Verification & Validation of High-Order Short-Characteristics-Based Deterministic Transport Methodology on Unstructured Grids

---

Reactor Concepts RD&D

Dr. Yousry Azmy

North Carolina State University

In collaboration with:

Idaho National Laboratory

Rob Versluis, Federal POC

David Nigg, Technical POC

Final Report NEUP Project 09-798

**Verification & Validation of High-Order Short-Characteristics-Based Deterministic
Transport Methodology on Unstructured Grids**

Submitted by Project Principal Investigator

Yousry Y. Azmy
Department of Nuclear Engineering
North Carolina State University

Contents

1.	Abstract.....	3
2.	Introduction	5
3.	Objective.....	6
4.	Detailed Report on Project Tasks	6
3.1	Task A – Verification Phase 1: Mesh refinement studies (INL)	6
3.1.1	Simple Test Problem Results	7
3.1.2	Numerical Benchmarks.....	16
3.2	Task B – Development of MMS Benchmark Suite (NCSU)	32
3.2.1	FEM Formulation	32
3.2.2	Development & Implementation of Three-Dimensional MMS for THOR	32
3.2.3	THOR Illustration of MMS.....	35
3.2.4	Development of a GUI for MMS3D	38
3.3	Task C – Grind Times Study (NCSU)	39
3.3.1	Objectives (i) & (ii)	39
3.3.2	Objective (iii).....	39
3.4	Task D – Novel Algorithm for Computing the Fundamental Eigenmode (NCSU).....	43
3.4.1	Implementation into THOR	43
3.4.2	The JFNK Methodology.....	43
3.4.3	Numerical Results.....	47
3.5	Task E – Verification Phase 3: Comparison to Monte Carlo (NCSU).....	50
3.5.1	ATR Description	51
3.5.2	Multigroup Cross Sections Generation.....	54
3.5.3	Computational Grid Generation.....	57
3.5.4	Current and Future Work on the ATR Assembly	60
3.6	Task F – Validation of THOR	63
5.	Bibliography.....	64
6.	List of Publications from the Project	65

1. Abstract

THOR is a radiation transport code that solves the steady-state, multigroup, discrete ordinates approximation of the linear Boltzmann equation in three-dimensional geometry on unstructured tetrahedral cells. The spatial approximation implemented in THOR is the Arbitrarily High Order Transport method of the Characteristic type, AHOTC, extended to Unstructured Grids, AHOTC-UG. The tasks of this project were designed to raise the production level of THOR by supplementing its capabilities then conducting a comprehensive Verification and Validation (V&V) exercise based on Idaho National Laboratory's (INL) Advanced Test Reactor (ATR) configuration and measured data.

The primary development work on the code commenced with a study of numerical stability of the underlying equations in the optically-thin cell limit that revealed the cause for the structural instability observed in earlier results. Basically the recursive algorithm used in evaluating the flux spatial moments in terms of lower-order moments accumulated the error to unacceptable magnitude for higher orders. This deficiency was addressed by reformulating the equations and subsequent solution algorithm into a non-recursive form that was found to be numerically stable with increasing spatial expansion order. Additionally we examined the numerical stability of the spatial weights associated with AHOTC-UG and constructed asymptotic expansions that are resilient in the optically thin and thick cell regimes. Many improvements intended to enhance THOR's robustness and computational efficiency were implemented, including a cycle-breaking algorithm that may be necessary in some complex automatically-generated unstructured grids.

One innovative development that was implemented in THOR is the JFNK approach to solving criticality problems. Essentially the non-linear (more precisely algebraically quadratic) k -eigenvalue problem is solved with Newton iterations where in each iterative step a linear system is solved with a Krylov method. However, instead of using the exact Jacobian per Krylov solve, only the effect of the Jacobian on an arbitrary vector is implemented thereby drastically improving the method's efficiency. This forms the theoretical foundation of JFNK, but actual implementation must deal with practical issues like storing the Krylov vectors and selecting proper parameters for controlling both the Newton and Krylov iterations. The selected approach was to use the original implementation of the standard Power Iterations as the kernel computation and wrap the JFNK solution algorithm around it. Three strategies for this are implemented and tested using various test problems, and two of these are found to reduce execution time by a factor of five to eight compared to Power Iterations. One thing that still requires further development (but that was not proposed within the scope of this project) is acceleration of the inner iterations and preconditioning of the JFNK's Krylov solver. These efforts will be pursued by the PI and his research group in the future in order to make THOR a more attractive option for practitioners.

In order to quantify the competitiveness of THOR with alternative radiation transport software we conducted a thorough grind-time analysis. Initially we used the results of this analysis to identify bottlenecks in THOR's performance and this led to (sometimes) major revisions of the code's sections/components resulting in substantial reductions in execution time. Once we were satisfied with THOR's computational efficiency we measured the final version's grind-time and compared it to two production-level transport codes, Denovo and TORT. We found that while the grind times for the three codes are of the same order of magnitude, THOR is five to ten times slower. This is expected because Denovo and TORT employ Cartesian grids whose mesh sweep is very efficient because its sequence maps trivially to cell indices for a given discrete ordinate. In contrast, unstructured mesh solvers, like THOR, must either pre-compute or compute on the fly the sequential order of sweeping the mesh for each discrete ordinate thereby adversely affecting computational efficiency. So unstructured grid solvers improve the fidelity of the geometric representation of complex configurations but suffer a heavier computational load. In our judgment the factor of ten penalty in grind-time is an acceptable compromise in applications like the ATR where a Cartesian grid is unlikely to provide a sufficiently accurate representation of the serpentine fuel elements and other core details.

The verification exercise for THOR comprised three stages. First, the standard mesh refinement studies using simple geometries that either possessed analytic solutions (non-scattering media) or high quality reference solutions (obtained on ultra-fine mesh with "trusted" code). This stage also included comparison of THOR solutions to standard Benchmarks with well-documented solutions. Second, as part of this project we developed a suite of three-dimensional Method of Manufactured Solutions (MMS) that provided analytic solutions in scattering media with pointwise resolution. The MMS permitted a convergence-order study for THOR that established agreement of the observed and theoretical rates of convergence. In the third stage we compared the solution obtained by THOR for the ATR configuration to Monte Carlo (specifically MCNP) solutions to the same ATR model. The first two stages of the Verification were completed successfully and we are now able to claim that the present version of THOR is verified. The third stage, and consequently the Validation exercise, was fraught with troubles due to the lack of a correct tetrahedral meshing of the ATR configuration. At the time of composing the proposal that yielded award of this project INL presumably possessed a tessellation of the ATR geometry produced by CUBIT and Attila's mesh generator. Our examination of this mesh revealed multiple problems with it that INL scientists were not able to help us fix, so we ended up fixing as much as we can and we continue to work towards a correct and robust tessellation of the ATR geometry. Our goal, even after the conclusion of this project, is to reap the benefits of our efforts on THOR with a successful validation that we hope to submit to PHYSOR 2014 in Kyoto, Japan. We are now close to achieving this goal because of a breakthrough in our ability to manipulate the geometry to eliminate unnecessary detail before tessellation.

2. Introduction

The solution of the linear particle transport equation is of great importance for many fields both in research and engineering. Probably the most prominent application that requires the solution of the neutron transport equation is nuclear reactor design because the exact knowledge of the neutron population distribution enables determination of the fission rate distribution in the fuel. This, in turn, provides the distributed heat source within the fuel that drives the thermalhydraulic analysis for normal and accident performance studies and fissile fuel consumption for burnup calculations.

The solution of the neutron transport equation is still a challenge even on today's leadership class machines because of the high dimensionality of the phase space, i.e. the independent variables that need to be discretized before implementing the solution onto a computer. This report details the initial steps of the development of the unstructured grid SN transport code THOR with special emphasis on the performed Verification & Validation (V&V) exercise. Within the remainder of this work a code will be used to refer to the implementation, in some programming language, of a collection of algorithms that solve a partial differential equation. Hence, a transport code is the implementation of algorithms to solve the neutral particle transport equation.

The algorithms implemented in THOR comprise standard methods for the discretization of the energy and particle's direction of motion variables (referred to as direction variable): energy is discretized using the multigroup formalism and direction is discretized using the SN method. Both are considered the most competitive schemes compared to alternative discretization methods for energy and direction. However, the novelty in THOR is the approach taken to discretize the spatial variables using the Arbitrarily High Order Transport method of the Characteristic type (AHOTC) on Unstructured tetrahedral Grids (AHOTC-UG). [The term high-order here refers to the local expansion order not the order of the solution accuracy.] Within this report, the AHOTC-UG method is sometimes referred to as the short characteristic method. In contrast to the standard approach in the radiation transport community of using structured grids, the spatial domain is discretized using tetrahedrons allowing for a much greater flexibility in the representation of intricate details in the problem's geometry to the desired fidelity. The AHOTC method allows for an arbitrary polynomial order representation of the source and solution distributions within each tetrahedron and of the solution on the tetrahedron's faces. Thus, the user has two approaches to improve the accuracy of the solution: refine the mesh or increase the polynomial expansion order. The latter is typically more efficient in problems with large homogeneous regions.

THOR is unique within the radiation transport community because it features an arbitrary order local expansion basis (referred to as high-order) implemented on an unstructured grids. Comparable codes that solve the first order form of the radiation transport equation are either low order on structured meshes: Denovo, or they feature a fixed order on an unstructured grid: ATTILA (Discontinuous Finite Element Method of order one on tetrahedrons). THOR supports two polynomial function space families of order Λ : The Lagrange family retains all cross-moments up to order Λ , while the complete family only retains polynomial cross terms whose individual powers in x , y and z add up to an integer less or equal Λ . In order to distinguish results from the Lagrange and complete function spaces, the respective expansion orders will be denoted Λ and Λ' , respectively.

Within the project subject of this report, the Verification & Validation (V&V) of THOR is performed. Verification & Validation is an essential piece of software quality assurance. Verification is a process ensuring that the code solves the underlying system of equations correctly. The most rigorous method of code verification is the order of convergence test utilizing a Manufactured Solution. Within this work this approach of verification is utilized and augmented by solving several other benchmark problems featuring known solutions, comparison with Monte-Carlo calculations and code-to-code comparison with other well established SN transport codes. Validation is to ensure that the solutions obtained from the code match reality sufficiently well. Therefore, validation requires comparison of the computed responses with experimentally measured data. Within the scope of this work the validation exercise is based on the Advanced Test Reactor (ATR) Core Internal Changeout (CIC) 94 benchmark configuration. In addition to the V&V procedure, the project comprises subtasks that (1) equip THOR with a new algorithm for the computation of the fundamental eigenmode and that (2) identify computational bottlenecks to improve THOR's performance.

3. Objective

The high-order discrete-ordinates short characteristics neutron transport code THOR has been developed for 3D unstructured tetrahedral grids. An essential prerequisite for its deployment in modeling reactor cores is comprehensive and successful V&V exercises against Monte Carlo simulation and experimentally measured data, respectively. Verification comprises: 1) spatial mesh and expansion order refinement studies monitoring convergence to reference solutions; 2) creation of code-independent suite of benchmarks based on the Method of Manufactured Solutions; 3) comparison against continuous-energy and -angle Monte Carlo to quantify the error due to the multi-group approximation and the specific cross section library deployed. The validation stage involves modeling and comparison of numerical results to experimental measurements of criticality parameters and power distribution in Idaho National Laboratory's (INL) Advanced Test Reactor (ATR) Zero-Power critical configuration measurements.

4. Detailed Report on Project Tasks

In this section we provide a detailed description of the accomplishment of each task using the same numbering sequence as in the awarded proposal. The responsible party for delivery of each task is listed parenthetically at the end of the task's title. For some tasks the responsible party listed here may be different than in the original proposal.

3.1 Task A – Verification Phase 1: Mesh refinement studies (INL)

In order to establish the overall reduction in error and convergence rate of the method, a set of preliminary numerical tests were devised and solved with THOR. This set of numerical tests is divided into two subsets: simple test problems, with either analytical or numerical reference solutions, and numerical benchmarks, which have been historically proposed by the NEA/OECD 3-D transport expert group, and whose solution is usually obtained via multi-group Monte Carlo techniques. The results of this task were published in the PhD dissertation of Dr. Rodolfo Ferrer [1], Co-PI on the proposal of this project. The following discussion is an excerpt from [1].

3.1.1 Simple Test Problem Results

The goal of this section is to present and discuss the results of a set of simplified, mono-energetic problems comprised of a homogeneous region with vacuum boundary conditions and a fixed distributed source. These fixed source problems are further subdivided into two cases: The first is a non-scattering medium, and the second case is a medium that supports scattering collisions. A set of angular directions and weights, consisting of the level symmetric quadratures, S_2 through S_8 , was tested and compared against the exact reference solution for the non-scattering case. In the second case, a set of different scattering ratios, defined as $c = \sigma_{sc}/\sigma_T$ were considered and the approximate solution for each instance, based on the S_2 quadrature set, was compared with fine mesh numerical solutions obtained with the TORT code [2].

3.1.1.1 Non-Scattering Medium

Previous research regarding the analytical solution with respect to the spatial variables of the discrete ordinates approximation to the transport equation has suggested the existence of spatial discontinuities in the exact solution of the angular flux within a non-scattering medium given a fixed incoming angular flux boundary condition. In the particular cases presented here, the boundary conditions are usually assumed to involve vacuum or zero incoming angular fluxes, and the distributed source is assumed to be flat or constant. In the absence of scattering, the exact solution of the angular flux exhibits several inflexion planes across a selected single angular direction. These inflexion planes suggest that the first derivatives of the angular flux as a function of the independent spatial variables are discontinuous. The meshing approach for the problem configuration involves an “unaligned” spatial grid, in which the cell faces are not parallel to the angular direction, hence some cells must be split or divided into constituent Characteristic Tetrahedrons or CTs (see Ref [1] for details).

3.1.1.1.1 S_2 Quadrature in a Non-Scattering Medium

In most realistic applications the spatial grid is rarely purposely generated in order to obtain any type of geometric consistency with respect to the angular quadrature set. Due to this practical consideration, it becomes necessary to test any spatial discretization of the transport equation in more general cases, particularly in situations where the grid is fixed and different quadrature sets are used to represent the angular variable integration.

In order to relax the requirements regarding the alignment between the mesh and the angular directions in the quadrature set, a cube tessellation identical to that proposed by Azmy and Barnett was adopted with analogous mesh generation and grid refinement. In particular, a unit cube was initially tessellated into five tetrahedrons, as sketched in Figure 1. Note that, unlike the previous cube tessellation, not all of the edges or faces of the selected five tetrahedrons are consistently aligned with respect to a particular angular direction in order to avoid the splitting of cells into CTs.

Furthermore, a set of five spatial grids, based on the cube tessellation presented above, were generated in order to spatially discretize each of the $3 \times 3 \times 3$ sub-cubes comprising the full domain of $3 \times 3 \times 3$ mean-free-paths (mfp) cube. Hence, a total of five meshes, ranging from 135 to 552,960 tetrahedral cells, were generated in order to test the convergence behavior of the AHOTC-UG methodology for cases with an unaligned grid. A table summarizing the total number of tetrahedrons for each mesh is shown in Table 1 and a sketch of the five ‘unaligned’ spatial

grids is present in Figure 2, where the grid becomes finer as the figures change from left to right and from top to bottom. Note that, due to the geometric nature of the tessellation, each cube must be surrounded by ‘rotated’ cubes of the same size at the same refinement level, in order to maintain coincidence between cell faces in adjacent sub-cubes.

Figure 1: Subcube tessellated into five tetrahedrons.

Table 1: Total number of tetrahedral cells present for each unaligned mesh.

