

Year End Progress Report on Rattlesnake Improvements

Yaqi Wang, Mark D. DeHart, Frederick N.
Gleicher, Javier Ortensi and Sebastian
Schunert
Idaho National Laboratory

September 2015



The INL is a U.S. Department of Energy National Laboratory
operated by Battelle Energy Alliance

DISCLAIMER

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

Year End Progress Report on Rattlesnake Improvements

**Yaqi Wang, Mark D. DeHart, Frederick N. Gleicher, Javier Ortensi
and Sebastian Schunert
Idaho National Laboratory**

September 2015

**Idaho National Laboratory
Reactor Physics Analysis and Design Dept.
Idaho Falls, Idaho 83415**

<http://www.inl.gov>

**Prepared for the
U.S. Department of Energy
Office of Nuclear Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517**

1. INTRODUCTION.

Rattlesnake is a MOOSE-based radiation transport application developed at INL to support modern multi-physics simulations. Rattlesnake is able to perform both steady state and transient calculations and both fixed source and eigenvalue calculations for the multigroup radiation transport equations. Various spatial discretization schemes are supported, including (1) finite element method (FEM) with discrete ordinates method (SN) and spherical harmonics expansion method (PN) for the self-adjoint angular flux (SAAF) formulation, (2) continuous FEM (CFEM) with SN for the least squares (LS) formulation, (3) diffusion solution with both CFEM and discontinuous FEM (DFEM). A separate toolkit, YAKXS, for multigroup cross section management has been developed to support Rattlesnake calculations with feedback both from changes in field variables, such as fuel temperature, coolant density, etc., and in isotopic inventories. The framework for doing nonlinear diffusion acceleration (NDA) within Rattlesnake has been implemented; NDA calculations have been performed with the SAAF-SN-CFEM scheme and using Monte Carlo with OpenMC [1]. Because it is part of the MOOSE system, Rattlesnake has also been coupled to BISON [29] and RELAP-7 [5] for full-core multi-physics simulations [8].

Rattlesnake development was migrated into the INL-internal GITLAB development environment at the end of CY14. Since then a total of 369 merge requests have been accepted within the Rattlesnake project. This provides a measure of the number of updates performed to the software. However, note that it is noted that because Rattlesnake is built on top of the MOOSE framework, additional improvements made in MOOSE resulted in improved Rattlesnake performance. Some of MOOSE updates were made at the request of the Rattlesnake team, while others were made to improve MOOSE structure and solvers, from which Rattlesnake benefitted.

During FY15, significant improvements have been made in Rattlesnake. This report summarizes some of the key development activities completed this year. Section 2 describes work to improve performance of the nonlinear diffusion acceleration scheme when used to accelerate SAAF-SN-CFEM transport calculations. Section 3 describes the status of implementation of a first order version of the discrete ordinates approximation of the transport equation using DFEM. In Section 4 the development of an improved quasi-static (IQS) approximation for acceleration of transient spatial kinetics calculations is presented. In Section 5, the development of a multi-scale solution capability is discussed. Finally, Section 6 describes improvements completed in least-squared finite element solutions.

2. Optimization of Rattlesnake Calculations using NDA with SAAF-SN-CFEM

Nonlinear diffusion acceleration is an efficient transport iterative scheme for accelerating the solution of the neutron transport equation. NDA is a diffusion model that has equivalent physics to a transport model, and can be used to accelerate more computationally expensive transport calculations by improving estimates of the source term, i.e., the scattering and fission sources that comprise the right hand side of the transport equation. When transport updates are performed using the SAAF-SN-CFEM transport method at each NDA iteration, both scattering and fission source are constructed from the previous NDA iteration and fixed. As a result, angular fluxes of all energy groups and all streaming directions are decoupled and can be updated independently. However, in neutron transport, the number of variables (angular fluxes for each energy group) is the number of energy groups times the number of

streaming directions, which could potentially add up to thousands of variables to solve for each element. This would require more memory and slower convergence to form the streaming-collision matrix for all angular fluxes in one single large matrix and perform the transport update all together. Instead, it is more appropriate to take advantage of the block-diagonal matrix structure, where each block corresponds to an angular flux, and solve all angular fluxes separately. Because all angular fluxes are independent and can be solved simultaneously, the approach can accommodate multiple threads in a natural manner. However, MOOSE interfaces only allow construction of one sub-matrix (corresponding to an individual variable) at a time, requiring sub-matrices to be built one at a time within a full sweep (evaluation of finite element shape functions, variables, material properties on each element in the mesh). This caused a data bottleneck that reduced the solution to a crawl. To alleviate this, a new interface was developed and added to MOOSE that allows construction multiple sub- matrices in one pass through the entire problem domain.

Rattlesnake was then found to begin to fail when the number of degrees of freedom (DoFs) in a calculation exceeded more than approximately 4.3×10^9 ; this many DoFs are not uncommon, especially for multigroup SN transport calculations. However, this exceeds the maximum value that an unsigned four-byte (32 bit) integer can represent (4,294,967,295). Thus it was determined that both Rattlesnake and all of the packages upon which it is built (MOOSE, LibMesh [12], PETSc [2]) needed to be compiled to support 64-bit indexing, allowing index values of up to 1.8×10^{19} . Multiple changes in MOOSE and Rattlesnake were required to allow for this higher magnitude indexing. Identification of all source code requiring this update was a significant undertaking. However, once implemented, larger problems solved with a detailed SAAF-SN-CFEM solve with NDA were able to run within Rattlesnake.

