Quantifying the Effect of Undersampling Biases in Monte Carlo Reaction Rate Tallies

B. T. Rearden, W. J. Marshall, and C. M. Perfetti

Oak Ridge National Laboratory

J. Miss and Y. Richet

Institut de Radioprotection et de Sûreté Nucléaire

June 2014

Abstract

The OECD/NEA Working Party on Nuclear Criticality Safety's Expert Group on Advanced Monte Carlo Techniques (AMCT) was formed to advance the knowledge base of information in regards to Monte Carlo criticality calculations and transfer that knowledge to end users. One item of interest in the community is the effective use of Monte Carlo depletion, especially as related to burnup credit. The AMCT group will investigate issues that affect the accuracy of Monte Carlo-calculated fluxes, reaction rate tallies, and cross sections. This paper outlines an approach for quantifying how undersampling induces biases in reaction rate tallies and produces inaccurate tally variance estimates. The goal of this benchmark exercise is to verify that Monte Carlo codes can calculate accurate steady-state reaction rates for complex Monte Carlo models involving burned fuel and to create a set of best practices to minimize the effects of undersampling in Monte Carlo calculations.

Background

Most Monte Carlo codes calculate the eigenvalue and fundamental mode of the source distribution of fissile systems by simulating multiple "generations" of fission neutrons, where the fission sites created during one generation serve as the birth sites for neutrons in the next generation. By simulating enough generations and neutrons, Monte Carlo codes can obtain accurate eigenvalue and reaction rate tally estimates. A phenomenon known as undersampling occurs when the neutrons in one generation do not interact with all regions in a problem. Undersampling results in generational eigenvalue estimates that have not incorporated information from all regions in the system, and reaction rate tallies that are "skipped" and not scored for an entire generation; because of these reasons, undersampling can lead to inaccurate eigenvalue and reaction rate tally and tally variance estimates.

In Reference 1 Brown examines how undersampling and cycle-to-cycle correlations introduce a bias in the eigenvalue for several difficult problems, including Whitesides' K-effective of the World problem and Nakagawa and Mori's 2D PWR quarter-core model [2] [3]. Brown found that using 1,000 particles per

generation instead of 5,000 particles per generation results in a 100 pcm bias in the eigenvalue of the K-effective of the World problem, and a 20 pcm bias in the eigenvalue of the PWR problem. The impact of undersampling was even more severe for fission reaction rate tallies, and Brown found that using less than 10,000 particles per generation caused the fission rates near the center and the edge of the PWR core to differ by between 2 and 3 standard deviations from their reference values.

In Reference 4 Mervin et. al. examine how undersampling and cycle-to-cycle correlations create poor variance estimates for flux tallies in full-core Monte Carlo simulations. In the worst-case scenario Mervin found that the uncertainties were under-predicted by a factor of 40.

Methodology

This study will continue and extend Brown's and Mervin's analysis on the impact of undersampling, and aims to develop a set of best practices for the Monte Carlo user community to ensure accurate reaction rate tally results in full-core Monte Carlo calculations. The impact of undersampling and cycle-to-cycle correlations will be quantified by performing several simulations for a group of models and varying the number of particles per generation across these simulations. Each simulation will skip a sufficient number of generations to ensure fission source convergence, and will simulate a very large number of generations to ensure that the tally statistical uncertainty is small relative to the expected size of the undersampling bias. The total number of active histories should be equal across all simulations. The tally bias will be quantified using "1/M" plots, which Brown has used to quantify the eigenvalue undersampling bias. Shown below in Figure 1, Brown's 1/M plots compare the calculated eigenvalue (or reaction rate tally in our study) as a function of 1/M, where M is the number of neutrons per generation [1]. Figure 1 suggests that at least 5,000 particles per generation should be used to minimize the eigenvalue bias in this problem.

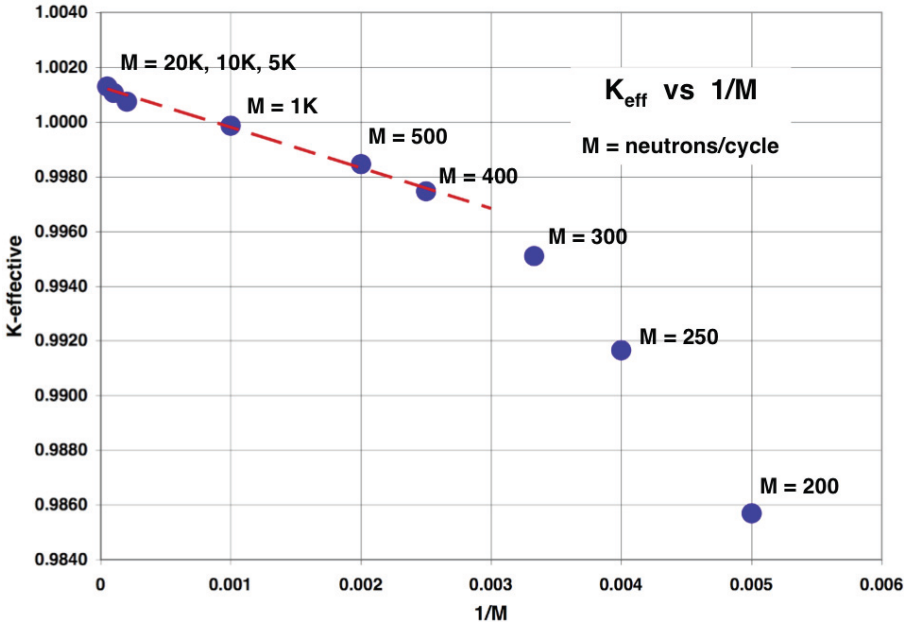

Figure 1. Sample 1/M plot for quantifying the undersampling bias in the eigenvalue for the K-effective of the World problem [1].

