




Enhancements to the Subplane Capability in MPACT

Shane Stimpson
Aaron Graham
Benjamin Collins
Oak Ridge National Laboratory

March 27, 2019



DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website <http://www.osti.gov>

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone 703-605-6000 (1-800-553-6847)
TDD 703-487-4639
Fax 703-605-6900
E-mail info@ntis.gov
Website <http://classic.ntis.gov/>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Website <http://www.osti.gov/contact.html>

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

REVISION LOG

Revision	Date	Affected Pages	Revision Description
0	12/21/2018	All	Original Document

Document pages that are:Export Controlled: NoneIP/Proprietary/NDA Controlled: NoneSensitive Controlled: NoneUnlimited: All

Approved for Public Release:

CONTENTS

FIGURES	vi
TABLES	vii
ACRONYMS	viii
1. INTRODUCTION	1
2. SUBPLANE BACKGROUND AND PREVIOUS WORK	1
3. OVERVIEW OF IMPROVEMENTS.....	2
3.1 Subgrid Treatment of Spacers.....	2
3.2 Spatial Multilevel CMFD Solver	3
3.3 Dynamic Subplane Mesh for Rod Movements.....	4
3.4 Usability Improvements.....	5
3.4.1 Automated Meshing	5
3.4.2 Flexible Axial Edit Specification	6
3.4.3 Easy to Use Input Specification	6
4. RESULTS	7
4.1 Overview of Cases Evaluated	7
4.2 Quarter Assembly (Progression Problem 3a)	7
4.2.1 Accuracy Assessment	10
4.2.2 Performance Improvements.....	11
4.3 Colorset (Progression Problem 4a)	12
4.3.1 Accuracy Assessment	12
4.3.2 Performance Improvements.....	13
4.3.3 Dynamic Meshing	14
4.4 Quarter Core (Progression Problem 5a).....	15
4.4.1 Accuracy Assessment	15
4.4.2 Performance Improvements.....	16
5. CONCLUSIONS.....	16
6. FUTURE WORK	17
6.1 Coupling Improvements.....	17
6.2 Improved Multilevel CMFD Convergence and Performance.....	17
6.3 Enhanced Axial Reflector Treatment.....	17
6.4 Depletion Improvements.....	17
7. ACKNOWLEDGMENTS	18
8. REFERENCES	18

FIGURES

Figure 2.1.1. Subplane Illustration [8].	2
Figure 3.1.1. Subgrid Illustration Applied to Spacer Grids.	3
Figure 3.2.1. Demonstration of Various Spatial Levels.....	4
Figure 3.3.1. Illustration of Multistate Variable Mesh.	5
Figure 4.2.1. Problem 3a Specification [12].	8
Figure 4.2.2. Axial Mesh Comparison Between Options.	9
Figure 4.2.3. Comparison of Various Grid Treatment Approaches.....	11
Figure 4.3.1. Problem 4a Colorset Radial Layout [12].....	12
Figure 4.4.1. Problem 5a Radial Layout [12].	15

TABLES

Table 4.2.1. Problem 3a Accuracy Comparisons.....	10
Table 4.2.2. Problem 3a Accuracy Comparisons (Tighter Tolerance)	10
Table 4.2.3. Problem 3a Performance Comparisons.....	12
Table 4.3.1. Problem 4a Accuracy Comparisons.....	13
Table 4.3.2. Problem 4a Accuracy Comparisons (Tighter Tolerance)	13
Table 4.3.3. Problem 4a Performance Comparisons.....	14
Table 4.3.4. Dynamic Meshing Results on 4a-IRW	14
Table 4.4.1. Problem 5a Accuracy Comparisons.....	16
Table 4.4.2. Problem 5a Performance Comparisons.....	16

ACRONYMS

AIC	silver-indium-cadmium
B ₄ C	boron carbide
BOC	beginning of cycle
CPM	collision probabilities method
CMFD	coarse mesh finite difference
CASL	Consortium for Advanced Simulation of Light Water Reactors
ML	multilevel
MOC	method of characteristics
ORNL	Oak Ridge National Laboratory
PWR	pressurized water reactor
VERA	Virtual Environment for Reactor Analysis

1. INTRODUCTION

In recent years, much of the development effort for the Virtual Environment for Reactor Applications (VERA) [1] has been devoted to improving performance and targeting industry-class computers. This includes reducing both runtime (improving efficiency) and core-count requirements to enable simulation on smaller machines. In MPACT [2], in particular, many aspects of the neutronic iteration strategy have been evaluated and improved, such as the subgroup [3] and method of characteristics (MOC) [4] runtimes, and many efforts have focused on improving the coarse mesh finite difference (CMFD) runtime [5,6,7,8], which has become a dominant component of the runtime as other aspects have been improved.

In the current default 2D/1D approach in MPACT, the axial meshes for the MOC plane and CMFD nodes have a one-to-one mapping (i.e., one CMFD node axially per MOC plane). However, in the subplane approach, multiple CMFD meshes are permitted axially within each MOC plane. This is attractive because it can help reduce the computational burden by reducing the number of MOC planes and could also impact stability by using larger MOC planes. Aaron Graham's recent PhD dissertation focused on reducing control rod decussing errors in MPACT, particularly within the context of the subplane method [9,10], so some initial foundation had already been laid to incorporate the subplane method. While Graham's efforts focused primarily on accuracy, this milestone addresses preparations to make subplane a production-level capability in MPACT.

2. SUBPLANE BACKGROUND AND PREVIOUS WORK

In the 2D/1D approach, 3D problems are decomposed into an axial stack of radial planes, where 2D-MOC is used to solve radially, and 1D-nodal methods operating on a homogenized pin-cell basis solve axially. The 2D and 1D solvers are coupled through transverse leakage, and CMFD is used to accelerate the global solve. To this point, MPACT has used a 1-to-1 mapping between the nodal and MOC solvers axially. However, with a target axial node size of ~8 cm, this can yield a substantial number of MOC planes (typically 55–60 in most cases).

