*M&C 2017 - International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering, Jeju, Korea, April 16-20, 2017, on USB (2017)*

# Accuracy and Performance of 3D MOC for Full-Core PWR Problems

Geoffrey Gunow, Samuel Shaner, William Boyd, Benoit Forget, and Kord Smith

*Massachusetts Institute of Technology*
*geogunow@mit.edu; shaner@mit.edu; wboyd@mit.edu bforget@mit.edu; kord@mit.edu*

**Abstract** - *The Method of Characteristics (MOC) has seen wide interest in reactor physics because of its accuracy and efficiency in computing lattice physics problems. While most of its use has been in solving 2D problems, there has been recent interest in extending MOC to 3D in order to more accurately calculate 3D power distributions in LWRs. While the method is naturally extensible to 3D, it presents significant computational difficulties. In this study we present PWR results from OpenMOC which mitigate the computational difficulties of 3D MOC by using domain decomposition, efficient track generation, axially extruded ray tracing, CMFD acceleration, and a linear source approximation. The BEAVRS benchmark is analyzed using OpenMC to generate multi-group cross-sections from Monte Carlo simulations for OpenMOC. First MOC and CMFD parameter studies are conducted on a single assembly to determine the MOC and CMFD parameters. Using optimal parameters to reduce computational burden and maintain solution accuracy, 3D MOC results are presented for the full core simulation of the BEAVRS benchmark.*

## I. INTRODUCTION

A common goal in multi-group deterministic neutron transport is to determine the spatial distribution of the neutron scalar flux for each energy group throughout the reactor core. This distribution can then be used to calculate reaction rates necessary for thermal analysis. There has recently been a shift to replacing the legacy diffusion theory based reactor physics codes with modern codes that more explicitly model the physics and yield higher accuracy solutions. The Method of Characteristics (MOC) has long been used in lattice physics analysis for cross section generation. Full-core analysis using MOC has typically followed a 2D/1D approach due to the computational costs of explicitly solving the full 3D MOC problem [1, 2, 3, 4]. This approach solves a series of 2D MOC problems in the *xy*-plane at different axial heights throughout the core and then couples these solutions together with a low-order diffusion or transport operator in the axial direction. This approach has yielded great success for many full-core problems, however such methods cannot be refined axially to produce the true solution to the 3D transport equation.

Recently OpenMOC [5], an open-source neutron transport code developed at MIT, has been extended to support direct 3D MOC calculations. This solver has been shown to be computationally competitive with 2D/1D methods in small 3D reactor problems while maintaining the rigorous accuracy and flexibility of a 3D MOC solver [6]. This is accomplished by using axially extruded ray tracing and a 3D linear source approximation [7, 8]. In this study, we extend this work to enable spatially-converged 3D MOC on a full-core PWR geometry.

We begin by detailing the memory and compute challenges of solving a full-core PWR problem in 3D MOC. In order to make the problem tractable on modern supercomputers, spatial domain decomposition (SDD), efficient track generation, coarse mesh finite difference (CMFD) acceleration with axial prolongation of CMFD fluxes, and higher order sources are essential. Our approaches to these aspects are highlighted with special attention to how they reduce the memory and compute challenges of solving large problems.

Even with a scalable and efficient solver, the accuracy of the results can be limited by the input cross sections. In this work, we used transport-corrected P0 cross sections generated with 3D full-core Monte Carlo simulations of the Benchmark for Evaluation and Validation of Reactor Simulations (BEAVRS) PWR model. Given the limited work in applying 3D MOC to large PWR problems, the exact spatial discretization of source regions, angular quadrature, and CMFD group structure required to efficiently converge a large PWR problem is still an open question. We present results on a full axial height 3D assembly to understand the effect of these parameters on solution convergence. Finally, results are presented for the full-core BEAVRS benchmark with 70-group cross sections.

## II. MEMORY AND COMPUTE CHALLENGES OF 3D MOC

The primary road blocks to performing 3D MOC with a purely shared-memory parallel design is the compute and memory requirements to solve large problems. Recent work has investigated these challenges using an MOC "mini-app" which mimics the computational performance of a 3D MOC solver but is unencumbered by the physics challenges of modeling a real geometry [9]. While using a slightly more refined track laydown than is used in this study, the previous results demonstrate the need for spatial domain decomposition to distribute the workload. Using 5,780 nodes to model the full BEAVRS benchmark geometry, each node contained 12.5 GB of boundary flux data, just below the 16 GB of on-node memory for the IBM Blue Gene/Q Supercomputer *Mira* at Argonne National Laboratory that the previous results and the full-core BEAVRS results in this study were obtained from. The results also highlighted that at this level of domain decomposition, communication of angular fluxes between neighboring nodes accounted for only 2.3% of the runtime for a wide variety of test cases, indicating that the 3D MOC algorithm can efficiently map to many node systems without concern of significant communication overhead.

## III. SPATIAL DOMAIN DECOMPOSITION

In order to reasonably solve problems larger than 3D C5G7, some decomposition of the problem is required. To accomplish this, domain decomposition was implemented in OpenMOC with the Spatial Domain Decomposition procedure outlined by Kelley *et. al.* except angular fluxes on domain boundaries were not updated during each transport sweep [10]. In this scheme, a Cartesian domain decomposition is implemented with modular domains. On each domain, the same track laydown is created using Modular Ray Tracing (MRT). Each domain operates independently of the other domains during a transport sweep, computing the change in angular fluxes across tracks and the local scalar fluxes in source regions.