Mesh ID	Mesh Size
1	135
2	1,080
3	8,640
4	69,120
5	552,960

A single reference solution was generated by numerically integrating the exact analytical solution over the set of 27 ‘subcube’ regions that compose the cubic domain. Once all the angular fluxes were obtained as region averages for each of the subcubes in each angular direction, the discrete ordinates quadrature rule was applied in order to obtain a reference region-averaged scalar flux Φ_{ijk} over each subcube (i, j, k).

The AHOTC-UG approximation of the scalar flux over each tetrahedral cell for each grid was averaged over each of the 27 subcube regions, hence a set of approximate solutions ϕ_{ijk} were obtained and compared against the reference solutions. The maximum (taken over the 27 subcubes) absolute value of the resulting errors, denoted by $e_{\max}$, of these sets of errors was chosen as the main error norm to compare solution accuracy with respect to spatial mesh refinement and spatial expansion orders. A plot of the maximum errors is shown in Figure 3 and

Figure 4 as a function of the longest cell's edge length for each of the spatial grids and various Λ and Λ' spatial expansion orders, respectively.

Figure 2: Spatial grids generated from successive refinement via subcubes tessellated into five tetrahedrons.

The curves of the maximum errors for this test problem are not increasingly spaced away from each other as a higher spatial expansion order is used in the computation. For example, the maximum error curve for the $\Lambda = 1$ and $\Lambda = 2$ spatial expansions plotted in Figure 3 shows the error decreasing as a higher order expansion is used. However, the slope of the two curves indicates that the convergence rate is similar between the two cases. In a similar trend, the maximum error curves plotted in Figure 4 show the error only slightly decreasing between the $\Lambda' = 2$ and $\Lambda' = 3$ spatial expansion orders. However, in this case it is possible to identify a slightly higher convergence rate for $\Lambda' = 3$ as the mesh is refined.

In order to further analyze the maximum error behavior, it is necessary to borrow some concepts from error analysis and *a priori* error estimation. Generally speaking, the convergence behavior of numerical discretizations in the asymptotic regime may be represented by the proportionality between the computed error and the convergence law given by $C h^p$, where C is a bounded constant that is independent of the spatial grid, h is a measure of the spatial discretization such as cell side length, and p is the order of convergence. Given the availability of two maximum error values obtained from the numerical approximation and the reference solution, such as $e_{1,max}$ and $e_{2,max}$, and the maximum cell side length over each spatial grid, such as h_1 and h_2 , it is possible estimate the convergence rate of the maximum error with respect to

mesh refinement by applying the following proportionality relation

$$p = \frac{\ln \frac{e_{1,max}}{e_{2,max}}}{\ln \frac{h_1}{h_2}}.$$

Figure 3: Maximum scalar flux error over 27 subcubes for various Λ spatial expansion orders and S_2 quadrature as a function of mesh refinement for unaligned semi-structured grids in a non-scattering medium.

In order to quantify the differences in the asymptotic behavior of the convergence order between the different expansion orders, the convergence rate as computed by the above formula is tabulated in Table 2 and Table 3 for the Λ and Λ' spatial expansion orders, respectively.

A few remarks are in order regarding the computation of the convergence rate shown in these tables. First, it is clear from these results that the asymptotic behavior of the convergence rate p is not achieved until very fine spatial grids are used in the computation. In particular, the convergence rate for $\Lambda' = 1$ between the first two spatial grids, shown in Table 3, indicates that the meshes are too coarse to accurately obtain the expected asymptotic behavior, hence the rather high convergence rate. In addition, it is worthwhile to note that a suboptimal convergence rate is observed, especially as higher order spatial expansions Λ and Λ' are used to approximate the angular flux. This is an indication that, while the error may be decreased with the use of a higher order spatial expansion, the convergence rate possesses a ‘ceiling’ or limit which depends on the spatial regularity of exact solution.

The behavior of the scalar flux as a function of higher angular quadratures will be presented and discussed in the next section. The S_2 quadrature possesses a single direction per octant, as dictated by the $M(N) = N(N+2)/8$ relation for the level symmetric Gauss-Legendre quadrature

sets. In order to further investigate the spatial dependence of the error, it is possible to obtain an approximate error by computing the difference between the cell-wise average scalar flux between two different expansions on exactly the same grid. In particular, the difference between the cell averaged scalar flux for the $\Lambda = 0$ and $\Lambda = 2$ spatial expansions can be computed and plotted over the same mesh as was done in Ref [1].

Figure 4: Maximum scalar flux error over 27 subcubes for various Λ' spatial expansion orders and S_2 quadrature as a function of mesh refinement for unaligned semi-structured grids in a non-scattering medium.

Table 2: Convergence rates for S_2 quadrature in a non-scattering medium and various Λ expansion orders.

h(cm)	$\Lambda=0$	$\Lambda=1$	$\Lambda=2$
1.4142	0.69	2.33	2.23
0.7071	0.79	2.29	2.57
0.3536	0.85	2.30	2.67
0.1768	0.90	2.36	2.92

3.1.1.1.2 S_4 Quadrature in a Non-Scattering Medium

In order to test the behavior of the maximum error as a function of the level symmetric Gauss-Legendre angular quadrature order, the same set of spatial grids and expansion orders were used in obtaining a set of solutions employing an S_4 quadrature. However, in order to decrease the necessary runtime, the finest spatial grid was not solved for the various Λ and Λ' expansion orders.

Table 3: Convergence rates for S_2 quadrature in a non-scattering medium and various Λ' expansion orders.

$h(\text{cm})$	$\Lambda'=0$	$\Lambda'=1$	$\Lambda'=2$	$\Lambda'=3$	$\Lambda'=4$
1.4142	0.69	2.71	1.98	2.24	2.29
0.7071	0.79	1.4	2.16	2.39	2.51
0.3536	0.85	1.8	2.28	2.47	2.58
0.1768	0.90	1.99	2.31	2.58	2.75

Applying analogous scalar flux and error definitions as presented in the previous section, the maximum error for various Λ and Λ' spatial expansion orders were obtained and are shown in Figure 5 and Figure 6, respectively. Unlike the previous problem, it is evident from the figures that the maximum error has a significantly different asymptotic behavior with respect to the convergence rate, even when very small errors are achieved via mesh refinement.

Figure 5: Maximum scalar flux error over 27 subcubes for various Λ spatial expansion orders and S_4 quadrature as a function of mesh refinement for unaligned semi-structured grids in a non-scattering medium.

In order to quantify the behavior of the maximum error with respect to mesh refinement, the convergence rates were tabulated and are shown in Table 4 and Table 5 for the various Λ and Λ' spatial expansion orders, respectively.

Inspection of these convergence rates reveals significant differences from those obtained in the

previous section. In particular, the convergence rate for $\Lambda = 1$, $\Lambda' = 2$, and $\Lambda' = 3$ are particularly erratic in their behavior. This behavior is caused by the fact that the errors due to solution irregularity at the characteristic planes, and from global and local truncation, no longer coincide in any particular spatial region of the subcube domains. Hence, two competing asymptotic convergence rates may be found as the mesh is refined. The first convergence rate corresponds to the error originating from the global and local approximations, which are located at the center subcube. The second convergence rate corresponds to the characteristic planes, which contain discontinuous first-order derivatives hence lower convergence rates are achieved in these regions. Due to the fact that the larger magnitude of the error is dominated in these spatial grids by the global and local truncation errors, lower convergence rates are not observed by simply inspecting the maximum error over the subcube. In order to obtain the lowest asymptotic convergence rate of the maximum error, in which AHOTC-UG attempts to resolve discontinuities in the first-order derivative of the analytical solution, much finer spatial grids would have to be considered.

Figure 6: Maximum scalar flux error over 27 subcubes for various Λ' spatial expansions and S_4 quadrature as a function of mesh refinement for unaligned semi-structured grids in a non-scattering medium.

Table 4: Convergence rates for S_4 quadrature in a non-scattering medium and various Λ expansions.

h(cm)	$\Lambda = 0$	$\Lambda = 1$	$\Lambda = 2$
1.4142	1.39	1.48	3.35
0.7071	1.47	2.95	3.77
0.3536	1.34	6.85	3.46

Table 5: Convergence rates for S_4 quadrature in a non-scattering medium and various Λ' expansions.

h(cm)	$\Lambda' = 0$	$\Lambda' = 1$	$\Lambda' = 2$	$\Lambda' = 3$	$\Lambda' = 4$
1.4142	1.39	2.88	1.71	5.53	3.53
0.7071	1.47	3.54	3.22	6.18	3.5
0.3536	1.34	3.4	5.9	2.52	3.59

Additional quadrature orders were tested and reported in Ref [1] but the two cases summarized above, namely S_2 and S_4 capture the main features of the observed behavior.

3.1.1.2 Scattering Medium

The test cases involving the S_2 quadrature set and multiple scattering ratios also assume the same unaligned spatial grid, hence reproducing the same identical set of problems originally considered by Azmy and Barnett in their original formulation of the AHOT-C-UG methodology.

The second set of problems is designed to further test the higher order spatial representation of the angular flux, via its influence on the scattering term. The reformulation of the AHOTC-UG approach as a Petrov-Galerkin projection requires the expansion of the source term, which includes the fixed and scattering sources, into a consistent polynomial basis expansion up to order Λ or Λ' . The inclusion of scattering, in turn, requires the evaluation of the spatial moments of the scalar flux based on the spatial moments of the angular flux solution obtained from some previous inner iteration. Hence, the inclusion of scattering is an important feature, which verifies the consistency and the conditioning of the linear system of discrete variable equations. A fine-mesh TORT [2] calculation, with approximately one billion computational cells, was used in each of these cases to obtain a reference solution. Due to the subcube approach adopted for the non-scattering and scattering problems, the structured Cartesian grid solution from TORT (confined to Cartesian grids) was easily integrated over the subcubes in order to obtain the reference values.

In order to reduce THOR's execution time for practicality reasons, only four spatial grids were used for the Λ and Λ' spatial expansion orders and the S_2 angular quadrature set. The convergence criteria applied to the inner iterations was set to 10^{-7} for all spatial grids and expansion orders. The maximum error in the resulting solutions was observed to decrease in the asymptotic regime at a convergence rate that was found to be similar, albeit slightly lower, in comparison to the convergence rate of the scalar flux solution for the S_2 quadrature set in a non-scattering medium. The next sets of problems involve three values of the scattering ratio and correspond exactly to the test cases proposed by Azmy and Barnett.

3.1.1.2.1 S_2 Quadrature and $c=0.1$ Scattering Ratio

In this test problem the scattering ratio is set to $c = 0.1$ and a set of four spatial grids are used to solve the set of AHOTC-UG discrete equations for various spatial expansion orders of type Λ and Λ' , respectively. In particular, Figure 7 and Figure 8 show the behavior of the maximum error over the 27 subcubes as the spatial grid is refined for various spatial expansion orders. The maximum error curves yield a similar behavior to those encountered in the first problem of

the last section, which also involved an S_2 angular quadrature.

Figure 7: Maximum scalar flux error over 27 subcubes for various Λ spatial expansion orders and S_2 quadrature as a function of mesh refinement for $c = 0.1$ scattering case.

Figure 8: Maximum scalar flux error over 27 subcubes for various Λ' spatial expansions and S_2 quadrature as a function of mesh refinement for $c = 0.1$ scattering case.

A more quantitative approach to the analysis of the maximum scalar flux error curves can be achieved by computing the convergence order p as a function of the cell side length h . An identical methodology to the approach presented in the previous section is applied to the numerical estimation of the convergence rate. The convergence rates, tabulated in Table 6 and

Table 7, show similar results to those obtained in first problem of the last section (non-scattering medium). However, the convergence rates for all Λ and Λ' spatial expansion orders are found to be slightly lower than in the non-scattering case. In addition, the convergence rate for the $\Lambda = 2$ case between the two finest grids is found to slightly decrease in comparison to rates based on coarser grids. This lower convergence rate may originate from the fact that the reference solution for the cases involving scattering medium has to be computed numerically and may not provide sufficient accuracy with respect to the $\Lambda = 2$ results computed by AHOTC-UG.

Table 6: Convergence rates for S-2 quadrature and $c=0.1$ scattering ratio and various Λ expansions.

$h(\text{cm})$	$\Lambda = 0$	$\Lambda = 1$	$\Lambda = 2$
1.4142	0.69	2.32	2.21
0.7071	0.79	2.25	2.31
0.3536	0.85	2.07	1.69

Table 7: Convergence rates for S-2 quadrature and $c=0.1$ scattering ratio and various Λ' expansions.

$h(\text{cm})$	$\Lambda' = 0$	$\Lambda' = 1$	$\Lambda' = 2$	$\Lambda' = 3$
1.4142	0.69	2.59	1.98	2.25
0.7071	0.79	1.5	2.13	2.33
0.3536	0.85	1.89	2.09	2.16

This concludes the mesh refinement studies. The results have shown good promise for THOR and the underlying AHOTC-UG methodology because convergence of the computed solution to the reference solution was obtained in all cases. In Dr. Rodolfo Ferrer's PhD dissertation [1] additional results are presented for both the scattering and non-scattering case, but these results do not contribute any more information to this report.

3.1.2 Numerical Benchmarks

Additional numerical experiments were also completed using 3 standard computational benchmarks: Godiva, Kobayashi, and Takeda. The results of this task are published in the PhD dissertation of Dr. Rodolfo Ferrer [1], Co-PI on the original proposal that resulted in award of this project. The following results are excerpted from Ref. [1]. Only results for the Kobayashi 1ii and 2ii configurations and Godiva results are presented here. A complete discussion of AHOTC-

UG's results for the numerical benchmarks tests can be found in Ref [1].

3.1.2.1 Kobayashi Benchmark

The previous two sections have been devoted to presenting and discussing the results for a set of simplified test problems used by previous investigators [3] as a test-bed for the original AHOTC-UG formulation. In addition to these original test cases, certain variations, such as those involving aligned grids, non-scattering media, and multiple angular quadrature orders, were analyzed in order to verify the convergence of the AHOTC-UG solutions to those of the discrete ordinates transport equation. While these test problems served the set goals of establishing consistency of the underlying discretization schemes and conditioning of the linear system of equations [1], and presented some interesting features of the asymptotic behavior of the convergence order for various Λ and Λ' expansion orders, they are not completely representative of all potential theoretical and practical issues that are encountered in realistic radiation transport applications.

In order to highlight certain problems that are of particular interest to the computational radiation transport community, several benchmarks have been proposed over the last few decades in order to compare the performance of different discretizations implemented into radiation transport production codes. In particular, the treatment of void regions has been an important issue due to the fact that discrete ordinates-based angular discretizations are known to suffer from ray effects in non-scattering regions. One of several benchmarks proposed by the Expert Group on 3-D Radiation Transport Benchmarks, under the auspices of the Nuclear Energy Agency (NEA) of the Organisation for Economic Co-operation and Development (OECD) is the set of Kobayashi benchmarks [3], which requires the treatment of internal voids within simple geometric configurations.

The Kobayashi benchmarks involve a set of three problems containing a fixed source surrounded by a void material, which in turn is surrounded by a non-scattering region. The goal of the original set of problems, for which analytical solutions can be obtained, was to study the accuracy of space and angular discretizations of the mono-energetic transport equation in a non-scattering medium with internal voids. This particular problem configuration is known to pose a significant challenge to discretizations based on the discrete ordinates methods. In essence, propagation of radiation into a void region surrounded by a non-scattering region gives rise to 'ray effects', which manifest themselves as the scalar flux solution having unphysical preferential propagation in the direction of the discrete directions comprising the angular quadrature.

Due to the difficulties stemming from ray effects, a second set of problems was proposed, and solved via the Monte Carlo approach, in which the non-scattering medium was replaced by a material with $c = 0.5$ in order to mitigate the unphysical oscillations in the obtained solutions originating from the discrete ordinates approximation. Unfortunately, ray effects may still severely affect the solution, and perhaps more importantly, the convergence of the numerical solution to the exact solution.