The subplane method was previously used in the 2D/1D method [11] and has the potential to offer substantial reductions in computational burden. In this approach, multiple axial nodes are allowed to exist within each MOC plane. This in turn allows the axial solvers to be refined independently. In the results shown herein, the subplane/nodal axial mesh is targeted to be very consistent with the axial meshes currently employed in MPACT but with fewer MOC planes, the number of which depends on which material boundaries are simulated more explicitly.

Figure 2.1.1 shows a representative illustration of 2D/1D with the subplane approach. The same solvers are used axially and radially as without the subplane, but some approximation is made when it comes to the transverse leakages and cross sections used in each.




Figure 2.1.1. Subplane Illustration [8].

Aaron Graham implemented the initial subplane capability as part of his dissertation work [9]. The bulk of his thesis focused on ways to treat the control rod decoupling problem which manifests itself when a control rod is partially inserted into an MOC plane. Graham also developed several “subgrid” solver variants that treat rodged and unrodged regions of the MOC plane with a partially inserted rod to inform the coarse mesh cross sections used in the CMFD and 1D-nodal axial sweepers. This work proved to be very successful, as it substantially reduced the errors that had been introduced from partially inserted rods.

3. OVERVIEW OF IMPROVEMENTS

3.1 Subgrid Treatment of Spacers

Building on Graham’s work, which focused on subgrid treatment of control rods, the subgrid approach was extended to apply to spacer grids, as well. In the subgrid approach, the cross sections used in the CMFD and axial nodal solves were modified to more accurately represent material changes that may be occurring within each MOC plane. For example, one of the approaches assessed in the results explicitly models the spacer grids with a separate MOC plane (and subplane in between the spacers), but another approach using subgrid incorporates the spacer grid into the top of the MOC plane (Figure 3.1.1). This allows for a further reduction of the number of MOC planes, but it introduces more approximation.

There are two primarily variants of subgrid solvers that are relevant to this work. The first is simply termed as a *subplane treatment*, in which axial cross section data within the MOC plane vary to

correctly capture the material changes (rodded vs. unrodded or spacer grid vs. no grid). The second is an extension of the subplane treatment which employs a 1D collision probabilities method (CPM) solver for each of the disparate regions of the subplane, using the flux solution to better inform the coarse mesh homogenized cross sections. The setup of the CPM solver cases is illustrated on the left side of Figure 3.1.1.

More details on the subplane and collision probabilities treatments can be found in Graham's dissertation. A similar concept was extended to spacer grids, though the underlying solvers are identical. Additionally, the results in this work focus exclusively on the subplane treatment.

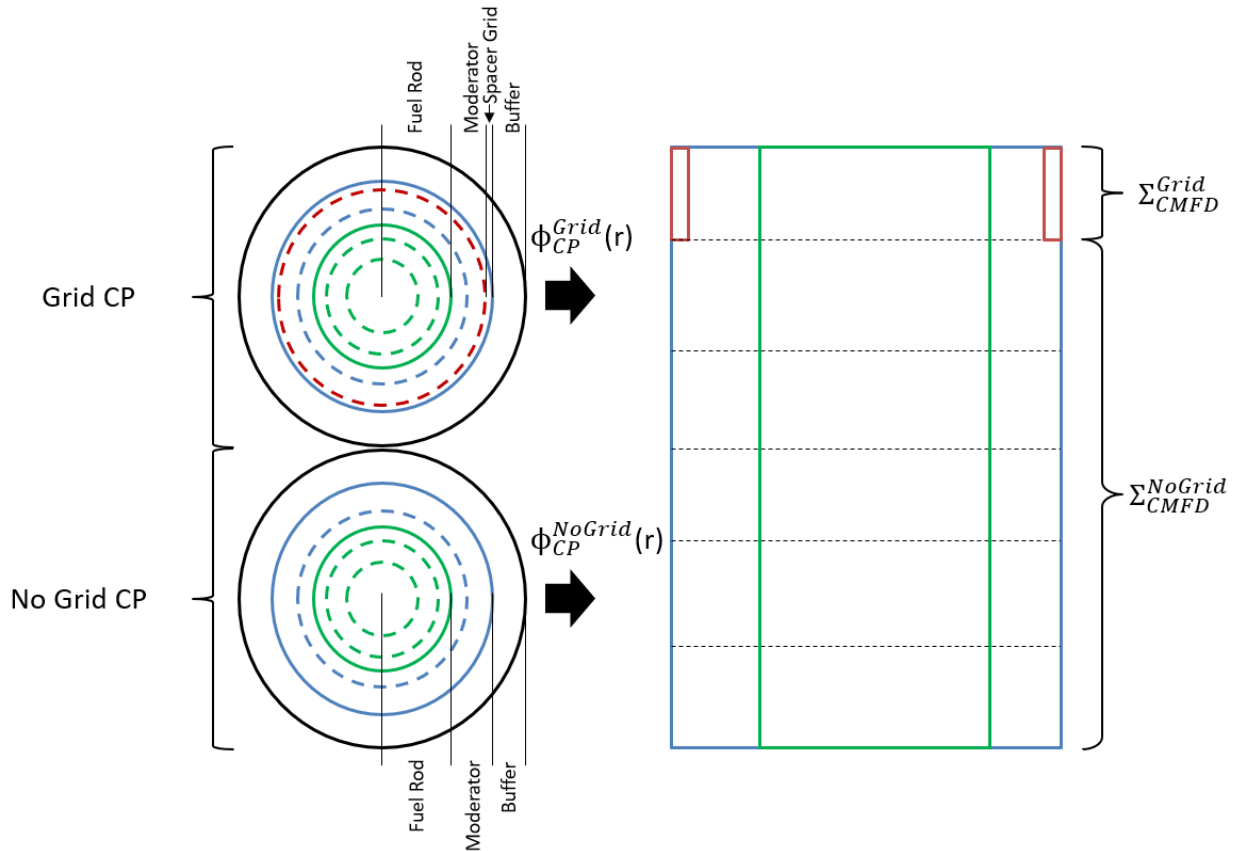


Figure 3.1.1. Subgrid Illustration Applied to Spacer Grids.

