Chapter 3

Methodology

The objective of this work is to develop a method that, like CADIS and FW-CADIS, automatically generates variance reduction parameters for fixed-source, deep-penetration radiation transport problems. In general, the variance reduction parameters generated by CADIS and FW-CADIS are not sufficient for problems that are strongly anisotropic with respect to the flux. This method will extend existing methods to generate importance maps–that in turn generate variance reduction parameters–that are informed by angle to remedy this issue. The first section in this chapter describes the mathematical foundation of this new method. A discussion on how the method's performance will be quantified follows. Finally, a description of the software being used and how the method is added to this software concludes this chapter.

3.1 Theory: Angle-Informed Importance Maps for CADIS and FW-CADIS

There exist methods to generate variance reduction parameters for deep penetration radiation transport problems with strong anisotropy in the flux. These methods have shown to have varying success, and may not be fully automated. The solution proposed in this dissertation is a formulation that we have named the Ω-CADIS-methods. This section will commence with a brief discussion of the foundational research on which the Ω-CADIS-methods are built. That discussion serves as a primer for the subsequent section, which is an introduction to the Ω-CADIS-methods and a discussion on how they differ from their predecessors.

3.1.1 Previous Work

As discussed in Sections 2.3.1 through 2.5, the existing gold standard for automatically generating variance reduction parameters for deep penetration fixed-source radiation transport problems are CADIS and FW-CADIS. Both of these methods are very effective at generating

variance reduction parameters for local and global solutions, respectively. However, CADIS and FW-CADIS have only been implemented to perform variance reduction in space and energy, not angle. As a result, solutions for problems with strong anisotropy in the flux are not always optimized with these methods, resulting in slow convergence times and low FOM values. Problems with strong anisotropies in the flux require more than just space- and energy- variance reduction techniques. A number of angle-informed variance reduction methods have been investigated, most notably AVATAR, LIFT, and a modified version of CADIS using AVATAR-type angular parameters.

LIFT, AVATAR, and Simple Angular CADIS all showed that by including angular information into Monte Carlo variance reduction parameters the FOM can be improved. However, none of these methods used the actual angular flux to calculate the variance reduction parameters for the problem they were optimizing. Without explicitly using the angular flux solutions they were limited in which types of problems they were applicable, because some assumption of the degree of anisotropy of the flux was made. Further, LIFT and Simple Angular CADIS showed that by including substantial angular biasing in the weight windows in problems where the approximation to the angular flux is not sufficient, the FOM can decrease not unsubstantially, defeating the purpose of using these methods.

3.1.2 The Ω Methods

The foundation of the Ω-methods is built upon CADIS and FW-CADIS. As with both methods, the Ω-methods will use a version of the adjoint scalar flux to consistently bias a Monte Carlo problem with the intention of reducing the variance. In Section 2.2.1 the concept of importance was introduced. Notably, it was shown that the adjoint flux is a good marker for the likelihood of particles to contribute to a tally, which is the particle's importance. It was also shown that the product of the forward and adjoint fluxes generates a pseudo-particle flux called the contributon flux, where contributons are "importance particles". These importance particles can be used to show preferential flow paths from a source to a tally or desired location.

By using a version of the adjoint scalar flux that has been formulated with the contributon flux, the direction of particle flow will be incorporated into the importance map and, consequently, the variance reduction parameters. By using this variant of the adjoint scalar flux, the method, like traditional CADIS, will show increasing importance as the particles travel near the adjoint source. However, because this variant of the adjoint flux incorporates directionality of the particle flow, not all regions near the adjoint source are equally important. In this way, the adjusted flux incorporates features from both the adjoint- and contributon- fluxes.