The objective of this section is to present the results from THOR for all three cases of the Kobayashi benchmark in which the non-scattering material has been replaced by a modestly scattering medium. These sets of problems, which are referred to in the benchmark description [3] as Problem 1ii, 2ii, and 3ii, were solved with THOR using the level-symmetric Gauss-Legendre

S_{16} angular quadrature on a variety of truly unstructured spatial grids, various Λ and Λ' spatial expansion orders, and a 10^{-6} inner iteration convergence tolerance. Due to the fact that the AHOTC-UG approach is based on the discrete ordinates approximation, an asymptotic convergence behavior was not identified as a desirable goal, since ray effects will prohibit the convergence to the correct results. In addition, the reference solutions for the Kobayashi benchmark is given in terms of point values, hence an accurate assessment of the spatial convergence of the AHOTC-UG approach for $\Lambda \geq 1$ or $\Lambda' \geq 1$ would require the reconstruction of the scalar flux over particular tetrahedral cells, given a set of spatial moments obtained from the direct calculation. In order to cope with this shortcoming, a post-processing algorithm was devised which identifies the closest tetrahedron to a particular point in space, with respect to the Euclidean norm, and assigns the average value of the scalar flux computed over that tetrahedron to the point value evaluated at that specified location. Evidently, this is a gross approximation, especially for high-order spatial expansion orders in which scalar fluxes possess detailed structures within their respective domain.

Aside from these limitations, the solutions of the Kobayashi benchmarks may still yield useful information regarding the AHOTC-UG methodology. In particular, the reformulation of the AHOTC-UG approach presented in Ref [1], which is based on the equivalence between the arbitrary-order balance equation and the arbitrary-order characteristic relation, allows for the stable treatment of internal voids, which is the motivation behind the Kobayashi Benchmark. Hence, the goal of this section is to provide an overview of the results from THOR for the $\Lambda = 0$, $\Lambda = 1$, and $\Lambda' = 1$ spatial expansion order solutions over the finest possible grid for each particular problem configuration.

3.1.1.2.1 Kobayashi Benchmark Problem 1ii

A schematic of the first Kobayashi benchmark problem is shown in Figure 9. The problem configuration involves a cubic domain with dimensions $100 \times 100 \times 100 \text{ cm}^3$ and two sets of boundary conditions: vacuum boundary conditions applied to the top and side surfaces facing away from the two inner regions (shown in yellow and red), and reflective boundary conditions applied to the inner side surfaces shared by all three material regions. The three 'concentric' material regions are defined as follows: the small center cube (red) with dimensions $10 \times 10 \times 10 \text{ cm}^3$ contains a unit distributed source $S = 1$ and a total and scattering cross-section values of $\sigma_T = 0.1 \text{ cm}^{-1}$ and $\sigma_{sc} = 0.05 \text{ cm}^{-1}$, respectively, a middle void region (yellow) with dimensions $50 \times 50 \times 50 \text{ cm}^3$ containing low total cross-section in the benchmark specification for spherical harmonics methods/codes (but set to zero in the THOR calculation), and an outer region (green) containing material with the same nuclear properties as in the inner source region.

The reference solution provided for the Kobayashi Benchmark Problem 1ii is given in terms of point values along certain lines traversing the problem domain. In particular, the scalar flux reference solution is provided along the $x = z = 5 \text{ cm}$ line starting from $y = 5 \text{ cm}$ and increasing by multiples of 5 cm along the same line up to $y = 95 \text{ cm}$. In a similar development, a second reference scalar flux solution is given along the main diagonal line, defined by $x = y = z$. Finally, a third scalar flux solution along the x -direction located 5 cm behind the void region is provided, which is defined by $y = 55 \text{ cm}$ and $z = 5 \text{ cm}$. Keeping in mind the previous discussions regarding ray effects and the discrete ordinates approximation, only values for the scalar flux which remained within the two innermost regions were compared, as shown in Figure 10, Figure 11, and Figure 12 for the three trajectories described above. This range of values avoids potential

difficulties brought upon by ray effects, even in the outer region, which contains a scattering medium.

Figure 9: Kobayashi Benchmark Problem 1ii geometry model.

Figure 10: Comparison between Monte Carlo reference solution and AHOTC-UG $\Lambda = 1$ and S_{16} quadrature results on finest grid for Kobayashi Benchmark Problem 1ii along y .

Figure 11: Comparison between Monte Carlo reference solution and AHOTC-UG $\Lambda = 1$ and S_{16} quadrature results on finest grid for Kobayashi Benchmark Problem 1ii along $x = y = z$.

Figure 12: Comparison between Monte Carlo reference solution and AHOTC-UG $\Lambda = 1$ and S_{16} quadrature results on finest grid for Kobayashi Benchmark Problem 1ii along x for $y = 55$ cm and $z = 5$ cm.

Generally speaking, good agreement is found in the first two sets of scalar fluxes obtained by THOR in comparison to the reference Kobayashi benchmark results, along the y -direction and the $x = y = z$ diagonal. While it may be possible to refine the spatial mesh further to obtain an improvement in the scalar flux solution, the use of a higher-order spatial expansion may become less favorable, since it becomes more expensive to solve (in terms of computational resources), and without the ability to reconstruct the scalar flux the improvement gained by a higher-order spatial expansion is lost due to the fact that the detailed intra-cell shape cannot be reconstructed, and hence, the approximate and exact scalar fluxes at the exact coordinate location cannot be compared. Still, the AHOTC-UG reformulation to accommodate internal voids yields solutions that show good agreement with the Monte Carlo solution, even throughout the void regions, though THOR appears to overestimate the scalar flux in this same region. However, this overestimation may originate from the fact that the THOR model assumes an exact void ($\sigma_T = 0$), whereas the Monte Carlo solution used in the benchmark assumed a very small value for the total cross-section ($\sigma_T = 10^{-4}$).

In addition, **Table 8** presents tabulated results for the scalar flux at the center of the source region obtained with THOR for $\Lambda=0$, $\Lambda' = 1$, and $\Lambda=1$ as a function of mesh refinement. The error percentages presented in these tables are computed with respect to the reference Monte Carlo solution. Additionally, a reference solution is provided in this Table from TORT in order to provide a complete comparison among the various solution methods. Generally speaking, the error in the scalar flux at location (5,5,5) decreases with mesh refinement, hence approaching the Monte Carlo reference solution. However, in comparison to the scalar flux obtained by TORT, the results obtained with THOR still require further refinement in order to improve the level of agreement with respect to a TORT reference solution. It is worthwhile to note that the TORT solution was generated by creating a mesh which contains 729,000 Cartesian computational cells, while the finest THOR solution presented in this work is limited to 34,638 tetrahedra. However, it may not be necessary to run finer grids with THOR for higher-order expansions if flux reconstruction is implemented, since this allows for an improved estimation of the point quantities prescribed by the benchmark exercise.

Table 8: Results for Kobayashi 1ii benchmark scalar flux value at $r = (5,5,5)$ obtained with THOR for $\Lambda = 0$, $\Lambda' = 1$, and $\Lambda = 1$ compared to TORT and Monte Carlo reference solutions.

Expansion Order	Mesh 1	% Error	Mesh 2	% Error	Mesh 3	% Error
Cells	11,357	-	20,693	-	34,638	-
$\Lambda = 0$	5.95510	28.2	5.23805	36.8	8.50083	2.5
$\Lambda' = 1$	5.93596	28.4	5.31473	35.9	8.93778	7.8
$\Lambda = 1$	5.91844	28.6	5.31702	35.9	8.92267	7.6
TORT	8.29063	-	8.29063	-	8.29063	-
Monte Carlo	8.2926	-	8.2926	-	8.2926	-

The third reference solution, defined as a line in the x -direction located 5 cm away from the void region, shows significant disagreements with the Monte Carlo solution. Unlike the previous cases, which more or less require the flux to be computed along the particle direction of motion,

the third set of reference solution highlights the influence of ray effects in the solution. Particularly striking is the agreement between the THOR solution and the reference solution at $r = (25,55,5)$. Inspection of the S_{16} angular quadrature reveals that, in fact, this location lies on a straight line emanating from the source region along one of the discrete directions in the quadrature set. Hence, the use of discrete directions for the solution of the angular flux can accidentally provide a reasonably close solution along the discrete directions, with respect to an exact reference solution.

In addition to the one-dimensional plots of the scalar flux along the lines on which the reference Monte Carlo solution was obtained, a three-dimensional plot of the scalar flux was generated in order to visually inspect THOR's solution, as shown in Figure 13. The presence of ray effects may be observed in the figure by noting that the shape of the scalar flux is preferentially higher in the direction of the discrete angles that comprise the S_{16} quadrature.

Figure 13: Three-dimensional plot of the $\Lambda = 1$ and S_{16} quadrature solution for Kobayashi Benchmark Problem 1ii over the finest spatial grid consisting of 34,638 tetrahedral cells.

3.1.1.2.2 Kobayashi Benchmark Problem 2ii

A similar schematic to the first Kobayashi benchmark problem is shown in **Figure 14** for Kobayashi Benchmark Problem 2ii. In this particular configuration, the same boundary conditions and material properties are applied to each of the three regions assigned the same color-code as in the first Kobayashi benchmark. However, problem geometry is adapted by decreasing the height in the z -direction from 100 cm to 60 cm and the side length along the x -direction from 100 cm to 60 cm. In addition, the void region is reconfigured in order to

simulate the effects of an elongated streaming path or channel parallel to the y -direction with dimensions $10 \times 10 \times 100 \text{ cm}^3$.

Figure 14: Kobayashi Benchmark Problem 2ii geometry model.

The reference solution provided for the Kobayashi Benchmark Problem 2ii is given in terms of scalar flux point values along certain three-dimensional lines. In this particular case, the scalar flux reference solution is provided along the streaming channel or void region (yellow) in the y -direction, namely $x = y = 5 \text{ cm}$ and $y_i = 5 + 10i$, where $i = 0, \dots, 9$. A comparison of the scalar flux at the first five points between THOR and the reference Monte Carlo solutions are shown in **Figure 15** that exhibits a similar agreement to that observed for Problem 1ii. Since the second reference solution in the x -direction is computed along the opposite end of the problem, and hence very far from the source, it is highly susceptible to ray effects detrimental to accuracy, so no comparison is performed between the results obtained by THOR and the reference solution for this trajectory.

Table 9 presents tabulated results for the scalar flux obtained with THOR for $\Lambda = 0$, $\Lambda' = 1$, and $\Lambda = 1$ as a function of mesh refinement. Generally speaking, the error in the scalar flux location (5,5,5) decreases with mesh refinement, hence approaching the Monte Carlo reference solution. An analogous trend in terms of percentage error as a function of mesh refinement is found in this problem as compared to the Kobayashi Benchmark Problem 1ii.

A three-dimensional plot of the scalar flux produced by THOR was generated in order to visually inspect the solution, as shown on the left in Figure 16. As expected, the presence of streaming along the void channel yields a higher flux level in the voided region. This streaming effect is verified by inspecting the opposite viewpoint of the three-dimensional configuration, which, as shown on the right in Figure 16, exhibits a peak in the scalar flux at the exiting area of the void region. The fact that this peak is not centered about the y -axis is evidence of the ray effects.

Finally, a two-dimensional view of the scalar flux generated by THOR is shown in Figure 17. This plot of the scalar flux was generated by restricting the mesh to the plane defined by $z = 0 \text{ cm}$, that is on the very bottom of the schematic for Problem 2ii, shown in Figure 17. This two-dimensional plot shows, in effect, the (x,y) distribution of the scalar flux across the very

geometric center of the problem, since reflective boundary conditions are applied across the planes defined by $x = 0$, $y = 0$, and $z = 0$. As expected, the scalar flux is observed to preferentially stream in the direction of the voided channel along the y -direction, a correct behavior in contrast to the unphysical ray effects.

Figure 15: Comparison between Monte Carlo reference solution and AHOTC-UG $\Lambda = 1$ and S_{16} quadrature results on finest grid for Kobayashi Benchmark Problem 2ii along y , $x = z = 5$ cm.

Table 9: Results for Kobayashi 2ii Benchmark scalar flux value at $r = (5,5,5)$ obtained with THOR for $\Lambda = 0$, $\Lambda' = 1$, and $\Lambda = 1$ compared to TORT and Monte Carlo reference solution.

Expansion Order	Mesh 1	% Error	Mesh 2	% Error	Mesh 3	% Error
Cells	12,687	-	20,944	-	39,331	-
$\Lambda = 0$	8.68244	0.8	6.21674	27.9	8.46783	1.7
$\Lambda' = 1$	9.21726	7.0	6.15575	28.6	8.64850	0.4
$\Lambda = 1$	9.18939	6.6	6.16986	28.4	8.64483	0.3
TORT	8.61209	-	8.61209	-	8.61209	-
Monte Carlo	8.61696	-	8.61696	-	8.61696	-

Figure 16: Three-dimensional plots of the $\Lambda = 1$ and S_{16} quadrature solution for Kobayashi Benchmark Problem 2ii over the finest spatial grid consisting of 39,331 tetrahedral cells.

Figure 17: Two-dimensional plot of the $\Lambda = 1$ and S_{16} quadrature fine grid solution for Kobayashi Benchmark Problem 2ii over the x, y plane for $z = 0$ cm.

3.1.2.2 Godiva Benchmark

The previous set of Kobayashi benchmarks tested the AHOTC-UG methodology, implemented into the THOR computer code, against Monte Carlo solutions for simple ‘shielding’ geometries involving internal voids. Certain difficulties were encountered in the comparison between the

THOR results and the reference solutions due to the discretization of the angular variables via the discrete ordinates approach. However, the results obtained via THOR were found to be in good agreement with the Monte Carlo solution, especially in light of the presence of ray effects. In addition, two important conclusions may be drawn from the exercise; first, the reformulation of the AHOTC-UG approach reported in Ref [1] does in fact allow for the treatment of pure void regions in a manner that is numerically stable and consistent with the previous formulation, and second, the original AHOT-C formalism has been correctly generalized to a method that can handle three-dimensional unstructured grids composed of arbitrary tetrahedrons.

In order to further test the AHOTC-UG reformulation, and showcase the advantage of unstructured tetrahedral grid spatial meshing, the Godiva benchmark [4] was solved with various spatial grids, Λ and Λ' spatial expansion orders, and several S_N angular quadratures. The Godiva benchmark is a criticality problem for a bare spherical fast neutron system (reactor) with a radius of 8.71 cm, which consists of a single material region composed of highly enriched uranium with reference six-group nuclear data consisting of total, scattering, and fission cross-sections [4]. This criticality benchmark problem can be solved with one-dimensional spherical-geometry transport codes with multi-group and criticality search capabilities. However, the simple geometry and material composition of the Godiva benchmark can be advantageous in the development and verification of general unstructured geometry transport codes and discretizations.

In addition to testing the unstructured grid capabilities of the AHOTC-UG formulation, the Godiva benchmark was used to test the multi-group and eigenvalue search algorithms implemented into THOR. A set of three spatial grids, shown in Figure 18 and consisting of 274 (left), 2,945 (middle) and 20,055 cells (right), were generated to model the Godiva benchmark geometry under the assumption of reflective boundary conditions across boundaries defined by the $x = 0$, $y = 0$, and $z = 0$ planes. In order to conserve the overall volume of the system, the solid body description of the sphere was artificially adjusted until the meshed volume became roughly equal to the real volume, or at least within 0.05 percent difference.

Figure 18: Spatial grids generated for the Godiva benchmark.

Unlike previous test problems and benchmarks cases, the reference solution for the Godiva benchmark is given by a single number, the multiplication factor k_{eff} , which is the eigenvalue belonging to the fundamental or dominant eigenmode. Several results for the Godiva benchmark eigenvalue are presented in [4], each of which depends on the particular numerical method used to solve the transport equation. For the purposes of benchmarking THOR, the

reference solution ($k_{eff} = 0.99597$) is assumed to be the eigenvalue computed by the DIRECT code, which is based on integral transport theory and, according to the authors of [4], is the best available solution to the benchmark.