Most of the communication is between neighbor domains during the transfer of angular fluxes. This is done in a bulk asynchronous communication at the end of each transport sweep with MPI. Track connections are obtained by computing the periodic track connections at interfaces between domains. Since tracks were generated with MRT, the periodic track connections yield the appropriate track index in the connecting domain.

Additionally there is communication between all nodes to compute global quantities such as the total fission source and $k_{eff}$. However, this communication cost is trivial in comparison to the communication cost required to transfer angular fluxes between domains.

## IV. EFFICIENT TRACK GENERATION

The memory requirements scale with the number of tracks and the compute requirements scale with the number of segments (which scales with the number of tracks). Clearly, there are incentives to minimize the number of tracks while not giving up on solution accuracy. When considering track generation methods it is important to keep in mind the physics aspects of the problem being solved. The material properties and flux in a PWR change gradually in the axial direction and sharply in the radial direction. This suggests that the polar (predominantly axial) spacing can be much coarser than the azimuthal (radial) spacing. In previous studies for the C5G7 3D Rodded B case we have shown that solution accuracy can be maintained with a polar $z$-spacing 30× larger than the azimuthal spacing [6]. Other track generation approaches require the polar spacing to be less than or equal to the azimuthal track spacing, resulting in far more tracks than are necessary to capture the relevant physics effects [11]. Furthermore, tracks with shallow polar angles capture more of the radial flux at different axial locations while steep polar angles capture more of the axial flux detail at different radial locations. This suggests that shallow polar angles can have a larger spacing than steep polar angles. These observations have led us to adopt the algorithm for modular track generation outlined in Algorithm 1 for angles in the first octant. The angles in other octants are computed using their complement in the first octant. Implementation details such as determining the track start and end points and connecting track indices are integral to the track generation process, but their description would require several pages and were thus not included in this paper.

Applying Algorithm 1 to the BEAVRS geometry with a spatial domain size of [21.42 cm, 21.42 cm, 20.0 cm] ($\Delta_x/D_x$, $\Delta_y/D_y$, $\Delta_z/D_z$), 32 azimuthal angles ($n_\phi$), 0.1 cm azimuthal spacing ($\delta_\phi$), 12 polar angles with Gauss-Legendre polar quadrature ($n_\theta$), and 1.5 cm polar $z$-spacing ($\delta_z$) yields the track parameters in Table I. In describing the spatial domains, $\Delta_{x/y/x}$ and $D_{x/y/z}$ refer to the geometry width and the number of domains in each dimension, respectively, so that $\Delta_{x/y/x}/D_{x/y/z}$ is the domain width in each dimension. Table I highlights that the track generation scheme can handle high ratios of the polar to azimuthal spacing with sufficiently low corrections to the quadrature angles. The relatively high max polar $z$-spacing correction is the result of the shallowest polar angle having very few crossings with the $xy$-plane and is not a cause for concern. The parameters also demonstrate the increasing polar spacing, $\delta_{\theta^{i,j}}$, as the polar angle gets shallower.

## V. AXIALLY EXTRUDED RAY TRACING

In addition to efficient track generation, OpenMOC also reduces computational cost by employing axially extruded ray tracing [7]. This method transforms the geometry into an axially extruded geometry over each domain in which the superposition of all radial detail is contained at every axial level. At the beginning of the simulation all axial levels are ray traced in the $xy$ plane. The calculated intersections implicitly form the superposition of all radial detail and are stored for use in the transport sweep. Since there are far fewer 2D tracks than full 3D tracks, the storage requirements are trivial in comparison with the the storage of angular fluxes for each 3D track. In each radial region, the geometry is mapped vertically to form an axial mesh for the radial region. During the transport sweep, the stored 2D intersections and axial mesh are used to calculate 3D intersections on-the-fly. This allows for dramatic reduction in storage requirements accompanied by a modest reduction in ray tracing time.

## VI. CMFD ACCELERATION

CMFD acceleration was implemented in OpenMOC as described by Smith [12] with damping on the current correction term. Of note, our CMFD solver incorporates a general group-condensation procedure for running CMFD with a reduced group structure, as illustrated in Figure 1. We have previously shown for the 2D C5G7 benchmark problem, rapid convergence can be achieved with a 2-group CMFD formulation [13]. Some have taken a multilevel CMFD approach of using energy multi-grid acceleration of the CMFD solve whereby a two-group CMFD formulation is used to accelerate a fine-group CMFD solve [2]. In this work, we explore several group structures for CMFD acceleration with 70-group cross sections for MOC to determine a group structure that has a good balance of rapid acceleration of the MOC solve while not inducing significant computational burden on the overall solve.

One other aspect of CMFD that has only recently been discussed in literature is alternative prolongation schemes for CMFD whereby the CMFD flux or update ratios from multiple cells are combined to form a source-region specific

**Algorithm 1** 3D modular track generation algorithm implemented in OpenMOC

User specifies $n_\phi$, $\boldsymbol{\delta_\phi}$, $n_\theta$, $\boldsymbol{\delta_z}$, $D_x$, $D_y$, $D_z$, and polar quadrature set.