3.2 Spatial Multilevel CMFD Solver

A multilevel-in-energy CMFD solver was previously implemented in MPACT [4] in which the standard multigroup structure (~51 groups) was collapsed into a user-controlled number of levels (default of 3), down to 2 groups at the coarsest solve used during eigenvalue iterations. In support of this milestone, a multilevel-in-space CMFD solver was added that employs three levels: (1) pin-homogenized nodes for each axial subplane level, (2) quarter assembly-collapsed nodes for each axial level, and (3) quarter assembly-collapsed nodes with full subplane collapse. This is briefly illustrated in Figure 3.2.1.

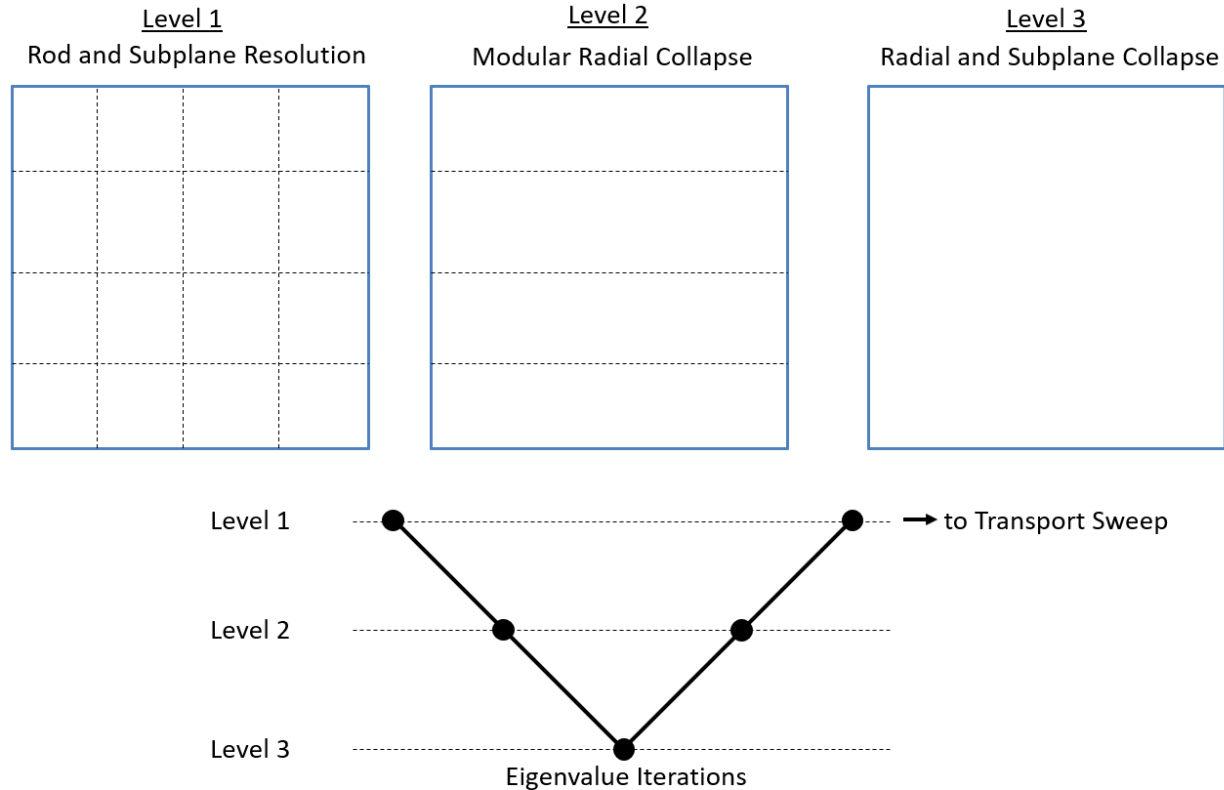


Figure 3.2.1. Demonstration of Various Spatial Levels.

Figure 3.2.1 shows a standard V-cycle in which levels 1 and 2 are used as smoothers which implement fixed source solves to quickly inform the coarsest solution from level 3, which is then used to solve the eigenvalue problem. Therefore, many solves are executed on level 3. Once the eigenvalue solution is converged, levels 2 and 1 are solved again to project the solution back, eventually to the fine mesh used in the transport sweeper. After the V-cycle is completed, the transport sweep (nodal and MOC) is executed to complete the outer iteration.

MPACT is currently limited to using a V-cycle, but eventual extensions may include W-cycles or more advanced treatments as necessary to improve the overall convergence. As shown in the results section, the V-cycle performs fairly well for both multilevel-in-space and energy (individually), but it often takes more outer iterations to full converge. This can potentially be resolved with a more complex iteration or with a combination of space and energy multilevels (more than 3 levels, for example).

3.3 Dynamic Subplane Mesh for Rod Movements

Throughout a cycle depletion, control rod banks will move to control the reactivity and power shape. In MPACT, the axial mesh has been held constant, relying heavily on decussing methods [9] to yield sufficient accuracy, as rods may be partially inserted into each plane. One of the benefits of the subplane method is the flexibility it can provide by more easily allowing for a dynamic axial mesh without completely reinitializing the problem. As long as the number of MOC planes remains constant, the subplane mesh can be adjusted to ensure that material changes are explicitly captured, at least with the subplane mesh. This requires reinitializing several objects

(coarse mesh, coarse ray trace, etc.) to account for the new mesh, but the overhead for this is minimal. This approach can help improve the accuracy of the subgrid solvers employed for control rods.