In the first set of results the level symmetric Gauss-Legendre S_4 angular quadrature and various Λ and Λ' spatial expansions were used to solve the eigenvalue problem over the three spatial grids depicted in Figure 18. The eigenvalue and source convergence tolerances were set to 10^{-6} and at most four inner iterations per outer iteration were allowed in all cases. Figure 19 depicts the absolute difference, or error, between the eigenvalue obtained by THOR and the reference solution for $\Lambda = 0$ and $\Lambda = 1$ as a function of the longest edge-length over all tetrahedral cells comprising each spatial grid, which was computed numerically. In addition, Figure 20 shows the error in the eigenvalue as a function of mesh refinement for $\Lambda' = 0, 1, 2$ spatial expansion orders.

Figure 19: Maximum eigenvalue error for various Λ spatial expansion orders and S_4 quadrature as a function of mesh refinement for the Godiva benchmark.

As expected, the error in the eigenvalue is observed to decrease monotonically as the mesh is refined and the spatial expansion order is increased. In fact, it is possible to estimate the convergence rate of the error in the eigenvalue in the same manner used to estimate the convergence rate of the maximum error for the problems involving the cube domain with either non-scattering or scattering material. These eigenvalue convergence rates are tabulated in

Figure 20: Maximum eigenvalue error for various Λ' spatial expansions and S_4 quadrature as a function of mesh refinement for the Godiva benchmark.

Table 10 for the various Λ and Λ' spatial expansion orders. It is worthwhile to note that, while the convergence rate is higher for $\Lambda = 1$ relative to the convergence rate of $\Lambda' = 1$ and $\Lambda' = 2$, the average normalized runtime between the former and the two latter expansion orders is roughly 24 and 2, respectively. In addition, if the degrees of freedom are defined as the number of unknowns based on the spatial expansion orders, and under the assumption that this quantity can be related to memory requirements, then the normalized computational burden in terms of memory requirement is roughly 2.4 between $\Lambda=1$ and $\Lambda' =1$ and 1.1 between $\Lambda=1$ and $\Lambda' =2$.

Figure 20: Maximum eigenvalue error for various Λ' spatial expansions and S_4 quadrature as a function of mesh refinement for the Godiva benchmark.

Table 10: Eigenvalue convergence rates for Godiva benchmark with respect to various Λ and Λ' expansions and the S_4 angular quadrature.

h (cm)	$\Lambda = 0$	$\Lambda = 1$	$\Lambda' = 1$	$\Lambda' = 2$
1.877468	1.07	3.99	2.98	3.40
1.051101	1.01	3.52	2.28	3.29

The geometry of the Godiva benchmark, which involves a perfectly spherical homogeneous region, makes it possible to solve the criticality benchmark as a one-dimensional problem in the radial, or r , direction. Conversely, it is worthwhile to note that, given a perfectly symmetric sphere centered about the origin of a global Cartesian system, the three-dimensional angular flux solution will be identical with respect to all angular directions, provided that symmetric rotations are applied so as to match the boundary conditions in each direction. This peculiarity is due to the fact that the spherical domain is independent of any particular angular direction with respect to the global coordinate system. Hence, the solution of the Godiva benchmark in three-dimensional geometry should be independent of the particular angular quadrature set under consideration. As depicted in Figure 21 and Figure 22, given a fixed mesh (in this case coarsest spatial grid), and various spatial approximations of the angular flux, the use of different angular quadrature will not change the solution to the transport problem. These results not only point to an interesting feature of the Godiva benchmark, but also confirm that the AHOTC-UG approach is correctly solving a multi-group eigenvalue problem in an arbitrary tetrahedral grid that is angle-independent. Any minor differences in the error as a function of angular

quadrature can be attributed to the fact that the perfectly spherical geometry is being approximated by a spatial grid consisting of tetrahedrons with piecewise linear faces.

Figure 21: Maximum eigenvalue error for various Λ spatial expansion orders and various S_N quadratures given a fixed mesh (coarsest) for the Godiva benchmark.

Figure 22: Maximum eigenvalue error for various Λ' spatial expansions and various S_N quadratures given a fixed mesh (coarsest) for the Godiva benchmark.

Finally, a two-dimensional plot of the scalar flux over the finest spatial grid for the thermal energy group $g = 6$ is shown in Figure 23. The $\Lambda = 1$ spatial expansion order was applied in this

particular test case along with the S_8 angular quadrature. Inspection of the spatial distribution of the scalar flux reveals that the solution only depends on the radial position, hence it is truly a one-dimensional problem. In addition, the coupling between the scattering and fission source distributions appears to have been implemented correctly, given the behavior of the eigenvalue convergence as a function of spatial mesh and angular quadrature, and the shape of the scalar flux for the thermal group, which depends on the down-scattering from higher energy groups into lower energy groups. Since any irregularity in the solution of the scalar flux for $g \leq 5$ produced by THOR would eventually find its way into the evaluation of the lower energy group source distribution through either scattering or fission, it may be safely concluded that the AHOTC-UG approach, along with the THOR implementation, has been verified via the Godiva benchmark.

In addition to the presented results, a ‘frozen’ version of the code has been established by Dr. Rodolfo Ferrer completed and submitted a draft of the THOR code manual thereby completing Task A of this project.

Figure 23: Two-dimensional plot of the thermal ($g = 6$) scalar flux computed with $\Lambda = 1$ and S_8 quadrature on the finest grid for the Godiva Benchmark.

3.2 Task B – Development of MMS Benchmark Suite (NCSU)

3.2.1 FEM Formulation

A Two-Dimensional version of the MMS benchmark suite based on variants of Larsens' benchmark has been implemented in the code MSBS-2D and verified versus *Mathematica* results. The code MSBS-2D computes sets of cell Legendre moments of the angular and scalar fluxes and of the distributed source up to arbitrary order via analytical integration. MSBS-2D allows for an arbitrary degree of solution regularity that can be controlled via the boundary conditions. A summary describing MSBS-2D was accepted for publication at the ANS Winter meeting in Las Vegas, Nevada [5], and a comprehensive paper was submitted to the ANS Math & Comp topical meeting in Rio de Janeiro, Brazil, 2011 [6].

The following spatial discretization methods for Cartesian meshes were implemented: The arbitrarily high order methods of the nodal type (AHOTN) and of the characteristic type (AHOTC), bi-polynomial Discontinuous Galerkin Finite Elements Method (DGFEM) and the Higher Order Diamond Difference scheme (HODD). A comprehensive error analysis based on test cases created with MSBS-2D was conducted for the aforementioned discretization schemes. The outcome of this study was that AHOTN and AHOTC feature superior accuracy for optically thick cells, while the difference among the four methods is marginal for sufficiently optically thin cells.

All four discretization schemes, AHOTN, AHOTC, HODD, and HODD, can be cast into a FEM framework, thus furnishing a theoretical foundation for comparison among the methods. Further, the particular trial spaces can be employed within the flux reconstruction procedure described later. The results described above were summarized in a paper and submitted to the ANS Math & Comp topical meeting in Rio de Janeiro, Brazil, 2011 [6].

A flux reconstruction method was devised that determines the trial function expansion coefficients of the within cell flux shape for AHOTN, DGFEM and HODD. Thus, for any of these discretization schemes we can provide interpolation formulas across spatial mesh cells either for obtaining point flux values at an arbitrary position within the mesh cell or for restricting and averaging the flux to some subset of the mesh cell (prolongation). In previous reports we described the flux reconstruction method as the tool for prolonging the flux from the numerical mesh to a fixed reference mesh. However, we realized that employing a continuous L_p norm is more rigorous than prolonging the flux to a fine mesh and then calculating the error and hence we abandoned the prolongation approach. In favor of the three-dimensional study, we did not implement the computation of the continuous L_p norm within the framework of the two-dimensional study. We emphasize here that the prolongation is only one potential area of application of the flux reconstruction capability and that we will pick up the results obtained so far if we opt to measure the accuracy for the three-dimensional study in a continuous L_p norm.

3.2.2 Development & Implementation of Three-Dimensional MMS for THOR

The Method of Manufactured Solutions (MMS) in conjunction with the order of convergence test is the state of the art method for computer code verification. Instead of attempting to obtain an analytical solution, the MMS prescribes the exact solution as a known, analytical

function f and then determines the corresponding source term S such that f solves the partial differential equation (in our case the discrete ordinates approximation of the one-speed, steady-state transport equation) given the source S . As this procedure typically only requires differentiating (and potentially integration operations) f it is usually straight forward to obtain S . Selection of the boundary conditions is critical to attaining the required regularity of the exact solution as explained below.

For the verification of THOR let a mono-energetic transport problem be given on the homogeneous cuboidal domain D with dimensions $[0, X] \times [0, Y] \times [0, Z]$. The selected analytical solution shall be the solution of the non-scattering auxiliary transport problem:

$$\begin{aligned}\hat{\Omega} \cdot \nabla \psi_n + \sigma_t \psi_n(\vec{r}) &= Q \\ \psi_n(\vec{r}) &= B_n(\vec{r}) \text{ for } \vec{r} \in \partial D^-\end{aligned}$$

where $\hat{\Omega}_n = (\mu_n, \eta_n, \zeta_n)$ is a single unit vector selected from a quadrature set $\{\hat{\Omega}_n, w_n\}_{n=1, \dots, N}$; σ_t and Q are the constant, positive total cross section and an angularly isotropic and spatially uniform, non-negative parameter, respectively. Using this auxiliary problem to define the manufactured solution ensures its physical meaningfulness and guarantees that parameters can be chosen to render the final distributed source positive. Prescribed inflow boundary conditions are given on the inflow boundary ∂D^- . Using the method of characteristics the solution for a discrete ordinate in the positive octant $\mu_n, \eta_n, \zeta_n > 0$ can be obtained as:

$$\psi_n = \begin{cases} \psi_w \left(y - \frac{\eta_n}{\mu_n} x, z - \frac{\zeta_n}{\mu_n} x \right) e^{-\frac{\sigma_t x}{\mu_n}} + \frac{Q}{\sigma_t} \left(1 - e^{-\frac{\sigma_t x}{\mu_n}} \right), & Or_y(\vec{r}) > 0, Or_z(\vec{r}) < 0 \\ \psi_s \left(z - \frac{\zeta_n}{\eta_n} y, x - \frac{\mu_n}{\eta_n} y \right) e^{-\frac{\sigma_t y}{\eta_n}} + \frac{Q}{\sigma_t} \left(1 - e^{-\frac{\sigma_t y}{\eta_n}} \right), & Or_x(\vec{r}) < 0, Or_z(\vec{r}) > 0 \\ \psi_b \left(x - \frac{\mu_n}{\zeta_n} z, y - \frac{\eta_n}{\zeta_n} z \right) e^{-\frac{\sigma_t z}{\zeta_n}} + \frac{Q}{\sigma_t} \left(1 - e^{-\frac{\sigma_t z}{\zeta_n}} \right), & Or_x(\vec{r}) > 0, Or_y(\vec{r}) < 0 \end{cases} \quad (1)$$

where ψ_w , ψ_s and ψ_b are the prescribed inflow boundary conditions on the West, South and Bottom domain face, respectively. As evident from Eq. 1 the solution is given by distinct expressions over the three subdomains illuminated by each of the incoming face boundary conditions; the subdomains are delineated by $Or_k(\vec{r})$ with respect to the k direction with $k = x, y, z$ is given by:

$$Or_k(\vec{r}) = \vec{r} \cdot (\hat{\Omega} \times \hat{e}_k).$$

Equivalent expressions apply to all other octants. Note that depending on the boundary conditions on the inflow faces, the solution can feature various degrees of smoothness because multiple expressions in Eq. 1 could apply when $Or_k(\vec{r}) = 0$ for one or more k .

For the general verification of THOR a manufactured solution of the transport problem with scattering is desired:

$$\begin{aligned}\hat{\Omega} \cdot \nabla \psi_n + \sigma_t \psi_n(\vec{r}) &= q(\vec{r}) + \sigma_s \phi(\vec{r}) \\ \phi(\vec{r}) &= \sum_{n=1}^N w_n \psi_n(\vec{r}) \\ \psi_n(\vec{r}) &= B_n(\vec{r}) \text{ for } \vec{r} \in \partial D^-. \end{aligned} \quad (2)$$

where $\sigma_s \leq \sigma_t$ is a non-negative isotropic and spatially uniform scattering cross section and q is the distributed source. By inspection Eq. (1) is the solution of Eq. (2) if the distributed source q is computed using:

$$q(\vec{r}) = Q - \sigma_s \phi(\vec{r}). \quad (3)$$

For the purpose of this verification exercise a smooth (i.e. infinitely differentiable) exact solution is desired such that potential non-smoothness does not interfere with the observed order of accuracy with mesh refinement and hence with the verification exercise. Selecting the following boundary conditions (again, but without loss of generality, for a discrete ordinate in the positive octant):

$$\begin{aligned}\psi_w &= C e^{-\frac{\sigma_t}{\eta_n} y + \frac{\sigma_t}{\zeta_n} z} + \frac{Q}{\sigma_t} \\ \psi_s &= C e^{-\frac{\sigma_t}{\mu_n} x + \frac{\sigma_t}{\zeta_n} z} + \frac{Q}{\sigma_t}, \\ \psi_b &= C e^{-\frac{\sigma_t}{\mu_n} x - \frac{\sigma_t}{\eta_n} y} + \frac{Q}{\sigma_t}\end{aligned}$$

leads to a manufactured solution featuring no unbounded partial derivatives:

$$\psi_n = C e^{-\frac{\sigma_t}{\mu_n} x - \frac{\sigma_t}{\eta_n} y + \frac{\sigma_t}{\zeta_n} z} + \frac{Q}{\sigma_t}. \quad (4)$$

In the above equations C is an arbitrary positive constant.

The MMS benchmark suite for tetrahedral meshes is implemented in the code MMS3D(UG). Through the MMS formalism the exact, pointwise solution everywhere in the domain is known, as well as corresponding distributed sources and boundary conditions. On input THOR requires monomial moments over a cell (tetrahedron) and face (triangle) of the source and boundary conditions, respectively. On output THOR computes cell monomial moments of the approximate solution. Therefore, MMS3D(UG) needs to compute the required monomial cell and face moments of the corresponding quantities. The algorithm performing these tasks needs to integrate the product of monomials and the angular flux Eq. (4) over cell volumes and boundary

faces. Instead of the analytical integration routines utilized in the purely Cartesian code MMS3D, MMS3D(UG) uses numerical integration routines comprised in the *cubpack* numerical software package [7].

3.2.3 THOR Illustration of MMS

Using the manufactured solution Eq. 4 two test cases are set up to facilitate the order of convergence verification test of THOR. The order of convergence test compares the observed order of accuracy r_o obtained from a mesh-refinement study based on the MMS test problems to the formal order of accuracy inherently associated with the discretization method itself. The formal orders of accuracy for the participating THOR discretization methods are listed in Table 11. If observed and formal order agree to within some reasonable tolerance correctness of the computer code is inferred. Within this verification exercise the step characteristics (Step), linear characteristics (LC), quadratic characteristics (QC) and mixed-trilinear characteristics (TLC) methods are verified. The basis functions used for the expansion of the source and inflow fluxes along with the resulting formal orders of accuracy for the four discretization methods are reported in Table 11.

Table 11: Comparison of THOR discretization methods and formal orders of accuracy.