The adjusted adjoint scalar flux quantity, or the Ω-adjoint scalar flux, is

$$\phi^{\dagger}_{\Omega}(\vec{r}, E) = \frac{\int_{\Omega} \psi^{\dagger}(\vec{r}, E, \hat{\Omega})\psi(\vec{r}, E, \hat{\Omega})d\hat{\Omega}}{\int_{\Omega} \psi(\vec{r}, E, \hat{\Omega})d\hat{\Omega}}. \tag{3.1}$$

The Ω-flux is a hybridization of the adjoint scalar flux and the contributon flux. It is both a normalized contributon flux and a forward-weighted adjoint flux. As a result, it should inherit some of the advantages of each of the traditional adjoint and the contributon fluxes. Because it maintains dimensionality of the traditional adjoint scalar flux, it can be used in place of the standard adjoint scalar flux in both CADIS and FW-CADIS variance reduction parameter generation. This means that the method can capitalize on existing infrastructure used to generate variance reduction parameters for CADIS and FW-CADIS, and only the software handling the transport and flux-generation requires modification.

3.1.2.1 CADIS-Ω

As with CADIS, CADIS-Ω consistently biases a problem's source and particle weights according to their importance. However, CADIS-Ω uses the Ω-adjoint scalar flux rather than the standard adjoint scalar flux to generate the biased source distribution, weight windows, and the particle birth weights. Furthermore, because $\phi^\dagger_\Omega$ is used to calculate these values in CADIS-Ω, the consistent-biasing hallmark for which CADIS is known is maintained. The adjusted formulation of CADIS using the Ω fluxes is given by Eqs. (3.2). The biased source distribution used by CADIS-Ω is formulated just as it is in CADIS, except the adjusted adjoint fluxes are used:

$$\begin{aligned}\hat{q}_\Omega &= \frac{\phi^\dagger_\Omega(\vec{r},E)q(\vec{r},E)}{\iint \phi^\dagger_\Omega(\vec{r},E)q(\vec{r},E)dEd\vec{r}} \\ &= \frac{\phi^\dagger_\Omega(\vec{r},E)q(\vec{r},E)}{R_\Omega}.\end{aligned} \tag{3.2a}$$

The starting weights of the particles sampled from the biased source distribution, $\hat{q}$ are given by

$$\begin{aligned}w_{0,\Omega} &= \frac{q}{\hat{q}_\Omega} \\ &= \frac{R_\Omega}{\phi^\dagger_\Omega(\vec{r},E)},\end{aligned} \tag{3.2b}$$

and the new target weights for the particle are

$$\hat{w}_\Omega = \frac{R_\Omega}{\phi^\dagger_\Omega(\vec{r},E)}. \tag{3.2c}$$

3.1.2.2 FW-CADIS-Ω

FW-CADIS differs from CADIS in that it requires a forward deterministic calculation to generate $q^\dagger$, which is used as the source distribution in the adjoint deterministic problem

(recall that CADIS sets $q^\dagger = \sigma_d$). Depending on the type of global response desired, FW-CADIS runs a deterministic forward calculation to approximate the global response in the problem. The inverse of these responses is then used to generate the biased adjoint source distribution for the adjoint deterministic run. Therefore, the behavior of FW-CADIS-Ω in the forward biasing portion of the calculation will remain unchanged from FW-CADIS. The generalized form for the adjoint source definition is given by the fraction of the response in a region of phase space, P, over the total response in the problem, or

$$q_\Omega^\dagger(P) = q^\dagger(P) = \frac{\sigma_d(P)}{R}.$$

When applied to the spatially-dependent global dose, $\int \phi(\vec{r}, E)\sigma_d(\vec{r}, E)dE$, the adjoint source will be

$$q_\Omega^\dagger(\vec{r}, E) = q^\dagger(\vec{r}, E) = \frac{\sigma_d(\vec{r}, E)}{\int \sigma_d(\vec{r}, E)\psi(\vec{r}, E,)dE}.$$

The adjoint source for the spatially-dependent total flux $\int \phi(\vec{r}, E)dE$ is

$$q_\Omega^\dagger(\vec{r}) = q^\dagger(\vec{r}) = \frac{1}{\int \phi(\vec{r}, E)dE}.$$