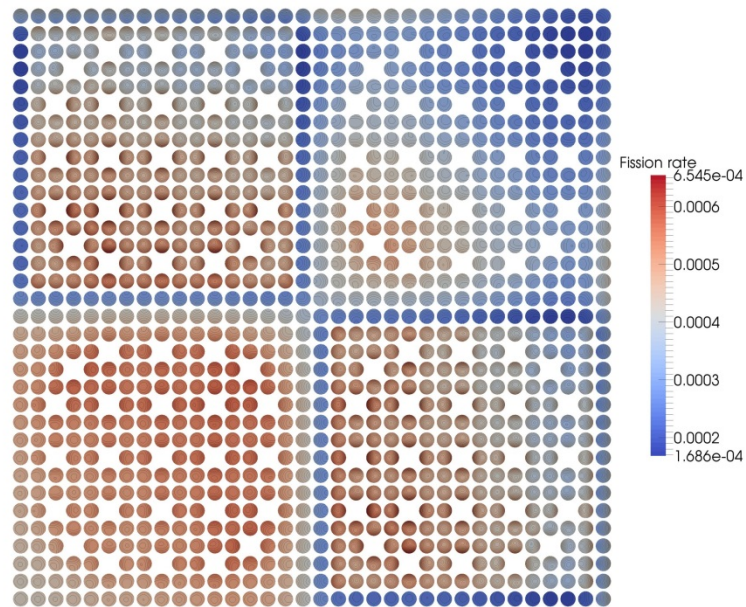


Fig. 1: Fission rate of fuel rods in C5G7 benchmark.

Figure 1 shows the configuration of the two-dimensional C5G7 benchmark problem [13], a 2x2 assembly configuration with a set of two UO₂ fueled and two mixed oxide (MOX) fueled Westinghouse 17x17 assemblies with a water reflector region on top and right sides, solved on a pin-resolved mesh. After the previously described code changes were implemented, the updated version of Rattlesnake was able to perform a convergence study for the C5G7

benchmark problem with SAAF-SN-CFEM; solution of this modest problem had not been possible previously. Convergence results with respect to angular and spatial refinement were reported in a paper published at M&C 2015 [26]. The computed eigenvalue and power distribution converge to the reference solution [16] to about 1 pcm and within 0.001% root mean square error. The most detailed calculation performed (refined mesh and quadrature order) required in excess of 45 billion degrees of freedom and executed in under 20 hours on 1200 CPUs. These results demonstrate the ability of Rattlesnake (and by inheritance, other MOOSE applications) to solve extremely large problems in an efficient manner. Scoping calculations, targeting the discretization error on the pin powers below 0.1%, can now be complete in a few minutes on one single computing node with 24 processors. Figure 1 illustrates the computed fission rate of all fuel rods in and reveals the azimuthal asymmetry of the fission rate in the fuel rods, especially the rods near water holes and reduced loading rods (MOX assembly). It will be necessary to consider such effects for multi-physics calculations for fuel performance analysis.

In the future, this work will attempt to extend to perform additional scalability studies for the multi-grid method on greater numbers of CPUs to further assess the computational efficiency of Rattlesnake. Similar convergence studies for the three-dimensional version of the C5G7 benchmark will also be conducted.

3. Discontinuous FEM with SN for the First-Order Transport Equation

Relative to continuous FEM based discretization schemes, the discontinuous FEM method of the SN solver with the first-order transport equation has advantages and disadvantages due to its sweeping algorithm in space. The drawbacks are the parallel scalability on supercomputers and difficulty in handling higher order isoparametric meshes. On the other hand, the sweeping algorithm is matrix-free, which is preferred for extremely large simulations. Because the transport solves happen at the element level, the sweeping algorithm does not have any specific requirements on the shape of the elements. It works reasonably with elements having poor aspect ratios. In addition, because the sweeping algorithm is based on the first-order transport equation, it does not require special void treatments, as is required for approaches based on the second-order form of the transport equation. Thus there is some value in addition of the DFEM-SN scheme to Rattlesnake. It is desirable to retain the matrix-free residual evaluation with MOOSE, which supports massively parallelism for unstructured higher order meshes. Implementation requires a dedicated sweeper to perform the transport updates and can be used either as a preconditioner for the direct transport solve or for NDA. Although this sweeper could be (and eventually will be) a general MOOSE-based capability, at this time a native sweeper has been implemented directly into Rattlesnake to test other aspects of the DFEM-SN implementation. This prototypic sweeper algorithm is somewhat crude at this time in the sense that it only supports a Jacobi-style sweep with multiple processors, i.e. the sweep can start from the processor interface without having updated information from the upwind processors. The sweeper currently works on unstructured higher order meshes with variable total cross sections within single elements, which makes its grind time relatively large (20 microseconds per solve) relative to a fully-optimized sweeper.

Once this capability was functional and tested, a flexible NDA implementation for the SN-DFEM scheme was added. This NDA method is able to potentially discretize the associated diffusion system with DFEM on a coarser mesh and with a coarser energy group structure than the SN equations. The only constraint is that the low-order mesh and energy groups must be a subset of the high-order or transport mesh and energy groups; i.e., element and energy boundaries in the coarse mesh must match boundaries of the high order mesh and energy structure, respectively. The benefit of this capability is that it provides an extension of

the stability region to optically thicker mesh cells.

The DFEM-SN method differs from current methods [28,20,7,14] by allowing the SN and diffusion system to be discretized with arbitrary FEM bases. It also differs from the current methods by using an additive update of the SN system: The high-order scalar fluxes Φ_g are updated with the prolongation of the low-order scalar fluxes φ_p with the following equation:

$$\Phi_g \leftarrow -\Phi_g + \Pi^{-1} [\varphi_p - \Pi[\Phi_g]]$$

where Π is the projection operator from the high-order solution space to the low-order space. The particular form of these equations ensures that no spurious information that is outside of the low-order function space spanned by the SN FEM basis is introduced by the prolongation.

This method introduces the notion of equivalence of the SN and diffusion solutions and derives suitable definitions of the cross sections as the following:

- Diffusion coefficient: $D_p = \frac{\Pi[\sum_{g \in p} \Phi_g]}{3 \Pi[(\sum_{g \in p} (\Sigma_{t,g} - \sum_{g'=1}^G \Sigma_{s,1}^{g \rightarrow g'}) \Phi_g)]}$.