Method	Source exp.	Face flux exp.	Formal accuracy
Step	$\{1\}$	$\{1\}$	1
LC	$\{1, x, y, z\}$	$\{1, x, y\}$	2
QC	$\{1, x, y, z, xy, yz, xz, x^2, y^2, z^2\}$	$\{1, x, y, xy, x^2, y^2\}$	3
TLC	$\{1, x, y, z, xy, yz, xz, xyz\}$	$\{1, x, x^2, y, y^2, xy, x^2y, xy^2, x^2y^2\}$	3

The parameters characterizing the two MMS test cases I and II are reported in Table 12. For both cases I and II, the S_4 level symmetric quadrature is utilized throughout. A sequence of 6 embedded meshes featuring 384, 3,072, 24,576, 196,608, 1,572,864 and 12,582,912 tetrahedrons (the last mesh only for LC and test case II) are employed. The first mesh is created using the *netgen* meshing tool for test case I then each mesh 2 through 6 is obtained from the previous mesh by octasection. The resulting 6 meshes are extended to case II by simply stretching the vertex coordinates by the ratio of the respective domain dimensions along each coordinate. In the framework of a mesh refinement study the solution, i.e. the cell average scalar flux in each tetrahedron, is obtained and compared to the reference computed by averaging the MMS Eq. 4 over each tetrahedral cell. For this purpose a code MMS3D(UG) was created that computes the integrals of the scalar flux, distributed source and inflow fluxes (potentially weighted with the polynomial expansion functions in Table I), over tetrahedrons and boundary triangles, respectively. The code MMS3D(UG) extends the functionality previously implemented and exercised in Cartesian geometry to tetrahedral geometry. In contrast to its Cartesian counterpart MMS3D(UG) uses numerical integration instead of analytical expressions for obtaining cell polynomial moments of the source and scalar fluxes utilizing the *cubpack* numerical software package. From the difference of the exact and numerical scalar fluxes the L_2 error norm defined as:

$$e_{2,h} = \left[\sum_i^{I_t} V_i (\phi_i^h - \phi_i)^2 \right]^{1/2},$$

is computed, where V_i is the volume of tetrahedron $i = 1, \dots, I_t$. Further, ϕ_i^h is the numerical estimate of the cell average scalar flux and ϕ_i is the exact cell average flux. In the framework of the mesh refinement study the L_2 error norm is computed for all meshes such that the rate of convergence estimated from the mesh h_1 to mesh h_2 solutions is given by:

$$r_o = \frac{\log \frac{e_{2,h_2}}{e_{2,h_1}}}{\log \frac{h_2}{h_1}}.$$

Table 12: Parameters for MMS test cases I and II.

Test Case	$X/Y/Z$	σ_t	σ_s	Q	C
I	1/1/1	1.0	0.2	0.8	0.8
II	10/1/0.2	0.5	0.05	0.8	0.8

Because the MMS solution in this case, i.e. Eq. 4, is analytic the order of convergence test stipulates that for each method $r_o \rightarrow r_f$ (the corresponding formal accuracy) listed in Table 11 verifies correct implementation of this method in THOR.

In Figure 24 the L_2 error is plotted versus the mesh parameter $h = (X \cdot Y \cdot Z / I_t)^{1/3}$ and in **Table 13** the L_2 error along with the associated observed accuracies are assembled for test case I. Note, that QC and TLC implementations are significantly slower to execute thus less computationally efficient than the SC and LC implementations such that not all meshes are utilized for these two expansion orders. Future versions of THOR will improve on the efficiency of these expansion orders. For this test case the observed accuracy orders approach the formal accuracies with mesh refinement thus indicating the correct implementation of the SC, LC, TLC and QC spatial discretizations within THOR. Analogous results for the second test case that features domain aspect ratios that deviate significantly from unity (about 50) are reported in Figure 25 and Table 14, respectively. Since the domain's tessellation is created by simply displacing the vertices obtained when meshing test case I's domain, the tetrahedrons comprising case II's mesh feature aspect ratios that deviate significantly from unity. As practical geometries for nuclear reactor cores often feature a very detail-rich x-y plane and are rather uniform along the z-axis, meshes tailored to these problems are likely to exhibit large aspect ratios as included in this test case. Again, the observed accuracies approach the formal orders of accuracy, even though finer meshes are required to attain the same proximity to the formal accuracy, suggesting that THOR's implementation retains its correctness, i.e. does not deteriorate, for meshes with aspect

ratios that deviate from unity significantly. The MMS benchmark exercise demonstrates that by applying the convergence order test conjecture the selected discretization methods: Step characteristics, linear characteristics, quadratic characteristics and tri-linear characteristics are implemented correctly.

Figure 24: L_2 error versus mesh parameter h for selected spatial discretization methods and MMS case I.

Table 13: L_2 error and observed accuracies r_o for test case I.

h	SC		LC		TLC		QC	
	$\ \epsilon\ _2$	r_o	$\ \epsilon\ _2$	r_o	$\ \epsilon\ _2$	r_o	$\ \epsilon\ _2$	r_o
1.38-1	6.41-2	-	7.96-3	-	3.24-4	-	3.97-4	-
6.88-2	3.73-2	0.78	2.31-3	1.78	4.01-5	3.01	4.87-5	3.03
3.44-2	2.02-2	0.88	6.31-4	1.87	4.78-6	3.07	5.75-6	3.08
1.72-2	1.06-2	0.93	1.66-4	1.93	5.66-7	3.08	-	-
8.60-3	5.43-3	0.97	4.25-5	1.96	-	-	-	-

Figure 25: L_2 error versus mesh parameter h for selected spatial discretization methods and MMS case II.

Table 14: L_2 error and observed accuracies r_o for test case II.

h	SC		LC		TLC		QC	
	$\ \epsilon\ _2$	r_o	$\ \epsilon\ _2$	r_o	$\ \epsilon\ _2$	r_o	$\ \epsilon\ _2$	r_o
1.38-1	3.61-2	-	5.92-3	-	4.73-4	-	5.61-4	-
6.88-2	2.53-2	0.51	2.66-3	1.15	1.87-4	1.34	2.09-4	1.42
3.44-2	1.47-2	0.78	1.00-3	1.41	4.42-5	2.08	4.71-5	2.15
1.72-2	7.97-3	0.89	3.42-4	1.55	6.87-6	2.69	7.20-6	2.71
8.60-3	4.23-3	0.92	1.06-4	1.70	-	-	-	-
4.30-3	-	-	2.88-5	1.87	-	-	-	-

3.2.4 Development of a GUI for MMS3D

In order to make use of the MMS3D and MMS3D(UG) utilities easier for the user, a GUI was created assisting in setting up desired test problems for three-dimensional Cartesian grids and three-dimensional tetrahedral grids, respectively. The GUI is written in the python scripting language using the Tkinter package.

3.3 Task C – Grind Times Study (NCSU)

The objectives of Task C are (i) to create stable and efficient AHOTC-UG kernels; (ii) to measure the grind time (execution time per iteration, per mesh cell, per discrete ordinate) and to compare these to grind times of available production level codes; and (iii) to conduct a study of the behavior of the execution time with increasing angular quadrature order and number of mesh cells to verify the theoretically expected linear profile.

3.3.1 Objectives (i) & (ii)

Constant and linear AHOTC-UG kernels were implemented in THOR, and these will be referred to as Step Characteristics (SC) and Linear Characteristics (LC) kernels. The original arbitrary expansion order implementation is still available in THOR but its grind times are too large to be practical for large problems at this stage. The SC and LC kernels underwent verification as outlined above in Tasks A and B. Numerical stability is ensured by computing volume and face integrals using asymptotic expansions for optically thin cells and analytical integrations above a certain threshold.

3.3.2 Objective (iii)

Transport solvers are in essence a set of loops wrapped around the kernel computation (per cell, per angle, per energy group). Therefore, the total execution time should be proportional to the number of traversals through these loops. For each energy group, one loop iterates over the number of angular directions which in turn is wrapped around the loop over all spatial mesh cells. Therefore, the execution time per inner iteration (single group by definition) should be proportional to the number of angular directions and the number of mesh cells.

3.3.2.1 Linearity with respect to number of angular directions

Within the execution of this task no modifications to the code were necessary to verify linearity of the execution (or grind) time with respect to number of angular directions comprising the employed angular quadrature.

Figure 26 illustrates the said linearity for the Takeda IV and Godiva test problems exercising the SC and LC discretization methods.

3.3.2.2 Linearity with respect to number of mesh cells

In Cartesian meshes the order in which mesh cells are visited in the course of a mesh sweep along a given discrete ordinate is fully specified by cell indices ordered in a Cartesian grid, and therefore the “mesh sweep” reduces algorithmically to a set of loops with fixed starting index, ending index and index increment. Thus, the mesh sweep's execution will naturally scale linearly with the number of mesh cells. In unstructured meshes the order in which mesh cells have to be visited is not generally fixed by the cell indices and depends on the angular direction. The algorithm that determines this order can make the mesh sweep scale super-linearly with the number of mesh cells. This super-linear scaling can make a transport code inefficient for large meshes. Older versions of THOR suffered from this super-linear scaling. A Breadth-First Search algorithm that pre-computes the sweep order of the mesh cells was implemented into THOR. Scaling of the execution time with number of mesh cells is depicted in **Figure 27**. The simple cube test case is a unit cube with a mesh created from subdividing an orthogonal grid of

subcubes into tetrahedrons. The netgen cube configuration is the same domain as the simple cube test problem but meshed with the netgen meshing tool. Takeda IV is a hexagonal fast reactor benchmark and the ATR is the Advanced Test Reactor at Idaho National Laboratory.

Figure 26: Execution time per inner iteration *versus* number of discrete ordinates for two test problems (Godiva and Takeda), various meshes and SC and LC discretization methods. For all cases the execution times scale linearly with number of discrete ordinates.

Figure 27 demonstrates that the execution time for a single mesh sweep scales linearly with the number of mesh cells. Further, extending the SC netgen cube results to about 3 million mesh cells the ATR results line up nicely with the results obtained for the significantly simpler netgen cube test case. This indicates that THOR's performance is scalable to large meshes.

3.3.2.3 Grind time study

The grind time for each case shown in

Figure 26 and **Figure 27** is determined by the vertical-axis intercept as described below. For the dependence on angle (

Figure 26) we have:

$$t_i = t_g \times n_c \times N ,$$

where t_i is the execution time for a single inner iteration, t_g is the grind time, n_c is the number of cells and N is the number of angular directions. Taking the logarithm gives:

$$\log(t_i) = \log(N) + \log(t_g) + \log(n_c)$$

Note the curves shown in

Figure 26 and **Figure 27** are parallel to the linear trend (indicated by the dashed line in these figures) thereby validating the coefficient of the N and $\log(n_c)$ terms in the above equation. The axis intercept, i.e. $\log(t_i)$ at $N = 1$, is the sum of the logarithms of the grind times and the number of mesh cells.

Figure 27: Execution time per inner iteration and discrete ordinate (execution time of one mesh sweep) *versus* number of mesh cells for the simple cube, netgen cube, Takeda and ATR test cases. For the ATR only a single SC data point is available.

For **Figure 27** (where t_i is divided by the number of angular directions) the final expressions become:

$$t_{i'} = t_i / N$$

$$\log(t_{i'}) = \log(t_g) + \log(n_c)$$

Writing t_g as the product of the average number of CTs per cell, κ , times the CT grind time

$t_{g,CT}$ the axes intercepts c_{ax} for

Figure 26 and **Figure 27**, become:

Figure 26:
$$c_{ax} = \log(\kappa) + \log(t_{g,CT}) + \log(n_c)$$

Figure 27:
$$c_{ax} = \log(\kappa) + \log(t_{g,CT}).$$

In Table 15 grind times of THOR's SC and LC methods per characteristic tetrahedron determined as described above from the measured results shown in

Figure 26 and **Figure 27** are compared to DENOVO and TORT grind times. [THOR's grind times were computed as the arithmetic average (separately for SC and LC) of the values of $t_{g,CT}$ computed as described above from the measured times for the simple cube, Godiva, and Takeda caes.] THOR's AHOTC-UG method slices each tetrahedral cell into characteristic tetrahedra (CT) and applies the characteristic equations to the slices. Depending on the case, a mesh cell can be divided into 2 to 4 tetrahedra and therefore depending on the particular case considered the cell's grind time can vary. Typically the cell's grind time is about four times as large as the CT grind time. The comparison in Table 15 is not entirely fair because DENOVO and TORT are orthogonal mesh codes which solve significantly easier equations per cell at the price of potentially using many more mesh cells to approximate a complicated geometry appropriately.

While the Cartesian-mesh production codes TORT and DENOVO feature much smaller grind times (note that LC and SC real grind times are four times higher) THOR's CT grind times are in the ball park of production level codes given that THOR uses tetrahedral meshes.

Table 15: Grind times per cell (per CT for SC, LC) in micro seconds for various methods. LD is the linear discontinuous method, TLD trilinear-discontinuous, TWD stands for Theta-weighted Diamond method.

Method	Grind Time (μs)
DENOVO SC	0.220
DENOVO LD	0.370
DENOVO TLD	2.900
TORT TWD	0.100
THOR SC	0.700
THOR LC	5.800

3.4 Task D – Novel Algorithm for Computing the Fundamental Eigenmode (NCSU)

This report section the implementation of the Jacobian-Free Newton Krylov (JFNK) method for the solution of the k -eigenvalue problem into the transport code THOR. We begin with a brief description of how the new algorithm is implemented into THOR, then we review the theoretical foundation of the JFNK method for solving eigenvalue problems, and finally we compare the performance of the implemented JFNK method to the standard Power Iteration (PI) method based on two test problems.

3.4.1 Implementation into THOR

The implementation of the transport kernel, i.e. the solution of the characteristic equations within each tetrahedron, into THOR is rather involved and has been described in earlier sections of this report as well as in Ref [1]. A standard inner-outer iteration scheme is wrapped around the transport kernel to facilitate the computation of k -eigenvalue problems using the power iteration method; additionally, though of little importance to this project, the code allows for fixed source problems to be solved. Neither the inner nor the outer iterations are accelerated and subsequently a high scattering ratio or dominance ratio can inhibit fast convergence. To date three versions of THOR have been equipped with the JFNK module.

1. THOR-C0: The spatial expansion is fixed to 0-th order (step characteristic) and the characteristic integrals are calculated on the fly. This version is stable for optically thick cells, and executes very fast.
2. THOR-CC: The spatial expansion order is arbitrary-order. The characteristic integrations are not calculated on the fly, but are given in terms of a pre-computed power series in terms of the optical thickness. Since the truncation error of this series becomes significant, the method becomes unstable for optical thicknesses greater than unity. However, execution is faster than for the next version THOR-CCE.
3. THOR-CCE: The spatial expansion order is arbitrary-order and the characteristic integrations are computed on the fly. This version executes very slowly but is stable for optically thick cells.

THOR features two different collections of spatial moments that are retained as the spatial order Λ is increased. Let i, j, k be the spatial moment indices associated with the x, y and z direction. Then:

$$\begin{aligned}\text{Option 1: } & 0 \leq i \leq \Lambda, 0 \leq j \leq \Lambda, 0 \leq k \leq \Lambda \\ \text{Option 2: } & 0 \leq i + j + k \leq \Lambda\end{aligned}$$

In the remainder of this section the second option is denoted by negative spatial order.

3.4.2 The JFNK Methodology

The fully discretized eigenvalue problem in neutron transport theory can conveniently be written in operator notation:

$$L\psi = MSD\psi + \frac{1}{k}MFD\psi, \quad (5)$$

where the operators L, M, S, F and D are the streaming plus removal, the moments-to-discrete, scattering, fission and discrete-to-moment operators, respectively. A precise definition of these operators is given in Ref. [8], so we will just heuristically describe the actions of various operators on the angular flux vector ψ .