The adjoint source for the energy- and spatially-dependent flux $\phi(\vec{r}, E)$ is

$$q_\Omega^\dagger(\vec{r}, E) = q^\dagger(\vec{r}, E) = \frac{1}{\phi(\vec{r}, E)}.$$

One advantage of FW-CADIS-Ω is that, from a transport perspective, the Ω-method is no more expensive than standard FW-CADIS. Because both versions require a forward and adjoint deterministic calculation, an extra transport step is not required as it is for CADIS-Ω. This is attractive, but the nature of FW-CADIS might not be the most well-suited for the Ω-methods. Because FW-CADIS attempts to evenly distribute particles throughout the problem using the forward-biased adjoint fluxes, the additional forward normalization with the Ω-methods will likely skew the particle distribution in the problem in the forward direction, and it may place too great of importance on the forward-moving particles in generating the variance reduction parameters.

3.2 Computational Success Metrics

3.2.1 Anisotropy Quantification

As the Ω-methods are analyzed, it is important to determine the types of problems in which the methods are successful. In addition to describing the physics that induce anisotropy in the flux, quantifying the degree of anisotropy of the problem is useful in characterizing the

method. In this section, a number of methods by which the anisotropy can be quantified in these problems are proposed. A brief description of how these methods capture anisotropy in the problem is also included. While each metric proposes an avenue by which the problem can be analyzed, there are certainly other methods that one may propose. The methods described in the following subsections are proposed because they use data generated from the existing method. The degree to which they impose a computational burden will be addressed in their analysis.

3.2.1.1 The Scalar Contributon Ratio

The hybrid methods software that will be used for this project is ADVANTG, developed at ORNL. Section 3.3 explains how the software used interacts with other pieces of software and how they were modified to execute this method. The standard release of ADVANTG provides the contributon flux as an output option, which can then be used to analyze problem physics by a user. If this option is selected as an output, a SILO file containing the contributon fluxes for each discretized cell in space and energy is created. This is useful for problem analysis as the user may see preferential streaming paths for particles in the problem using this metric. The contributon flux generated in this process is given by the product of the scalar adjoint and forward fluxes (Eq. (3.4)).

As mentioned in Section 2.2.3, the contributon flux can be calculated by using the product of the forward and adjoint fluxes. In standard software packages that calculate the contributon flux, like ADVANTG, the scalar contributon flux is calculated by the product of the scalar adjoint and forward fluxes. This can be written as

$$\phi^c(\vec{r}, E) = \phi^\dagger(\vec{r}, E)\phi(\vec{r}, E). \tag{3.4}$$

A more precise calculation of the contributon flux could be generated from integrating the angular contributon flux over all angle, as

$$\begin{aligned} \Phi^c(\vec{r}, E) &= \int_\Omega \psi^c(\vec{r}, E, \hat{\Omega}) d\hat{\Omega} \\ &= \int_\Omega \psi^\dagger(\vec{r}, E, \hat{\Omega})\psi(\vec{r}, E, \hat{\Omega}) d\hat{\Omega}. \end{aligned} \tag{3.5}$$

Both Eqs. (3.4) and (3.5) calculate the contributon flux as a function of space and energy, but the differences in their calculation is addressed in their notation, namely using ϕ^c or Φ^c. The standard release of ADVANTG only has access to the scalar fluxes, so Eq. (3.5) is not an accessible option for a user. Because the Ω calculations require full angular flux map, the scalar contributon flux can be calculated with the latter formulation, rather than the former in the modified version developed to support this work.

