

Hybrid and Parallel Domain-Decomposition Methods Development to Enable Monte Carlo for Reactor Analyses*

JOHN C. WAGNER,** SCOTT W. MOSHER, THOMAS M. EVANS, DOUGLAS E. PELOW and JOHN A. TURNER

Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA

This paper describes code and methods development at the Oak Ridge National Laboratory focused on enabling high-fidelity, large-scale reactor analyses with Monte Carlo (MC). Current state-of-the-art tools and methods used to perform “real” commercial reactor analyses have several undesirable features, the most significant of which is the non-rigorous spatial decomposition scheme. Monte Carlo methods, which allow detailed and accurate modeling of the full geometry and are considered the “gold standard” for radiation transport solutions, are playing an ever-increasing role in correcting and/or verifying the deterministic, multi-level spatial decomposition methodology in current practice. However, the prohibitive computational requirements associated with obtaining fully converged, system-wide solutions restrict the role of MC to benchmarking deterministic results at a limited number of state-points for a limited number of relevant quantities. The goal of this research is to change this paradigm by enabling direct use of MC for full-core reactor analyses. The most significant of the many technical challenges that must be overcome are the slow, non-uniform convergence of system-wide MC estimates and the memory requirements associated with detailed solutions throughout a reactor (problems involving hundreds of millions of different material and tally regions due to fuel irradiation, temperature distributions, and the needs associated with multi-physics code coupling). To address these challenges, our research has focused on the development and implementation of (1) a novel hybrid deterministic/MC method for determining high-precision fluxes throughout the problem space in k -eigenvalue problems and (2) an efficient MC domain-decomposition (DD) algorithm that partitions the problem phase space onto multiple processors for massively parallel systems, with statistical uncertainty estimation. The hybrid method development is based on an extension of the FW-CADIS method, which attempts to achieve uniform statistical uncertainty throughout a designated problem space. The MC DD development is being implemented in conjunction with the Denovo deterministic radiation transport package to have direct access to the 3-D, massively parallel discrete-ordinates solver (to support the hybrid method) and the associated parallel routines and structure. This paper describes the hybrid method, its implementation, and initial testing results for a realistic 2-D quarter core pressurized-water reactor model and also describes the MC DD algorithm and its implementation.

KEYWORDS: *hybrid methods, FW-CADIS, variance reduction, domain-decomposed Monte Carlo*

I. Introduction

Current state-of-the-art tools and methods used to perform “real” commercial reactor analyses use high-fidelity deterministic transport codes to produce few-group parameters at the assembly level for use in low-order methods applied at the core level. Monte Carlo (MC) methods allow detailed and accurate modeling of the full geometry and energy details, and are considered the “gold standard” for radiation transport solutions. These are playing an ever-increasing role in correcting and/or verifying the several-decades-old methodology used in current practice. However, the prohibitive computational requirements

associated with obtaining fully converged system-wide solutions restrict the role of MC to benchmarking deterministic results at a limited number of state-points for a limited number of relevant quantities. A goal of current research at Oak Ridge National Laboratory (ORNL) is to change this paradigm by enabling direct use of MC for full-core reactor analyses. The most significant of the many technical challenges that must be overcome are the slow, non-uniform convergence of system-wide MC estimates and the memory requirements associated with detailed solutions for reactor problems involving hundreds of millions of material and tally regions due to fuel irradiation, temperature distributions, and the need to couple with multiple physics packages. To address these challenges, our research has focused on development in the following two areas: (1) a hybrid deterministic/MC method for determining high-precision fluxes throughout the problem space in k -eigenvalue problems and (2) an efficient MC domain-decomposition (DD) algorithm that partitions the problem phase space onto multiple processors for massively parallel systems, with statistical uncertainty estimation. This

*Notice: This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

**Corresponding author, E-mail: wagnerjc@ornl.gov

paper describes the extension of the FW-CADIS¹⁾ method to variance reduction of MC reactor analyses and provides initial results for calculating group-wise fluxes throughout a generic 2-D pressurized water reactor (PWR) quarter core model. This paper also describes a new MC DD algorithm developed specifically for this application to reduce communication cost and enable variance estimates.