(It should be noted here that the definition of the diffusion coefficient does not affect the consistency of the method and it is therefore a free parameter. The described choice is the most straightforward definition, but might not be the best definition in terms of the convergence properties.)

- Removal Cross section: $\sigma_{r,p} = \frac{\Pi[(\sum_{g \in p} (\Sigma_{t,g} - \sum_{s=1}^G \Sigma_s^{g \rightarrow g'}) \Phi_g)]}{\Pi[\sum_{g \in p} \Phi_g]}$.
- Scattering cross section: $\sigma_s^{p' \rightarrow p} = \frac{\Pi[\sum_{g' \in p'} \Phi_{g'} (\sum_{g \in p} \Sigma_s^{g' \rightarrow g})]}{\Pi[\sum_{g' \in p'} \Phi_{g'}]}$.
- Fission cross section: $\sigma_{f,p} = \frac{\Pi[\sum_{g \in p} \nu \Sigma_{f,g} \Phi_g]}{\Pi[\sum_{g \in p} \Phi_g]}$.
- Fission spectrum: $\chi_p = \frac{\Pi[\sum_{g \in p} \chi_g \sum_{g'=1}^G \nu \Sigma_{f,g'} \Phi_{g'}]}{\Pi[\sum_{g'=1}^G \nu \Sigma_{f,g'} \Phi_{g'}]}$.

All cross sections are consistent with standard definitions. The closure used for reproducing scalar fluxes with the low-order equation is though drift terms with drift vectors and element face closures.

The NDA equations are solved using a Picard iteration scheme. Fourier analysis for this iteration scheme with low-order diffusion system having the same mesh and energy group structure were performed [23]. Unconditionally stability of the iteration was observed and further confirmed with numerical results [22]. The implementation was then tested with the C5G7 benchmark problem. For discretizing the SN problem the triangular mesh for a single MOX assembly is shown in Fig. 2 and contained roughly 120,000 elements. In contrast, the diffusion equation was discretized either on the same mesh or on a rectangular mesh coincident with the pin cell boundaries. Furthermore, seven energy groups were used for the SN problem, while the diffusion system employed either the same energy group structure or a two group structure created by collapsing the first four groups into coarse group one and the last three groups into coarse group two. The SN system used linear bases functions, while the

diffusion system used either constant shape functions (equivalent to coarse-mesh finite difference, CMFD), or linear shape functions (NDA). Convergence of the fission source is measured as the absolute L2 difference between subsequent Picard iterations. The fission source error is plotted versus the iteration number in Fig. 3.




Fig. 2: Triangular mesh for a single C5G7 MOX assembly.

CMFD and NDA appear to be stable for all combinations of diffusion meshes and energy group structures. For the NDA coarse group/coarse mesh case, convergence stalls at an iterative error of 10^{-5} . In general, reducing the complexity of the diffusion system reduces the effectiveness of the NDA method. When using a fine mesh and seven energy groups for the diffusion system, the NDA and CMFD methods are stable and effective (8 iterations for convergence to 10^{-6} fission source difference) and virtually no advantage is gained when the diffusion system is discretized using linear shape functions. However, when coarsening the diffusion mesh, the NDA method is significantly faster than CMFD. Using coarse diffusion meshes while retaining a stable and effective acceleration is the target application of the

presented NDA method. Reducing the number of groups in the diffusion system leads to a larger impact on the effectiveness of acceleration than using a coarse mesh. Interestingly, using a coarser diffusion mesh or varying the diffusion FEM basis in addition to using the coarse energy group structure has little effect on convergence of the NDA method with the exception of the convergence stall observed for NDA coarse mesh and two energy groups.




Fig. 3: Fission source error plotted versus iteration number for C5G7. The diffusion system's discretization is varied using fine/coarse meshes (fm, cm) and fine/coarse energy group structures (G=7, G=2) with various combinations thereof.

The results indicate that the presented method is stable for all combinations of diffusion shape functions, meshes, and energy group structures. In the future, testing of this implementation is planned using more practical calculations, such as analysis of the Transient Test Reactor (TREAT) and perhaps even the Advanced Test Reactor (ATR). Development is planned for a rigorous sweeper that will allow sweeping through disconnected domains or incorporate a third-party sweeper for massively parallelization.

4. Improved Quasi-Static (IQS) Kinetics Capabilities

The improved quasi-static (IQS) method is a transient spatial kinetics approach that factors

flux into space- and time-dependent components. These components reflect the power and shape of the flux. The magnitude of the power is time-dependent, while the shape is both space- and time-dependent. However, the impetus of the method is the assumption that the shape is only weakly dependent on time; therefore, the shape may not require computation at every time step, invoking the quasi-static nature of reactor kinetics. Between flux shape updates, the following PKE (point kinetics equation) is solved for the power p and delayed neutron precursors ξ_i , $i = 1, \dots, I$:

$$\begin{aligned}\frac{dp}{dt} &= \left[\frac{\rho - \bar{\beta}}{\Lambda} \right] p + \sum_{i=1}^I \lambda_i \xi_i, \\ \frac{d\xi_i}{dt} &= \frac{\bar{\beta}_i}{\Lambda} p - \bar{\lambda}_i \xi_i, \quad 1 \leq i \leq I,\end{aligned}$$

where the effective reactivity ρ , delay-neutron fraction $\bar{\beta}_i$ and the delayed-neutron precursor decay constant $\bar{\lambda}_i$ are defined as:

$$\begin{aligned}\frac{\rho}{\Lambda}(t) &= \frac{(\phi^*, P[\varphi] - L[\varphi])}{(\phi^*, \frac{1}{v}\varphi)}, \\ \frac{\bar{\beta}}{\Lambda} &= \sum_{i=1}^I \frac{\bar{\beta}_i}{\Lambda} = \frac{1}{k_{eff}} \sum_{i=1}^I \frac{(\phi^*, \bar{\beta}_i P[\varphi])}{(\phi^*, \frac{1}{v}\varphi)}, \\ \lambda_i &= \frac{(\phi^*, \bar{\lambda}_i C_i)}{(\phi^*, C_i)},\end{aligned}$$

where P and L are the production and loss operators, φ is the shape and ϕ^* is the adjoint flux. The equation is not specified with a rigorous FEM notation because the boundary conditions necessary for the loss operator L are not expressed here. The weak form for the shape function is almost the same as the original weak form but with an extra term:

$$\left(\phi^*, \frac{1}{v} \frac{dp}{dt} \varphi \right).$$

Additionally, to improve consistency and accuracy, each macro (shape) time step can be iterated so the best shape is used to compute power at the micro (power) time steps. This iteration process must converge the shape such that the uniqueness condition:

$$\left(\frac{d}{dt} \left(\phi^*, \frac{1}{v} \varphi \right) \right)$$

is preserved. Currently, IQS has been implemented for the multigroup diffusion equations. Preliminary results verified the correctness of the implementation by comparing with a separately developed MATLAB code. The implementation of IQS put pressure on many features of MOOSE that revealed potential for improvement. The pressure on the *save-in* feature in MOOSE propagated its application to boundary conditions and initial solves. MOOSE also updated its ability to restart dense vector data and to set the *MooseApp* executioner immediately after the executioner is created. The framework of Rattlesnake enabled the fast rapid (one month) development of this preliminary capability. The next step will be testing this implementation with more benchmark problems and ultimately with TREAT models. If full transport transient

simulations are deemed necessary, implementation IQS for SAAF-SN-CFEM with other transport schemes will be given priority. Such a capability is likely to be adopted eventually – the priority will depend on needs.

5. Multi-Scale Transport

Rigorous neutron transport calculations for simulation of neutron behavior in a reactor core at a fine resolution in seven-dimensional phase space (one in time, three in space, two in angle and one in energy) will remain a challenge in a foreseeable future. However, uniform fine resolution over a large domain of interest is rarely required or necessary. Typically one can apply higher resolutions in regions of interest; for example, an experiment bundle loaded within TREAT, with a lower resolution treatment and various homogenization levels representing the core outside of the experiment. Although the solution accuracy of the regions of interest can be limited by the accuracy of the lower-resolution regions, especially those close to the interface, it is possible that accuracy can be significantly improved by properly taking into account the environmental conditions imposed by lower-resolution regions. This approach is defined here as a multi-scale transport solution. Discrete ordinates methods (SN) are suitable for treatment of heterogeneous problems due to the decoupling feature of streaming operators, while spherical harmonics expansion methods (PN) typically generate more accurate results with the same number of unknowns for problems with significant homogenization. It is believed that regions of different levels of homogenization can be treated most efficiently with hybrid SN-PN calculations. Such hybrid calculations have been studied in the past [18, 3, 9,10,1,15]. References [18,3,1,15] describe one-dimensional hybrid transport methods, while Refs. [9,10] detail the coupling between Variational Nodal Methods (VNM) with PN and the SN nodal method (SNM), both of which are locally conservative. Current research seeks to develop a multi-scale transport capability to leverage all of the available transport schemes within Rattlesnake.

Implementation of multi-scale transport capabilities in an existing code, even a developmental code, is a non-trivial task. Although the interface conditions for coupling SN and PN are unique, they contain common spatial coupling components, which can be useful to other non-transport physics in the MOOSE herd and drawn from other FEM implementations outside of INL. For example, Mortar FEM (MFEM) for coupling two regions or subdomains has been investigated and applied intensively outside of the transport community [4,6,30,11] and is being implemented within the MOOSE framework at present. A one-dimensional MFEM method was adapted within MOOSE and tested for feasibility; a three-dimensional approach has been added, but continues to be tailored for the needs of MOOSE applications. Within Rattlesnake, developers are responsible for development and implementation of appropriate interface conditions, but by design of the MOOSE concept must work closely with framework developers to address special needs or capabilities.

Earlier work [1] has studied adaptive domain partitioning based on a posterior estimation techniques, where partitioning describes different transport scales and levels of discretization. In this work, the regions of interest will be known a priori by the physics of the problem being evaluated, and the problem domain would be partitioned into subdomains by the user before the calculation, with discretization schemes applied on each subdomain. Spatial interface conditions are implemented via the general mortar FEM framework available in MOOSE. Leveraging the flexibility of the mortar FEM capability, meshes on different subdomains may conform at the mortar interface but they are not required to do so. A custom MOOSE mesh modifier is used to split conforming meshes

within the code and interface conditions are added automatically during a simulation. From a users' perspective, the only change from a single-scale approach is that a separate transport method has to be specified for each subdomain. In this work CPU times are not reported because more efficient solving techniques handling the so-called saddle point problem resulting from multi-scale coupling are under development. The work reported here considers only multi-group approximation in energy variable. Embedded energy structures are assumed across two neighboring scales, i.e. any group in these two scales is either a subset or completely outside of another group. Although users are allowed to amend the interface conditions with the properties defined on interfaces, such as discontinuity factors from the generalized homogenization theory [24], this paper will only talk about the nominal conditions.

