

# LA-UR-12-24277

Approved for public release; distribution is unlimited.

Title: Reactor Physics Verification of the MCNP6 Unstructured Mesh Capability

Author(s): Burke, Timothy P.  
Kiedrowski, Brian C.  
Martz, Roger L.  
Martin, William R.

Intended for: International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, 2013-05-05/2013-05-09 (Sun Valley, Idaho, United States)



**Disclaimer:**

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

## **REACTOR PHYSICS VERIFICATION OF THE MCNP6 UNSTRUCTURED MESH CAPABILITY**

**Timothy P. Burke**

Department of Nuclear Engineering and Radiological Sciences  
University of Michigan  
2355 Bonisteel Boulevard, Ann Arbor, MI 48109, USA  
tpburke@umich.edu

**Brian C. Kiedrowski and Roger L. Martz**

X-Computational Physics Division, Monte Carlo Codes Group  
Los Alamos National Laboratory  
P.O. Box 1663, MS A143 Los Alamos, NM 87544  
bckiedro@lanl.gov; martz@lanl.gov

**William R. Martin**

Department of Nuclear Engineering and Radiological Sciences  
University of Michigan  
2355 Bonisteel Boulevard, Ann Arbor, MI 48109, USA  
wrm@umich.edu

### **ABSTRACT**

The Monte Carlo software package MCNP6 has the ability to transport particles on unstructured meshes generated from the Computed-Aided Engineering software Abaqus. Verification is performed using benchmarks with features relevant to reactor physics – Big Ten and the C5G7 computational benchmark. Various meshing strategies are tested and results are compared to reference solutions. Computational performance results are also given. The conclusions show MCNP6 is capable of producing accurate calculations for reactor physics geometries and the computational requirements for small lattice benchmarks are reasonable on modern computing platforms.

*Key Words:* Monte Carlo, Big Ten, C5G7, Geometry

### **1. INTRODUCTION**

MCNP6 [1], a Los Alamos National Laboratory Monte Carlo radiation transport software package, has the capability to perform particle tracking on unstructured meshes generated by Abaqus/CAE [2]. The driver for this development has been related to fixed-source problems, and significant verification has been performed for those types of calculations [3]; however, less attention has been given to eigenvalue problems, which are important for reactor physics and criticality safety. Preliminary eigenvalue verification has been performed [4] using the simple Godiva sphere. To show this capability is useful for a broader class of eigenvalue problems, verification using more complicated systems is needed. This work specifically focuses on problems with geometric features typical of reactor physics problems.

To accomplish this, a criticality benchmark and a fuel assembly benchmark were used for calculations in

MCNP6 using both the Constructive Solid Geometry (CSG) native to MCNP6 and the unstructured mesh geometry generated using Abaqus/CAE. Specifically, the Big Ten criticality benchmark [5] and the C5G7 computational benchmark [6] were selected. The Big Ten criticality benchmark was selected because it has a relatively simple geometry containing multiple cylindrical components. The C5G7 3-D Mixed Oxide (MOX) Fuel Assembly Benchmark was chosen to test the unstructured mesh capabilities on a semi-realistic small reactor problem.

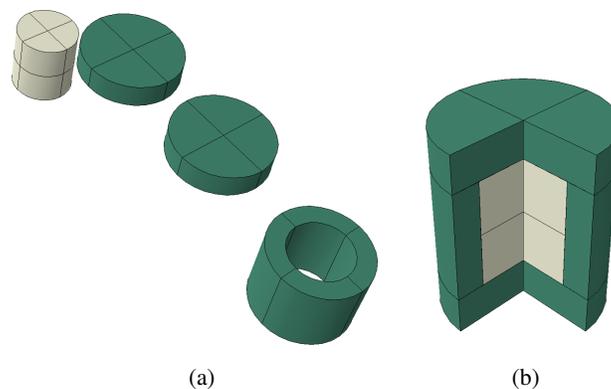
The results show that the MCNP6 unstructured mesh capability can match CSG results ( $k_{\text{eff}}$  and local pin powers) if the models are prepared with appropriate meshes. Accurate results can be achieved in the C5G7 benchmark with reasonable computational resources by modern standards.