Figure 3.3.1 illustrates how the subplane mesh can vary, focusing on two planes of a representative problem.

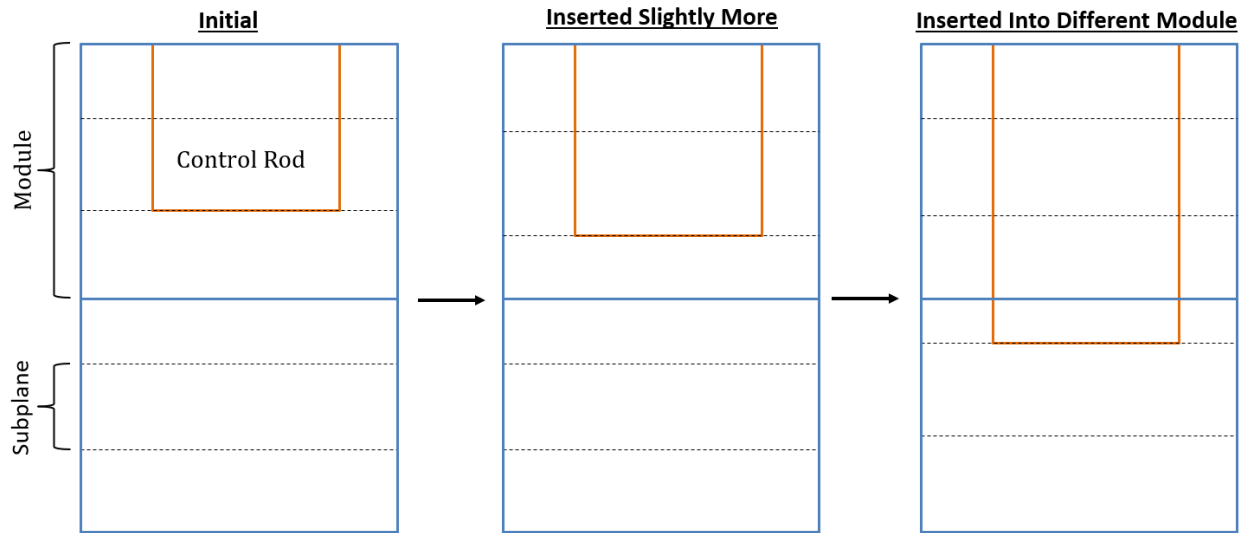


Figure 3.3.1. Illustration of Multistate Variable Mesh.

As can be seen in this simple illustration, as the rod moves through different states, the subplane mesh adapts to ensure that the material changes in the rod are sufficiently meshed. This includes material variations within control rods, such as stainless steel tips, as well as silver-indium-cadmium (AIC) and boron carbide (B_4C) combination rods.

3.4 Usability Improvements

3.4.1 Automated Meshing

Currently in MPACT, axial meshing in the fuel is governed by the `axial_edit_bounds` flag (truncated below):

```
[EDITS]
  axial_edit_bounds
    11.951
    15.817
    24.028
    32.239
    40.45
    48.662
    56.87
    ...
```

This can be problematic for subplane meshing because the bounds are used to define the boundaries of the MOC planes. Previous subplane runs used a reduced `axial_edit_bounds` array to specify fewer MOC planes, but they resultantly sacrificed axial resolution in the pin power and other output distributions.

In current efforts, when subplane is being used, it automatically meshes the fuel based on the target subplane thickness specified, and it automatically registers spacer grids. If using subgrid solvers for spacer grids, it will only impose the upper bound of each spacer grid as a MOC plane boundary, treating the grid as part of the plane. Otherwise, the spacer grid will be explicitly represented with a MOC plane. Comparisons of these approaches are shown in the results section.

3.4.2 Flexible Axial Edit Specification

When using subplane, since the axial mesh is automatically determined from important material changes (i.e., spacer grids and control rods), the `axial_edits_bounds` are used exclusively for their intended purpose of controlling the output data mesh. A new module, *DataShapers*, has been added to MPACT. DataShapers uses the Legendre expansion coefficients from the 1D-axial nodal solver solution to accurately evaluate the subplane mesh solution to the desired resolution per the edit bounds. In the progression problem cases used in this report, the edit bounds align very well with the subplane mesh produced from a 9 cm target mesh (see next subsection). However, the DataShapers object is very flexible and can operate over an arbitrary mesh, even if it does not perfectly align with the subplane mesh.

3.4.3 Easy to Use Input Specification

A primary goal of this milestone is to ensure that the method is easy to use. All of the improvements discussed previously can be specified using the following flags. All of the progression problem cases in the next section were run exclusively with these changes. No adjustments to geometry or axial edit bounds were necessary.

```
[MPACT]
num_space 16
subplane_target 9.0
cmfd mlcmfd
multilevel space
subgrid_spacers true
rod_treatment subplane
```

4. RESULTS

4.1 Overview of Cases Evaluated

Results for three progression problems will be evaluated here: (1) **Problem 3a**, Quarter Assembly, (2) **Problem 4a**, 3×3 Assembly Cluster, and (3) **Problem 5a**, Quarter Core. To focus on the accuracy of the approximations imposed by the subplane method and the subgrid treatment of the spacer grids, the control rods have been removed from all cases.