The number of particles per generation needed to minimize the undersampling bias will be problem-specific, and this study will determine how many particles per generation should be used in a variety of models with increasing complexity. This study will focus on PWR models because the goal of this benchmark is to verify the end-use of Monte Carlo codes for reaction rate and depletion calculations related to burnup credit. Additionally, PWRs are sufficiently complex to develop a general set of best practices. It is certainly possible to conceive of a more loosely-coupled Monte Carlo problem where the undersampling bias has a more significant impact, but the scope and complexity of this exercise is limited to PWRs.

In addition to allowing a clear depiction of the undersampling bias, 1/M plots have the added advantage of avoiding the need for delegating a code to produce benchmark results. Code-to-code differences for tally and eigenvalue responses could be the same order of magnitude, or larger, than the undersampling bias, thus making tally and eigenvalue comparisons difficult for even similar Monte Carlo codes. 1/M plots only examine the accuracy of tallies within *one* code as a function of the number of particles per generation, and thus allow participants to isolate the effect of the undersampling bias without introducing interference from code-to-code biases.

1/M plots should be developed for two global quantities, k_{eff} and the Energy of the Average Lethargy of neutrons causing Fission (EALF), and local flux tallies at various locations in the models in this study; a list of the locations of these flux tallies is given later with the model descriptions. This list is designed to sample areas in which the flux would be expected to be changing quickly as a function of position and places where the flux should be fairly constant. Future expansions of this benchmark may include local reaction rates in addition to the flux tallies desired for this benchmark exercise. The flux will be tallied over only the fuel within the rod; gap and cladding regions are excluded from the tally volume.

The accuracy of the k_{eff}, EALF, and flux tally uncertainty estimates will be measured by comparing the calculated tally variance with the true tally variance. The true tally variance can be calculated by performing N repeated Monte Carlo simulations with different random number seeds and calculating the variance of this group of N tally results using the equation

$$\sigma_x^2 = \frac{1}{N-1}\sum_i^N (x_i - \bar{x})^2 .$$

In Reference 4 Mervin plots the ratio of the real to calculated tally variance, or f_σ, for flux tallies. An example of these plots is given in Figure 2 for a pin cell model for differing numbers of initial generations skipped and differing numbers of histories per generation. As seen in the figure, all of the cases reveal a significant (~factor of 10) underestimation of the variance in flux tallies near the bottom of the fuel pin.

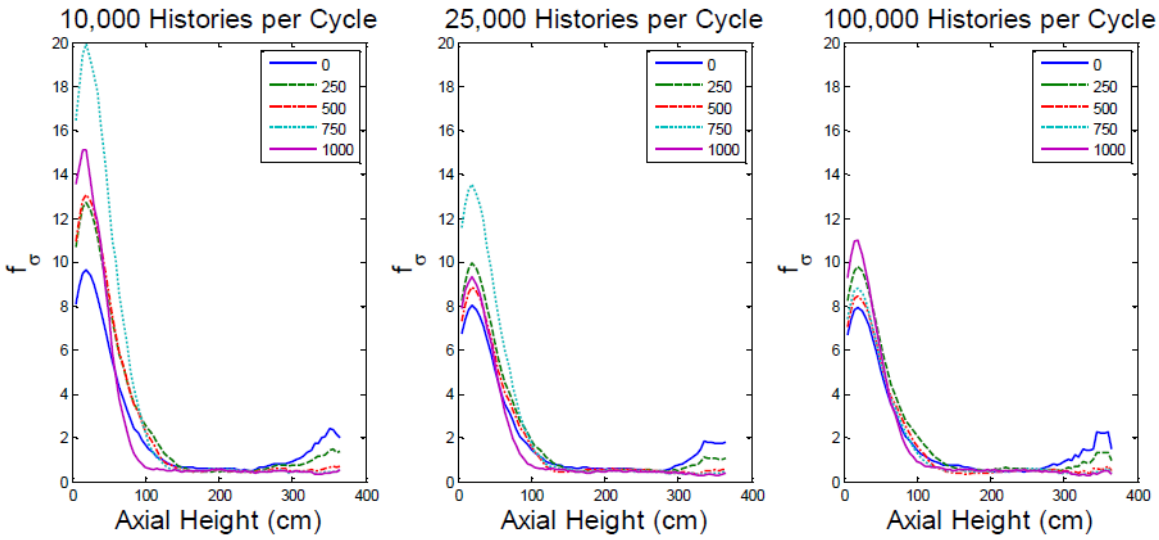

Figure 2. Comparison of f_σ with five different values for the number of initial cycles skipped and three different values for the number of histories per cycle.

Benchmark Exercises

The goal of this investigation is to determine a set of easy-to-follow guidelines (i.e. the minimum number of particles per generation that should be used) to minimize the impact of undersampling and cycle-to-cycle correlations in Monte Carlo eigenvalue calculations.

This benchmark study seeks to determine how spatial and isotopic complexity induces a bias in reaction rate tallies and it is important to isolate the effects of the spatial and isotopic components of undersampling. Thus, the benchmark will consist of several stages of increasing complexity to isolate the effect of 2D, 3D, and isotopic complexity on flux tallies. A summary of the six benchmark models is provided in Table 1, where there are three models for reactor (R) configurations and three models for storage (S) configurations. Details of each configuration are provided below. The detailed tally location descriptions are provided below.

Table 1. Configurations for Benchmark Exercises.