One of the key elements required by multi-scale transport calculations is the SN-PN interface condition, in particular, is the coupling between SAAF-SN-CFEM and SAAF-PN-CFEM from the viewpoint of application developers. The following condition has been developed on an SN-PN domain interface $\Gamma = \partial\text{DPN} \cap \partial\text{DSN}$:

$$\int_{\Gamma} (([\Psi]), \Lambda^*)_{\Gamma} + ([\Psi^*], \Lambda)_{\Gamma} d\bar{\Omega}$$

The same angular space for the SN side is used for the Lagrange multiplier, i.e. $\{\Lambda_m, \Lambda_m^*, m = 1, \dots, M\}$. On the SN side,

$$\sum_{m=1}^M w_m ((\Psi_m, \Lambda_m^*)_{\Gamma} + (\Psi_m^*, \Lambda_m)_{\Gamma}).$$

On the PN side,

$$- \sum_{m=1}^M w_m (\vec{R}^T(\bar{\Omega}_m) \vec{\Phi}, \Lambda_m^*)_{\Gamma} - (\vec{\Phi}, \vec{P} \bar{\Lambda})_{\Gamma},$$

where

$$\bar{\Lambda} \equiv \sum_{m=1}^M w_m \vec{Y}(\bar{\Omega}_m) \Lambda_m.$$

It is noted that if the angular quadrature n integrate spherical harmonics up to order $NSn > N$, the first term in Eq. (10) and Eq. (11) will ensure the continuity of the angular flux moments. It also will make all the flux moments evaluated on the SN side with order being greater than N and less equal to NSn to zero. Unfortunately, the above interface condition renders non-smooth PN solutions close to the interface based on numerical observations.

It is noted that the angular function space on the SN side is not clear defined so far. Different interpolation schemes can be applied to construct the space, or to construct the solution over the sphere, with M discrete angular fluxes. One way of achieving this is through Lagrange interpolation. We let the constructed flux in a space spanned by the spherical harmonics with the dimension being the number of directions M . We will denote the set of these basis spherical harmonics as \vec{Y} . It is noted that the spherical harmonics on

the P_N side will always be a subspace of the above space. Then we can select some spherical harmonics in $L > N$ space to make the dimension equal to M . Ref. [17] suggested that the following space odd and $0 < m < N^-$ with $l = N^-$. for two-dimensional calculations if level-symmetric quadrature is used. Here N^- is the level-symmetric SN order.

We then denote the sphere nodal function associated with direction m as

We expand the nodal functions as

and form the following linear equation from the conditions,

(15)

where

If we define we have

The constructed angular flux

(17)

From the constructed angular flux, we can evaluate the angular moments

(18)

Typically the angular flux moments evaluated in this is different from the one directly evaluated from the angular quadrature. We want to use

where A^m is the m -th row of the matrix A . We will use Eq. (19) for the interface condition in Eq. (11). This procedure fixes the non-smoothness connection along the PN -PN interfaces.


it is worthwhile to mention the the interface conditions for PN -PN with different orders can be implemented by making the Lagrange multiplier angular space the same as the angular space with the higher PN order. The Lagrange multiplier essentially make the all angular moments of the lower PN order continuous across the interface and the higher angular moments on the higher PN order side zero. SAAF-CFEM-P0 is equivalent with the diffusion approximation if the boundary condition is imposed in a way of decoupling the even and odd parities. Both SAAF-CFEM-SN and SAAF-CFEM-PN interface condition with diffusion continuous FEM are tested and working.

We tested the implemented interface condition with a one-group homogeneous fixed-source problem. The domain is $8\text{cm} \times 8\text{cm}$ discretized by a 8-by-8 uniform Cartesian grid. Total and scattering cross sections are uniform, 1cm^{-1} and 0.9cm^{-1} respectively, and the external source is isotropic and uniform with $1\text{cm}^{-3}\text{s}^{-1}$.

Left, right and top boundaries are reflecting and the bottom boundary is vacuum. The problem is illustrated in Fig.4. It should be pointed out that the exact scalar flux is constant along lines with $y = \text{const}$. We divide

Reflecting

(8,8)



(0,8)

Vacuum
(8,0)

Figure 4: A one-group homogeneous problem.

the domain into two subdomains by the line starting from (0, 6)cm to (6, 6)cm and apply
SAAF-CFEM-SN

at the lower part and SAAF-CFEM-PN at the upper part. Level-symmetric angular quadratures of orders ranging from 8 to 18 are used for the angular discretization within the SAAF-CFEM-SN subdomain. Within the SAAF-CFEM-PN subdomain, PN orders ranging from 0 to 6 are used. For the solution of the equation system resulting from the FEM discretization, we exclusively employ the PJFNK (preconditioned Jacobian- free Newton-Krylov) solver in this study. The preconditioning matrix comprises contributions from the streaming, collision, vacuum boundary condition, and the interface terms. We gathered the relative errors of integration of scalar flux in the lower part and upper part with the corresponding SN solution over the entire domain into Table1and Table2. The last row of these two tables show the the relative error between the SN solution and the reference solution 379.046cm/s and 156.31848cm/s generated with S24.

Table 1: Relative error of the lower integrated scalar flux.

PN	SN order						
	order	8	10	12	14	16	18
	0	1.93E-04	1.93E-04	1.93E-04	1.93E-04	1.93E-04	1.93E-04
	1	2.07E-04	2.07E-04	2.07E-04	2.07E-04	2.07E-04	2.07E-04
	2	3.68E-06	3.62E-06	3.60E-06	3.59E-06	3.58E-06	3.58E-06
	3	-1.71E-06	-1.77E-06	-1.80E-06	-1.81E-06	-1.81E-06	-1.82E-06
	4	4.59E-09	1.67E-08	1.78E-08	1.68E-08	1.76E-08	1.74E-08
	5	7.82E-09	1.82E-08	1.97E-08	1.86E-08	2.00E-08	1.98E-08
	6	-5.63E-09	1.55E-09	8.04E-10	-1.09E-09	-2.22E-10	-8.34E-10
		-8.68E-04	-5.78E-04	-3.83E-04	-2.70E-04	-1.80E-04	-1.21E-04

Table 2: Relative error of the upper integrated scalar flux.