The first measure of anisotropy quantification that will be evaluated is the ratio between these two quantities, as described by Eq. (3.6). The ratio between these two values is evaluated for every cell, x, y, z, and energy group, E_g. If the adjoint or forward angular

flux is significantly peaked in Ω, this will result in a deviation between ϕ^c and Φ^c, because there will be a multiplicative effect in the angular flux captured in Φ^c but not ϕ^c. The more isotropic the flux in $\vec{r}$ and E, the closer these values will be and the quantity will approach unity.

$$M_1 = \left.\frac{\phi^c}{\Phi^c}\right|_{x,y,z,E_g} \tag{3.6}$$

3.2.1.2 The Ratio of Adjoint Fluxes

As discussed in previous sections, the Ω-methods use the Ω-scalar flux in place of the standard adjoint scalar flux. Therefore the ratio between these two quantities would also provide a useful metric for comparing which regions have significantly differing bias parameters in standard-adjoint and Ω-adjoint situations. This metric will deviate from unity if the forward flux is anisotropic. This metric is calculated for every cell and every energy group in the problem, as shown in Eq. (3.7).

$$M_2 = \left.\frac{\phi^\dagger_\Omega}{\phi^\dagger}\right|_{x,y,z,E_g} \tag{3.7}$$

Metrics one and two both reasonably appear to compute the anisotropy in the flux using versions of the contributon and adjoint fluxes, respectively. However, by expanding the Ω-adjoint scalar flux in metric two,

$$\begin{aligned} M_2 &= \left.\frac{\phi^\dagger_\Omega}{\phi^\dagger}\right|_{x,y,z,E_g} \\ &= \left.\frac{\int_\Omega \psi^\dagger(\hat{\Omega})\psi(\hat{\Omega})d\hat{\Omega}}{\int_\Omega \psi(\hat{\Omega})d\hat{\Omega}}\frac{1}{\phi^\dagger}\right|_{x,y,z,E_g}, \end{aligned}$$

integrating the forward angular flux over all angle,

$$= \left.\frac{\int_\Omega \psi^\dagger(\hat{\Omega})\psi(\hat{\Omega})d\hat{\Omega}}{\phi}\frac{1}{\phi^\dagger}\right|_{x,y,z,E_g},$$

and rearranging the terms,

$$\begin{aligned} &= \frac{\int_\Omega \psi^\dagger(\hat{\Omega})\psi(\hat{\Omega})d\hat{\Omega}}{\phi\phi^\dagger}\Bigg|_{x,y,z,E_g} \\ &= \frac{\Phi^c}{\phi^c}\Bigg|_{x,y,z,E_g} \\ &= \frac{1}{M_1}\Bigg|_{x,y,z,E_g}, \end{aligned}$$

it becomes evident that the ratio of adjoint fluxes is the inverse of the scalar contributon ratio. As a result, metric one will not be used in the analyses of the characterization problems.

3.2.1.3 The Maximum to Average Flux Ratio

An alternative metric to quantify anisotropy is to calculate the ratio between the maximum and average angular contributon flux in each $\vec{r}, E$ voxel. The higher this quantity, the more peaked the contributon flux is in Ω. Note that while using the Ω-flux would seem like the natural choice, no angular information is directly accessible once the Ω scalar flux has been calculated. One can compare the standard adjoint scalar flux and the Ω-adjoint scalar flux and infer how anisotropic the flux in the cell might be, but due to the normalization that occurs in Eq. (3.1), the variation of angular Ω fluxes throughout Ω for a cell in x, y, z, E_g is not calculated. As such, the contributon flux must be relied upon as a next-best evaluator of that metric:

$$M_3 = \frac{\psi^c_{Max}}{\psi^c_{Avg}}\Bigg|_{x,y,z,E_g}. \tag{3.8}$$

While Eq. (3.8) directly measures the anisotropy in the problem using the angular contributon fluxes, it doesn't compare the difference between the fluxes used in the Ω- and the standard adjoint methods. Metric three can be reformulated to incorporate this information using

$$\begin{aligned} M_4 &= \frac{\frac{\psi^c_{Max}}{\psi^c_{Avg}}}{\frac{\psi^\dagger_{Max}}{\psi^\dagger_{Avg}}}\Bigg|_{x,y,z,E_g} \\ &= \frac{M_3}{\frac{\psi^\dagger_{Max}}{\psi^\dagger_{Avg}}}\Bigg|_{x,y,z,E_g}, \end{aligned} \tag{3.9}$$