Nine different solver configurations will be evaluated for each case to assess the impact of relevant improvements:

1. **Master:** default mesh and solver options (explicit MOC planes for each axial level [58 total planes] and a single level CMFD solver for all 51 groups and pin-wise coarse mesh cells in the system)
2. **ML Space:** same axial mesh as Master but using a 3-level multilevel in space CMFD solver per Section 3.2
3. **ML Energy:** same axial mesh as Master but using a 3-level multilevel in energy CMFD solver; with 3 levels, Level 1 has 51 groups, Level 2 has 10 groups, and Level 3 has 2 groups
4. **Subplane:** same CMFD option as Master and with explicit MOC planes for spacer grid but has a single MOC plane between spacer grids for a total of 23 MOC planes.
5. **Subplane+ML Space:** same mesh as Subplane and ML Space CMFD solver
6. **Subplane+ML Energy:** same mesh as Subplane and ML Energy CMFD solver
7. **Subplane+Subgrid:** same CMFD option as Master but has a single MOC plane between the tops of each spacer grid for a total of 16 MOC planes, using the subgrid approach for spacer representation
8. **Subplane+Subgrid+ML Space:** same mesh as Subplane+Subgrid and ML Space CMFD solver
9. **Subplane+Subgrid+ML Energy:** same mesh as Subplane+Subgrid and ML Energy CMFD solver

4.2 Quarter Assembly (Progression Problem 3a)

Figure 4.2.1 shows the radial and axial layouts of the quarter assembly model, *Progression Problem 3* [12]. It is a standard Westinghouse 17×17 assembly with 6 Zircaloy spacer grids in the active fuel and 2 Inconel grids, one at the bottom of the active fuel and another in the upper reflector region.




Figure 4.2.1. Problem 3a Specification [12].

Figure 4.2.2 shows the axial mesh used in each of the cases. The first three have 58 planes, the second three have 23 planes, and the last three have 16 planes. Full axial decomposition is used in Problems 3a and 4a (58, 23, and 16 procs), and radial and axial decompositions are used in Problem 5a.

		Plane Upper Boundaries		
		Master	Subplane	Subgrid
Core Plate		418.937	418.937	418.937
Nozzle		411.337	411.337	411.337
Gap		402.51	402.51	402.51
Plenum		398.711	398.711	398.711
Inconel Grid		395.133	395.133	395.133
Plenum		391.267	391.267	
		382.711	382.711	382.711
		374.7898		
		366.8686		
		358.9474		
		351.0262		
Zircaloy Grid		343.105	343.105	343.105
		339.295	339.295	
		331.23		
		323.165		
		315.1		
		307.035		
		298.97		
Zircaloy Grid		290.905	290.905	290.905
		287.095	287.095	
		279.03		
		270.965		
		262.9		
		254.835		
		246.77		
Zircaloy Grid		238.705	238.705	238.705
		234.895	234.895	
		226.83		
		218.765		
		210.7		
		202.635		
		194.57		
Zircaloy Grid		186.505	186.505	186.505
		182.695	182.695	
		174.63		
		166.565		
		158.5		
		150.435		
		142.37		
Zircaloy Grid		134.305	134.305	134.305
		130.495	130.495	
		122.43		
		114.365		
		106.3		
		98.235		
		90.17		
Zircaloy Grid		82.105	82.105	82.105
		78.295	78.295	
		70.084		
		61.873		
		53.662		
		45.45		
		37.239		
		29.028		
Inconel Grid		20.817	20.817	20.817
Gap		16.951	16.951	16.951
Nozzle		11.053	11.053	11.053
Core Plate		5	5	5

Figure 4.2.2. Axial Mesh Comparison Between Options.

4.2.1 Accuracy Assessment

Table 4.2.1 shows the solution differences for the 9 cases, with Master as the reference. In general, subplane does a very good job of sustaining the Master solution. While the incorporation of subgrid solvers introduces some error, it is reasonable for the method. It can also be seen that the ML Space and ML Energy CMFD solvers yield similar results but with some variation, particularly in pin powers. This is related to the convergence criteria used by default in MPACT. The multilevel CMFD solvers change the path to solution, so there is potential that the converged solution slightly deviates from other cases in which the solution should be identical.

Table 4.2.1. Problem 3a Accuracy Comparisons

Case	Eigenvalue	dk (pcm)	dP (RMS %)	dP (MAX %)
Master	1.17461	---	---	---
ML Space	1.17461	0.04	0.045	0.078
ML Energy	1.17461	0.13	0.040	0.065
Subplane	1.17461	0.14	0.054	0.099
Subplane+ML Space	1.17461	0.15	0.028	0.103
Subplane+ML Energy	1.17461	0.23	0.025	0.116
Subplane+Subgrid	1.17461	0.33	0.052	0.322
Subplane+Subgrid+ML Space	1.17461	0.06	0.072	0.365
Subplane+Subgrid+ML Energy	1.17462	0.38	0.092	0.362

For comparison, the default fission source tolerance in MPACT was dropped from $5e-5$ to $1e-5$, yielding the following results in Table 4.2.2. As can be seen, there is much better consistency between the different CMFD solvers.

Table 4.2.2. Problem 3a Accuracy Comparisons (Tighter Tolerance)

Case	Eigenvalue	dk (pcm)	dP (RMS %)	dP (MAX %)
Master	1.17461	---	---	---
ML Space	1.17461	0.02	0.018	0.048
ML Energy	1.17461	0.02	0.017	0.030
Subplane	1.17461	0.08	0.013	0.042
Subplane+ML Space	1.17461	0.08	0.014	0.044
Subplane+ML Energy	1.17461	0.10	0.014	0.048
Subplane+Subgrid	1.17462	0.27	0.105	0.352
Subplane+Subgrid+ML Space	1.17461	0.16	0.069	0.321
Subplane+Subgrid+ML Energy	1.17462	0.33	0.062	0.313

To illustrate how well the subplane and subgrid solvers perform, Figure 4.2.3 shows an axial power plot comparing Master, Subplane with Explicit Grids, and Subplane with Subgrid Solvers. On the whole, the distributions agree well. The blue line denotes the differences between Master and Subplane with explicit grids, where some error is introduced by using subplane between the grids, but the axial power difference is less than 0.1%. When using subgrid solvers, the solution is still good, with a maximum axial power difference of $\sim 0.15\%$, but it is clearer that there is greater error in the spacer grid locations, as evidenced by the larger downward spikes.