## 2. BIG TEN CRITICALITY BENCHMARK

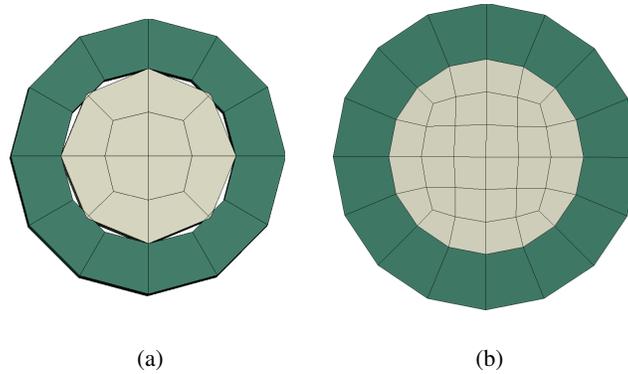
### 2.1. Geometry and Meshing

The Big Ten criticality benchmark has a fairly simple geometric specification: a low-enriched uranium cylindrical core region surrounded by an axisymmetric cylindrical reflector consisting of depleted uranium. Big Ten was selected because it features cylindrical components, which are geometrically similar to fuel pins found in reactors, although the length scales are quite different. Additionally, the simplicity of this model – featuring only two components – allows for easier verification of the tracking algorithms handling nested cylindrical bodies and provides some insight as to optimal strategies for developing similar, but more complicated, models.

In Abaqus/CAE, users have a variety of methods available to create their models, and some of those methods lead to suboptimal performance for radiation transport calculations. Users can create separate parts for various sections of their model that are meshed independently and MCNP6 can perform transport on those multiple parts. The case using this meshing method is referred to as multi-part model. The multi-part model has three components: a cylinder for the fuel, an annulus for the reflector surrounding the fuel, and one part for the top and bottom of the cylindrical reflector. These independently meshed components are then assembled together in Abaqus/CAE to create the final representation of the model. The parts comprising the multi-part model are shown in Fig. 1a, and a cutaway view of the assembled



**Figure 1: (a) Individual components and (b) assembled multi-part Big Ten model.**



**Figure 2: Top view of meshes for the (a) multi-part and (b) merged-part Big Ten models.**

model is displayed in Fig. 1b. Alternatively, users can merge parts together and mesh the merged part to create what is referred to here as the merged-part model.

The impact of these two modeling techniques on transport calculations was tested using both models. The differences arise when the parts are meshed. The merged-part model has a contiguous mesh, meaning the mesh of one material region shares nodes with the meshes of adjacent regions. This is not a necessary condition in the model composed of multiple parts, where there can be gaps and overlaps in the geometry. The difference can be visualized when a coarse mesh is imposed on the models; Fig. 2 shows radial cross sections of the multi-part and merged-part Big Ten models meshed with first-order hexahedron elements with an 18 cm global seed size – the seed size gives Abaqus/CAE an element edge size that it attempts to match, but the edges will usually not be this size exactly because of geometric conformity issues. As seen in Fig. 2, there are no gaps or overlaps in the merged-part model while there are a significant amount of gaps and overlaps in the multi-part model even though both models use the same mesh seed size. The user can specify a smaller seed to reduce gaps and overlaps.

The accuracy of the unstructured mesh for both models was tested with various meshing strategies. A total of 16 different unstructured mesh models were created using Abaqus/CAE with each model comprised of either first- or second-order tetrahedra or hexahedra. The cases considered for the simple cylindrical model are described in Table I. Four different global seed sizes were used per element type: 1, 3, 6, and 18 cm. All parts of the model have an axial seed size of 10 cm. This was done to prevent prohibitive runtimes with smaller seed sizes. Without the 10 cm axial seed size, element counts are on the order of hundreds of thousands for global seed sizes of interest. This reduction in element count allowed for reasonable computer times (less than a CPU day) and comparison of computing statistics for the seed sizes chosen.

## 2.2. Results

The calculations for the Big Ten model were done using 10 inactive cycles followed by 150 active cycles with 20,000 histories per cycle using the sequential version of MCNP6. Table I shows resulting eigenvalues of the various models and their 1- $\sigma$  statistical uncertainties in parentheses following the value. The “ $k_{\text{eff}}$  % Error” is the difference of  $k_{\text{eff}}$  from the CSG reference case given at the bottom of Table I.