II. Hybrid Method

Hybrid methods that utilize approximate forward and/or adjoint fluxes from deterministic transport calculations to generate variance reduction (biasing) parameters for accelerating fixed-source MC simulations have been developed and implemented at ORNL for use in SCALE²⁾ and MCNP.³⁾

The two principal hybrid methods used at ORNL are (1) Consistent Adjoint Driven Importance Sampling (CADIS)^{4,5)} for optimization of a localized detector (tally) region (e.g., flux, dose, or reaction rate at a particular location) and (2) Forward-Weighted CADIS (FW-CADIS) for optimizing distributions (e.g., mesh tallies over all or part of the problem space) or multiple localized detector regions (e.g., simultaneous optimization of two or more localized tally regions). FW-CADIS is an extension of the CADIS method, which has been used for more than a decade to dramatically improve the efficiency of MC simulations for a variety of fixed-source source-detector-type problems. See Refs. 4–10 for examples. The basis of the FW-CADIS method is the development of a function that represents the importance of particles to the objective of uniform MC particle density in the desired tally regions; the details of this method are described in Ref. 1. This method is significant in that it enables high-fidelity MC results for distributions (e.g., spatial dose distribution) in large phase space regions, a capability typically attributed only to deterministic methods.

Implementation of the FW-CADIS method for fixed-source simulations¹⁾ involves two deterministic transport calculations: a forward fixed-source calculation consistent with the problem being solved, followed by a fixed-source adjoint calculation in which the adjoint source is weighted by the inverse of the forward distribution being sought. The adjoint function from the second calculation is used to calculate consistent space- and energy-dependent source and transport (weight windows) biasing parameters for use in the MC calculation. In contrast to the fixed-source problems for which CADIS and FW-CADIS were originally developed and applied, reactor analyses involve k -eigenvalue calculations to calculate space- and energy-dependent fluxes that may be coupled with depletion/decay calculations to simulate the burnup of nuclear fuel. Enabling MC reactor analyses requires the ability to perform eigenvalue simulations that obtain sufficiently precise scalar fluxes for computing reaction rates throughout the reactor core. Thus, as with previous fixed-source applications, the goal of the FW-CADIS method for reactor analysis is to determine space- and energy-dependent distributions with acceptable

statistical uncertainties throughout the entire problem space.

With the FW-CADIS method, the adjoint source is weighted by the inverse of the forward distribution being sought, which in this case is the space-dependent, group-wise flux. Extension of the FW-CADIS method to eigenvalue problems requires a forward eigenvalue deterministic calculation, as opposed to a fixed-source deterministic calculation, to compute approximate scalar fluxes for weighting the adjoint source in a subsequent fixed-source deterministic adjoint calculation in which fission is treated as capture. The adjoint function is then used to generate source and transport biasing parameters with the CADIS methodology for the eigenvalue MC calculation, in which the optimization objective is not the eigenvalue, but rather the multigroup fluxes in fissionable regions. For this last step, it is necessary to use the fission source from the forward deterministic calculation to calculate the consistent source and transport biasing parameters. The steps to calculate the FW-CADIS weight windows for a 2-D quarter-core PWR model are visually illustrated in **Fig. 1**. In the initial implementation, rather than biasing the selection of the source particles as is done in fixed-source problems, the source particles are sampled normally (i.e., as in an analog simulation). The source particle weight is subsequently modified by the weight window. Investigations to evaluate this aspect of the implementation are planned for future work.