Figure 4.2.3. Comparison of Various Grid Treatment Approaches.

4.2.2 Performance Improvements

Table 4.2.3 shows performance results for the various options when using the default tolerances. As discussed previously, the number of processors/cores used for each case depends on the axial mesh used, and it decreases as more approximations are applied and as the number of MOC planes is reduced. In turn, this also has substantial implications for the total memory necessary to solve the problem. For example, with 58 planes, the requirement is roughly 10 GB of memory, but this is reduced to 3.5 GB when running with 16 planes and subgrid treatment. Even with explicit grids in subplane, it is reduced to 4.5 GB.

As compared to Master, Subplane with default CMFD actually requires more core-hours to solve. Much of this is attributable to the load imbalance present between the planes with explicit spacer grids and the planes between them. Since the planes in between will have more CMFD/Nodal work because they have more subplanes, some procs will have to wait idle, thus decreasing efficiency. Much of this is resolved when incorporating ML CMFD, as CMFD is lighter, so the imbalance is much less noticeable.


Overall, Subplane+ML and Subplane+Subgrid+ML reduce the runtime down to 0.57 core-hours. This is significant savings over current Master (almost 3×) and still has substantial savings over default axial mesh with ML (1.6×).

Table 4.2.3. Problem 3a Performance Comparisons

Case	Iterations	Total Memory (GB)	Walltime (min)	Cores	Time (core-hrs)
Master	9	10.52	1.54	58	1.49
ML Space	9	10.67	0.96	58	0.92
ML Energy	9	10.67	0.93	58	0.90
Subplane	13	4.53	4.39	23	1.68
Subplane+ML Space	10	4.56	1.48	23	0.57
Subplane+ML Energy	10	4.56	1.49	23	0.57
Subplane+Subgrid	16	3.52	5.25	16	1.40
Subplane+Subgrid+ML Space	12	3.52	2.16	16	0.58
Subplane+Subgrid+ML Energy	12	3.52	2.15	16	0.57

4.3 Colorset (Progression Problem 4a)

Problem 4a is a step up from the Problem 3a in that it is a 3×3 assembly colorset featuring two different enrichments, and it also includes burnable poison rods (per Figure 4.3.1). The axial layout is identical to that in Problem 3a. Similarly, all cases have 58, 23, or 16 planes and procs.

**Figure 4.3.1. Problem 4a Colorset Radial Layout [12].**

4.3.1 Accuracy Assessment

Table 4.3.1 shows the solution comparison for Problem 4a. This time, more noticeable solution differences are seen when moving to Subplane and Subplane+Subgrid. The maximum pin power errors approaching 1% are somewhat concerning, but they may be tolerable given the performance gains.

Table 4.3.1. Problem 4a Accuracy Comparisons

Case	Eigenvalue	dk (pcm)	dP (RMS %)	dP (MAX %)
Master	1.00090	---	---	---
ML Space	1.00090	0.02	0.037	0.104
ML Energy	1.00090	0.14	0.023	0.151
Subplane	1.00091	1.14	0.387	0.744
Subplane+ML Space	1.00092	1.24	0.238	0.588
Subplane+ML Energy	1.00091	1.19	0.298	0.657
Subplane+Subgrid	1.00097	6.91	0.329	0.680
Subplane+Subgrid+ML Space	1.00097	6.76	0.233	0.577
Subplane+Subgrid+ML Energy	1.00097	6.98	0.272	0.614

Table 4.3.2 shows similar results with a tighter fission source tolerance of $1e-5$. As before, there is much better consistency, but additional iterations to reach convergence are incurred.

Table 4.3.2. Problem 4a Accuracy Comparisons (Tighter Tolerance)

Case	Eigenvalue	dk (pcm)	dP (RMS %)	dP (MAX %)
Master	1.00090	---	---	---
ML Space	1.00090	0.05	0.014	0.042
ML Energy	1.00090	0.08	0.009	0.035
Subplane	1.00091	1.20	0.269	0.630
Subplane+ML Space	1.00092	1.22	0.240	0.600
Subplane+ML Energy	1.00091	1.21	0.252	0.615
Subplane+Subgrid	1.00097	6.90	0.175	0.515
Subplane+Subgrid+ML Space	1.00097	6.84	0.216	0.560
Subplane+Subgrid+ML Energy	1.00097	6.91	0.215	0.563

4.3.2 Performance Improvements

Table 4.3.2 shows the performance comparisons, which are very similar to those seen in Problem 3a. As before, there is a substantial reduction in memory footprint, with $3\times$ core-hours reduction over Master, and roughly $1.5\times$ reduction over ML Space. Subgrid seems to offer some savings in core-hours and memory, but it is relatively small.

Table 4.3.3. Problem 4a Performance Comparisons

Case	Iterations	Total Memory (GB)	Walltime (min)	Cores	Time (core-hrs)
Master	12	47.92	16.93	58	16.36
ML Space	11	47.75	8.08	58	7.81
ML Energy	11	47.75	8.14	58	7.87
Subplane	12	22.25	34.91	23	13.38
Subplane+ML Space	17	21.97	14.60	23	5.60
Subplane+ML Energy	11	21.97	14.75	23	5.66
Subplane+Subgrid	14	18.83	44.92	16	11.98
Subplane+Subgrid+ML Space	19	18.53	19.58	16	5.22
Subplane+Subgrid+ML Energy	11	18.53	19.60	16	5.23

4.3.3 Dynamic Meshing

To test the dynamic meshing capability, the integral rod worth version of Problem 4a (now with the control rod included) was assessed. All cases here use subplane with explicit grid representation (23 planes). The case models a central control rod withdrawing at 10% per state. The reference solution in this case is each of the rod positions modeled separately as a single state. With the automated meshing, the initial subplane mesh conforms perfectly to the rod location of the first state. For comparison, a stacked case (all 11 positions modeled in a single file) is used both with and without subplane remeshing. Without remeshing, the mesh will remain what it was after the first state (0 steps withdrawn). With remeshing, the mesh will be reset to conform to the new rod position.