Several trends are seen in the results. The eigenvalue strongly depends upon upon the mass of the fissile

**Table I: Comparison of Big Ten CSG and Unstructured Mesh Results.**

Element Type	Mesh Seed (cm)	Number of Elements	Fuel Volume % Error	Refl. Volume % Error	$k_{\text{eff}}$	$k_{\text{eff}}$ % Error	runtime (min)
1 <sup>st</sup> -Order Hex	1	36912	0.024	0.225	0.99276(29)	-0.215	139.00
	3	4848	0.210	0.277	0.99327(24)	-0.164	82.79
	6	1944	0.837	0.456	0.99187(30)	-0.305	68.35
	18	984	9.969	2.330	0.97400(29)	-2.101	59.52
1 <sup>st</sup> -Order Tet	1	139526	0.024	0.179	0.99327(28)	-0.164	178.07
	3	24233	0.210	0.237	0.99227(29)	-0.264	107.34
	6	10492	0.837	0.409	0.99243(30)	-0.248	84.67
	18	4993	9.969	2.284	0.97397(28)	-2.104	68.35
Multi-Part 2 <sup>nd</sup> -Order Hex	1	36912	0.000	-0.001	0.96533(29)	-2.972	1057.20
	3	4848	0.000	-0.001	0.96297(29)	-3.209	586.99
	6	1944	0.000	0.000	0.97364(31)	-2.137	549.78
	18	984	0.005	0.007	0.98499(29)	-0.996	698.45
2 <sup>nd</sup> -Order Tet	1	139526	0.000	-0.001	0.98038(27)	-1.459	916.91
	3	24233	0.000	-0.001	0.99086(28)	-0.406	606.54
	6	10492	0.000	0.000	0.99230(29)	-0.261	474.02
	8	4993	0.078	0.007	0.98639(29)	-0.855	450.69
1 <sup>st</sup> -Order Hex	1	47520	0.032	0.031	0.99499(30)	0.009	120.89
	3	5840	0.183	0.182	0.99504(29)	0.014	71.22
	6	1920	0.642	0.641	0.99404(28)	-0.086	61.69
	18	480	2.550	2.550	0.99101(32)	-0.391	55.95
1 <sup>st</sup> -Order Tet	1	329718	0.034	0.033	0.99531(27)	0.041	243.32
	3	42395	0.160	0.160	0.99521(29)	0.031	83.98
	6	12038	0.642	0.641	0.99456(29)	-0.034	67.09
	18	3501	2.550	2.550	0.99078(33)	-0.414	57.50
Merged Part 2 <sup>nd</sup> -Order Hex	1	47520	0.000	-0.001	0.99570(31)	0.080	837.71
	3	5840	0.000	-0.001	0.99507(30)	0.017	436.88
	6	1920	0.000	0.000	0.99463(29)	-0.027	392.89
	18	480	0.005	0.007	0.99433(29)	-0.057	424.63
2 <sup>nd</sup> -Order Tet	1	329718	0.000	0.000	0.99535(32)	0.045	978.93
	3	42395	0.000	0.000	0.99505(29)	0.015	562.02
	6	12038	0.000	0.000	0.99456(31)	-0.034	440.71
	18	3501	0.078	0.005	0.99500(31)	0.010	387.98
CSG	—	—	—	—	0.99490(29)	—	—

material in the problem, and as such the percent error in the volumes is a predictor of how well the unstructured mesh represents the actual problem. The general trend in the first-order element models is that as the number of elements increases, or the error in the representation of the volume decreases,  $k_{\text{eff}}$  converges on the value of the CSG model. This trend is less applicable for models meshed with second-order elements because of the element's higher degree of curvature, enabling them to accurately reproduce volumes even with a very coarse mesh. Even so, the coarse mesh with an 18 cm seed does not produce the correct eigenvalue. The larger second-order elements may preserve volume, however they may not preserve the shape of the problem. Because this assembly is small relative to the neutron mean-free paths (which are long because Big Ten is a fast assembly), the neutron leakage in the coarse mesh differs from the CSG model and thus produces a different  $k_{\text{eff}}$ .