**for all** $i \in n_\phi/4$ **do** {Loop over azimuthal angles in first octant}

$$\phi^i = \frac{2\pi * (i + 0.5)}{n_\phi}$$ Compute equally spaced $\phi^i$

$$n_x^i = \left\lceil \frac{\Delta_x * \sin(\phi^i)}{D_x * \delta_\phi} \right\rceil \quad n_y^i = \left\lceil \frac{\Delta_y * \cos(\phi^i)}{D_y * \delta_\phi} \right\rceil$$ Compute the # of $x$ and $y$ intersections, $n_x^i$ and $n_y^i$

$$\phi^i = \text{atan}\left(\frac{\Delta_y * D_x * n_x^i}{\Delta_x * D_y * n_y^i}\right)$$ Correct $\phi^i$ with $n_x^i$ and $n_y^i$

$$\delta_{\phi^i} = \frac{\Delta_x * \sin(\phi^i)}{D_x * n_x^i}$$ Correct $\delta_{\phi^i}$ with $n_x^i$ and $\phi^i$

**for all** $j \in n_\theta/2$ **do** {Loop over polar angles in first octant}

$\theta^{i,j}$ Obtain $\theta^{i,j}$ from desired polar quadrature set

$$n_{xy}^{i,j} = \left\lceil \frac{\Delta_y * \tan(\frac{\pi}{2} - \theta^{i,j})}{D_y * \sin(\phi^i) * \delta_z} \right\rceil$$ Compute the # of $xy$ intersections, $n_{xy}^{i,j}$

$$n_z^{i,j} = \left\lceil \frac{\Delta_z * D_y * \sin(\phi^i) * n_{xy}^{i,j} * \tan(\theta^{i,j})}{D_z * \Delta_y} \right\rceil$$ Compute the # of $z$ intersections, $n_z^{i,j}$

$$\theta^{i,j} = \text{atan}\left(\frac{\Delta_y * D_z * n_z^{i,j}}{\Delta_z * D_y * n_{xy}^{i,j} * \sin(\phi^i)}\right)$$ Correct $\theta^{i,j}$ with $n_{xy}^{i,j}$ and $n_z^{i,j}$

$$\delta_{\theta^{i,j}} = \frac{\Delta_z * \sin(\theta^{i,j})}{D_z * n_z^{i,j}}$$ Correct $\delta_{\theta^{i,j}}$ with $n_z^{i,j}$ and $\theta^{i,j}$

**end for**

**end for**

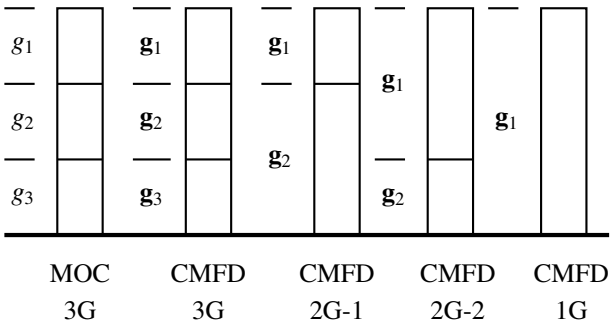


Fig. 1. Illustration of possible CMFD energy group structures for a three group MOC calculation.

prolongation value. In 2015 Li proposed a 2/3–1/3 scheme for updating sectorized pin cell regions in 2D MOC with a unique prolongation value [14]. We have further extended this approach to a general form that uses the distance-weighted flux ratio from the k-nearest neighbor (KNN) CMFD cells to a source region to update the MOC flux [15]. In this work, we use the KNN scheme in the radial direction and investigate a quadratic flux prolongation scheme in the axial direction where the CMFD cell flux between three neighboring axially-stacked CMFD cells is fit with a quadratic function and the interpolated flux at the source region centroid location is used in forming the update ratio.

## VII. LINEAR SOURCE APPROXIMATION

In most MOC implementations a flat source approximation is used. This stems from the rigorous definition of the angular neutron flux $\psi$ for a particular energy group through a region with constant material properties and cross sections as

$$\frac{d\psi(s,\Omega)}{ds} + \Sigma_t \psi(s,\Omega) = q(s,\Omega) \tag{1}$$

where

$\Omega$ = The direction

$s$ = The distance traversed along direction $\Omega$

$\Sigma_t$ = The transport cross-section

$q$ = The neutron source.

The usual flat source approximation assumes that the neutron source $q(s,\Omega)$ is constant over the domain. Alternatively, a higher order approximation could be used.

| **Parameter** | **Value(s)** |
|---|---|
| average $\delta_{\theta^{i,j}}$ by polar angle (cm) | [0.29, 0.63, 0.92, 1.16, 1.30, 1.27] |
| tracks per polar angle (%) | [46, 21, 13, 9, 6, 5] |
| number of tracks | 14,956 |
| average azimuthal angle correction (%) | 0.27 |
| max azimuthal angle correction (%) | 1.23 |
| average azimuthal spacing correction (%) | 0.18 |
| max azimuthal spacing correction (%) | 0.39 |
| average polar angle correction (%) | 2.16 |
| max polar angle correction (%) | 4.85 |
| average polar $z$-spacing correction (%) | 7.10 |
| max polar $z$-spacing correction (%) | 33.33 |

TABLE I. Track laydown parameters for first octant tracks in the BEAVRS geometry. The average $\delta_{\theta^{i,j}}$ by polar angle and tracks by polar angle are arranged from the steepest to the shallowest polar angle.