Table 4.3.4 shows the results of this comparison. Without remeshing, there are reasonably large errors for a number of the rod positions. It is worth noting that the default P_3 axial solver is understood for some of the meshes generated when using the remesh capability, so these results were generated using the two-node NEM axial solver. These stability issues can arise when thin planes are generated with modeling the rod more explicitly, which is not surprising given previous experience, but worth further investigation to resolve.

Table 4.3.4. Dynamic Meshing Results on 4a-IRW

Steps Withdrawn	Individual			Stacked - No Remesh			Stacked Remesh		
	dk (pcm)	dP (RMS)	dP (MAX)	dk (pcm)	dP (RMS)	dP (MAX)	dk (pcm)	dP (RMS)	dP (MAX)
0	---	---	---	0	0.009	0.211	0	0.000	0.000
23	---	---	---	0	0.020	0.207	0	0.000	0.000
46	---	---	---	0	0.066	1.295	0	0.000	0.000
69	---	---	---	0	0.026	0.298	0	0.000	0.000
92	---	---	---	0	0.004	0.109	0	0.000	0.000
115	---	---	---	0	0.003	0.030	0	0.000	0.000
138	---	---	---	0	0.082	1.999	0	0.000	0.000
161	---	---	---	0	0.013	0.202	0	0.000	0.000
184	---	---	---	0	0.006	0.124	0	0.000	0.000
207	---	---	---	0	0.020	0.409	0	0.000	0.000
230	---	---	---	0	0.000	0.000	0	0.000	0.000

4.4 Quarter Core (Progression Problem 5a)

Figure 4.4.1 shows the core layout for the quarter core Problem 5a, including a third enrichment value and a much larger core than 4a. This problem also includes a radial reflector region not shown in the figure. A radial decomposition of 16 cores was used when running this problem. In combination with the 58, 23, and 16 axial cores, totals of 928, 368, and 256 cores were used for the respective axial mesh variants. These cases were all run on the Panacea cluster at Oak Ridge National Laboratory (ORNL).

	H	G	F	E	D	C	B	A
8	2.1 20	2.6 20	2.1 20	2.6 20	2.1 20	2.6 20	2.1 20	3.1 12
9	2.6 20	2.1 24	2.6 20	2.1 20	2.6 20	2.1 24	3.1 16	3.1
10	2.1 24	2.6 20	2.1 20	2.6 20	2.1 20	2.6 16	2.1 16	3.1 8
11	2.6 20	2.1 20	2.6 20	2.1 20	2.6 20	2.1 16	3.1 16	3.1
12	2.1 20	2.6 20	2.1 20	2.6 20	2.6 24	2.1 12	3.1	
13	2.6 20	2.1 16	2.6 24	2.1 12	2.6 12	3.1	3.1	
14	2.1 24	3.1 16	2.1 16	3.1 16	3.1	3.1		
15	3.1 12	3.1 8	3.1 8	3.1 8	Enrichment Number of Pyrex Rods			

Figure 4.4.1. Problem 5a Radial Layout [12].

4.4.1 Accuracy Assessment

Table 4.4.1 shows the accuracy comparisons for 5a. When compared to 4a, the pin power solution differences when incorporating subplane and subgrid are comparable and did not increase substantially with the problem size. The eigenvalue difference with subgrid is more pronounced, but it is still very low at only 7 pcm.

Table 4.4.1. Problem 5a Accuracy Comparisons

Case	Eigenvalue	dk (pcm)	dP (RMS %)	dP (MAX %)
Master	1.00063	---	---	---
ML Space	1.00063	0.02	0.093	0.307
ML Energy	1.00063	0.03	0.086	0.307
Subplane	1.00064	1.31	0.347	0.901
Subplane+ML Space	1.00064	1.34	0.260	0.710
Subplane+ML Energy	1.00064	1.26	0.270	0.775
Subplane+Subgrid	1.00071	7.94	0.232	0.747
Subplane+Subgrid+ML Space	1.00070	7.76	0.217	0.664
Subplane+Subgrid+ML Energy	1.00071	7.99	0.125	0.604

4.4.2 Performance Improvements

Due to its larger size and the potential use of radial decomposition, the performance comparisons for Problem 5a are substantially more interesting than those in Problems 3a and 4a. The memory reduction with subplane is still very sizable, but Master heavily struggles in CMFD to solve the problem effectively. Just by incorporating ML Space CMFD, the time is reduced to ~350 core-hours (~3.5×). With ML Energy CMFD, the improvement is less pronounced at ~550 core-hours (2.25×).

In this case, going to Subplane with the default CMFD actually shows a considerable improvement, whereas the benefit was less clear in Problems 3a and 4a. Adding on ML Space or ML Energy CMFD to Subplane drops the runtime even lower to ~220–260 core-hours (~5.6× over Master). Incorporating subgrid does help reduce memory, but the overall runtime comparisons are fairly comparable to subplane with explicit grids. This is potentially attributable to the increased iterations that were required in those cases.