ID	Configuration	Geometry	Isotopics	Temperature	Reaction Tally Locations
R1 (Optional)	2D Quarter Core	17x17 quarter core radial slice	Uniform 20 GWD/MTU with equilibrium Xenon	Reactor – Uniform Mid-Plane	Center and edge bundles
R2	3D Infinite Lattice	17x17 bundle in infinite lattice	18 axial zones; varying 20 GWD/MTU with equilibrium Xenon	Reactor – 18 Axial zones	Top, mid-plane, and bottom
R3 (Optional)	3D Quarter Core	17x17 quarter core	18 axial zones; 20 GWD/MTU with equilibrium Xenon; Uniform radially	Reactor – Uniform radially, 18 axial zones	Center and edge bundles Top, mid-plane, and bottom
S1 (Optional)	2D Storage Cask	17x17 in cask geometry radial slice	Uniform 40 GWD/MTU with 5 year cooling time	Uniform storage temperature	Center and edge bundles
S2	3D Infinite Lattice	17x17 bundle in infinite lattice	18 axial zones; 40 GWD/MTU with 5 year cooling time	Uniform storage temperature	Top, mid-plane, and bottom
S3 (Optional)	3D Cask	17x17 in full cask	18 axial zones; 40 GWD/MTU with 5 year cooling time; Uniform radially	Uniform storage temperature	Center and edge bundles Top, mid-plane, and bottom

Model descriptions

The models to be examined in this study are based on the Commercial Reactor Critical (CRC) benchmarks and the GBC-32 spent fuel storage and transportation canister model [5] [6]. A simplified CRC PWR model was chosen because it is a realistic model of a PWR without being overly cumbersome. The GBC-32 problem was chosen both because it is an important problem for criticality safety and shielding analysis. Furthermore, much of the fission source in the GBC-32 problem is concentrated near the top of the shipping cask, which may lead to undersampling in the lower regions of the cask.

PWR Core Model

The simplified CRC PWR model is based on the Sequoyah Unit 2 Cycle 3 middle of cycle (MOC) statepoint from Reference 5. This statepoint was selected as it was near the desired cycle burnup of 20 GWd/MTU, as shown in Table 1, but the temperatures, moderator densities, and ^{135}Xe number densities are selected to approximate a full power condition. Sequoyah Unit 2 is a Westinghouse 4-loop PWR fueled with 17×17 fuel assemblies. Cycle 3 used standard fuel with a rod outer diameter of approximately 0.95 cm (0.374 in.). A radial slice of the model is shown in Figure 3. The full core is a truncated 15×15 array of fuel assemblies, with baffle plates adjacent to the outer faces of the assemblies on the core periphery. The water reflector outside the baffle plates is modeled as is the core barrel. The

outer diameter of the core barrel defines the radial boundary for the benchmark problem. Both the core baffle and barrel are modeled as 304 stainless steel.

The simplified model developed for this benchmark exercise includes no radial fuel variation but does include axial variations in fuel isotopic number densities and temperatures. The moderator density and temperature also varies as a function of elevation but not radial position in heated channels. The cladding temperature is interpolated between the fuel and moderator temperatures such that 15% of the temperature change occurs from the cladding to the moderator. The soluble boron concentration is set to yield a calculated k_{eff} value near 1.0 in combination with an axial moderator density and temperature profile designed to give a somewhat realistic, flat core average axial power profile.

A uniform 18-zone axial structure is used to discretize the axial variations in 20.32 cm (8 in.) segments. The axial mesh is provided in Table 1 of Appendix A. Isotopic number densities are also provided in Appendix A for the fuel in Tables 2 – 19 and for the moderator in Tables 20 – 37. Isotopic number densities for the water reflector, cladding, stainless steel, and smeared axial reflectors are provided in Appendix A Tables 38 – 42. Temperatures are specified for axially varying components in Appendix A Table 43 and for non-varying components in Table 44.

Fuel assembly dimensions are provided in Appendix A Table 45, and other reactor dimensions are provided in Table 46. The lower plenum and smeared reflector region use the temperature and moderator density for axial zone 1, while the upper plenum and smeared reflector region use the conditions from axial zone 18. Both the upper and lower fuel rod plenums are modeled as void filled cladding tubes, while the smeared reflectors are intended as approximate representation of fuel assembly end fittings and core internals. The lengths of the plenums and smeared reflectors are given in Appendix A Table 46. The smeared reflector lengths are determined such that a total of 30 cm is included in the model above and below the active fuel. An axial representation of the quarter core is sketched in Figure 4.

The 2D quarter core slice identified as Case R1 in Table 1 should use the compositions from axial zone 10, representing the core slightly above the midplane. The baffle plates, water reflector, and core barrel will use the same temperature as the 3D model in Case R3.

The 3D infinite lattice of fuel assemblies recommended in Case R2 in Table 1 uses a single fuel assembly including plenums and smeared axial reflectors. The top and bottom surfaces of the model have leakage boundary conditions applied and the radial boundaries are reflective so that the model represents an infinite radial array of a single fuel assembly.

The optional 3D quarter core model recommended in Case R3 in Table 1 uses the full radial and axial description of the core model. Because there is no variation in the core compositions, the X and Y axes can be modeled as reflective or periodic about the origin. The outer radius of the core barrel is a leakage boundary condition.

Each assembly is assigned a core horizontal identifier (CHid) which consists of the X and Y coordinates of the assembly location. CHid1 is the X coordinate, numbered from left to right as shown in Figure 3, and CHid 2 is the Y coordinate, numbered from bottom to top as shown in Figure 3. The tally locations selected for this benchmark are in the 6 assemblies identified in Figure 3. The order of the assemblies, to maintain uniform tally numbering to simplify comparing results, is shown in Table 2.