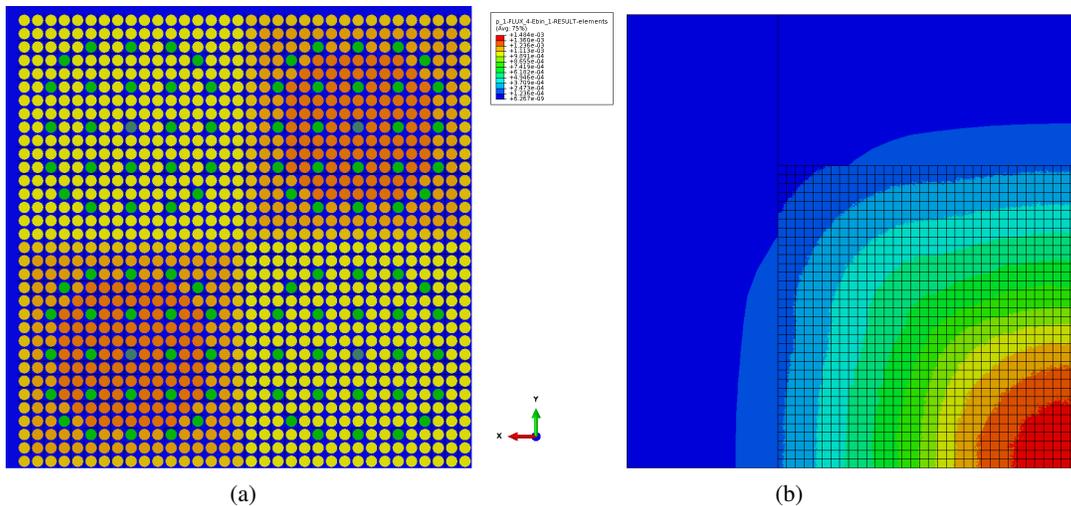
There is a large difference between the resulting  $k_{\text{eff}}$  of the merged-part and multi-part models. The merged-part model with both first- and second-order elements converges to the same eigenvalue as the CSG model (within  $1-\sigma$  uncertainty) when the mesh seed is decreased. The first-order elements for the multi-part model also appear to converge, but at a much slower rate.

Even though gaps and overlaps are handled by the element-to-element particle tracking routines, the two models produce varying results because of a difference in the resulting geometries. Differences greater than 200 pcm are seen even when using fine meshes with over 30,000 elements. The merged-part model better represents the true geometry than the multi-part model, even though they both use the same mesh seeding. The contiguous mesh resulting from merging the various parts into a single part better represents the material interfaces. This demonstrates that users should consider merging parts that have curved interfaces because failing to do so may lead to significant inaccuracies.

At the time these calculations were performed, second-order elements have undergone less rigorous testing than the first-order elements. Currently, the results for the second-order elements in the multi-part model diverge from the correct value. This issue does not exist in the merged-part model, and indicates that there could be a tracking issue between curved surfaces of parts when using second-order elements. Furthermore, the runtimes for second-order hexahedra show erratic behavior for both the merged-part and the multi-part model; every other element type shows increased runtimes corresponding to an increase in element count. As MCNP6 development continues, the second-order element tracking methodologies will improve, addressing these issues.

Table I also shows that the merged-part model has faster runtimes than the multi-part model for similarly-meshed problems. In the merged-part model, particles tracking element-to-element take advantage of more efficient nearest neighbor lists. Whenever a particle reaches the edge of a part, other search routines determine if there is an element in another part to which it can transition. This makes tracking on multi-part models less efficient. This effect becomes less pronounced when the number of elements per part is increased.

Some conclusions can be drawn from the Big Ten simulations: First, merging parts yields shorter runtimes as long as the number of elements per part remains manageable. Furthermore, merging parts with curved interfaces yields more accurate results because of volume conservation and the elimination of gaps and overlaps. Additionally, first-order hexahedra provide the best runtimes, but first-order tetrahedra are comparable. Finally, volume conservation is not the only parameter to consider when doing eigenvalue calculations when leakage is important – shape preservation also matters.



**Figure 3: Abaqus/CAE visualizations of (a) the C5G7 core geometry and (b) scalar flux contours.**

### 3. C5G7 COMPUTATIONAL BENCHMARK

#### 3.1. Geometry and Meshing

The C5G7 benchmark is a small reactor physics benchmark designed in the early 2000's to primarily test the ability of deterministic transport software to handle detailed (no homogenization) lattice physics geometries. The benchmark consists of four assemblies: two UO<sub>2</sub> and two MOX. The MOX assembly consists of three different blends of MOX fuel. An Abaqus/CAE visualization of a slice showing the detailed lattice geometry is shown in Fig. 3a.