In OpenMOC, linear sources were implemented using track-based spatial moments over source regions to obtain linear source expansion coefficients, as described in [8]. In this formulation the neutron source $q(s, \Omega)$ is approximated as

$$q(s, \Omega) = \bar{q}(\Omega) + \hat{q}(\Omega)\,(s - s_c) \tag{2}$$

where $\bar{q}(\Omega)$ and $\hat{q}(\Omega)$ are pre-calculated linear source expansion coefficients and $s_c = s_T/2$ where $s_T$ is the total track length through the region.

The methodology was presented for 2D MOC. It showed that the converged solution to the 2D C5G7 benchmark requires an order of magnitude less memory and 5× less computational time. Using a linear source approximation greatly reduces the required computation to converge a problem by allowing for a much coarser source region discretization. While this methodology was presented for 2D MOC, it can easily be extended to three dimensions [16] showing similar improvements over flat sources for the 3D C5G7 Rodded-B benchmark [6].

## VIII. THE BEAVRS BENCHMARK

The BEAVRS benchmark is a highly-detailed, full-core PWR geometry that is intended for validation of high-fidelity reactor analysis codes [17]. The BEAVRS geometry provides an excellent test bed for new reactor analysis codes and has seen widespread adoption since its introduction in 2013 as a problem that encompasses many of the challenges present in industrial analysis. In this study, we tried to adhere to the benchmark specification as close as possible while making only a few minor modifications. These modifications include:

- The radial extent was truncated at one assembly-width outside the core fuel assemblies, resulting in a $17 \times 17$ lattice of fuel assembly-sized regions.
- The lower axial extent was truncated at the bottom support plate and the upper axial extent was truncated at the top of the fuel rods in order to make the geometry 400 cm and evenly fit into an evenly spaced domain decomposition mesh.
- The bottom and top spacers (Grids 1 and 8) were altered from 3.3579 cm to 2.0 cm in height and the middle spacers (Grids 2-7) were altered from 5.715 cm to 6.0 cm in height such that an integral number of evenly spaced source regions can be used in the spacer regions.

## IX. MULTI-GROUP CROSS-SECTION GENERATION

This work used the OpenMC Monte Carlo code [18] to generate multi-group cross sections (MGXS) for the BEAVRS model. In particular, the `openmc.mgxs` module was employed to generate XML tally input files, as well as to process the tally data output by OpenMC, to compute MGXS for use by OpenMOC. The `openmc.mgxs` module leverages OpenMC's fully-featured Python Application Programming Interface (API) [19] and is presented in detail in [20].

The MGXS were generated from a single full-core Monte Carlo simulation of the identical BEAVRS geometry modeled by OpenMOC. The single-step Monte Carlo-based approach to MGXS generation used by this analysis is thoroughly described and rigorously validated in [20]. The cross sections were tallied for each unique material, including each of the three fuel enrichments (1.6%, 2.4% and 3.1%), zircaloy cladding, helium gap, borated water moderator, and borosilicate glass burnable poisons. The flux and reaction rate tallies used to compute MGXS were spatially homogenized across all heterogeneous zones comprised of each material throughout the entire core geometry. This approach to spatial homogenization is referred to as *null spatial homogenization* in [20].

The MGXS were generated in the 70 energy group structure used by the CASMO-4 lattice physics code [21]. The `openmc.mgxs` module was used to tally the 70-group reaction rates and fluxes necessary to compute macroscopic total ($\Sigma_{t,g}$), fission ($\Sigma_{f,g}$), and fission production ($\nu\Sigma_{f,g}$) cross sections, along with group-to-group scattering matrices ($\Sigma_{s,g'\to g}$) and normalized fission spectra ($\chi_g$) for each material. The neutron multiplication due to $(n, xn)$ reactions was directly

embedded into the scattering matrix. The in-scatter form of the transport correction [22] was tallied and subtracted from both the total MGXS and the diagonals of the scattering matrices to approximate the effect of anisotropic scattering with the isotropic-in-lab scattering kernel used by OpenMOC. This transport correction approximation is commonly referred to as the in-scatter or micro-reversibility approximation with P0 flux.

The single-step Monte Carlo-based MGXS generation scheme with 70 energy groups has been previously shown to enable accurate deterministic multi-group transport solutions with OpenMOC for a 2D version of the BEAVRS model [20]. In particular, the eigenvalues computed by OpenMOC with 70-group MGXS was within 150 pcm of the OpenMC reference value. The majority of the pin-wise reaction rate bias between OpenMC and OpenMOC was shown to be due to the spatial homogenization model used to account for spatial self-shielding effects. The null spatial homogenization scheme used here to tabulate MGXS for each material is unable to account for these effects. As a result, we expect that the errors for 3D OpenMOC calculations will be of a similar magnitude to those cited for null homogenization in the previous study of the 2D BEAVRS model. This paper, however, focuses on solving the multi-group transport problem efficiently - regardless of the source of the multi-group cross sections.