as a measure between the anisotropies of the standard and contributon fluxes. This equation is a logical progression from metric two and metric three. This metric contains more information on how perturbed the contributon flux is when compared to the original adjoint

flux that is normally used in CADIS and FW-CADIS. In the case of a strongly anisotropic forward flux, the forward flux would significantly change the distribution of the contributon fluxes in a cell, but it would not affect the flux distribution of the standard adjoint angular fluxes. By comparing the anisotropy in the contributon fluxes to those in the standard adjoint, the perturbation of the Ω flux by the forward flux in the cell can be evaluated. In regions where the forward flux is not anisotropic, then the contributon anisotropy ratio should be approximately the same as the standard adjoint anisotropy ratio.

Further, because the contributon flux incorporates directionality of the forward and adjoint fluxes, the maximum to average ratio of the contributon flux can differ from the adjoint flux. In regions where the adjoint angular flux and the forward angular flux are traveling in the same direction, the contributon ratio should be greater than the adjoint ratio, and this metric will be greater than one. In regions where they are travelling in opposite or perpendicular directions, the contributon flux will evaluate to a more isotropic state, and metric four will be less than unity. This metric provides substantially more information than metric two because it compares the behavior of the directional contributon and adjoint fluxes, rather than comparing the overall behavior of the flux in the cell.

Both Eqs. (3.8) and (3.9) compare the maximum angular flux in a cell to the average flux in the same cell. Because the average angular flux is the normalization factor, the maximum flux in the cell is compared to some relative measure of the total flux behavior in that cell. If, for example, the flux has several directional peaks, the average will reflect that. The fact that Eq. (3.9) contains information on the global behavior in the contributon and average cell, the directionality of the fluxes, and the degree of isotropy of the forward flux is attractive. However, this is also a fairly computationally expensive calculation and it may not be worth the computational cost when compared to metrics two and three.

3.2.1.4 The Maximum to Minimum Flux Ratio

An additional metric to quantify anisotropy in the contributon flux distribution is to calculate the ratio between the maximum and minimum angular fluxes for each region of x, y, z, E_g phase-space, as described in metric five, or Eq. (3.10). This quantity incorporates information about the behavior of the local maximum relative to the local minimum angular flux in each cell.

$$M_5 = \left.\frac{\psi^c_{Max}}{\psi^c_{Min}}\right|_{x,y,z,E_g} \tag{3.10}$$

This metric may be more appropriate to describe the anisotropy of the flux in cells where the distribution of flux values in the cell are not well reflected by the average flux in the cell. As with metric three (Eq. (3.8)), metric five (Eq. (3.10)) only quantifies the anisotropy of the contributon flux in the cell. There is no comparison or normalization to compare the anisotropy with respect to another method. To compare it to the anisotropy of the flux in the standard adjoint problem, a ratio similar to that of Eq. (3.9) may be formulated:

$$M_6 = \left. \frac{\frac{\psi^c_{Max}}{\psi^c_{Min}}}{\frac{\psi^\dagger_{Max}}{\psi^\dagger_{Min}}} \right|_{x,y,z,E_g} = \left. \frac{M_5}{\frac{\psi^\dagger_{Max}}{\psi^\dagger_{Min}}} \right|_{x,y,z,E_g}. \tag{3.11}$$

As with Eq. (3.9), Eq. (3.11) uses a ratio from the standard adjoint formulation to normalize the anisotropy of the contributon flux. Equation (3.11) is consistent with Eq. (3.10) and normalizes using the maximum to minimum ratio of angular fluxes of the adjoint. These two metrics will show the relative behavior of the flux in the cell, but because neither incorporates information about the total flux behavior within the cell, they may be very sensitive to the variance of the angular flux within the cell. Using the ratio of both the contributon and adjoint fluxes may help to smooth this if the variance of flux distributions within the contributon and standard adjoint is similar in a particular cell. However, if these two differ significantly, then metric six (Eq. (3.11)) may have a synergistic effect and will over-emphasize the variance when quantifying the anisotropy of the cell.