1. The vector of scalar flux moments can be obtained from the vector of angular fluxes by

$$\phi = D\psi.$$

2. Further, the inversion of the operator L is what is usually referred to as a transport sweep. For a fixed source problem without scattering, i.e. $L\psi = q$, the solution can be found by sweeping the mesh only once which can be denoted by $\psi = L^{-1}q$.
3. The scattering and fission sources can be constructed by:

$$Q_s = MSD\psi$$

$$Q_f = MFD\psi.$$

4. The eigenvalue problem Eq. (5) is often recast into the more convenient form:

$$(I - DL^{-1}MS)\phi = \frac{1}{k}DL^{-1}MF\phi,$$

which does not contain angular fluxes any more. By inspection we can identify this as a generalized eigenvalue problem of the form

$$A\phi = \frac{1}{k}B\phi,$$

with

$$A = (I - DL^{-1}MS)$$

$$B = DL^{-1}MF.$$

5. The action of the inverse of the operator A is equivalent to a single outer iteration, which comprises a sweep through energy groups from highest to lowest energy and potentially thermal iterations which are usually performed using Gauss-Seidel fixed point iterations. At each group a block matrix has to be inverted, which is usually done using Richardson iterations commonly referred to as inner iterations.
6. Inner iterations in group g are represented by the action of the operator

$A_g^{-1} = (I - DL^{-1}MS_{gg})^{-1}$ on the sum of the fission and scattering from other groups into group g .

Given a vector of unknowns u comprised of all the scalar flux moments (spatial as well as angular moments) and the multiplication factor k

$$u = (\vec{\phi}, k)^T$$

the JFNK method is based on writing a nonlinear system of equations

$$F(u) = 0$$

whose roots are a solution of the k -eigenvalue problem. Thus, we recast the eigenvalue problem as a nonlinear system of equations, which can be solved using Newton's method.

The standard Newton's method uses the following iteration strategy:

$$\begin{aligned} J(u_n)\delta u_n &= -F(u_n) \\ u_{n+1} &= u_n + \delta u_n \\ \text{until } F(u_n) &\leq \text{tol.} \end{aligned} \tag{6}$$

However, we do not want to construct the Jacobian due to its large size and complexity and hence resort to an approximation of its action on the arbitrary vector v as follows:

$$J(u)v = \frac{F(u+\varepsilon v) - F(u)}{\varepsilon} + O(\varepsilon).$$

Within a single Newton iteration a linear system of equations has to be solved (see first line in Eq. (6)), but the Jacobian $J(u_n)$ is unknown except for his action on an arbitrary vector. Therefore, we use GMRES to solve the above linear system since the only information required is $J(u_n)\delta u_n$. Within the framework of the JFNK method an important question is how tightly to converge the GMRES iterations. One possible stopping criterion for the GMRES iterations is:

$$\|F(u_n) + J(u_n)\delta u_n\| \leq \eta \|F(u_n)\|.$$

The parameter η is called forcing factor and has a strong impact on the efficiency of the JFNK method. If the forcing factor is chosen too restrictive, meaning that the GMRES iterations are converged too tightly, then the cost in both execution time and memory consumption of the JFNK method increases dramatically due to the inherent features of GMRES. It has been found that an optimal choice of the forcing factor is problem dependent and hence hard to make *a priori*. However, it was also found that a forcing factor of 0.01 works sufficiently well for a variety of test problems. In the presented implementation $\eta = 0.01$ is adopted throughout.

We will now introduce three different methods for formulating the nonlinear problem $F(u) = 0$, which have been implemented in THOR. In the remainder of this section these 3 formulations are referred to as JFNK- n , where $n=1, 2$ or 3 . For all three formulations, it is advantageous to consider $F(u)$ to be a vector-valued nonlinear function depending on the vector u . Within the Newton iteration all we have to do is to evaluate this nonlinear function. The various methods then differ in the way $F(u)$ is evaluated.

1. JFNK-1: $F(u)$ is given by

$$F\left(u = \begin{pmatrix} \phi \\ k \end{pmatrix}\right) = \begin{pmatrix} \phi - A_L^{-1} \left(\frac{1}{k} B + A_U\right) \phi \\ k - k \frac{\|FP\phi\|}{\|F\phi\|} \end{pmatrix},$$

where the operators A_L and A_U are the lower triangular and strictly upper triangular parts of the operator A . The application of $A_L^{-1} \left(\frac{1}{k} B + A_U\right)$ on the scalar flux vector is equivalent to a single outer iteration without thermal iterations (no iterations on the upscattering).

In algorithmic form:

- a) Start from group $g=1$
- b) Do inner iterations in group g
- c) Increment group $g=g+1$
- d) If $g = G + 1$, go to e), else go to b)
- e) End

The last component of the nonlinear function F is computed using a closure relation which is somewhat arbitrary. With the definition of $P = A_L^{-1} \left(\frac{1}{k} B + A_U\right)$, we see that this relation is a standard update formula used commonly in power iterations in its residual form. For the methods 2 and 3, P is going to be whatever operator the scalar flux vector ϕ is subjected to. This method is also referred to as “single outer iteration with lagged upscattering”.

2. JFNK-2: $F(u)$ is given by

$$F\left(u = \begin{pmatrix} \phi \\ k \end{pmatrix}\right) = \begin{pmatrix} \phi - DL^{-1}M \left(\frac{1}{k} F + S\right) \phi \\ k - k \frac{\|FP\phi\|}{\|F\phi\|} \end{pmatrix}.$$

The operator $\left(\frac{1}{k} F + S\right)$ computes a lumped, fixed source comprising scattering and fission, which is then subjected to a transport sweep. In algorithmic form:

- a) Compute sources in groups $g=1, \dots, G$
- b) Go to group $g=1$
- c) Perform mesh sweep group g
- d) Increment $g = g + 1$
- e) If $g = G + 1$, go to f), else go to c)
- f) End

Note, that the evaluation of F does not require any iterative procedure, it only comprises building a fixed source and sweeping once on it. Hence, this method is also referred to as the “flat iteration scheme”.

3. JFNK-3: $F(u)$ is given by

$$F\left(u = \begin{pmatrix} \phi \\ k \end{pmatrix}\right) = \begin{pmatrix} \phi - (I - DL^{-1}MS_L)DL^{-1}M\left(\frac{1}{k}F + S_U\right)\phi \\ k - k \frac{\|FP\phi\|}{\|F\phi\|} \end{pmatrix},$$

where S_L and S_U are the strictly lower triangular and upper triangular parts of S , respectively. Although the operator looks more complicated, its action on the scalar fluxes is not difficult to implement as becomes obvious through the algorithmic form:

- a) Compute fission, self-scattering and up-scattering sources in groups $g = 1, \dots, G$
- b) Go to group $g=1$
- c) Compute down-scattering into group g from groups $g' > g$. Perform mesh sweep on group g
- d) Increment $g = g + 1$
- e) If $g = G + 1$, go to f), else go to c)
- f) End

Hence method 3 is very similar to method 2 with the difference that the most up-to-date information is used for the calculation of the downscattering source. JFNK-3 might also be referred to as “Flat iteration scheme with updated downscattering”.

So far, neither the Newton iteration nor GMRES are “optimized”. In terms of the Newton iteration the update-formula in Eq. (6) could be optimized using under-relaxation or over-relaxation possibly in conjunction with a line-search algorithm. The GMRES iterations are to date not preconditioned, but it is well known that GMRES realizes its full potential only when it is properly preconditioned.

3.4.3 Numerical Results

The performance of the JFNK method is compared to the unaccelerated PI-SI iterations. Since Newton iterations generally use a different stopping criterion than PI iterations, we adjusted the Newton iteration to be consistent with the PI. The following stopping criteria are used:

$$\frac{|k^{(p+1)} - k^{(p)}|}{|k^{(p+1)}|} \leq \varepsilon_k$$

$$\max_i \frac{|(F\phi^{(p+1)})_{i,l=0} - (F\phi^{(p)})_{i,l=0}|}{|(F\phi^{(p+1)})_{i,g,l=0}|} \leq \varepsilon_F$$

where p is the PI or Newton iteration index, i is an index that runs over all spatial cells and l is the spatial moment index with $l = 0$ corresponding to the flux cell-average. For all test cases we employ $\varepsilon_k = 10^{-7}$, $\varepsilon_F = 10^{-6}$. Further, for PI and JFNK-1, the inner iteration stopping criterion is set to $\varepsilon_{SI} = 10^{-10}$ with a maximum of 4 executed source iterations. Hence, source iterations are, at least in the early stages, never run to convergence.

The Godiva test problem and the Takeda-4 test problem (configuration: rods fully inserted) are used to demonstrate the performance of JFNK with respect to PI. Since the employed mesh for the Takeda-4 problem exhibits several cells with optical thicknesses greater than one, THOR-CC could not be used for this test problem. Table 16 shows a summary of all the performed calculations and Table 17 shows the number of mesh cells for the meshes referenced in Table 16.

Table 16: Summary of the performed benchmark calculations.

Test Case	Godiva									Takeda-4					
Quadrature	S4									S8			S4		
Mesh	1			2			3			1	2	3	1		
Spatial Order	0	-1	-2	0	-1	-2	0	-1	-2	0	0	0	0	-1	-2
C0	x			x			x			x	x	x	x		
CC	x	x	x	x	x	x	x	x	x						
CCE													x	x	

Table 17: Number of mesh cells of the employed meshes.

	Mesh 1	Mesh 2	Mesh 3
Godiva	274	2945	20055
Takeda-4	15733		

Calculations that are marked with an x in Table 16 were executed for all JFNK implementations and PI. For each calculation, the total execution time and the total number of mesh sweeps is tracked. Both are good indicators for the performance of an iteration scheme, which is obvious for the execution time. The sweeps are a good indicator for the total amount of work required for obtaining the results, as most of the execution time is spent within the sweeps or daughter subroutines called within the mesh sweep; additionally computations comprising the mesh sweep are identical for all methods under consideration.

Results for the test cases are presented in Table 18 through Table 21. Note further, that for each result it was checked that the final eigenvalue was consistent with the PI eigenvalue. Maximum differences in the eigenvalues were all smaller than the stopping criterion and hence we can conclude that all the methods converged to the same eigenvalue.

The results draw a clear picture. First of all, the influence of the quadrature, the level of mesh-refinement (i.e. the number of mesh cells) and the spatial order on the number of sweeps is noticeable but rather small compared to the difference in number of sweeps across the methods. Hence, we need not discuss the outcome of the test cases separately since the same

conclusions will hold for all of them. However, we need to distinguish between the two test cases Godiva and Takeda.

For the Godiva test case, JFNK-2 and JFNK-3 do better than PI, the number of sweeps for PI is at least twice as large as for JFNK-2 and JFNK-3. Among JFNK-2 and JFNK-3 no clear winner emerges. Unfortunately, however, the JFNK-1 performs slightly worse than PI. We ran additional tests using a tighter stopping criterion and found that eventually JFNK-1 becomes superior to PI. However, we believe that these stopping criteria are tighter than usually chosen in practice and hence the result stands that JFNK-1 performs worse than PI.

Table 18: Results for the THOR-C0 calculations for the Godiva test case.

		S4 Quadrature			S8 Quadrature		
		Mesh 1	Mesh 2	Mesh 3	Mesh 1	Mesh 2	Mesh 3
JFNK-1	# Sweeps	597	498	541	597	498	543
	Ex. Time [s]	12	149	2207	43	572	6522
JFNK-2	# Sweeps	210	168	192	210	174	192
	Ex. Time [s]	5	55	834	18	208	2824
JFNK-3	# Sweeps	204	198	186	204	198	186
	Ex. Time [s]	4	66	821	16	227	2814
PI	# Sweeps	456	456	456	456	456	456
	Ex. Time [s]	10	154	1936	35	538	5552

Table 19: Results for the THOR-C0 calculations for the Godiva test case.

		S4 Quadrature	S8 Quadrature
JFNK-1	# Sweeps	848	848
	Ex. Time [s]	2389	7776
JFNK-2	# Sweeps	456	372
	Ex. Time [s]	1428	3839
JFNK-3	# Sweeps	360	360
	Ex. Time [s]	1070	3805
PI	# Sweeps	2624	2624
	Ex. Time [s]	7197	22375

Table 20: Results for the THOR-CC calculations for the Godiva test case.

		Lambda =0			Lambda=-1			Lambda=-2		
	Mesh	1	2	3	1	2	3	1	2	3
JFNK-1	# Sweeps	597	497	524	519	501	547	524	501	543
	Ex Time, s	83	823	7062	313	3396	26241	4023	38172	281991

JFNK-2	# Sweeps	210	168	192	222	210	234	210	210	228
	Ex Time, s	29	276	2723	136	1415	12199	1459	16145	119973
JFNK-3	# Sweeps	204	198	222	210	168	222	210	168	222
	Ex Time, s	28	325	3128	130	1125	10699	1466	12869	116558
PI	# Sweeps	456	456	456	480	456	456	480	456	456
	Ex Time, s	60	707	5881	301	3177	22833	3769	38492	261988

Table 21: Results for the THOR-CCE calculations for the Takeda-4 test case.

		Lambda=0	Lambda=-1	Lambda=-2
JFNK-1	# Sweeps	848	960	
	Ex. Time [s]	25054	803007	
JFNK-2	# Sweeps	372	412	
	Ex. Time [s]	10991	342750	
JFNK-3	# Sweeps	360	404	
	Ex. Time [s]	10667	337800	
PI	# Sweeps	2032		
	Ex. Time [s]	59792		

For the Takeda-4 test case, all JFNK implementations perform better than PI. Among the JFNK implementations JFNK-3 performs best (slightly better than JFNK-2 and much better than JFNK-1). Compared to PI JFNK-3 reduces the number of sweeps by a factor of 7. However, even JFNK-1 reduces the number of sweeps by about a factor of 3 compared to PI.

3.5 Task E – Verification Phase 3: Comparison to Monte Carlo (NCSU)

In this section efforts for the verification of THOR based on the Advanced Test Reactor (ATR) 94 core internal changeout (CIC-94) benchmark are described. The ATR benchmark was selected because it (1) poses a challenging problem for deterministic S_N transport codes because of its complicated geometry, (2) the CIC-94 configuration is described accurately in [9], and (3) Ref. [9] comprises an MCNP model of the ATR which enables code-to-code comparison of THOR results *versus* MCNP as a verification step before entering the validation phase (Task F).

As a general comment on INL's involvement and contributions to Tasks E and F of this project we quote the following text received by email from our INL collaborators on December 15, 2013:

Support for THOR Validation

In 2012, INL supported THOR validation in modeling of the Advanced Test Reactor (ATR). A significant level of effort has been invested for modern physics methods upgrades for ATR analysis since 2010. As a result of this work, INL has validated a number of computer codes for ATR core analysis. Building on this work, attempts have been made to validate THOR by comparison to experimental measurements and computational results for a clean core configuration.

As with any deterministic transport solver, THOR requires both multigroup cross sections for materials in the core and specification of a computational grid with materials defined with the grid. INL was asked to assist in providing these data. Using a one-dimensional approximation of the ATR, simulating a single core lobe, SCALE was used to generate a set of 238 group cross sections. It was subsequently determined that such a fine group structure, coupled with a three-dimensional (3D) core mesh, would result in excessive computational requirements; the cross section library was collapsed to 10 energy groups weighted by the ATR core spectrum as computed from a two-dimensional (2D) NEWT model of the core.

Generation of a computational grid was done in parallel with cross section work but proved to be more problematic. INL has previously generated 3D meshes using SolidWorks and Attila. Solidworks was used to generate a CAD model for the Attila transport solver. Attila has the ability to generate its own mesh (with user specifications) within the bodies defined in the CAD model. The mesh is saved as a text file in RTT format. INL has previously processed this mesh file to generate 2D meshes for NEWT, but has never attempted to generate a 3D mesh in other formats from the mesh file. Thus, NCSU took the lead in conversion of the Attila mesh file to the format used by THOR. However, in experimenting with the meshed representation of the core, it was determined that the original SolidWorks model contained flaws.

Unfortunately, by this time the original staff member who had generated the SolidWorks model had left INL. INL staff attempted to correct the flaws in the model, but because of the steep learning curve for the software and the lack of familiarity with the model, this proved to be excessively time consuming and beyond what INL could support. Hence, the original SolidWorks models were provided to NCSU so that they could pursue modeling efforts independently.