## X. ASSEMBLY-LEVEL PARAMETRIC STUDY OF MOC AND CMFD PROBLEM REQUIREMENTS

With the generated cross sections from our Monte Carlo simulations, a series of 3D PWR single assembly problems were solved. In this study, we make several assumptions about the parameters required for convergence in the radial plane. These assumptions are motivated by extensive experience with parameters required for 2D MOC convergence and include:

- 32 azimuthal angles and 0.1 azimuthal ray spacing is sufficient for the radial MOC track mesh.
- 8 sectors in the moderator and 4 in the fuel will be sufficient when using the linear source approximation.
- A uniform radial mesh with mesh size equal to the pin pitch is sufficient to achieve efficient CMFD acceleration.

With these assumed parameters, there are other MOC parameters which remain to be determined to achieve the proper accuracy and efficiency of 3D MOC. These include the axial ray spacing, the axial height of source regions, the number of polar angles, the CMFD axial mesh, and the CMFD group structure. Therefore we first conduct computational experiments on a full 3D assembly.

For these tests we use the BEAVRS assembly of 1.6% enriched fuel with no burnable poisons. For investigating the MOC parameters such as ray spacing, source height, and number of polar angles, we aim to find the coarsest set of parameters that maintains solution accuracy. In the CMFD analysis, we aim to determine the coarsest parameters (*i.e.*, fewest groups and coarsest axial mesh) that maintains efficient solution convergence. The combination of these parameters will form the parameters used in full-core analysis.

### 1. Axial Ray Spacing and Axial Source Height

Axial ray spacing and axial source height form the first set of parameters investigated in this study. The two parameters are closely coupled since sufficient ray crossings must intersect each source region to form reliable linear source components. Additionally, the 3D MOC implementation in OpenMOC requires that for each region intersected in the radial plane, all axial zones must be traversed by at least one track. This necessitates that the axial ray spacing must not exceed the smallest source region height. From our BEAVRS model assumptions mentioned in section VIII., the smallest axial material region is 2.0 cm. This mandates a 2.0 cm limit for axial source height as we do not wish to fit sources across material boundaries.

For simplicity, we vary the axial source height by laying down a uniform axial mesh across the geometry, though the code allows for arbitrary axial mesh in every 2D region. The parameter sweep results are shown in Table II. For each axial ray spacing and axial source height pair, the eigenvalue and RMS pellet fission rate fractional error are computed relative to the finest case. We define the RMS fractional pellet fission rate error to be the root mean square fractional error of fission rates within each 2.0 cm fissionable region within every fuel pin. All cases shown use 6 polar angles with a Gauss-Legendre quadrature.

From these results, 0.75 cm ray spacing and 2.0 cm source height are sufficient to reduce the RMS error in fractional pellet fission rate error to below 1% relative to the finest case (0.1875 cm axial ray spacing and 0.5 cm source height). It is important to remind the reader that this parameter study was conducted with full axial detail, including grid spacers. Figure 2 shows the radially integrated normalized fission rate distribution as a function of axial height. Notice that the profile is nearly a cosine with depressions around grid spacers. A parameter study of axial ray spacing and source height without this full detail may lead to alternative conclusions with respect to the parameters necessary for convergence.

In addition we present the axially integrated fission rate distribution in Figure 3 for the same case plotted for each radial pin cell. These results demonstrate that the axial ray spacing sensitivity is much larger than the axial source heights sensitivity. Consequently, linear sources allow source heights of 2.0 cm without significant loss of accuracy, but axial ray spacing of 0.75 cm is required to reduce pellet-wise r.m.s. errors below 1%.

### 2. Polar Angles

The next MOC parameter we study is the number of polar angles. In OpenMOC there are many polar quadrature options, but all simulations performed in this analysis utilize the Gauss-Legendre polar quadrature. OpenMOC has the ability to generate a Gauss-Legendre quadrature for an arbitrary even number of requested polar angles and correct the generated polar angles to ensure linking tracks at domain boundaries. The results, presented in Table III, demonstrate that a Gauss-Legendre quadrature with 10 polar angles is sufficient to reduce RMS fractional pellet fission rate error below

TABLE II. Parameter study of axial ray spacing and axial source height required to converge a 1.6% enrich BEAVRS assembly.

| **Axial Ray Spacing (cm)** | **Axial Source Height (cm)** | $k_{eff}$ | $\Delta k_{eff}$ (pcm) | **RMS Fractional Pellet Fission Rate Error** |
|---|---|---|---|---|
| 1.5 | 2.0 | 1.003230 | 3.7 | 3.35% |
| 0.75 | 2.0 | 1.003214 | 2.1 | 0.70% |
| 0.75 | 1.0 | 1.003199 | 0.6 | 0.70% |
| 0.375 | 0.5 | 1.003219 | 2.6 | 0.25% |
| 0.1875 | 0.5 | 1.003193 | – | – |