A total of three models of the C5G7 problem were created: a multi-part model where every pincell is composed of two parts, one for the fuel and one for the water surrounding the fuel, and two merged-part models with different mesh seeds where every pincell is composed of one part with two material regions. The first merged-part pincell (merged-pincell 1) model has a global seed of 4 cm, an edge seed of 0.15 cm, and a circumferential seed of 0.075 cm. The second merged-part pincell (merged-pincell 2) model has a global seed of 4 cm, an edge seed of 0.10 cm, and a circumferential seed of 0.05 cm. The fuel in the multi-part pincell model has a global seed of 0.15 cm, an axial seed of 4 cm, and a circumferential seed of 0.075 cm to obtain the same mesh as the merged-pincell 1 model. The water surrounding the fuel in the multi-part pincell model has a global seed of 0.15 cm, an axial seed of 4 cm, a circumferential seed of 0.05 cm, and an edge seed of 0.1 cm. Only the unrodded configuration (all of the control rods are suspended above the core) has been modeled.

#### 3.2. Results

Eigenvalue calculations were performed using 50,000 particles per cycle with 50 inactive cycles and 2,500 active cycles. All simulations were conducted on Intel Xeon E5-2670, 2.6 GHz processors with 64 MPI processes using two threads each. Results of the eigenvalue calculations for the three unstructured mesh models and the CSG model are presented in Table II. The reference CSG values were obtained from previously calculated benchmarks using MCNP CSG [6].

**Table II: Geometry, Eigenvalue Results, and Performance Data for C5G7.**

Model	Elms per Pincell	Fuel Volume % Error	Number of Elms	$k_{\text{eff}}$	Runtime (min per cycle)	Memory per MPI Proc. (GB)
Multi-Part Pincell	3594	-0.357	4.1 M	1.13033(18)	3.6	3.1
Merged-Pincell 1	3072	-0.357	3.7 M	1.14432(8)	1.06	2.8
Merged-Pincell 2	5952	-0.153	7.1 M	1.14350(8)	1.18	5.2
CSG	–	–	–	1.14308(3)	–	0.03

**Table III: Assembly Power Comparisons for the Various Axial Sections of C5G7.**

Section	Model	Inner UO2	Percent Error	MOX	Percent Error	Outer UO2	Percent Error
Whole Assembly	CSG	491.2	–	212.7	–	139.4	–
	UM	491.2	0.00	212.7	0.02	139.3	-0.08
Slice 1	CSG	219.0	–	94.5	–	62.1	–
	UM	219.0	-0.01	94.5	-0.05	62.1	0.04
Slice 2	CSG	174.2	–	75.2	–	49.5	–
	UM	174.2	0.02	75.3	0.10	49.3	-0.32
Slice 3	CSG	97.9	–	42.9	–	27.8	–
	UM	98.0	0.10	43.0	0.28	27.8	0.08

As expected from the results from the Big Ten unstructured mesh simulations, the multi-part pincell model produces less accurate results with  $k_{\text{eff}}$  being approximately 1,000 pcm lower than the CSG model. When the pincell is merged into a single part, the results improve dramatically with the merged-pincell 2 model producing results that disagree by only 42 pcm. While this is still well outside the (very tight) 3 pcm  $1-\sigma$  statistical uncertainty band of the reference  $k_{\text{eff}}$ , this difference could be reduced through improved meshing.

Performance parameters for the various simulations are also presented in Table II. They show that merging the pincell into one part has a significant runtime reduction, agreeing with what is seen from the Big Ten simulations. Merging multiple pincells (e.g., a cluster of 9 pincells) into one part may further improve performance. Furthermore, runtime does not significantly increase with the number of elements; a doubling of the number of elements in the model only produces a modest increase in runtimes. As such, the runtimes required to increase the number of elements to reduce the discrepancy in  $k_{\text{eff}}$  should not be prohibitive. However, unstructured mesh representations of assemblies require hundreds of megabytes to gigabytes of memory, and, as expected, increasing the number of elements corresponds to a roughly linear increase in memory requirements. Thus, the memory requirements pose a limitation on the number of elements capable of being modeled in the unstructured mesh. Note that in the MCNP6 version used, memory has not yet been optimized, but this will improve with future development.