Within the assembly, each rod is assigned an assembly horizontal identifier (AHid) which consists of the X and Y coordinates of the rod location within the assembly. AHid1 is the X coordinate of the rod in the

assembly, numbered from left to right, and AHid2 is the Y coordinate of the rod in the assembly, numbered from bottom to top. The number conventions are shown in Figure 5. The tally locations selected for this benchmark are in the 5 rods identified in Figure 5 for each of the 6 assemblies. The order of the rods, to maintain uniform tally numbering to simplify comparing results, is shown in Table 2.

Order of Assembly Tally Locations

Order of Tally Location	CHid1	CHid2
1	2	2
2	8	2
3	4	4
4	7	6
5	6	7
6	2	8

Table 3.

For the three-dimensional cases, R2 and R3, tallies are defined for multiple, specific axial elevations. The vertical identifier (VID) corresponds to the axial node in the discretized 18-zone axial model. The zones are numbered from bottom to top, as shown in Figure 4. The elevations that will be tallied are 1, 5, 10, 15, and 18.

The axial fission density averaged over all radial regions in an axial zone will also be reported for the R2 and R3 cases. This will add 18 tallies for each of these two cases. The fission rate tallies should converge more quickly than the flux tallies, so this should not significantly increase the number of particles needed in each generation to achieve convergence of the flux tallies.

A spreadsheet for reporting results is provided along with this benchmark specification. An example of the format of the results table in the spreadsheet for a portion of the R1 case is provided in Table 4. A pair of columns is added to the right of the table to record the flux and its standard deviation for each number of particles per generation simulated. The first rows of data are the calculated k_{eff} and EALF for each simulation and their uncertainty even though the columns are labeled for flux tallies. Additional rows are added on the bottom of the table for the total of 30 tallies in R1, the 25 tallies in R2, and the 150 tallies in R3. As mentioned above, the flux is to be tallied over the fuel volume only in the desired locations. The axial fission density results are included for R2 and R3 with one blank line separating the flux tallies from the fission density tallies.

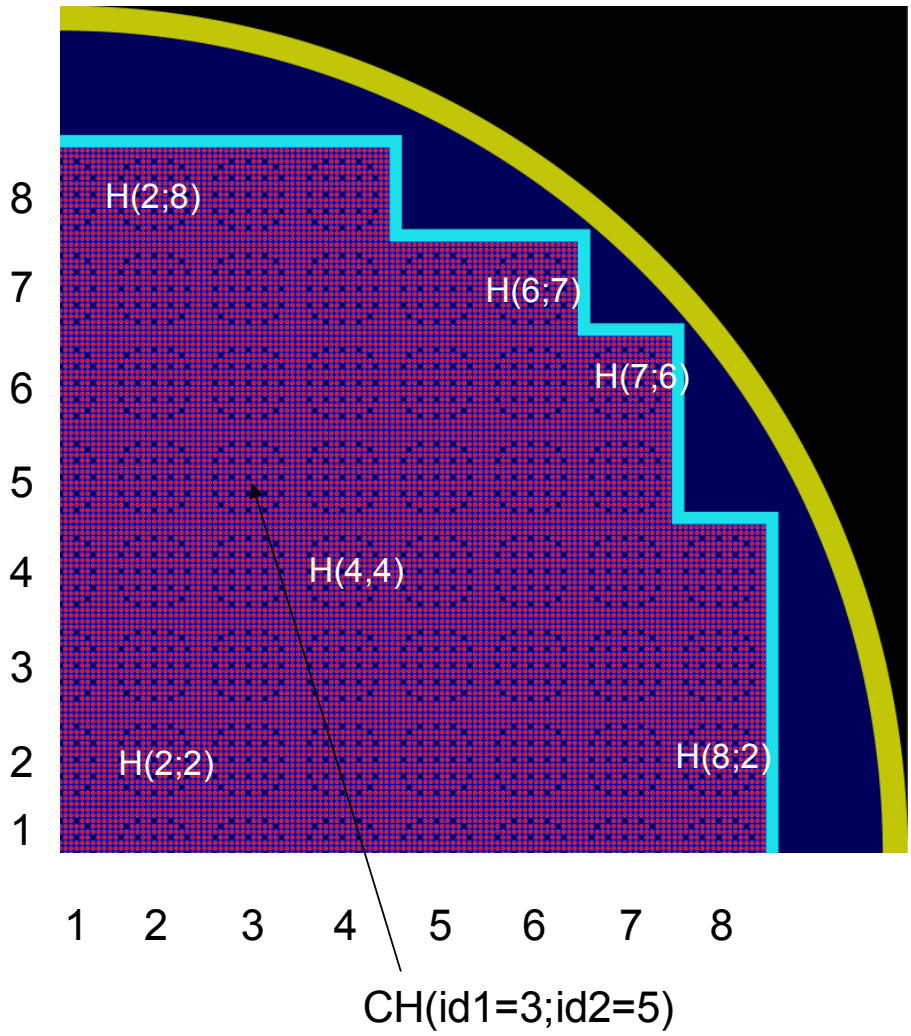

Figure 3. Sample PWR quarter core model showing assemblies containing tally locations.

