carlo calculations, Transactions of the American Nuclear Society 113 (2015) 649–651.

[4] S. Hamilton, G. Davidson, T. Evans, K. Banerjee, Accelerated Monte Carlo fission source convergence with fission matrix and kernel density

estimators, Transactions of the American Nuclear Society 114 (2016) 385–387.

[5] A. L. Lund, P. K. Romano, A. R. Siegel, Accelerating source convergence in Monte Carlo criticality calculation using a particle ramp-up technique, in: M&C 2017 - International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering, Jeju, Korea, 2017.

[6] M. Nowak, J. Miao, E. Dumonteil, B. Forget, A. Onillon, K. S. Smith, A. Zoia, Monte Carlo power iteration: Entropy and spatial correlations, Annals of Nuclear Energy 94 (2016) 856 – 868.

[7] T. Yamamoto, Y. Miyoshi, Reliable method for fission source convergence of Monte Carlo criticality calculation with Wielandt's method, Journal of Nuclear Science and Technology 41 (2) (2004) 99–107.

[8] A. M. Ibrahim, D. E. Peplow, K. B. Bekar, C. Celik, J. M. Scaglione, D. Ilas, J. C. Wagner, Hybrid technique in SCALE for fission source convergence applied to used nuclear fuel analysis, in: Topical Meeting held by the ANS Nuclear Criticality Safety Division (NCSD 2013), Wilmington, NC, USA, 2013.

[9] B. Rearden, M. A. Jessee, Eds., SCALE code system, version 6.2, Tech. Rep. ORNL/TM-2005/39, Oak Ridge National Laboratory, Oak Ridge, TN (2016).

[10] T. M. Evans, A. S. Stafford, R. N. Slaybaugh, K. T. Clarno, Denovo: A new three-dimensional parallel discrete ordinates code in SCALE, Nuclear Technology 171 (2010) 171–200.

[11] J. Jarrell, T. Evans, G. Davidson, A. Godfry, Full core reactor analysis: running Denovo on Jaguar, Nuclear Science and Engineering 175 (3) (2013) 283–291.

[12] S. P. Hamilton, T. M. Evans, Efficient solution of the simplifed P_N equations, Journal of Computational Physics 284 (2015) 155–170.

[13] S. Johnson, T. Evans, G. Davidson, S. Hamilton, T. Pandya, Exnihilo documentation, Tech. Rep. Release 6.2.0 (Dev), Oak Ridge National Laboratory (2017).

[14] T. M. Pandya, S. R. Johnson, T. M. Evans, G. G. Davidson, S. P. Hamilton, A. T. Godfrey, Implementation, capabilities, and benchmarking of Shift, a massively parallel Monte Carlo radiation transport code, Journal of Computational Physics 308 (2016) 239–272.

[15] S. R. Johnson, Omnibus: a new front end to Denovo and Shift, Transactions of the American Nuclear Society 117.

[16] T. Ueki, F. B. Brown, Stationary modeling and informatics-based diagnostics in Monte Carlo criticality calculations, Nuclear Science and Engineering 149 (1) (2005) 38–50.

[17] M. A. Smith, E. E. Lewis, B. C. Na, Benchmark on deterministic transport calculations without spatial homogenization — a 2-d/3-d MOX fuel assembly benchmark (C5G7 MOX benchmark), Tech. Rep. NEA/NSC/DOC(2003)16, OECD/NEA (2003).

[18] L. W. Newman, W. A. Wittkopf, W. G. Pettus, M. N. Baldwin, H. A. Hassan, Urania-Gadolinia: Nuclear model development and critical experiment benchmark, Tech. Rep. DOE/ET/34212-41 (BAW-1810), Babcock and Wilcox Co. (1984).

[19] J. Gehin, A. Godfrey, F. Franceschini, T. Evans, B. Collins, S. Hamilton, Operational reactor model demonstration with VERA: Watts Bar unit 1 cycle 1 zero power physics tests, Tech. Rep. CASL-U-2013-0105-001, Consortium for Advanced Simulation of LWRs (2013).