The spatial fission distribution from the forward deterministic calculation also can be used as the starting source distribution in the MC calculation to reduce the number of initial skipped cycles and help ensure reliable convergence.¹¹⁾

1. Implementation

Once the importance function is determined in the FW-CADIS method, the CADIS methodology is used to calculate consistent space- and energy-dependent biasing parameters as has been described elsewhere.^{4,5)} This capability is implemented and automated for fixed-source applications in the SCALE MAVRIC sequence which uses the SCALE Monaco code for multigroup MC calculations, and in the ADVANTG code¹²⁾ which uses MCNP for continuous-energy MC calculations. Both MAVRIC and ADVANTG use the Denovo 3-D discrete ordinates code for the deterministic calculations.^{13,14)} The FW-CADIS method for eigenvalue calculations described above has initially been implemented and tested in the ADVANTG code.

2. Results

To evaluate the effectiveness of the FW-CADIS method, it was applied to a generic 2-D PWR quarter core model (see **Fig. 2**). First, resonance-shielded and pin-cell homogenized cross sections over eight energy groups (with boundaries at 20 MeV; 820 and 17 keV; and 3, 0.625, 0.275, 0.15, 0.05, and 10^{-5} eV) were generated with SCALE.²⁾ These cross sections were used in forward eigenvalue and adjoint fixed-source calculations with Denovo, using an S_4

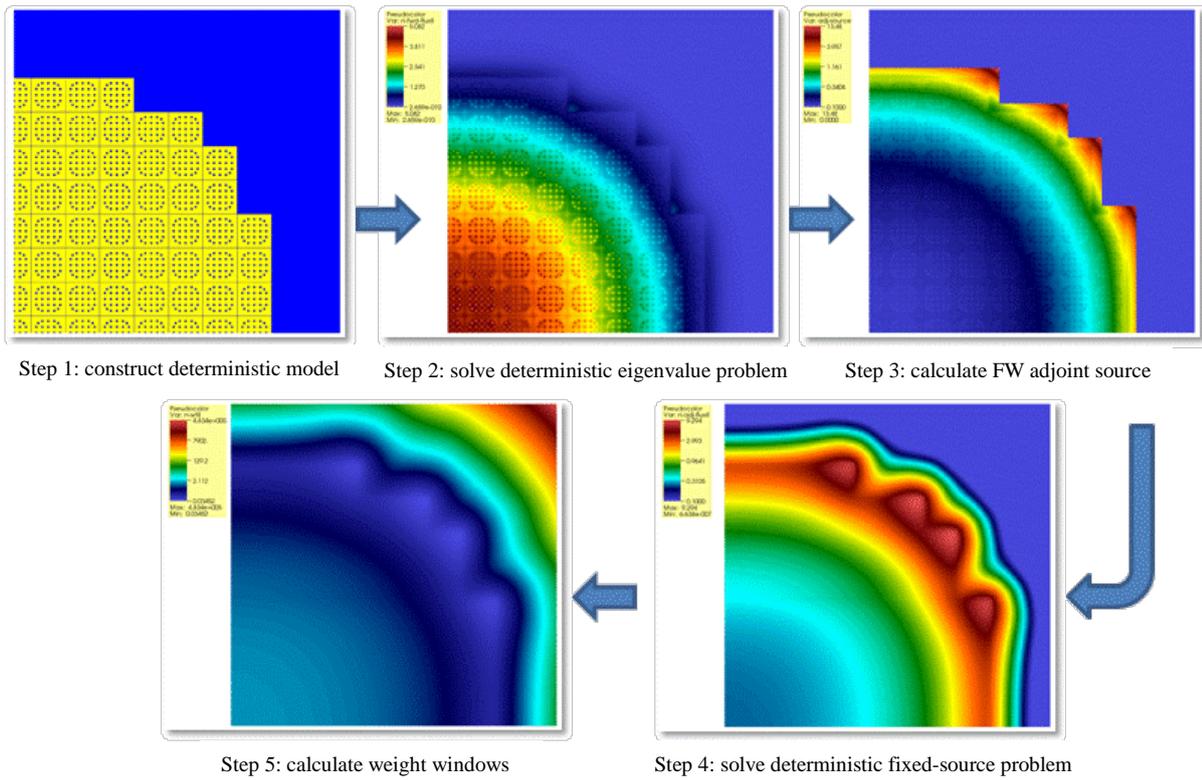


Fig. 1 Visual representation of the steps for calculating the FW-CADIS weight windows for the 2-D PWR quarter core model.