Metrics one through six quantify anisotropy in the problem solved by using different parameters to capture the problem physics. These metrics will be compared to one another to determine which is the most consistently correlated with predicting the Ω-method's success. A user may want to know if the Ω-method will effectively generate variance reduction parameters for a Monte Carlo simulation, and this may be a prescriptive solution for that issue. However, all of these metrics do require full angular flux solutions for both the forward- and adjoint- problem, so some computational burden will be required. The analysis of using these metrics will include some information of benefit to burden, which likely will come at the cost of time. That said, because the Monte Carlo solution is more computationally demanding, generating these metrics from the deterministic solution should be substantially less of an obstacle.

3.2.2 Figure of Merit

The FOM is a commonly used metric to measure Monte Carlo runtimes and to gauge the effectiveness of various hybrid methods. As discussed in Section 2.1, the FOM relates the relative error of a solution to the time required to achieve that variance. This was introduced in Eq. (2.20) as:

$$\text{FOM} = \frac{1}{R^2 T},$$

where T is the time and R^2 is the square of the relative error.

3.2.2.1 Relative Error

In tallies with multiple regions and/or energy bins, the FOM is usually calculated from the tally average relative error, or R_{avg}. This value is meaningful as it reflects the overall tally behavior. However, it is often desirable that all portions of the tally lie below a desired relative error threshold. A region with very low particle contribution may have a much higher relative error than the tally average, and may also converge much slower to a desired relative error. This results in a substantially different FOM than the tally average. In the results presented in later chapters, both relative errors will be used to calculate different FOMs, respectively

$$\text{FOM}_{avg} = \frac{1}{R_{avg}^2 T}, \tag{3.12a}$$

and

$$\text{FOM}_{max} = \frac{1}{R_{max}^2 T}. \tag{3.12b}$$

In addition to reporting both FOMs for the entire problem, comparing the distribution of values of the relative error for problems will be a useful metric in method characterization. If, for example, FW-CADIS acquires desirable results in a calculation, then the problem should have a relatively even uncertainty distribution for all cells. Comparing the distribution of relative errors between the analog case and the hybrid case reveals whether the method is effectively generating variance reduction parameters for the entire problem or if it is more effective in particular regions.

3.2.2.2 Timing

The previous section described two different means by which the FOM could be calculated using different relative errors. The question that one must now consider is: what time should be used to calculate the FOM? In an analog Monte Carlo simulation, this time is the runtime of the Monte Carlo simulation, $T = T_{MC}$. In a hybrid method, one could choose either

$$T_{Hybrid} = T_{MC} + T_{Deterministic}, \tag{3.13}$$

or

$$T_{Hybrid} = T_{MC}. \tag{3.14}$$

The FOM should remain a constant–with the exception of very early on in an MC calculation where statistics are very poor–for a problem. The issue with using Eq. (3.13) to calculate the FOM is that the deterministic runtime does not change the relative error of the Monte Carlo simulation. Thus, the FOM is not a constant throughout the Monte Carlo simulation when using Eq. (3.13) as the time. However, it would be disingenuous to not include the deterministic runtime into reports for the hybrid method, as the total computational time required to achieve some desired relative error is ultimately what the user is

seeking. As such, two reports of the FOM are included with the results for each simulation:

$$\text{FOM}_{MC} = \frac{1}{R^2 T_{MC}}, \tag{3.15a}$$

and

$$\text{FOM}_{Hybrid} = \frac{1}{R^2 (T_{MC} + T_{Deterministc})}. \tag{3.15b}$$

Note that the deterministic time used in Eq. (3.15b) is the time to run the transport and generate source biasing and weight window values for each problem. It will not include the time used to quantify the anisotropy as outlined in Section 3.2.1, as those parameters will be computationally demanding but not normally included in a hybrid method computation.