[20] F. Franceschini, A. Godfrey, J. Kulesza, R. Oelrich, Westinghouse VERA test stand: zero power physics test simulations for the AP1000 PWR, Tech. Rep. CASL-U-2014-0012-000, CASL (2014).

[21] J. Peterson, et al., Used Nuclear Fuel Storage Transportation & Disposal Analysis Resource and Data System (UNF-ST&DARDS), Tech. Rep. FCRD-NFST-2013-000117 Rev. 0, U.S. Department of Energy Nuclear Fuels Storage and Transportation Planning Project (March 2013).

[22] F. Brown, A review of best practices for Monte Carlo criticality calculations, in: American Nuclear Society Nuclear Criticality Safety Topical Meeting, Richland, WA, 2009.

[23] M. Wenner, A. Haghighat, Study of methods of stationarity detection for Monte Carlo criticality analysis with kenov.a, Transactions of the American Nuclear Society 97 (2007) 647–650.

[24] F. Brown, B. Nease, J. Cheatham, Convergence testing for MCNP5 Monte Carlo eigenvalue calculations, in: Joint International Topical Meeting on Mathematics & Computation and Supercomputing in Nuclear Applications (M&C + SNA 2007), Monterey, California, 2007.

[25] G. I. Bell, S. Glasstone, Nuclear Reactor Theory, Van Nostrand Reinhold Company, New York, NY, 1970.

[26] K. Azekura, H. Maruyama, T. Ikehara, M. Yamamoto, V. Mills, T. Marcille, Development of a BWR lattice analysis code LANCR based on an improved CCCP method, in: Advances in Nuclear Fuel Management III, Hilton Head Island, South Carolina, 2003.

[27] M. Williams, Private Communication (January 2017).

[28] S. P. Hamilton, T. M. Evans, G. G. Davidson, S. R. Johnson, T. M. Pandya, A. T. Godfrey, Hot zero power reactor calculations using the Insilico code, Journal of Computational Physics 314 (Supplement C) (2016) 700 – 711.

[29] R. Montgomery, VERA tools and workflows, Technical Report CASL-U-2014-0054-001, Oak Ridge National Laboratory (March 2014).

[30] CASL: Improving reactor performance with predictive science-based simulation technology that harnesses world-class computational power, Tech. Rep. CASL-U-2014-0089-000, Oak Ridge National Laboratory (2014).

[31] T. Pandya, T. Evans, G. Davidson, S. Johnson, A. Godfrey, Shift verification and validation, Tech. Rep. CASL-U-2016-1186-000, Consortium for Advanced Simulation of LWRs (2016).

[32] S. P. Hamilton, T. M. Evans, Efficient solution of the simplified P_N equations, Tech. Rep. CASL-U-2014-0352-000, Consortium for Advanced Simulation of LWRs (2014).

[33] S. Cathalau, J. Lefebvre, J. West, Proposal for a second stage of the benchmark on power distributions within assemblies, Tech. Rep. NEA/NSC/DOC(96)2 Revision 2, OECD/NEA (April 1996).

[34] R. Lefebvre, Private Communication (November 2017).

Appendix A. Energy Group Bounds for the 24-Group Structure

The 24-group structure is based on the 23-group LANCR structure, with an additional group bound at 5 eV. This structure is shown in Table A.8.

Table A.8: Energy group bounds for the 24-group structure

Energy group bound (eV)
2.000×10^{7}
8.208×10^{5}
1.111×10^{5}
5.531×10^{3}
1.864×10^{2}
3.761×10^{1}
3.538×10^{1}
2.770×10^{1}
2.168×10^{1}
2.040×10^{1}
1.597×10^{1}
7.150×10^{0}
6.700×10^{0}
6.300×10^{0}
5.000×10^{0}
1.097×10^{0}
1.045×10^{0}
9.500×10^{-1}
3.500×10^{-1}
2.060×10^{-1}
1.070×10^{-1}
5.800×10^{-2}
2.500×10^{-2}
1.000×10^{-2}
1.000×10^{-5}