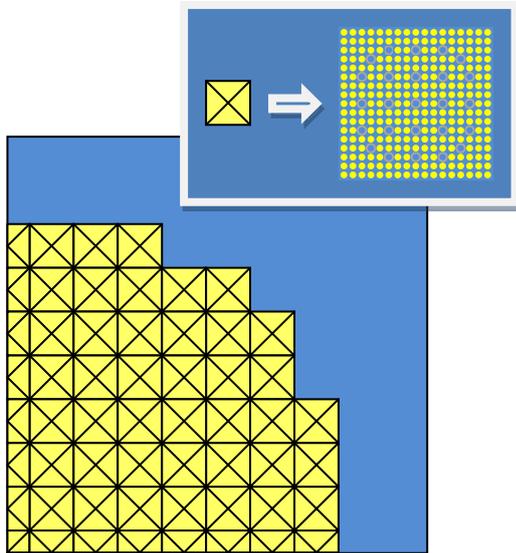


Fig. 2 2-D PWR quarter core model.

level-symmetric quadrature set and a $181 \times 181 \times 1$ spatial grid corresponding to pin-cells. With an inner iteration criterion of 10^{-3} on the flux residual, the total time for the forward and adjoint deterministic calculations was 50 minutes. Continuous-energy MCNP5 simulations were then performed with and without biasing parameters from the FW-CADIS method with 15,000 histories/cycle and 50 inactive cycles starting from an initial fission source

distributed uniformly across fuel pin centers. With the FW-CADIS parameters, 600 active cycles were completed in 249 minutes. To generate results for a fair comparison, a conventional MCNP5 simulation was performed with a time limit of 299 minutes (the total time consumed by the FW-CADIS transport calculations) and completed 1,068 active cycles. Results were generated using a mesh tally with pin-cell-sized spatial cells and eight energy bins. All calculations were performed using a single core of a machine with two 2.21 GHz quad-core AMD OpteronTM and 16 GB of shared memory.

In **Fig. 3**, the relative uncertainties of the tallied flux values in one of the thermal energy tally bins are plotted. The more uniform convergence obtained with the FW-CADIS method is clearly illustrated. The distribution of relative uncertainties for tallied fluxes in all eight energy bins is shown in **Table 1** and **Fig. 4**. (Only the tally cells within fuel assemblies are included in the results.) The conventional MCNP5 simulation does produce a significant number of tally cells with relative uncertainties of less than 1%, whereas the FW-CADIS method produced no flux estimates converged to this level. However, as is the intention of the method, a much narrower range of uncertainties is produced with FW-CADIS (1.0–6.6% versus 0.6–16.2%). Consequently, it would take ~ 322.8 hours for a conventional MCNP5 simulation to produce tally results in which all fuel-pin-bearing cells have relative uncertainties less than 2%. With the FW-CADIS method, this objective