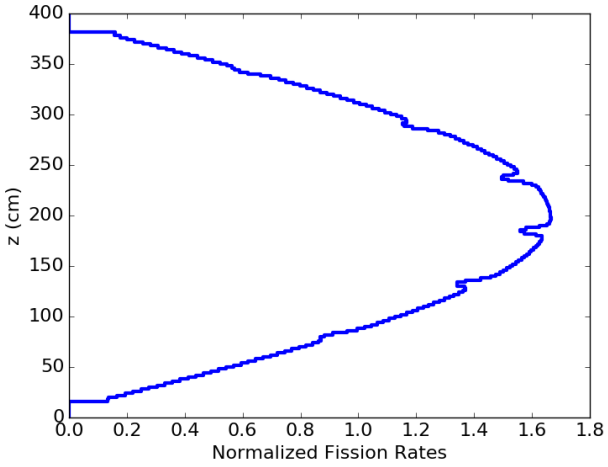


Fig. 2. Normalized radially integrated fission rates for a 1.6% enriched BEAVRS assembly with 0.1875 cm axial ray spacing, 0.5 cm source height, and 6 polar angles as a function of axial height ($z$).

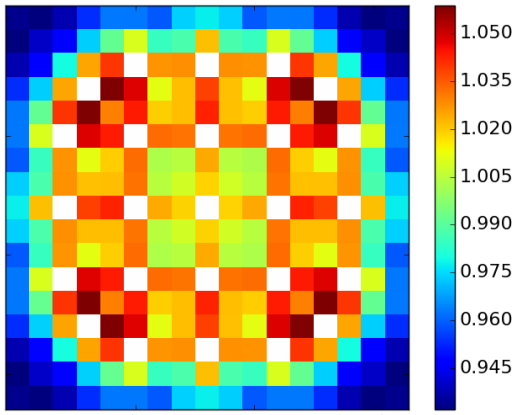


Fig. 3. Normalized axially integrated fission rates for a 1.6% enriched BEAVRS assembly with 0.1875 cm axial ray spacing, 0.5 cm source height, and 6 polar angles by pin cell.

1% relative to the finest case (18 polar angles). All cases use 0.75 cm axial ray spacing and 2.0 cm source height.

## 3. CMFD Prolongation

Now that the MOC parameters have been investigated, we turn to determining the CMFD parameters required for efficient convergence. It is important to note that in all the MOC parameter studies a $17 \times 17 \times 200$ CMFD mesh with 25 CMFD groups and a 0.7 CMFD relaxation factor are used. All cases converged to a criteria of $10^{-4}$ RMS fractional fission source difference and $\Delta k_{eff}$ less than 1 pcm within 14–16 iterations. This of course includes the cases from the axial ray spacing and axial source height study which included cases with source heights as low as 0.5 cm. Since we used a fixed CMFD mesh with 2.0 cm CMFD cell height this required prolonging over a few source regions axially. As mentioned earlier, we use a simple quadratic fit over neighboring axial zones to attempt to reconstruct the flux. This method uses the two axial neighboring CMFD cells to construct a quadratic fit that preserves the average flux in all three regions. Since we observed no significant change in the required number of iterations for convergence, this prolongation seems to have worked well within material boundaries.

To reduce the computational costs of the CMFD acceleration, we would like to coarsen the CMFD mesh axially. Bearing in mind that the materials can change over 2 cm intervals (particularly spacers) coarsening the CMFD mesh axially necessitates fitting the flux across material boundaries. This doesn't seem to be a particularly difficult challenge as we already form CMFD cells radially that encompass material boundaries. Specifically for the $17\times17\times200$ CMFD mesh successfully used in the previous MOC parameter studies, a pin cell CMFD mesh was used which incorporates such disparate materials as fuel, helium, zirconium, and water. However, we have observed that coarsening the CMFD mesh axially significantly hinders the convergence characteristics. For 8.0 cm and coarser axial CMFD mesh, the solution failed to converge. The results are presented in Table IV and all cases used 0.75 cm axial ray spacing with 2.0 cm source height and 6 polar angles. These results suggest maintaining a CMFD mesh of 2.0 cm axially to preserve the convergence rate. This conclusion may be sensitive to the crude approximation of fitting fluxes across neighboring cells, and the conclusion might be different if a low order nodal approximation were to be employed axially to produce more localized quadratic flux shapes.

## 4. CMFD Group Structure

The last parameter study conducted on the single assembly is the CMFD group structure. As noted, a 25 group CMFD

TABLE III. Parameter study of polar angles required to converge a 1.6% enrich BEAVRS assembly.

| **Number of Polar Angles** | $k_{eff}$ | $\Delta k_{eff}$ (pcm) | **RMS Pellet Fission Rate Fractional Error** |
|---|---|---|---|
| 6 | 1.003214 | 2.8 | 1.06% |
| 10 | 1.003215 | 2.9 | 0.79% |
| 14 | 1.003213 | 2.7 | 0.73% |
| 16 | 1.003181 | -0.5 | 0.62% |
| 18 | 1.003186 | – | – |

TABLE IV. CMFD prolongation study on a 400 cm tall 1.6% enriched BEAVRS assembly.

| **CMFD Cell Axial Height** | **Converged?** | **Required Iterations for Convergence** |
|---|---|---|
| 2.0 | Yes | 15 |
| 4.0 | Yes | 27 |
| 8.0 | No | – |
| 16.0 | No | – |

structure was used on all previous cases. However, we would like to use a coarser group structure to lessen the computational requirements of CMFD while maintaining convergence rate. A coarser group structure also significantly reduces the memory required for CMFD as currents need to be stored for every CMFD group on ever CMFD surface. A variety of group structures were tested and the results are presented in Table V. These group structures were chosen from the CASMO-4 group structures for the associated number of groups [21]. The results show that moving to lower CMFD group numbers comes with the cost of increased MOC iterations. If the cost of CMFD is significantly low, the higher CMFD group numbers might be preferred. But if the communication or storage requirements become prohibitively expensive for a large CMFD matrix, then the lower CMFD group numbers would be preferred. This high CMFD overhead becomes significant for full core problems if the CMFD solver is not domain decomposed.