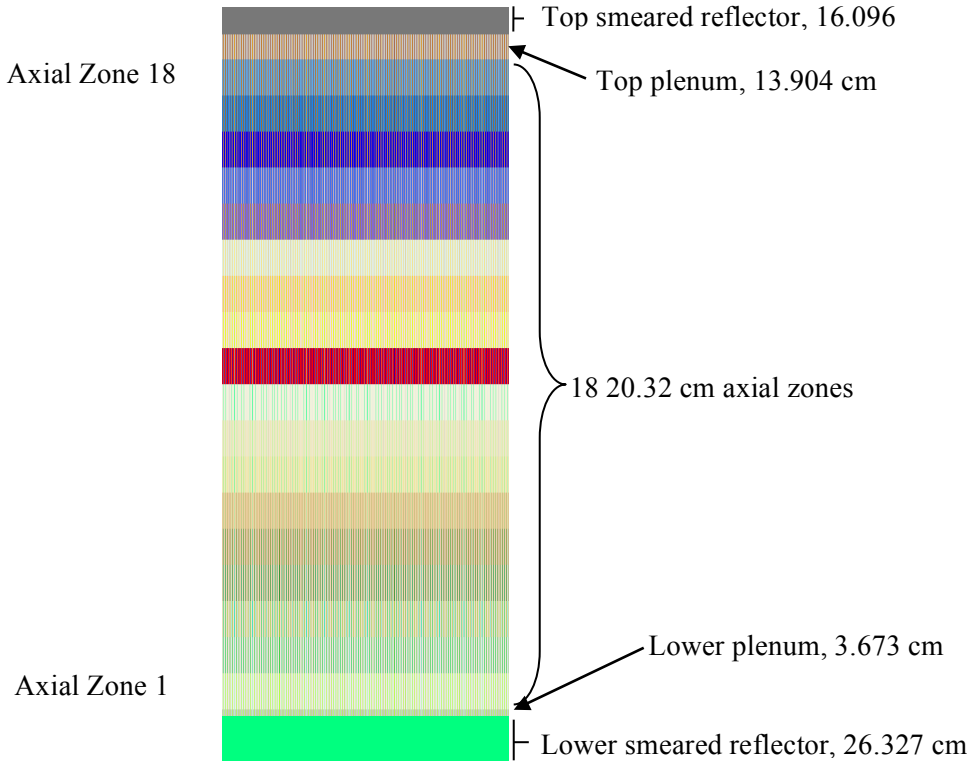

Figure 4. Axial representation of fuel assemblies.

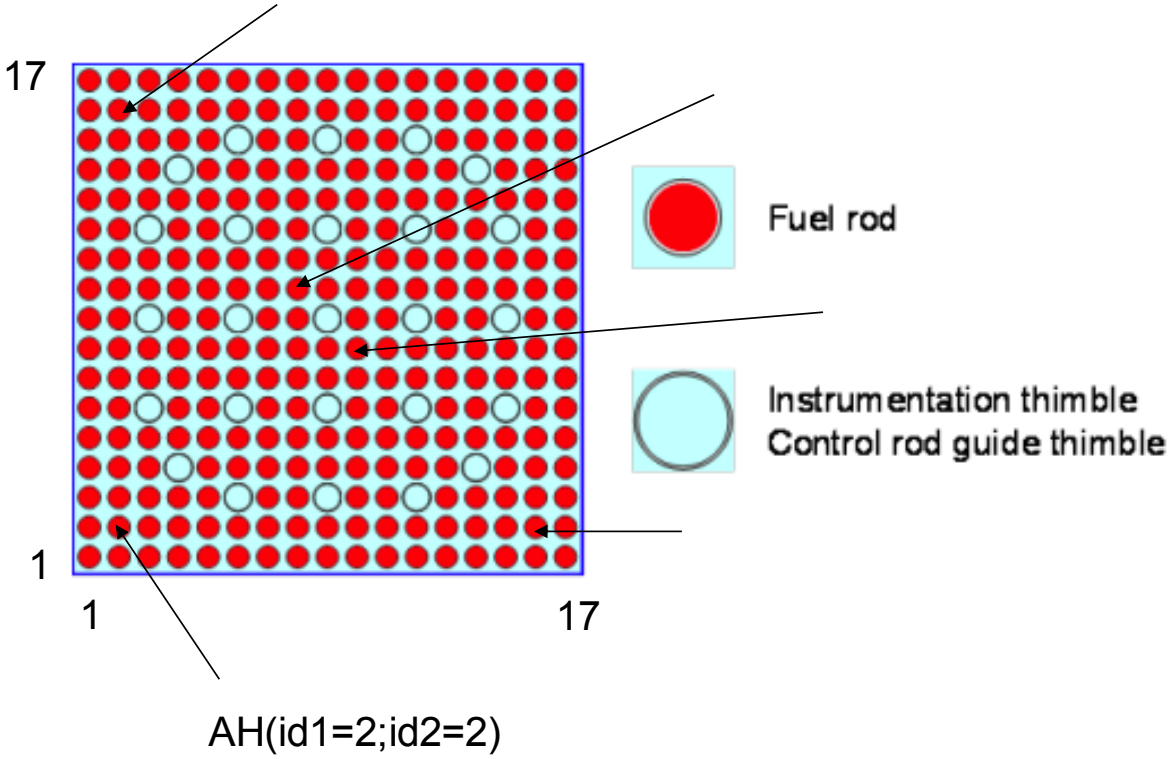

Figure 5. Assembly sketch showing rods containing tally locations.