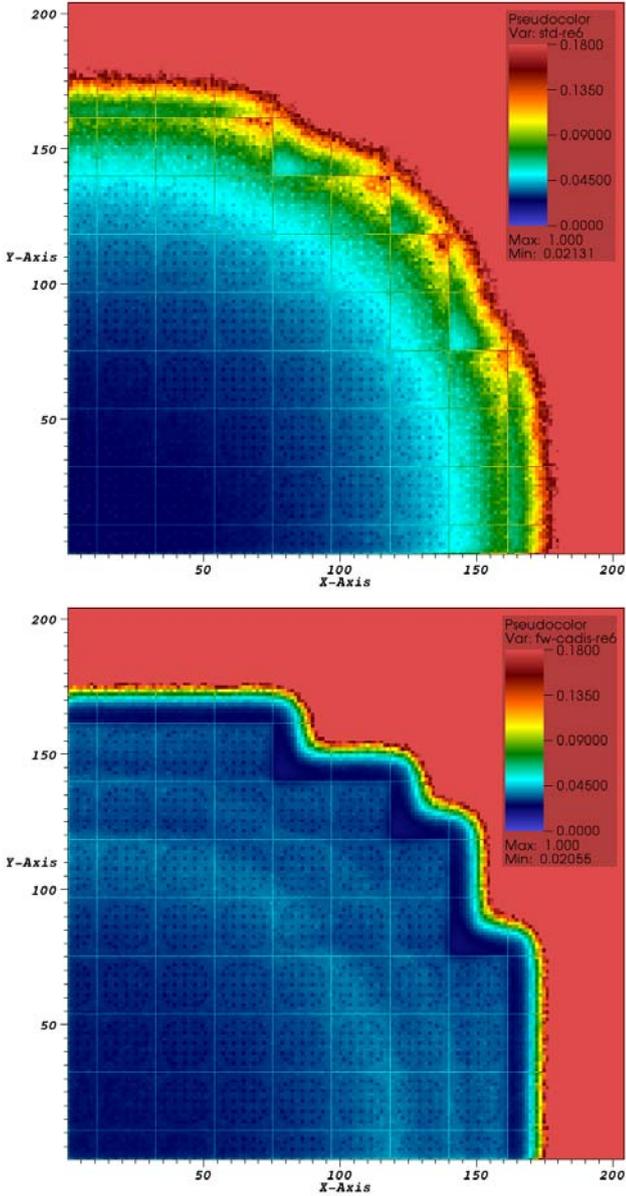


Fig. 3 Relative uncertainties in energy bin 5 (0.15 to 0.275 eV) computed using MCNP5 with (*bottom*) and without (*top*) FW-CADIS method.

Table 1 Number of space-energy flux tally cells by relative uncertainty (RU) interval using MCNP5 with and without the FW-CADIS method

RU %	MCNP	FW-CADIS	RU %	MCNP	FW-CADIS
0 - 1	11,127	0	9 - 10	397	0
1 - 2	31,670	38,233	10 - 11	212	0
2 - 3	34,229	25,571	11 - 12	113	0
3 - 4	16,382	41,652	12 - 13	55	0
4 - 5	8,301	6,872	13 - 14	24	0
5 - 6	4,674	227	14 - 15	13	0
6 - 7	2,911	21	15 - 16	1	0
7 - 8	1,651	0	16 - 17	1	0
8 - 9	815	0			

can be achieved in ~ 45.4 hours, including the deterministic run time (a speed-up of 7.1). Analyses were performed with different numbers of histories/cycle, different numbers of active cycles, and a less-strict convergence criterion for the Denovo calculations. The speed-up was observed to be between 6 and 10 in all cases.

Based on our experience with fixed-source problems, where speed-ups of $O(10^{2-4})$ are common,⁴⁻¹⁰ we expect the speed-up to be larger for problems that are computationally more difficult. Hence, for problems with a larger variation in flux values, e.g., smaller or less homogeneous systems, and/or finer resolution in tallies, either in terms of space or energy, we expect speed-ups greater than those observed with this test problem.

III. Monte Carlo Domain-Decomposition Algorithm

The new MC DD algorithm extends an algorithm,^{15,16} developed for the National Nuclear Security Administration's Advanced Simulation and Computing program, to address the challenges and exploit the characteristics of light-water reactor (LWR) simulations. The algorithm decomposes the entire problem domain over a set of processor cores (e.g., 256 to 512 cores) that is small enough to be localized within a single physical cabinet (or rack), and then replicates the problem across multiple sets so that each set simulates an independent batch of particles. The benefit of this multi-level scheme is threefold. First, there is no communication (message passing) between sets within a transport cycle. Second, the within-cycle communication that does occur takes advantage of the lower latency of within-cabinet communications (versus inter-cabinet communications). Third, we can estimate statistical uncertainties based on the variance of the independent batches, a requirement not handled by existing domain decomposition schemes. This will eliminate the under-prediction of statistical uncertainties due to cycle-to-cycle correlations between fission generations.