## XI. FULL-CORE BEAVRS PWR RESULTS

The single assembly results now motivate our choice for MOC and CMFD parameters to accurately and efficiently converge the 3D full core solution. Before attempting the 3D solution, we first solve a simplistic "2D" problem using the 3D MOC solver whereby a full radial section of height 10 cm is cut out from a non-spacer bearing region of the core and modeled with reflective boundary conditions and coarse parameters axially. Solving this problem and viewing the result allows confidence that the setup of the problem is correct. The normalized axially integrated fission rate distribution is plotted in Figure 4 where the fission rates are normalized such that the average fission rate across all fuel pins is 1.0.

Now, we turn to solving the BEAVRS benchmark with full detail over the 400 cm core. We choose our MOC and CMFD parameters to be the values given in Table VI which are based on the previous parameter studies. The final results of the full core 3D analysis are quite computationally intense. Just prior to submission of this paper, we determined that

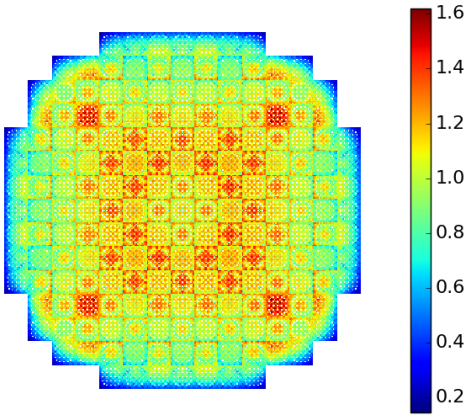


Fig. 4. Normalized axially integrated fission rates for the BEAVRS core modeled with axially reflective boundary conditions and a height of 10 cm.

domain decomposition of the CMFD solver was needed to allow 25-group acceleration of the 70 group MOC and on-node memory of the *Mira* IBM Blue Gene/Q was insufficient to store the full CMFD solutions matrix. With insufficient time before publication, this paper was submitted without results from the full-core 3D BEAVRS problem. Those results will be presented in the oral session.

## XII. CONCLUSIONS

Fully resolving solutions to a full PWR reactor core problems in three dimensions has long been a goal of the reactor physics community. In this investigation we have shown that it is computationally feasible to solve these problems using current technology and proper algorithmic choices. Specifically, OpenMOC is able to solve complex problems such as BEAVRS by using domain decomposition to allow for scalabil-

TABLE V. CMFD group structure study on a 70-group 1.6% enriched BEAVRS assembly.

| **CMFD Groups** | **Required Iterations** |
|---|---|
| 25 | 15 |
| 11 | 17 |
| 8 | 21 |
| 4 | 25 |

TABLE VI. Final parameters used in the full core MOC solution of the BEAVRS benchmark in OpenMOC.

| **Parameter** | **Value** |
|---|---|
| Azimuthal Angles | 32 |
| Radial Ray Spacing | 0.1 cm |
| Polar Angles | 10 |
| Axial Ray Spacing | 0.75 cm |
| Axial Source Height | 2.0 cm |
| CMFD Mesh | (1.26 cm, 1.26 cm, 2.0 cm) |
| CMFD Groups | 25 |

TABLE VII. Computational results of the full core BEAVRS benchmark simulations using 3D MOC.

| | BEAVRS (Extruded 2D) |
|---|---|
| Computed $k_{eff}$ | 1.023290 |
| Source Regions | $1.69 \times 10^6$ |
| Tracks | $4.476 \times 10^7$ |
| Segments | $2.058 \times 10^9$ |
| Transport Sweeps | 14 |
| Computational Nodes | 34 |
| Compute Cores per Node | 24 |
| Processor Type | Haswell chipset 2.50 GHz |
| Total Compute Time | 9.06 m |
| Transport Sweep Time | 3.28 m |
| MPI Communication Time | 37.08 s |
| CMFD Solver Time | 3.68 m |
| Time per Integration per core | 39.89 ns |
| Total Core Hours | 123.22 core-hr |

ity onto large computing platforms, efficient track generation to reduce the number of total integrations required, axially extruded ray tracing to reduce the storage requirements, CMFD acceleration to reduce the number of transport sweeps required for convergence, and a linear source approximation to reduce the number of source regions relative to the traditional flat source approximation. The assembly results have shown that the axial ray spacing of MOC tracks can be much coarser than the radial ray spacing while maintaining solution accuracy and the linear source in three dimensions allows for source discretization on a similar scale as material boundaries. While the number of polar angles required is significantly larger than that of 2D MOC, the total 3D MOC computational complexity remains within reason and enables realistic 3D transport solutions of full core LWR reactor configurations.