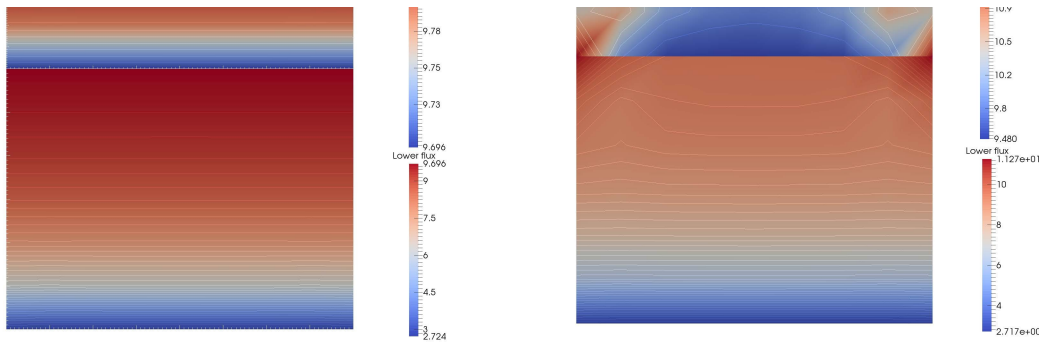
PN	SN order						
	order	8	10	12	14	16	18
	0	-5.30E-04	-5.31E-04	-5.31E-04	-5.31E-04	-5.31E-04	-5.31E-04
	1	-5.80E-04	-5.81E-04	-5.81E-04	-5.81E-04	-5.81E-04	-5.81E-04
	2	-1.01E-05	-1.00E-05	-9.95E-06	-9.93E-06	-9.93E-06	-9.93E-06
	3	3.57E-06	3.72E-06	3.79E-06	3.81E-06	3.82E-06	3.82E-06
	4	-2.88E-08	-5.93E-08	-6.15E-08	-5.89E-08	-6.11E-08	-6.04E-08
	5	-2.10E-08	-4.67E-08	-4.93E-08	-4.64E-08	-4.99E-08	-4.92E-08
	6	1.35E-08	-4.84E-09	-1.79E-09	3.28E-09	1.09E-09	2.88E-09
		-4.55E-05	-3.07E-05	-2.05E-05	-1.46E-05	-9.77E-06	-6.57E-06

We can see that the integrated flux at the lower part has larger angular discretization error than the one at the upper part because of the stronger transport effect close to the

vacuum boundary. With a fixed SN order, errors drop with the increased PN order. For all SN orders, the errors with different PN orders are very close except the PN order 6 or PN order larger than 4 with S8, possibly due to the contamination with the iterative error. We see significant error drop for N to $N + 1$ when N is odd, which can be explained by the weak coupling from odd parity to even parity at the boundaries [27].


The flux of S16-P4 are plotted in Fig.5a along with its contour lines. We can see the scalar flux are continuous across the mortar interface and the solution is relatively flat along the x direction even though it is not exactly constant. The same solution but with the interface condition in Eq. (11) are plotted in Fig.5b. We can clearly see the non-smoothness of the solution along the interface. The solution is indeed symmetric with respect to the line at (4, 0)cm to (4, 8)cm, which is not the case indicated by the color rendering during visualization.

We then rotated the mesh with 45° to test if the rotation can severely impact the validity of the implemented interface condition. 45° makes all the reflecting directions on the left, right and top boundaries are still in the quadrature. The solution is indeed affected by the rotation, for example, the integrated



(a) With Eq. (19).

(b) With Eq. (11).



(c) With 45 degree rotation of the mesh.

Figure 5: Flux with S16 and P4.

fluxes at the lower part and upper part are 379.245510cm/s and 156.323143cm/s from those un-rotated 378.977864cm/s and 156.316943cm/s. The flux are plotted in Fig.5c. Color bars show the the same limits of the fluxes as in Fig.5a. We can see that the interface condition is still working properly, i.e., the solution remains continuous across and smooth along the mortar interface.

We would like to test the hybrid calculations with more sample problems, either multigroup eigenvalue or transient, with or without group collapsing in the future. We will also show the CPU-time results with the fine-tuned mortar FEM framework and solvers of the saddle point problems. Ultimately, we will apply the multiscale transport capability in Rattlesnake for applications to the real reactors like ATR (advanced test reactor), TREAT at INL.

1 Improvements on Least-Square Methods (participants: Hans Hammer, Vincent Laboure, Jim Morel (Texas A& M University), Yaqi Wang)

Hans Hammer took over the least-squares (LS) Sn with continuous FEM (LS-SN-CFEM) in Rattlesnake from Jake Peterson last fall when Jake decided to leave the Ph.D. program at TAMU and take a job. In the process of making detailed comparisons between the LS and SAAF schemes in Rattlesnake for a 2-D k-eigenvalue benchmark, Hans found and fixed several bugs in the LS scheme. Results from the comparison study indicated that the two schemes are quite comparable in terms of both accuracy and cost for the problems considered. Results from an early version of this work were presented at SIAM-CSE15. The latest results were presented at M&C 2015 and a full paper was published in the conference proceedings [21].

A effective algebraic multigrid (AMG) solver is the key of success of LS scheme. Thomas Manteuffel and collaborators at the University of Colorado (UC) and LLNL have developed a multigrid method for solving the source iteration equations associated with the least-squares (LS) Sn equations [25]. Before he left the program, Jake Peterson showed that this method was very effective for the LS source iteration matrices. He did this by extracting the matrices from Rattlesnake for a single direction and reading them into a stand-alone research code written in Python by Jacob Schroder of LLNL. Recent spatial convergence studies made by Hans Hammer for the LS and SAAF Sn packages appear to indicate that the high-resolution calculations may not be properly converged even though the convergence tolerance is being satisfied. An effective multigrid solver should also alleviate this difficulty. Hans confirmed that the algorithm is effective for the simple test problems. He will continue working on some more realistic problem to make sure the algorithm is indeed effective. Then Hans will consider more efficient longer-term integration strategies. One that is particularly attractive is to rewrite the research code in C and implement it in HyPre. This would be a good investment because Tom and his collaborators would be willing to abandon the PYTHON version and continue all development of the multigrid solver in HyPre since they are formally connected with the HyPre team. It might also be possible to rewrite the research code in C and implement it in PETSc. This upcoming work might also provide a better alternative for SAAF schemes.