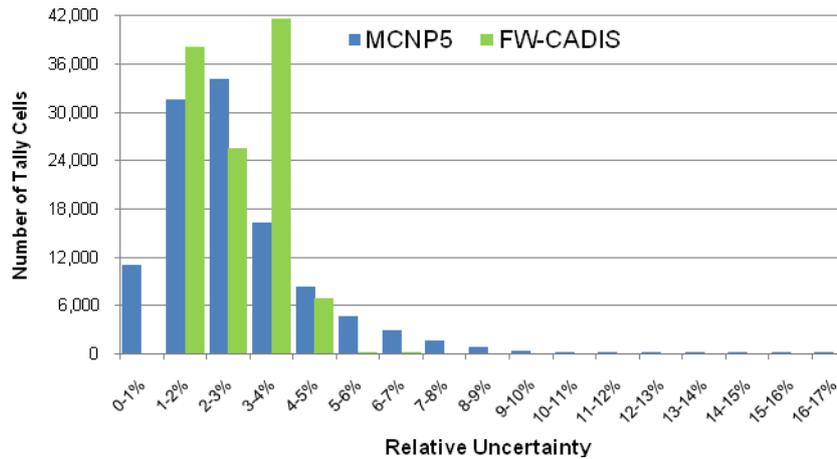


Fig. 4 Relative uncertainty distribution for tallied fluxes using MCNP5 with and without the FW-CADIS method.

Another important aspect of this new MC DD algorithm is that, within a set, the decomposed domains, referred to as blocks, are defined to overlap with one another to a significant degree (e.g., by one-half to one assembly width) to reduce the number of particles that scatter back and forth between blocks, thereby reducing communication cost. The block size is defined to cover a portion of the problem space (e.g., 4 to 9 assemblies in the radial dimension) and the simulated particle (fission) source is restricted to the central portion of the block (e.g., the center assembly in the block). To provide insight into the potential benefits of using overlapping regions, a study to evaluate block-to-block communication in a 2-D PWR core problem was performed. **Table 2** provides the results of the study and demonstrates that overlapping regions can be effectively used to minimize the amount of block-to-block communication during a transport cycle. Hence, for this type of application, the use of overlapping regions can virtually eliminate the costs associated with particles that scatter back and forth between blocks when no overlap is used.

Table 2 Block-to-block particle communication in a 2-D PWR core problem (Source is confined to the center assembly within the block.)

Width of over-lapping region (assemblies)	Total size of block (assemblies)	Source particles in the central block region that leave the block (percentage)
No overlap	1	76.12
0.5	2 × 2	21.03
1.0	3 × 3	4.54
1.5	4 × 4	1.05
2	5 × 5	0.27
2.5	6 × 6	0.07
3.0	7 × 7	0.02

This new multi-level overlapping-region (MLOR) DD algorithm is illustrated in **Fig. 5** and is currently being implemented and tested within the Denovo radiation

transport package. This permits direct access to the 3-D, massively parallel discrete-ordinates solver and its associated parallel routines and structure (to support the hybrid method and automated determination of domain boundaries to optimize load-balancing). Although implementation and testing are still under way, based on initial scaling studies this new MLOR DD algorithm is expected to significantly outperform existing algorithms.

IV. Summary

The FW-CADIS method for optimizing global distributions has been extended to k -eigenvalue calculations for reactor analyses and initially tested on a generic 2-D PWR core model. The method requires one forward and one adjoint discrete ordinates calculation to generate consistent source biasing and weight-window parameters for the subsequent MC simulation. It does not require any modifications to existing MC codes. Although further testing and analysis of results are still needed, the potential of this method to improve the efficiency and reliability of MC calculations for reactor analyses appears promising.

Additionally, a new MC DD algorithm has been developed that uses a multi-level set/block decomposition for load balancing and to reduce communication time/cost and overlapping regions to reduce communications. Initial implementation and testing indicate that this new MLOR DD algorithm should outperform existing MC DD algorithms for LWR analyses, and will enable variance estimation based on the variance of individual batches across sets.