Table 2. Order of Assembly Tally Locations

Order of Tally Location	CHid1	CHid2
1	2	2
2	8	2
3	4	4
4	7	6
5	6	7
6	2	8

Table 3. Order of Tally Locations Within Each Assembly

Order of Tally Location	AHid1	AHid2
1	2	2
2	16	2
3	2	16
4	10	8
5	8	10

Table 4. Portion of the Results Table for Case R1

	Core		Assembly		Vertical Position	100 Part./Gen.		200 Part./Gen.	
	CHid1	CHid2	AHid1	AHid2	VID	Flux	St. Dev.	Flux	St. Dev.
k_{eff}	N/A								
EALF	N/A								
Tally 1	2	2	2	2	0				
Tally 2	2	2	16	2	0				
Tally 3	2	2	2	16	0				
Tally 4	2	2	10	8	0				
Tally 5	2	2	8	10	0				

GBC-32 Canister Model

The GBC-32 canister model is built the match the description provided in [6], the report which first defined the cask for burnup credit studies. The canister design incorporates a fuel basket consisting of 32 fuel storage cells with neutron absorber panels sandwiched between the walls of adjacent cells. The panels are an aluminum clad B_4C/Al cermet material much like Boral. The fuel storage basket is surrounded by a steel cask body with a thickness of 20 cm and a lid and baseplate, each of which are 30 cm thick. The fuel modeled in the canister is the Westinghouse 17×17 Optimized Fuel Assembly (OFA) with a rod outer diameter of 0.9144 cm (0.360 in.) All assemblies are modeled with an assembly average burnup of 40 GWd/MTU and 5 years of post-irradiation cooling time. A radial cross section of the canister model is provided in Figure 6.

The model for the benchmark exercise contains no radial fuel variation but does include axial variation of isotopic number densities in the spent fuel. All materials in the model are assumed to be at a uniform temperature of 293K. The entire inner cavity of the cask is flooded with water with a density of 0.998

g/cm^3. The spent fuel is represented using the same 18 zone axial structure as the PWR core benchmark. The axial description is provided again in Table 1 of Appendix B, and the isotopic number densities for the 18 zones are provided in Tables 2 – 19 in Appendix B. Isotopic number densities for water, cladding, stainless steel, the absorber panel core material, and the absorber panel cladding are provided in Appendix B, Tables 20 – 24, respectively.

The dimensions for the fuel assembly are provided in Appendix B Table 25. The fuel storage basket dimensions are provided in Table 26 of Appendix B and the cask dimensions are in Table 27 of Appendix B. The fuel basket and neutron absorber panels are set to be the same height as the active fuel. This is not representative of actual cask designs, but is effective for the Monte Carlo analyses for which the benchmark cask is intended. The axial regions above and below the active fuel, storage basket, and absorber panels are modeled as water; no modeling of the assembly end fittings or cask fuel assembly spacers is performed. An axial cross section of the model is shown in Figure 7.

The 2D storage cask slice recommended as Case S1 in Table 1 should use the compositions from axial zone 10, representing the fuel and cask slightly above the midplane. The entire radial extent of the cask is included in the model.

The 3D infinite lattice of fuel storage cells recommended in Case S2 in Table 1 uses a single fuel assembly in a storage cell. The axial extent of the model includes only the active fuel, storage basket, and neutron absorber panels. The top and bottom surfaces of the model have leakage boundary conditions applied and the radial boundaries are reflective so that the model represents an infinite radial array of a single fuel storage cell.

The optional 3D cask model recommended in Case S3 in Table 1 uses the full radial and axial description of the cask model. Leakage boundary conditions are applied to all boundaries of the model so that it models a single, isolated cask.

Each storage cell is assigned an identifier (SCid), which is a number from 1 to 32, as shown in Figure 6. The tally locations selected for this benchmark are in the 6 cells labeled 4, 9, 14, 16, 22, and 29. The order of the assemblies, to maintain uniform tally numbering to simplify comparing results, is shown in Table 5.

The same fuel rods identified for the core benchmark with AHids as shown in Table 2. Order of Assembly Tally Locations

Order of Tally Location	CHid1	CHid2
1	2	2
2	8	2
3	4	4
4	7	6
5	6	7
6	2	8

Table 3 and Figure 5 will be used for SCids 9, 14, 16. Different AHids will be used for SCids 4, 22, and 29 so that symmetric tally locations are defined. The AHids for these three sets of tallies are provided in Table 6. The 3D infinite array model in S2 should use the fuel rods from SCid 14, which are shown in Table 3.

Six axial elevations will be tallied for the cask benchmark. The VIDs for the 6 elevations are 1, 2, 10, 16, 17, and 18. The emphasis on the top portion of the assembly is important because of the axial "end-effect" which causes overall assembly reactivity to be dominated by the reactivity of the top few nodes of the assembly, as discussed in [7].

A spreadsheet for reporting results is provided along with this benchmark specification. An example of the format of the results table in the spreadsheet for a portion of the S1 case is provided in Table 7. A pair of columns is added to the right of the table to record the flux and its standard deviation for each number of particles per generation simulated. The first rows present the calculated k_{eff} and EALF for each simulation and their uncertainty even though the columns are labeled for flux tallies. Additional rows are added on the bottom of the table for the total of 30 tallies in S1, the 30 tallies in S2, and the 180 tallies in S3. As mentioned above, the flux is to be tallied over the fuel volume only in the desired locations. The axial fission density results are included for S2 and S3 with one blank line separating the flux tallies from the fission density tallies.