Fourier analysis is needed for investigating better options on evaluating diffusion coefficients at the void or near-void regions for NDA. Hans Hammer performed a discrete linearized Fourier analysis for each of two NDA schemes for the least-squares Sn equations. The two NDA schemes differ in that one has a consistent non-conservative low-order acceleration equation that converges to the transport solution and the other has an inconsistent conservative low-order equation that converges to the transport solution only in the limit of infinite spatial refinement. Interestingly, the inconsistent conservative scheme is the more accurate of the two for k-eigenvalue calculations. Hans results indicate that both schemes are unconditionally effective with a linearized spectral radius of roughly one-quarter. This means that both schemes are about as effective as standard diffusion-synthetic acceleration, which is an excellent result. Typically the inconsistent conservative scheme is preferred because it has the global conservation and renders better results in k-eigenvalue calculations. Vincent also implemented LS-PN scheme and agreement between LS-SN and LS-PN is observed with fine angular resolution in both. Vincent also confirmed that the stand-alone LS calculations in k-eigenvalue calculations yield unsatisfactory results due to lack of the global conservation.

We have been investigating a plethora of weak boundary conditions for the LS Sn equations. Weak boundary conditions are desirable because Dirichlet corner and edge boundary conditions are ill-defined with vertex flux unknowns whenever a vertex shares multiple faces with different boundary conditions. The boundary functional for our LS equation which was suggested to us by Tom Manteuffel of the University of Colorado, is

$$\Gamma_b = c \int_{\partial D} \vec{\Omega} \cdot \vec{n} (\Psi - \Psi^{inc})^2 ds, \quad (20)$$

where c is an arbitrary constant. Tom indicated that c should not be less than 1. By comparing the boundary term with rest of the functional, we realized that c has units of inverse length. We decided to making c a function of material or mesh parameters. The first thing we tried was $c = \sigma_t$,

which was generally suggested by a thick diffusion-limit analysis. We developed several variations about this form to achieve good behavior with an unresolved layer at the boundary in the thick diffusion limit, but none were entirely satisfactory in

that they yielded the correct behavior to leading order, but were somewhat ill-behaved at first order. Hans Hammer generated some results which clearly indicated that $c = \sigma_t$ was too weak outside the diffusion limit with a highly refined mesh. So we decided to try $c = \max(1/h, \sigma_t)$. This was one of the best performing forms out of a total of 7 investigated. There is no ambiguity about the boundary conditions for the SAAF Sn equations. Yaqi has derived a modified form of the SAAF Sn equations that accommodates voids at the cost of an asymmetric rather than symmetric system. This form has a parameter, τ , in each cell that varies with the cell optical thickness. Vincent recently noticed that for certain simplified types of problems, both the LS and modified SAAF equations yielded the same solutions. He was later able to demonstrate that the two schemes are equivalent for a simplified class of problems if we set the constant in the LS weak boundary conditions as follows: $c = \max(1/(2h), \sigma_t)$. This is only slightly different from the definition given previously and gives some theoretical justification for what was previously an ad hoc definition.

References

1. Dmitriy Y. Anistratov and Nicholas D. Stehle. Computational transport methodology based on de-composition of a problem domain into transport and diffusive subdomains. *Journal of Computational Physics*, 231:8009–8028, 2012.
2. S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. Gropp, D. Karpeyev, D. Kaushik, M. Knepley, L. Curfman McInnes, K. Rupp, B. Smith, S. Zampini, and H. Zhang. *PETSc Users Manual*. Mathematics and Computer Science Division, Argonne National Laboratory, 3.6 edition, June 2015.
3. Ricardo C. Barros, Musa Yavuz, Marcos P. de Abreu, Hermes Alves Filho, and Jos'e Ant'onio M. Mello. Progress in spectral nodal methods applied to discrete ordinates transport problems. *Progress in Nuclear Energy*, 33(1/2):117–154, 1998.
4. Faker Ben Belgacem. The mortar finite element method with lagrange multipliers. *Numerische Mathematik*, 84:173–197, 1999.
5. Ray A. Berry, John W. Peterson, Hongbin Zhang, Richard C. Martineau, Haihua Zhao, Ling Zou, and David Andrs. *RELAP-7 Theory Manual*. INL, December 2014.
6. Dietrich Braess, Wolfgang Dahmen, and Christian Wieners. A multigrid algorithm for the mortar finite element method. *SIAM Journal on Numerical Analysis*, 37(1):48–69, 1999.
7. Y. A. Chao. Coarse mesh finite difference methods and applications. In *Proceeding Physor 2000*, Pittsburg, PA, May 2000. ANS.
8. Derek R. Gaston, Cody J. Permann, John W. Peterson, Andrew E. Slaughter, David Andrs, Yaqi Wang, Michael P. Short, Danielle M. Perez, Michael R. Tonks, Javier Ortensi, Ling Zou, and Richard C. Martineau. Physics-based multiscale coupling for full core nuclear reactor simulation. *Annals of Nuclear Energy*, 84:45–54, October 2015.
9. Enrico Girardi and Jean-Michel Ruggieri. Mixed first- and second-order transport method using domain decomposition techniques for reactor core calculations. In *Proc. International Conference on Supercomputing in Nuclear Application*, Paris, France, 2003.
10. Enrico Girardi, Jean-Michel Ruggieri, Patricia Sireta, and Guillaume Ritter. A new method for the treatment of local strong heterogeneities and its application to the phebus experimental facility. In *PHYSOR 2004 - The Physics of Fuel Cycles and Advanced Nuclear Systems : Global Developments*, Chicago, Illinois, April 25–29 2004. American Nuclear Society, Lagrange Park, IL.
11. Chisup Kim, Raytcho D. Lazarov, Joseph E. Pasciak, and Panayot S. Vassilevski. Multiplier spaces for the mortar finite element method in three dimensions. *SIAM Journal on Numerical Analysis*,

39(2):519– 538, 2001.