Planned next steps related to the hybrid method include (1) application to a 3-D core model, (2) implementation to automatically use the converged fission source distribution from the forward Denovo calculation as the starting source for the MC calculation, (3) optimization for directly calculating reaction rates, and (4) more thorough results analysis and testing. Planned next steps related to the MLOR DD algorithm include completing the initial implementation and thorough testing.

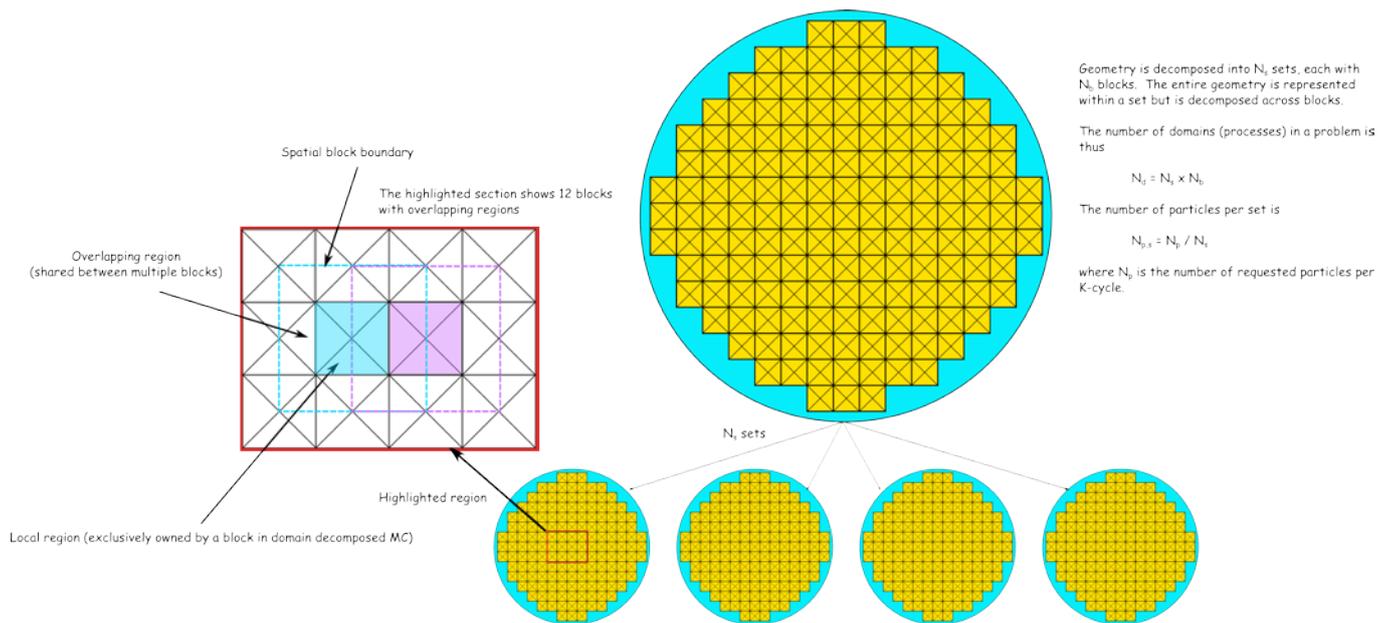


Fig. 5 Pictorial description of the multi-level overlapping-region (MLOR) parallel decomposition. Here the core geometry is decomposed into $N_s = 4$ sets. Each set has N_b blocks with overlapping regions such that the total number of parallel domains is $N_b \times N_s$. Particles are decomposed across sets where the number of particles per set is $N_{p,s} = N_p / N_s$. Each block can define an overlapping domain (shown in inset) to reduce block-to-block communication for particles that scatter at the interfaces between blocks. Optimal block sizes will be a compromise between computational and parallel efficiency and problem size.

Acknowledgments

Research sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U.S. Department of Energy.