Pin powers were obtained using an volume-averaged flux mesh tally (FMESH) in MCNP6 with a tally multiplier (FM card) to tally the absorption reaction rate in each axial pincell region. Using analysis tools provided in the reference benchmark [6], the reaction rates were then normalized to obtain pin powers where the total power is equal to the number of fuel pins. Comparisons of assembly powers for the various axial sections of the model are shown in Table III. The results all agree well within 1%, and the statistical uncertainties are all under 0.3% so it is unlikely this agreement is by chance. Only two assembly powers have errors larger than the statistical uncertainty reported by the reference CSG model, occurring in the low-power outer UO<sub>2</sub> assemblies. A contour plot generated by Abaqus/CAE of the MCNP6 calculated scalar flux profile on the mid-plane slice is displayed in Fig. 3b.

Further verification was performed on the consistency between the calculated volume-averaged fluxes with a mesh tally (FMESH) and unstructured mesh edits. This was done in a single calculation using the merged-pincell 2 model where the mesh tally and unstructured mesh edits conform geometrically. As expected, the results of both agree within four digits of precision, verifying that MCNP6 computes fluxes on the unstructured mesh edits correctly.

#### 4. CONCLUSIONS AND FUTURE WORK

MCNP6's unstructured mesh capability can accurately calculate eigenvalues and localized quantities for problems of interest to the reactor physics community. The ability to do this depends on the user's willingness to appropriately refine their models. First, a suitable mesh must be chosen to ensure preservation of fissile mass of each component, and, in some cases, preservation of shape to ensure correct leakage. Furthermore, merging parts with curved interfaces into a single part is important. Inaccurate eigenvalue results can be obtained because of gaps and overlaps.

Currently, second-order elements have significantly greater runtime, although less total memory is needed because their curvature allows fewer of them to achieve the same fidelity. Increasing the number of elements does not lead to significantly increased runtimes, but greater memory usage. Because of this, there are tradeoffs that a user must consider between runtime and memory limitations.

Modern computing resources are capable of handling the unstructured mesh representation of the C5G7, which is a simplistic representation of a small reactor. In the near future, memory availability is likely going to be the primary limitation. The C5G7 requires a few gigabytes of storage per MPI process, which is starting to approach memory limitations of typical modern computing platforms. Shared memory processing or threading with OpenMP available in MCNP6 helps reduce this memory requirement. Nonetheless, scaling this to full-core analysis of commercial reactors will almost certainly present challenges on memory storage and further research into various memory decomposition strategies is needed.

Additionally, the Abaqus/CAE (and other finite element analysis tools) offer the possibility of performing multiphysics analyses linking neutronics and thermal hydraulics. Some preliminary work on this has already been performed, and research and development down this avenue will continue.

#### ACKNOWLEDGMENTS

Funding for this work was provided by the U.S. DOE/NNSA Advanced Scientific Computing program.

## REFERENCES

- [1] MCNP6 Development Team, "Initial MCNP6 Release Overview - MCNP6 Beta 2," Los Alamos National Laboratory, LA-UR-11-07082 (2011).
- [2] Dessault Systems Simulia, Inc., "ABAQUS USER MANUALS, Version 6.11," Providence, RI (2011).
- [3] R. L. Martz, "MCNP6 Unstructured Mesh Initial Validation And Performance Results," *Nucl. Technol.*, accepted for publication (2012).
- [4] K. C. Kelley, R. L. Martz, D. L. Crane, "Riding Bare-Back on Unstructured Mesh for 21st Century Criticality Calculations," *Proc. PHYSOR 2010*, Pittsburgh, PA, USA (2010).
- [5] *International Handbook of Evaluated Critical Safety Benchmark Experiments*, Nuclear Energy Agency, NEA/NSC/DOC(95) 03 Vol III, September 2010 Edition (2010).
- [6] *Benchmark on Deterministic Transport Calculations Without Spatial Homogenization, MOX Fuel Assembly 3-D Extension Case*, Nuclear Energy Agency, NEA/NSC/DOC(2005)16 (2005).