12. Benjamin S. Kirk, John W. Peterson, Roy H. Stogner, and Graham F. Carey. libMesh : a C++ library for parallel adaptive mesh refinement/coarsening simulations. *Engineering with Computers*, 22(3):237– 254, December 2006.
13. E. E. Lewis, M. A. Smith, G. Palmiotti, T. A. Taiwo, and N. Tsoul-Fanidis. Benchmark specification for deterministic 2-D/3-D MOX fuel assembly transport calculations without spatial homogenisation (C5G7 MOX). Technical Report NEA/NSC/DOC(2001)4, OECD/NEA Expert Group on 3-D Radiation Transport Benchmarks, 2001.
14. L. Lee, K. Smith, and B. Forget. Techniques for stabilizing coarse-mesh finite difference (cmfd) in methods of characteristics (moc). In *MC2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) method*, Nashville, TN, April 2015. ANS.
15. Sergiy Manolov, Jim E. Morel, and Cristian Rabiti. Hybrid Sn-diffusion and Sn-P3 transport calculations. In *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, pages 2748–2758, Sun Valley, Idaho, May 5–9 2013.
16. Carolyn N. McGraw, Marvin L. Adams, W. Daryl Hawkins, Michael P. Adams, and Timmie Smith. Accuracy of the linear discontinuous galerkin method for reactor analyses with resolved fuel pins. In *PHYSOR 2014 - The Role of Reactor Physics toward a Sustainable Future*, The Westin Miyako, Kyoto, Japan, September 28 - October 3, 2014 2014.
17. W. F. Miller, Jr. and Wm. H. Reed. Ray-effect mitigation methods for two-dimensional neutron transport theory. *Nuclear Science and Engineering*, 62:391–411, 1977.
18. M. M. Nanneh and M. M. R. Williams. A diffusion-transport theory hybrid method for calculating neutron flux distributions in slab lattices. *Atomkernenergie Kerntechnik*, 47:221–224, 1985.
- 1 (19) Paul K. Romano and Nicholas E. Horelik, Bryan R. Herman, Adam G. Nelson, Benoit Forget, and Kord Smith. OpenMC: A state-of-the-art Monte Carlo code for research and development. *Annals of Nuclear Energy*, 82:90–97, 2015. Was 19
- [20]H. Park, D. A. Knoll, and C. K. Newman. Nonlinear acceleration of transport criticality problems. *Nuclear Science and Engineering*, 172:52–65, 2012.
21. Jacob R. Peterson, Hans R. Hammer, Jim E. Morel, Jean C. Ragusa, and Yaqi Wang. Conservative nonlinear diffusion acceleration applied to the unweighted least-squares transport equation in MOOSE. In *Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method*, Nashville, Tennessee, April 19–23 2015. ANS.
22. Sebastian Schunert, Yaqi Wang, Javier Ortensi, Frederick Gleicher, Mark DeHart, and Richard Martineau. A high-order nonlinear diffusion acceleration for the SN equations discretized with the discontinuous FEM II: Fourier analysis. *Transactions of the American Nuclear Society*, 113, 2015. Accepted.
23. Sebastian Schunert, Yaqi Wang, Javier Ortensi, Frederick Gleicher, Mark DeHart, and Richard Martineau. A high-order nonlinear diffusion acceleration for the SN equations discretized with the discontinuous FEM I: Theory and numerical results. *Transactions of the American Nuclear Society*, 113, 2015. Accepted.
24. K. S. Smith. Assembly homogenization techniques for light water reactor analysis. *Progress in Nuclear Energy*, 17,3:303–335, July 1986.
25. H. De Sterck, Thomas A. Manteuffel, Stephen F. McCormick, and Luke Olson. Least-squares finite element methods and algebraic multigrid solvers for linear hyperbolic PDEs. *SIAM Journal on Scientific Computing*, 26(1):31–54, 2005.

26. Yaqi Wang, Mark D. DeHart, Derek R. Gaston, Frederick N. Gleicher, Richard C. Martineau, Javier Ortensi, John W. Peterson, and Sebastian Schunert. Convergence study of Rattlesnake solutions for the two-dimensional C5G7 MOX benchmark. In *Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method*, Nashville, Tennessee, April 1–23 2015.
27. Yaqi Wang and Frederick N. Gleicher. Revisit boundary conditions for the self-adjoint angular flux formulation. In *PHYSOR 2014 - The Role of Reactor Physics toward a Sustainable Future*, The Westin Miyako, Kyoto, Japan, September 28 - October 3 2014.
28. Jeffrey Willert, H. Park, and D.A. Knoll. A comparison of acceleration methods for solving the neutron transport k-eigenvalue problem. *Journal of Computational Physics*, 274:681–694, 2014.
29. R. L. Williamson, J. D. Hales, S. R. Novascone, M. R. Tonks, D. R. Gaston, C. J. Permann, D. Andrs, and R. C. Martineau. Multidimensional multiphysics simulation of nuclear fuel behavior. *Journal of Nuclear Materials*, 423:149–163, 2012.
30. Barbara I. Wohlmuth. A mortar finite element method using dual spaces for the lagrange multiplier. *SIAM Journal on Numerical Analysis*, 38(3):989–1012, 2000.