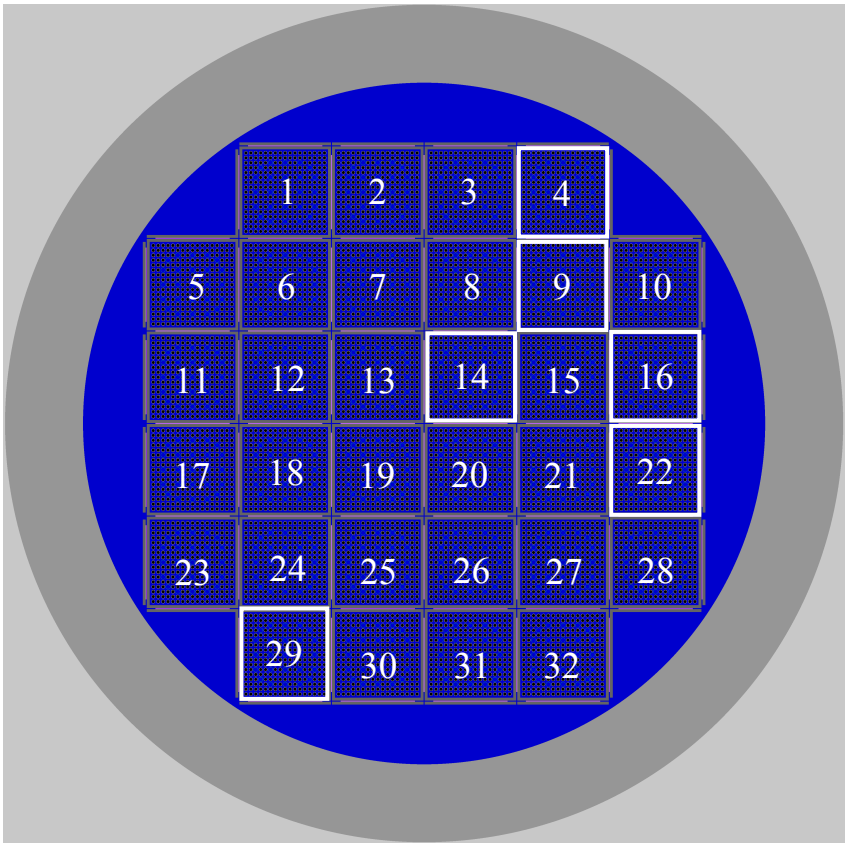

Figure 6. Cross section of the GBC-32 cask model, including cell numbers.

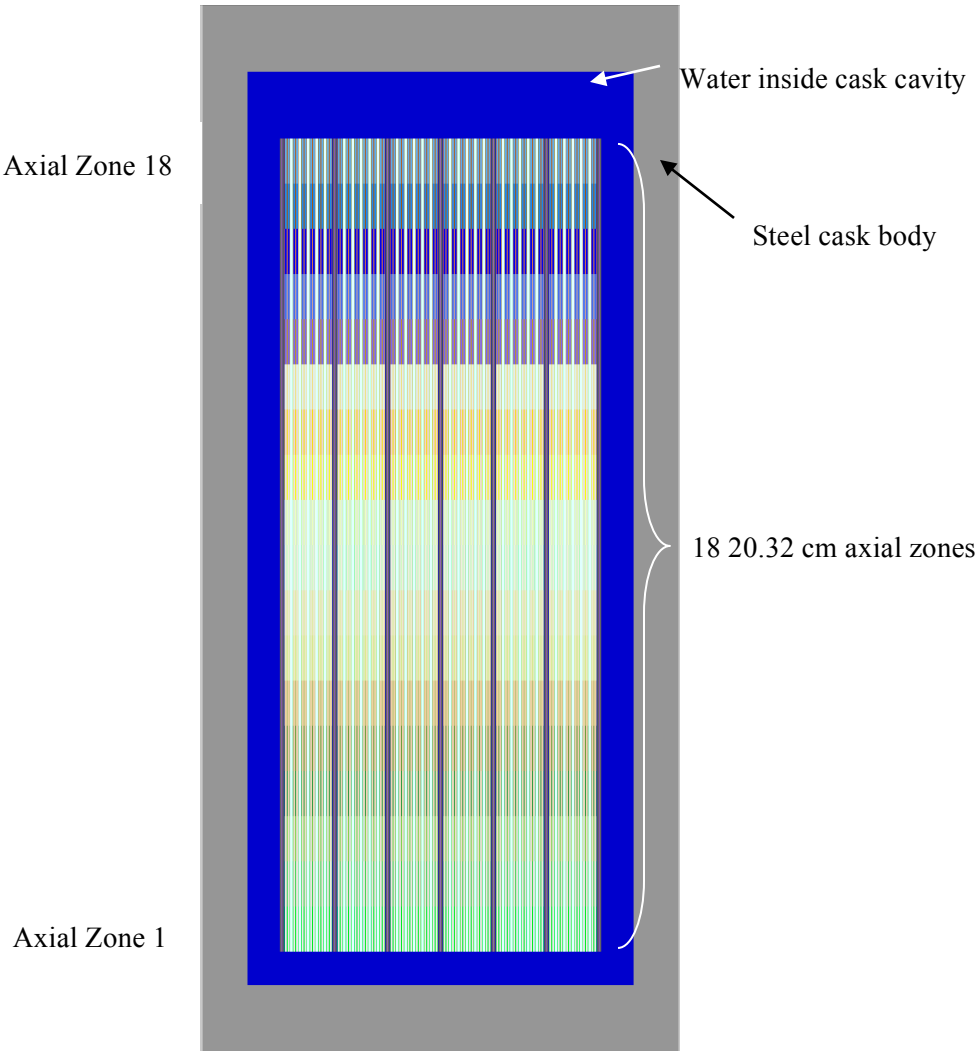

Figure 7. Axial cross section of GBC-32 3D model.