Table 4.4.2. Problem 5a Performance Comparisons

Case	Iterations	Total Memory (GB)	Walltime (min)	Cores	Time (core-hrs)
Master	13	1400.16	79.47	928	1229.14
ML Space	11	1407.41	22.35	928	345.70
ML Energy	12	1437.31	35.33	928	546.37
Subplane	16	610.58	91.87	368	563.45
Subplane+ML Space	15	614.53	35.69	368	218.92
Subplane+ML Energy	12	641.84	42.45	368	260.37
Subplane+Subgrid	14	544.35	100.32	256	428.02
Subplane+Subgrid+ML Space	18	546.78	54.46	256	232.36
Subplane+Subgrid+ML Energy	13	574.45	52.47	256	223.89

5. CONCLUSIONS

Overall, this milestone has been successful in improving the usability and performance of the subplane method in MPACT. All of the enhancements shown here can be activated with the addition of a few flags in the MPACT block. This includes the multilevel-in-space CMFD, which uses radial and axial collapsing, automated axial meshing, subgrid treatment of spacer grids, and

the dynamic subplane mesh capability for control rod movement. Additionally, consistent comparisons of all configurations were made possible through the new DataShapers editing capability.

In the problems tested, considerable reductions were seen in both memory and runtime through minor approximations (subplane and subgrid), at least with respect to solution data. For example, Problem 5a saw a nearly $3\times$ reduction in overall memory and a $5\text{--}6\times$ reduction in total core-hours. These approaches also make running on smaller machines more tractable, which is important, as VERA and MPACT are deployed to industry users. Looking forward, as additional work is done on the subplane method, it would seem that these improvements have a good chance of becoming default. At this point in time, it would be recommended to use Subplane + ML Space + Subgrid, but ongoing work to combine ML Space and Energy would change that recommendation once completed.

6. FUTURE WORK

While this milestone was successful, there is still plenty of work ahead before this capability is ready to be deployed, potentially as a default. Some of this work is already planned under FY19 milestones.

6.1 Coupling Improvements

At present, running cases with thermal hydraulic feedback through CTF and subplane is problematic. Modifications need to be made to the coupling mesh to ensure that the same resolution that is currently available in the coupling is preserved. Additionally, coupling with VeraShift and Tiamat likely need similar updates.

6.2 Improved Multilevel CMFD Convergence and Performance

Currently, the multilevel CMFD either collapses in space or energy. Future work should focus on combining the space and energy collapses as appropriate, probably with more than three levels to formulate a faster solver with better convergence properties.

6.3 Enhanced Axial Reflector Treatment

Of the 58, 23, or 16 planes used in the different configurations, 8 are used exclusively for the lower and upper axial reflector regions. If this could be reduced to 2 (one lower and one upper) using a subgrid representation similar to that used for spacers and control rods, additional savings could be gained. Additionally, thin planes, particularly in regions without fuel, can be particularly problematic to overall 2D/1D convergence, so addressing this and generating thicker MOC planes in the reflector could help improve stability.

6.4 Depletion Improvements

The initial testing with depletion has highlighted that additional improvements will likely be necessary to retain accuracy during depletion. With subplanes in the fuel, the current depletion methodology operates on the MOC plane basis, missing a substantial part of the axial shape. Modifications need to be made to capture this shape, potentially through further subgrid extensions or a more explicit approach.

7. ACKNOWLEDGMENTS

This research was supported by the Consortium for Advanced Simulation of Light Water Reactors (www.casl.gov), an Energy Innovation Hub (<http://www.energy.gov/hubs>) for Modeling and Simulation of Nuclear Reactors under US Department of Energy Contract No. DE-AC05-00OR22725.

This research made use of the resources of the Compute and Data Environment for Science (CADES) at ORNL, which is supported in part by the Office of Nuclear Energy of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

8. REFERENCES

1. J. Turner et al., “The Virtual Environment for Reactor Applications (VERA): Design and Architecture,” *Journal of Computational Physics*, **326**, 544 (2016).
2. B. Collins et al., “Stability and Accuracy of 3D Neutron Transport Simulations Using the 2D/1D Method in MPACT,” *J. Comput. Phys.*, **326**, 612 (2016).
3. S. Stimpson et al., “A Lumped Parameter MOC Approach and Multigroup Kernels Applied to Subgroup Self-Shielding in MPACT,” *Nucl. Eng. Technol.*, **49**, 1240 (2017).
4. S. Stimpson et al., “Improvement of Transport-Corrected Scattering Stability and Performance Using a Jacobi Inscatter Algorithm for 2D-MOC,” *Ann. Nucl. Energy*, **105**, 1 (2017).
5. B. Collins, S. Stimpson, “Acceleration Methods for Whole Core Reactor Simulations using VERA,” *Trans. Am. Nucl. Soc.*, **118**, 929 (2018).
6. B. Collins, S. Hamilton, S. Stimpson, “Use of Generalized Davidson Eigenvalue Solver for Coarse Mesh Finite Difference Acceleration,” *Proc. M&C 2017*, Jeju, Korea (April 16-20, 2017).
7. S. Stimpson, B. Collins, “Implementation of a Red-Black SOR CMFD Solver in MPACT,” *Trans. Am. Nucl. Soc.*, 115, pp. 1252-1255 (2016).
8. B. Yee, “A Multilevel in Space and Energy Solver for Multigroup Diffusion and Coarse Mesh Finite Difference Eigenvalue Problems,” Ph.D. Dissertation, University of Michigan (2015).
9. A. Graham, “Subgrid Methods for Resolving Axial Heterogeneity in Planar Synthesis Solutions for the Boltzmann Transport Equation,” Ph.D. Dissertation, University of Michigan (2017).
10. A. Graham et al., “Subplane Collision Probabilities Method Applied to Control Rod Cusping in 2D/1D,” *Ann. Nucl. Energy*, **118**, pp. 1-14 (2018).
11. J. Y. Cho et al., “Sub-plane scheme for a radial transport and axial diffusion code,” *Proc ICAPP 2007*.
12. A. Godfrey, “VERA Core Physics Benchmark Progression Problem Specifications,” Revision 4, CASL-U-2012-0131-004, Revision 4, CASL, August 29, 2014. <http://www.casl.gov/docs/CASL-U-2012-0131-004.pdf>