This ends the quoted text from INL. As stated above NCSU picked up the effort required to complete Tasks E and F to the best possible state given limitation of time and funding and the prohibitive challenges of tessellating the ATR's geometry from scratch.

3.5.1 ATR Description

The ATR configuration is depicted in Figure 28 and **Figure 29**. The ATR comprises 40 fuel assemblies each of which contains 19 fuel plates. The fuel plate geometry is depicted in **Figure 30**. Overall dimensions can be inferred from these figures, but for convenience they are repeated here:

- Total diameter: 200 cm.
- Outer diameter core reflector tank: 137.16 cm.
- Height (overall): 168.275 cm.
- Fuel height: 121.92 cm.

The fuel assemblies are arranged in a serpentine around the flux traps, which are high-flux regions within the reactor for conducting various experiments, for example irradiating material specimens. On the outside of the fuel serpentine are the control drums whose circumference is partially plated with Hafnium and by rotating them neutron absorption can be increased. For the ATR CIC-94 benchmark the control drums are fixed in a fixed position that allows for a

critical reactor core. Outside of the control drum region is the water reflector. In addition to the control drums, several Hafnium control rods are partially inserted into the core.

The geometric detail in the xy -planes, Figure 28, is much richer than the geometric detail along the z -axis, **Figure 29**. However, it must be pointed out that the ATR has a three-dimensional structure to it, i.e. it is not simply an extrusion of a plane, and additionally is of finite height.

Figure 28: An XY view at $x = 0$, $y = 0$ of the ATR benchmark model. Image courtesy Ref. [9].

The inherent difficulty of the ATR model is the intricate geometry, especially the fuel plates depicted in **Figure 30**, in conjunction with large homogeneous regions such as the flux traps and reflector region. The fuel meat has a thickness of about 0.05 cm and the water gap in between fuel plates has a thickness of about 0.2 cm and therefore the mesh in the fuel assembly regions needs to be very fine, i.e. the tetrahedrons used for approximating the geometry need to be very small in size. As a direct consequence many mesh cells are necessary to correctly model the fuel assemblies.

In contrast, the reflector and flux trap regions feature relatively little detail and therefore larger and fewer cells can be used to represent these regions within the spatial mesh. It is the disparity of regions rich in detail and low in detail, and the curved shape of the fuel plates that precludes Cartesian mesh S_N codes to be utilized for modeling the ATR configuration.

Reference [9] describes the ATR configuration in critical, zero power mode. In this configuration the ATR is critical $k_{eff} = 1$ but the power is sufficiently small that it does not change the temperature of the ATR materials. Accurate knowledge of the material temperature is essential to generate appropriate multigroup cross sections and accurate multigroup cross sections are required for an accurate solution of the ATR CIC-94 benchmark problem. In a potential critical, full-power situation a coupled thermo-hydraulics computation would have to be added to determine the temperature everywhere. The zero power configuration, however, allows validation of the neutronics part of reactor simulation independent from the usually coupled multiphysics framework.

Figure 29: An X-Z view at $x = 0$, $y = 0$ of the ATR benchmark model. Image courtesy Ref. [9].

The ATR CIC-94 configuration is a validation exercise for three-dimensional transport codes. Solving it using the diffusion approximation leads to unsatisfactory results due to the inherent anisotropy of the angular flux that is not well captured by the linear flux approximation inherent in the diffusion approximation.

In the framework of this project's verification and validation (V&V) exercises the computed multiplication factor k_{eff} and the assembly power distribution (total normalized to unity) will be compared to MCNP results and experimentally measured values, respectively.

Figure 30: ATR fuel assembly. Image courtesy Ref. [9].

3.5.2 Multigroup Cross Sections Generation

The generation of multigroup cross sections is performed using the *csas-mg* sequence from the SCALE code package. The *csas-mg* sequence is a variant of the *csas-5* sequences for criticality studies with the particular objective to provide self-shielded microscopic multigroup cross section sets. Self-shielding calculations are performed using *Bonami* for the resolved resonances and *Centrm* for the unresolved resonances. Several other modules that perform mixing and cross section utilities are also executed within the employed sequence.

In the particular case of the ATR, cross sections are self-shielded assuming an infinite homogeneous mixture of all constituent nuclides. It is typical for self-shielding calculations, even for LWR lattices, to represent the geometry in a rudimentary fashion, but in contrast to the ATR typical LWR computational methods utilize the periodic nature of the LWR lattices to introduce a notion of the geometric complexity into the self-shielding computation. For the ATR such a wealth of experience does not exist due to the unique design of the ATR. Thus, *csas-5* does not include a tailored self-shielding method encompassed in the *celldata* card tailored to the ATR geometry. Because of this lack of alternatives, the infinite homogeneous option in *csas-mg* is used for self-shielding calculations.

At its current state of development, THOR is a serial code, which imposes limitations on the maximum size of the problem that can be completed within a reasonable execution time. It is understood that three-dimensional S_N computations on tetrahedral grids are infeasible without proper parallelization of the executing code. Therefore, the current state of THOR can only be a stepping stone awaiting further development into a more efficient, parallelized version. Arising from the limitation on the problem size is a limitation on the number of energy groups that can be utilized for completion of the THOR V&V exercise. In this work, *csas-mg* is used to prepare a

seven group library. Clearly, a seven group multigroup library cannot provide high-fidelity answers but given the described limitations, a seven group library seems to be a good compromise between reasonable accuracy and computational cost.

Within the *csas-mg* sequence the native SCALE 238 group library is utilized as the starting point for preparing the seven group cross section library. The spectrum depicted in **Figure 31** is used to collapse this multigroup library into the seven group form before commencing the self-shielding computation. The spectrum was kindly provided by Dr. Mark DeHart of INL. Note that it makes a difference whether self-shielding or group collapsing is performed first, i.e. the two operations do not commute.

Figure 31: Spectrum utilized for collapsing the 238-group SCALE library into the seven group library (blue) utilized in THOR. The red shaded areas indicate the extent of each of the seven groups and the accumulated spectrum across this group.

For determining the necessary scattering expansion order, the *csas-mg* sequence was used to prepare multigroup libraries featuring P_0 , P_1 , and P_3 scattering expansion orders. These cross section libraries are immediately applicable in the SCALE Monte-Carlo code KENO. The MCNP code typically uses internal, continuous energy libraries, which feature virtually no approximation in energy. However, for comparison with KENO the CRSRD code was used to create an MCNP compatible seven group P_0 cross section library.

In **Figure 32** the normalized assembly power distributions obtained with MCNP continuous energy cross sections, and KENO 7-group, 44-group, and 238-group libraries are plotted versus the assembly numbers (left); in addition assembly powers exclusively obtained with the seven group library but featuring various scattering expansion orders are depicted in the right subplot. Several notable conclusions can be drawn from **Figure 32** are:

- Left subplot: KENO results using 7, 44 and 238 groups are all very similar to one another, while MCNP continuous energy results are quite different for a few assemblies. This indicates that the main discrepancy between the prepared multigroup and continuous energy libraries is not the small number of groups but the cross section preparation process (self-shielding) in general.

Figure 32: Left: Comparison of KENO- (multigroup) and MCNP- (continuous energy) computed normalized assembly power distributions for different KENO multigroup libraries with P_5 scattering order. Right: Comparison of the normalized assembly power distributions computed with seven group libraries featuring different orders of scattering expansions. The MCNP results feature P_0 scattering expansions.

- Right subplot: At least linearly anisotropic scattering P_1 must be retained to obtain reasonable results for the ATR benchmark problem.
- Right subplot: Keno and MCNP P_0 (isotropic scattering) results should be all but identical (within statistical variance). However, a significant discrepancy is observed. The reason for this discrepancy is that absorption and scattering in the cross section sets prepared by SCALE add-up to a value larger than the computed total cross section for some isotopes (highest energy group Beryllium for example). As MCNP has issues with this inconsistency a fixup is applied that scales the scattering and absorption cross sections in such instances so that their sum equals the total cross section. The underlying reason for the inconsistency of the SCALE cross section sets is that $(n, 2n)/(n, 3n)$ reactions are added twice/thrice to the scattering matrices, respectively. *The issue of how to obtain consistent cross sections from SCALE is not resolved at this point. It remains an open issue.*

In addition to the normalized power distributions depicted in **Figure 32**, **Table 22** shows differences (in pcm) between the k_{eff} obtained with MCNP continuous energy cross sections and KENO- and MCNP-computed k_{eff} values corresponding to the cases depicted in **Figure 32**. The difference in pcm is defined as:

$$\Delta k_{eff} = |k_{eff} - k_{eff}^{ref}| \cdot 10^5$$

Table 22: Difference (pcm) between MCNP continuous energy and KENO- and MCNP-computed k_{eff} for the ATR with varying multigroup cross section sets.

KENO P_5			KENO 7-group			MCNP, P_0 , 7-group
7-group	44-group	238-group	P_0	P_1	P_3	
348	71	753	709	573	7445	6513

In summary, the created seven group cross section set is sufficiently accurate for our purposes while keeping the computational load at an acceptable level. Clearly, a P_1 scattering expansion is necessary to obtain reasonable results. However, the unresolved issue pertaining to the $(n,2n)$ and $(n,3n)$ reactions and the resulting inconsistency of the cross section set is of concern for the future application of the described cross section set. It should be noted that because of problems with obtaining a suitable tetrahedral mesh for the ATR, the cross section set has not yet been applied beyond scoping calculations whose purpose was to check whether THOR's results are in the same "ballpark" as the KENO/MCNP results.

3.5.3 Computational Grid Generation

Obtaining a suitable tetrahedral mesh for the ATR CIC-94 configuration proved to be the most difficult task encompassed in the V&V exercise. To a minor extent this is attributed to the limitation that THOR is currently a serial code and the mesh cell count therefore needs to be limited to a few million tetrahedra. However, with a geometry as detailed as the ATR, a full three-dimensional model may require in excess of ten million tetrahedra. However, the primary obstacle to satisfactorily completing the V&V exercise is the unavailability of robust mesh generators and the multiple deficiencies discovered in the course of this research in available tetrahedral meshes of the ATR. Even simplified ATR models considered in this work that will be discussed within this section were problematic to be meshed using the software available today. At the time of composing the proposal that resulted in this project there was a functioning Attila mesh that we planned to employ in Tasks E and F but in the process of completing the described work limitations and errors in these meshes emerged and were beyond fixing.

The suggested validation exercise started under the assumption that an accurate tetrahedral mesh of the ATR CIC-94 is available from Idaho National Laboratory (INL). In fact, work at INL modeling the ATR using the tetrahedral S_N code ATTILA has produced a tetrahedral mesh with 2,459,099 tetrahedra. However, the mesh quality of the tetrahedra comprised in this mesh was

found to be very poor leading to several hundred thousand cyclic dependencies among these cells that need to be resolved using an algorithm implemented in THOR's grid-sweep routines. However, the "cycle-breaking" algorithm is designed for resolving a few cyclic dependencies and is thus incapable of dealing with cycles numbering in the hundreds of thousands. As a direct result the runtime of the sweep-order computation increased from about one second per discrete ordinate to twelve hours per discrete ordinate (approx. a factor of 40,000 increase). In addition, results obtained on the said mesh may suffer from the poor quality of the tetrahedral mesh both from deterioration of the characteristic discretization method and potential lack of iterative convergence.

Therefore, we made the request to the scientists at INL to re-mesh the ATR geometry using the mesh generator in ATTILA ensuring proper mesh quality and thus reducing the number of cyclic dependencies within the mesh. The resulting mesh featured 3,035,221 tetrahedrons but the number of cycles is significantly reduced compared to the previous mesh. This mesh was used for extensive scoping calculations to determine if the THOR solutions are of the same order as the results obtained with MCNP. Initially, the obtained multiplication factor was $k_{eff} < 0.9$ (MCNP for the same configuration yielded $k_{eff} = 0.99919$) constituting a difference in excess of 10,000 pcm. Our goal of successful verification is to compute a multiplication factor within 100 pcm for a well-tailored multigroup cross sections set. However, given the small number of employed energy groups 1,000 pcm would have sufficed for a scoping calculation.

Discrepancies in the multiplication factor of the described magnitude are often caused by a discrepant content of control material or fuel. Checking the total volume of the fuel and Hafnium in the mesh described above revealed that about 50% too much Hafnium was present. The reason for the presence of excess Hafnium was a simplification applied when the scientists at INL set up the solid geometry model that is later meshed using the mesh generator tool in ATTILA. The INL solid geometry model is an extruded 2D simplification of the ATR: The xy-layer of the ATR depicted in Figure 28 is extruded along the z-axis to an axial extent of 121.92 cm. Then sufficient reflector is added on the top and bottom to remedy excess leakage from the exposed fuel meat and moderator region at the top and bottom. The reflector is homogeneous water. Having such a quasi-two-dimensional model allows taking advantage of a feature in the ATTILA mesh generator to distribute vertices only on a small number of z levels. This reduces the tetrahedron count significantly but causes the tetrahedra to have skewed aspect ratios. However, the ATR geometry is not simply an extruded two-dimensional geometry, as manifested in partially inserted Hafnium control rods that extend only about 3 inches into the core (from above). In the model obtained from INL these control rods extend all the way through the core in the utilized mesh leading to a large amount of excess Hafnium causing the observed significant reduction in the reactor's criticality.

This problem was fixed by assigning 'water' to every tetrahedron, whose centroid is below 3 inches of the upper edge of the fuel region, and that was (incorrectly) assigned Hafnium as material in the INL model. Note that this fix does not exactly assign water/Hafnium to every location where water/Hafnium belongs because of the "granularity" of the tetrahedra. Thus, a residual discrepancy in the Hafnium control rod volumes persists that is dealt with by using density factors. Density factors are a well-known method to ameliorate discrepancies in material contents between the real configuration and the mesh it is represented by. Typically, these discrepancies arise from the approximation of curved surfaces by straight lines and/or planes, for example when tetrahedral meshes are used. Let V_r be the exact volume of a region, V_r^t the

sum of the volumes of all tetrahedra making up that region, and N_r be the total number density of nuclei in this region. Then the density correction adjusts the number density N_r to N_r' by:

$$N_r' \leftarrow N_r \frac{V_r}{V_r'}$$

The density correction is not only applied to the Hafnium control rod regions, but in general to all regions within the ATR. This is especially important for the curved fuel plates whose volume is not perfectly matched by the equivalent volume of all tetrahedra making up the fuel plates.

With the corrected ATR model, results were obtained using the SC discretization, P_0 scattering expansion, and relative stopping criteria set to 10^{-3} for the convergence of the eigenvalue and 10^{-2} for the convergence of the cell-wise flux between power iterations. Note, that standard stopping criteria are about a factor of ten tighter, but for the purpose of scoping calculation the utilized stopping criteria should suffice.

The obtained results were still disappointing: The computed multiplication factor is $k_{eff} = 0.97685$ constituting a difference of several thousand pcm compared to the KENO/MCNP seven group, isotropic scattering computations. Even though the large magnitude of the eigenvalue difference indicates a deficiency of the utilized cross section and/or mesh, it might still be possible to explain it by the interaction of multiple error sources: cross section set, discretization of angle and space, iterative convergence. However, by examining the computed, normalized assembly power depicted in **Figure 33**, it becomes clear that the mesh is deficient. The obtained power for all assemblies in the fourth lobe (numbers 31-40) is very low. Since the total reactor power is normalized to unity, all the other assemblies' power is increased.