Table 5. Order of Storage Cell Tally Locations

Order of Tally Location	SCid
1	4
2	9
3	14
4	16
5	22
6	29

Table 6. Tally Locations In Storage Cells 22 and 29

Order of Tally Location	AHid1	AHid2
Storage Cell 4		
1	17	17
2	16	2
3	2	16
4	10	8
5	8	10
Storage Cell 22		
1	2	16
2	16	16
3	2	2
4	10	10
5	8	8
Storage Cell 29		
1	1	1
2	2	16
3	16	2
4	8	10
5	10	8

Table 7. Portion of the Results Table for Case S1

	Storage Cell	Assembly		Vertical Position	100 Part./Gen.		200 Part./Gen.	
	SCid1	AHid1	AHid2	VID	Flux	St. Dev.	Flux	St. Dev.
Tally 1	4	2	2	0				
Tally 2	4	16	2	0				
Tally 3	4	2	16	0				
Tally 4	4	10	8	0				
Tally 5	4	8	10	0				

Calculations

The tasks of the benchmark exercise for each of the six models listed in Table 1 are provided below.

Task 1: Sufficiency of sample size

Each of these cases will be simulated using various numbers of particles per generation to generate 1/M plots. A recommend set of particles per generation to use for these simulations is: 100, 200, 500, 1k, 2k, 5k, 10k, 20k, 50k, 100k, etc… These simulations will use a large number of active generations so that the statistical uncertainty in the tally responses is small compared to the expected size of the undersampling bias. The number of skipped generations should be large enough to assure source convergence and should be constant for all simulations. The total number of active histories should be constant across all simulations. 1/M plots should be generated for the designated locations in each model to capture the impact of undersampling as a function of the radial location within the model. It has been demonstrated

that a significant amount of undersampling can occur for flux tallies in the R2 and S2 cases and benchmark **participants are recommended to focus their efforts on these cases for the initial stages of the benchmark study** [8].

This study seeks to understand how undersampling leads to biases and poor uncertainty estimates in Monte Carlo tallies during active generations and will not explore best practices for ensuring fission source convergence during inactive generations; therefore, all simulations must ensure that the problem fission source is sufficiently converged before active tallies begin. Studies by Brown [1] have indicated that fission source convergence occurs in large reactor simulations within 200 skipped generations, but participants should use as many inactive generations/histories as they deem necessary to ensure fission source convergence; alternatively, participants may wish to run one baseline calculation with a large number of inactive histories and generations to generate a starting fission source for their simulations, or may use deterministic transport methods to accelerate and ensure fission source convergence.

Task 2: Accuracy of variance estimators (Optional)

This benchmark study also seeks to investigate the conditions that lead to under- and over-estimation of tally statistical uncertainty by comparing the real versus calculated variance for the tallies examined in Task 1. The calculated tally variance will be obtained from any of the simulations in Task 1 and the real variance will be determined by simulating each case repeatedly using different random number seeds. These calculations should be performed only for cases with an acceptable number of particles per generation to generate reliable tallies. The real variance calculations should use on the order of 20 - 50 simulations for each selected simulation, and may impose a significant computational burden. It is expected that the real-to-apparent variance ratio will be significantly farther from unity for the local flux tallies than for the overall system k_{eff} and EALF.

Task 3: Development of metrics to predict undersampling (Optional)

Benchmark participants may wish to investigate and develop metrics for predicting the occurrence of undersampling in order to detect and prevent its occurrence in Monte Carlo simulations. Participants should focus on examining metrics that can identify tally biases and inaccurate tally uncertainty estimates while Monte Carlo simulations are being performed and can instruct Monte Carlo codes to increase the number of particles being simulated in each generation if undersampling is detected.

Results Reporting

A spreadsheet is provided with this benchmark specification. The results shall be reported in the format specified in the spreadsheet. Each participant can select the number of particles per generation and the number of steps in each 1/M plot. The axial fission density data will be collected in the same spreadsheet for the 3D cases.

References

1. F. B. Brown, "'K-effective of the World' and Other Concerns for Monte Carlo Eigenvalue Calculations," *Progress in NUCLEAR SCIENCE and TECHNOLOGY*, **2**, 738-742 (2011).
2. G. E. Whitesides, "A Difficulty in Computing the k-effective of the World," *Trans. Am. Nucl. Soc.,* **14**[2], 680 (1971).
3. M. Nakagawa, T. Mori, "Whole core calculations of power reactors by use of Monte Carlo method," *J. Nuc. Sci. Techol*., **30** [7], 692-701 (1993).
4. B. T. Mervin, S. W. Mosher, J. C. Wagner, G. I. Maldonado, "Uncertainty Under-Prediction in Monte Carlo Eigenvalue Calculations," *Nucl. Sci. & Eng.*, **173**, 276–292 (2013).
5. G. Radulescu, D. E. Mueller, and J. C. Wagner, "Sensitivity and Uncertainty Analysis of Commercial Reactor Criticals for Burnup Credit," *Nuclear Technology* **167**(2), pp.268-287 (2009).
6. J. C. Wagner, *Computational Benchmark for Estimation of Reactivity Margin from Fission Products and Minor Actinides in PWR Burnup Credit*, NUREG/CR-6747, ORNL/TM-2000/306, U.S. Nuclear Regulatory Commission, Washington, D.C. (2001).
7. J. C. Wagner, M. D. DeHart, and C. V. Parks, *Recommendations for Addressing Axial Burnup in PWR Burnup Credit Analyses*, NUREG/CR-6801, ORNL/TM-2001/273, U.S. Nuclear Regulatory Commission, Washington, D.C. (2003).
8. C. M. Perfetti, B. T. Rearden, "Quantifying the Effect of Undersampling in Monte Carlo Simulations Using SCALE," in *Proc. of the 2014 International Conference on the Physics of Reactors (PHYSOR 2014),* Kyoto, Japan, September 28 – October 3, 2014.