References

- 1) J. C. Wagner, E. D. Blakeman, D. E. Peplow, "Forward-Weighted CADIS Method for Variance Reduction of Monte Carlo Calculations of Distributions and Multiple Localized Quantities," M&C 2009, Saratoga Springs, N.Y. (May 3–7, 2009).
- 2) *SCALE: A Modular Code System for Performing Standardized Computer Analyses for Licensing Evaluation*, ORNL/TM-2005/39, Version 6, Vols. I–III. Available from Radiation Safety Information Computational Center at Oak Ridge National Laboratory as CCC-750 (January 2009).
- 3) X-5 MONTE CARLO TEAM, *MCNP—A General Monte Carlo N-Particle Transport Code, Version 5. Volume I: Overview and Theory*, LA-UR-03-1987, Los Alamos National Laboratory (2003).
- 4) J. C. Wagner, A. Haghghat, "Automated Variance Reduction of Monte Carlo Shielding Calculations Using the Discrete Ordinates Adjoint Function," *Nucl. Sci. Eng.*, **128**, 186 (1998).
- 5) A. Haghghat, J. C. Wagner, "Monte Carlo Variance Reduction with Deterministic Importance Functions," *Prog. Nucl. Energy*, **42**(1), 25–53 (2003).
- 6) A. Haghghat, H. Hiruta, B. Petrovic, and J.C. Wagner, "Performance of the Automated Adjoint Accelerated MCNP (A³MCNP) for Simulation of a BWR Core Shroud Problem," *Proc. International Conference on Mathematics and Computation, Reactor Physics, and Environmental Analysis in Nuclear Applications*, Madrid, Spain, September 27–30, 1999.
- 7) H. P. Smith and J. C. Wagner, "A Case Study in Manual and Automated Monte Carlo Variance Reduction with a Deep Penetration Reactor Shielding Problem," *Nucl. Sci. Eng.*, **149**, 23–37 (2005).
- 8) S. W. Mosher, T. M. Miller, T. M. Evans, and J. C. Wagner, "Automated Weight-Window Generation for Threat Detection Applications using ADVANTG," M&C 2009, Saratoga Springs, N.Y. (May 3–7, 2009).
- 9) J. C. Wagner, D. E. Pelow, T.M. Evans, "Automated Variance Reduction Applied to Nuclear Well-Logging Problems," *Nucl. Technol.* **168**, 799–809, December 2009.
- 10) A. M. Ibrahim, S. W. Mosher, T. M. Evans, D. E. Peplow, M. E. Sawan, P. P. H. Wilson, J. C. Wagner, "ITER Neutronics Modeling Using Hybrid Monte Carlo/Deterministic and CAD-based Monte Carlo Methods," accepted for publication *Nucl. Technol.*
- 11) M. T. Wenner, J. C. Wagner, L. M. Petrie, "A New Source Initialization Capability for the SCALE Burnup Credit Sequence," *Trans. Am. Nucl. Soc.*, **95**, 310–311 (2006).
- 12) S. W. Mosher, "A New Version of the ADVANTG Variance Reduction Generator," presented at the American Nuclear Society Radiation Protection and Shielding Division 2010 Topical Meeting, Las Vegas, Nev., April 18–23 (2010).
- 13) T. M. Evans, A. S. Stafford, K. T. Clarno, "Denovo—A New Three-Dimensional Parallel Discrete Ordinates Code in SCALE," *Nucl. Technol.*, **171**, 171–200 (2010).
- 14) T.M. Evans, G.G. Davidson, and R.N. Slaybaugh, "Three-Dimensional Full Core Power Calculations for Pressurized Water Reactors," *Journal of Physics: Conference Series*, accepted (2010).
- 15) T. A. Brunner, T. J. Urbatsch, T. M. Evans, N.A. Gentile, "Comparison of Four Parallel Algorithms for Domain Decomposed Implicit Monte Carlo," *Journal of Computational Physics* **212**, 527–539 (2006).
- 16) T. A. Brunner, P.S. Brantley, "An Efficient, Robust, Domain-Decomposition Algorithm for Particle Monte Carlo," *Journal of Computational Physics* **228**, 3882–3890 (2009)