## XIII. ACKNOWLEDGMENTS

The first and second authors are recipients of the DOE Office of Nuclear Energy's Nuclear Energy University Programs Fellowship. This research is being performed,in part, using funding received from the DOE Office of Nuclear Energy's Nuclear Energy University Programs (contract number: DE-NE0008578). This research made use of the resources of the High Performance Computing Center at Idaho National Laboratory, which is supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517.

## REFERENCES

1. Y. S. JUNG, ET AL., "Practical numerical reactor employing direct whole core neutron transport and subchannel

thermal/hydraulic solvers," *Annals of Nuclear Energy*, **62**, 357–374 (2013).
2. J. Y. CHO, ET AL., "Cell Based CMFD Formulation for Acceleration of Whole-core Method of Characteristics Calculations," *Journal of the Korean Nuclear Society*, **34**, 250–258 (2002).
3. B. KELLEY AND E. LARSEN, "2D/1D approximations to the 3D neutron transport equation. I: Theory," in "International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering," Sun Valley, ID, USA (May 2013).
4. N. Z. CHO, G. S. LEE, AND C. J. PARK, "A Fusion Technique of 2-D/1-D Methods for Three-Dimensional Whole-Core Transport Calculations," in "Korean Nuclear Society Spring Meeting," Kwangju, Korea (May 2002).
5. W. BOYD, S. SHANER, L. LI, B. FORGET, and K. SMITH, "The OpenMOC Method of Characteristics Neutral Particle Transport Code," *Annals of Nuclear Energy*, **68**, 43–52 (2014).
6. S. SHANER, ET AL., "Verification of the 3D Method of Characteristics Solver in OpenMOC," in "PHYSOR 2016," Sun Valley, ID, USA (May 2016).
7. G. GUNOW, ET AL., "Reducing 3D MOC Storage Requirements with Axial On-the-fly Ray Tracing," in "PHYSOR 2016," Sun Valley, ID, USA (May 2016).
8. R. FERRER J. RHODES III, AND K. SMITH, "Linear Source Approximations in CASMO 5," in "ANS Reactor Physics Topical Meeting (PHYSOR)," Knoxville, TN (April 2012).
9. G. GUNOW, ET AL., "SimpleMOC - A Performance Abstraction for 3D MOC," in "ANS MC2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method," Nashville, TN, USA (April 2015).
10. B. KELLEY, ET AL., "CMFD Acceleration of Spatial Domain-decomposition Neutron Transport Problems," in "PHYSOR 2016," Knoxville, TN, USA (April 2012).
11. B. KOCHUNAS, "A Hybrid Parallel Algorithm for the 3-D Method of Characteristics Solution of the Boltzmann Transport Equation on High Performance Computing Clusters," Ph.D. Thesis, University of Michigan, Department of Nuclear Engineering and Radiological Sciences (2013).
12. K. SMITH AND J. RHODES III, "Full-core, 2-D, LWR Core Calculations with CASMO-4E," in "ANS Reactor Physics Topical Meeting (PHYSOR)," Seoul, Korea (October 2002).
13. S. SHANER, "Transient Method of Characteristics via the Adiabatic, Theta, and Multigrid Amplitude Function Methods," Ph.D. Thesis, Massachusetts Institute of Technology, Department of Nuclear Science and Engineering (2014).
14. L. LI, ET AL., "Techniques for Stabilizing Coarse-Mesh Finite Difference (CMFD) in Methods of Characteristics (MOC)," in "ANS MC2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method," Nashville, TN, USA (April 2015).
15. S. SHANER, B. FORGET, AND K. SMITH, "K-nearest Neighbor Flux Updating Scheme for CMFD Acceleration," *Annals of Nuclear Energy, manuscript in preparation*.
16. S. SHANER, "In preparation," Ph.D. Thesis, Massachusetts Institute of Technology, Department of Nuclear Science and Engineering (2017).
17. N. HORELIK, B. HERMAN, B. FORGET, and K. SMITH, "Benchmark for Evaluation and Validation of Reactor Simulations (BEAVRS), v1.0.1," in "Proc. Int. Conf. Math. and Comp. Methods Applied to Nuc. Sci. & Eng.", Sun Valley, Idaho, USA (2013).
18. P. K. ROMANO and B. FORGET, "The OpenMC Monte Carlo Particle Transport Code," *Annals of Nuclear Energy*, **51**, 274–281 (2013).
19. W. BOYD, P. K. ROMANO, and S. HARPER, "Equipping OpenMC for the Big Data Era," in "PHYSOR," Sun Valley, ID, USA (2016).
20. W. R. D. BOYD, *Reactor Agnostic Multi-Group Cross Section Generation for Fine-Mesh Deterministic Neutron Transport Simulations*, Ph.D. thesis, Massachusetts Institute of Technology (2017).
21. M. EDENIUS, K. EKBERG, B. H. FORSSÉN, and D. KNOTT, "CASMO-4, A Fuel Assembly Burnup Program, User's Manual," *Studsvik0SOA-9501, Studsvik of America, Inc.* (1995).
22. A. YAMAMOTO, Y. KITAMURA, and Y. YAMANE, "Simplified Treatments of Anisotropic Scattering in LWR Core Calculations," *Journal of Nuclear Science and Technology*, **45**, *3*, 217–229 (2008).