Investigating the low power in the fourth serpentine lobe we found that in fact the mesh was again deficient. A ring around the Northwest flux trap was left unmeshed in the ATTILA meshing procedure leading the mesh file converter from the ATTILA mesh to THOR's to apply vacuum boundary conditions on the interior of all the tetrahedra bordering this ring. The ATTILA to THOR mesh file converter identifies boundary tetrahedra as those that do not have a neighboring tetrahedron, which is unfortunately true for those bordering the erroneously unmeshed ring. The additional leakage originating from the incorrect vacuum boundary conditions yields the reduced assembly powers in the fourth lobe as well as the large subcriticality discussed earlier. It is important to emphasize that this error is not a THOR error and originated in the mesh obtained from INL that we had no way of fixing.

Given the questionable fidelity of the solid geometry model provided by INL for this task, it was decided to create a high-fidelity, truly three-dimensional ATR solid geometry model. Two approaches were followed separately: First, a SolidWorks model for the ATR was developed, and second a MCNP-to-CAD converter developed at the University of Wisconsin, Madison, was utilized to convert the ATR MCNP model to the common `"*.sat"` CAD format. We planned to mesh both SolidWorks and/or the converted MCNP geometry using the general-purpose mesh generator CUBIT.

The creation of a high-fidelity SolidWorks model did not yield any measurable progress because of the complexity of the ATR geometry. Therefore, the remainder of this subsection will discuss progress made on converting the MCNP ATR model to a solid geometry model. The code

mcnp2cad developed at the University of Wisconsin, Madison, is used to convert the MCNP geometry into “*.sat” format which can then be imported into CUBIT.

Figure 33: Normalized assembly power computed by THOR , MCNP and KENO using the seven group P_0 cross sections set.

In a first attempt the full three-dimensional ATR geometry was meshed without regard to potential problems with misaligned surfaces. This initial attempt led to meshing failures in roughly half the solid geometry volumes and successful meshing in the other half. Despite the fact that only half of the volumes were meshed, in excess of ten million tetrahedra were utilized. Discussion with the CUBIT developers revealed that any tetrahedron count smaller than ten million would be unrealistic for the full three-dimensional ATR solid geometry model meshed using the CUBIT mesh generator. This cell-count is beyond the practical limits of THOR at present.

3.5.4 Current and Future Work on the ATR Assembly

Despite the mentioned issues, current meshing efforts are focused on using the converted “*.sat” file to generate a tetrahedral mesh using CUBIT. The current efforts include measures to apply reasonable simplifications to the ATR geometry representation, for example reducing the radial extent of the reflector and eliminating finely detailed objects that are sufficiently far from the core to impact its criticality and power distribution. Initial work in this direction has so far

been successful and progressing well; we plan to report its full results in a paper to submit to PHYSOR 2014, Kyoto, Japan. A brief summary of this effort follows.

As the number of tetrahedrons for meshing the full three-dimensional ATR geometry is anticipated to be in excess of ten million, we decided to simplify the ATR geometry. Completed initial work focused on a quasi two-dimensional model of the ATR that is obtained by extracting a thin slice out of the (axially) central region of the ATR. Reflective boundary conditions are then applied on the top and bottom of this slice mimicking a plane two-dimensional geometry. Due to the small thickness of the slice, the tetrahedron count comprising this mesh is reasonable for THOR's current capabilities and stands at 1,024,875.

The solid body geometry, that the mesh generation will be based on, is obtained using Boolean geometry operations provided by the CUBIT software package. Each volume within the three-dimensional ATR geometry is intersected with a *brick* comprising the full extent of the ATR model in the x-y plane; this *brick* is 0.2 cm thick and is centered at $z=63.1825$ cm. All necessary operations in CUBIT are thereby performed using the convenient python interface that is integrated in CUBIT and allows access to all members and attributes of the solid geometry model and the computational grid once created.

Inspecting the created quasi-2D ATR model, two initial flaws were identified. First, an annulus, according to Ref. [9] filled with water, surrounding a cobalt target was missing. The error probably originates from converting the MCNP geometry to the solid geometry format suitable for CUBIT. Using the python interface, we manually corrected this error by filling the affected volume with water.

Second, for creating a tetrahedral mesh, it is essential to know which material should be assigned to each volume and by implication to each tetrahedron belonging to this volume. After completing the meshing, the tetrahedrons' material tag is used to assign region identification numbers to each tetrahedron and create a mapping from regions to materials. This information is instrumental for enabling the THOR solution to the ATR problem. As the ATR model is created from the MCNP input listed in [9], the material assignment of each MCNP cell needs to be translated to material assignments of each CUBIT volume. This is facilitated by creating CUBIT groups, each of which corresponds to a single MCNP material definition, and then assigning all volumes to their corresponding groups. However, in the conversion process some volumes (about ten in total) are not assigned to a "material" group. Comparison with the MCNP model revealed that all of these 10 volumes feature Beryllium content. Hence we manually implemented the appropriate corrections to the CUBIT model.

For a fair comparison the quasi-2D ATR model cannot be compared to the three-dimensional MCNP results for verification purposes. Even though the ATR geometry is partially extruded in the central regions, it features some Hafnium control rods that are inserted 7.62 cm into the core from the top. In addition, the three-dimensional ATR model features axial leakage, but the two-dimensional slice model is reflected at the top and bottom. Therefore, comparisons are conducted with respect to a two-dimensional MCNP model that can be easily created by adding two planes at $z=63.1825\pm0.1$ cm and restricting each cell to be within these two planes (basically cropping the MCNP model to the desired slice shape).

Solutions for the quasi-2D ATR model are obtained with THOR (step characteristics) using the

described seven group cross section set with P_0 and P_1 scattering expansions. MCNP results are obtained with both seven group P_0 cross sections and continuous cross sections. Statistical uncertainties from MCNP are reduced sufficiently to not influence comparison of the results. In Figure 34 the assembly power distribution is plotted and in Table 23 the obtained k_{eff} for the described scenarios is listed. Comparing THOR(SC), P_0 and the MCNP multigroup, P_0 results on the one hand and the THOR(SC), P_1 and the MCNP continuous energy results, shows decent agreement between THOR and MCNP. The THOR and MCNP computed eigenvalues differ by 700 and 500 pcm, respectively, and the assembly power distribution is agrees reasonably well. The exception is the first lobe (assemblies 1-10) where differences between MCNP and THOR are in excess of 10%. These errors as well as the difference in k_{eff} are driven by the relatively coarse mesh, low quadrature order (S_4 level symmetric) and the relatively coarse cross section library when comparing MCNP continuous energy results and THOR multigroup P_1 results. A model refinement study is under way that attempts to quantify the fraction of the observed errors caused by the spatial and angular discretizations separately.

Figure 34: Comparison of MCNP continuous energy, MCNP-multigroup and THOR-SC P_0 and P_1 computed assembly powers for the quasi two-dimensional ATR model (ATR “slice”).

Future research will focus on simplification of the three-dimensional ATR model. The idea behind this approach is to remove unnecessary detail from the full three-dimensional ATR model. As the THOR validation exercise will only use the computed eigenvalue and the assembly power distribution as responses of interest, the detailed modeling of all the small holes used for experimentation is unnecessary. The plan is to merge small holes to the material body surrounding them as long as these holes do not contain a material that significantly perturbs the

neutron economy, for example highly absorbing materials. In addition, the region outside the core's radial extent outside the control drums' centers may be subject of justifiable significant geometric simplification. The control drums are coated with Hafnium, a strong neutron absorber, only along a partial arc. The drums are rotated so that the Hafnium arc is brought closer or farther from the core to effect higher or smaller absorption, respectively, to maintain the reactor's criticality. It is reasonable to expect that when the drums are "parked" with the Hafnium beyond the center-radius they possess negligible reactivity, otherwise they would be unnecessarily taxing the neutron economy even when not in use. This conjecture is further justified by the fact that the mean free path of thermal neutrons (that are primarily responsible for sustaining the chain reaction) in water is about one-third of a cm while the radius of the drums is ~10 cm. It is highly unlikely that thermal neutrons traveling farther than the drums' center-radius will reflect back and significantly influence the core's critical state or contribute appreciably to the assembly power distribution. Following this argument it becomes evident that modeling the full width of the water reflector, 30 cm or about 90 mean free paths, is an unnecessary overkill. It so happens also that due to its larger radial extent this region occupies a large volume that requires a large number of tetrahedral to tessellate while contributing little to the model's physics fidelity. Therefore a likely simplification to the model that we have started exploring is to truncate the radial extent after a few cm of the water reflector, or even on the outside surface of the control drums. To verify the reasonableness of this approximation we are conducting a sensitivity study with MCNP to determine the change in k_{eff} and the assembly power distribution with changing water reflector thickness.

Code and Cross Section Data	k_{eff}
MCNP continuous energy XS	1.03129
MCNP 7 group, P_0	1.08675
THOR(SC), 7 group, P_0	1.07979
THOR(SC), 7 group, P_1	1.03657

Table 23: Comparison of MCNP continuous energy, MCNP-multigroup and THOR-SC P_0 and P_1 computed k_{eff} for the quasi two-dimensional ATR model (ATR "slice").

The goal of the geometric simplification is to reduce the complexity of the ATR model in an attempt to capture the important physics with a grid that features a reasonable number of tetrahedra. As mentioned before, the meaning of "reasonable tetrahedron count" depends on the considered code (especially its parallelization). For THOR, a grid with no more than five million tetrahedra can be considered reasonable.

3.6 Task F – Validation of THOR

Due to personnel changes on INL's side of the collaboration and due to the delays caused by the inability to obtain a suitable tetrahedral mesh of the ATR CIC-94 configuration this task was not completed. Up to this point no validation of THOR versus the measured ATR CIC-94 power distribution was completed. However, we believe we identified many deficiencies that must be

resolved for any future validation of modeling and simulation tools to succeed not just for THOR. We reiterate that the meshing difficulties we encountered originate in earlier models of the ATR that preceded and were outside the scope of this project. Once these meshing deficiencies are fixed comparison of the computed multiplication factor and assembly powers to the measured results should be rather straightforward. Even though this project has ended at this point we continue to invest personal and NCSU resources on reaching this goal with the objective of submitting a paper on our validation results to PHYSOR 2014, Kyoto, Japan.

One additional important observation that we made in the course of this project regarding validation of the computed ATR assembly power distribution is the fact that the actual experimental measured quantities on ATR CIC-94 comprised a number of fission foils placed at selected points in the core. Conversion of the fission rates at these points to assembly powers is an added step that incurs its own modeling assumptions and errors that will taint the validation exercise. Clearly a straight comparison of the computed and measured fission rate in each individual fission foil provides a stronger validation result. We checked with INL staff who are well versed on the ATR CIC-94 experiments to try to obtain this “raw data” but unfortunately there is no record of it at this point. It appears that once this data was processed and the assembly powers were “inferred” the actual measured data was discarded. This is very unfortunate because it separates the measured and computed values by an intervening “model” and because translation of the experimental uncertainty in the actual measured data (fission foil fission rates) to the compared data (assembly power) is compromised. It will be very valuable for future efforts on this project, but also other validation exercises using the ATR CIC-94 results to search harder for the raw data and make it available to the community.

5. Bibliography

- [1] R.M. Ferrer, “An Arbitrarily High Order Transport Method of the Characteristic Type for Unstructured Tetrahedral Grids,” PhD Dissertation (2010).
- [2] W.A. Rhoades and D.B. Simpson, “The TORT Three- dimensional Discrete Ordinates Neutron/Photon Trans- port Code,” ORNL/TM-13221.
- [3] Kobayashi, K., Sugimura, N., and Nagaya, Y. 3-D Radiation Transport Benchmark Problems and Results for Simple Geometries with Void Regions. Tech. rep., NEA/OECD, 2000.
- [4] Asaoka, T., Asano, N., Nakamura, H., Mizuta, H., Chichiwa, H., Ohnishi, T., Miyasaka, S., Zukeran, A., Tsutsui, T., Fujimura, T., and Katsuragi, S. “Benchmark tests of radiation transport computer codes for reactor core and shield calculations,” *Journal of Nuclear Science and Technology* **15**, 1 (1978), 56–71.
- [5] Sebastian Schunert, Yousry Y. Azmy, “A Two-Dimensional MMS Benchmark Suite for Cartesian SN Transport Methods with Escalating Order of Non-smoothness of the Exact Solution”, in the *Transactions of the American Nuclear Society*, Winter Meeting, Las Vegas, Nevada, November 7-11, 2010.
- [6] Sebastian Schunert, Yousry Y. Azmy, “A Two-Dimensional Method of Manufactured Solutions Benchmark Suite based on Variations of Larsen’s Benchmark with escalating Order of Smoothness of the Exact Solution”, in the proceedings of the *International Conference on Mathematics and Computational Methods applied to Nuclear Science and Engineering (M&C 2011)*, May 8-12, 2011, Rio de Janeiro, Brazil.
- [7] A. Genz and R. Cools, “An Adaptive Cubature Algo-rithm for Simplices,” *ACM Transactions on Mathematical Software* **29**, pp. 297-308 (2003).

- [8] T.M. Evans and A.S. Stafford and R.N. Slaybaugh and K.T. Clarno, "DENOVO: A NEW THREE-DIMENSIONAL PARALLEL DISCRETE ORDINATES CODE IN SCALE," *Nuclear Technology* **171**, Number 2, pp. 171-200 (2010).
- [9] S.S. Kim and B.G. Schnitzler, "Advanced Test Reactor: Serpentine Arrangement of highly-enriched water-moderated Uranium-Aluminide Fuel Plates reflected by Beryllium," NEA/NSC/DOC(95)05/II, Volume II, HEU-MET-THERM-022.

6. List of Publications from the Project

1. Rodolfo M. Ferrer, Yousry Y. Azmy, "A Robust Arbitrarily High Order Transport Method of the Characteristic Type for Unstructured Tetrahedral Grids: Theory and Applications", in the proceedings of the *International Conference on Advances in Mathematics, Computational Methods, and Reactor Physics*, Saratoga Springs, New York, May 3-7, 2009, on CD-ROM, American Nuclear Society, La Grange Park, IL (2009).
2. D. F. Gill, Y. Y. Azmy, "Jacobian-Free Newton-Krylov as an Alternative to Power Iterations for the k-Eigenvalue Transport Problem", *Transactions of the American Nuclear Society* **100**, 291 (2009).
3. R.M. Ferrer, "An Arbitrarily High Order Transport Method of the Characteristic Type for Unstructured Tetrahedral Grids," PhD Dissertation (2010).
4. Sebastian Schunert, Yousry Y. Azmy, "A Two-Dimensional MMS Benchmark Suite for Cartesian SN Transport Methods with Escalating Order of Non-Smoothness of the Exact Solution", *Transactions of the American Nuclear Society* **103**, 359 (2010).
5. Daniel F. Gill, Yousry Y. Azmy, James S. Warsa and Jeffery D. Densmore, "Newton's Method for the Computation of k-Eigenvalues in SN Transport Applications", *Nuclear Science & Engineering* **168**, 37 (2011).
6. Sebastian Schunert & Yousry Y. Azmy, "A Two-Dimensional Method of Manufactured Solutions Benchmark Suite Based on Variations of Larsen's Benchmark with Escalating Order of Smoothness of the Exact Solution", in the proceedings of *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2011)*, Rio de Janeiro, RJ, Brazil, May 8-12, 2011, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2011).
7. Sebastian Schunert, Rodolfo M. Ferrer, Daniel F. Gill, Yousry Y. Azmy, "The JFNK Method for the Solution of the K-Eigenvalue Transport Problem in Tetrahedral Geometry", *Transactions of the American Nuclear Society* **104**, 49 (2011).
8. Rodolfo M. Ferrer & Yousry Y. Azmy, "A Robust Arbitrarily High-Order Transport Method of the Characteristic Type for Unstructured Grids", *Nuclear Science & Engineering* **172**, 33 (2012).
9. Sebastian Schunert, Yousry Azmy, and Rodolfo Ferrer "Verification of the Three-Dimensional Tetrahedral Grid SN Code THOR", in the proceedings of *International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering (M&C 2013)*, Sun Valley, Idaho, USA, May 5-9, 2013, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2013).