In this section, four different equations to calculate the FOM were presented: two using different relative errors, and two using different quantities for time. In analyzing the method, all four will be presented: $\text{FOM}_{MC,avg}$, $\text{FOM}_{MC,max}$, $\text{FOM}_{Det,avg}$, and $\text{FOM}_{Det,max}$. Further, the improvement in the FOM for each problem will be reported as those values normalized by $\text{FOM}_{analog,avg}$ for the two FOMs calculated with the tally average relative error and $\text{FOM}_{analog,max}$ for the FOMs calculated with the tally maximum relative error. The success of the Ω-method will depend on its ability to improve each one of these FOM values.

3.3 Software

In this section, the software in which the methods presented in Section 3.1.2 are implemented is described. A brief summary of each piece of software and what was added in each is discussed. While the details of the inner-workings of the software will not be described here, both pieces of software have rich documentation and user guides which an interested reader may reference.

3.3.1 Denovo

Denovo S_N is a three-dimensional discrete ordinates transport solver developed at Oak Ridge National Laboratory [62]. Denovo is a module in the larger Exnihilo massively-parallel radiation trasnsport code suite. There exist several other modules in Exnihilo. In addition to Denovo, the most pertinent package being Omnibus, a frontend pre- and post-processing module. The Ω-fluxes are generated by running two independent (a forward and an adjoint) determinstic solves in Denovo. The setup and generation of each simulation input is automated through ADVANTG (see Section 3.3.2). After the calculation has reached the desired convergence criteria, the full angular flux maps for the forward and adjoint solves are saved to an HDF5 [72] file. Denovo was modified to output the full angular flux maps for a simulation. The Ω-fluxes are then generated by passing the angular flux maps through the postprocessing module in Omnibus. Using this module, the integration described in Eq. (3.1) is performed, the scalar Ω-fluxes are saved to a SILO file, and the scalar fluxes are

passed to ADVANTG for variance reduction parameter generation. Appendix A.1 contains the code added to Omnibus to perform this calculation.

3.3.2 ADVANTG

ADVANTG [70] is a software package originally designed to automatically generate variance reduction parameters for the Monte Carlo radiation transport solver MCNP [9] using the CADIS and FW-CADIS methods. For this project, the ADVANTG functionality was extended to process the Ω-fluxes provided by Denovo through Omnibus and to generate variance reduction parameters for CADIS and FW-CADIS using said fluxes. In addition to the modifications required to perform CADIS and FW-CADIS, ADVANTG was further modified to generate anisotropy quantification metrics and a modified version of the scalar contributon flux, both of which were summarized in Section 3.2.1. The piece of code used to reroute the Ω-fluxes through CADIS and FW-CADIS as well as to generate the anisotropy metrics in ADVANTG is included in Appendix A.2.

Summary

In summary, this chapter presented the novel theory behind the Ω-methods; the metrics by which the Ω-methods will be compared with existing hybrid methods; and the software that was modified to implement the Ω-methods into an existing codebase. Two variants of the Ω-methods were presented: CADIS-Ω and FW-CADIS-Ω, which are referred to together as FW/CADIS-Ω. CADIS-Ω is a modification of CADIS, and is designed to generate VR parameters for local solutions in problems with strong anisotropy. FW-CADIS-Ω is a modification of FW-CADIS, and is designed for generating VR parameters for global solutions in problems with strong anisotropy.

Both CADIS-Ω and FW-CADIS-Ω are implemented in well-used, well-documented, massively-parallel, state-of-the-art radiation transport and hybrid methods software. The radiation transport code suite Exnihilo is modified to generate the Ω-fluxes. The hybrid methods package ADVANTG is modified to generate VR parameters for the Ω-methods using the Ω-fluxes.

To understand the performance of the Ω-methods and compare it consistently to existing methods, several performance metrics were proposed. First, a few variants of the FOM were described. They include: $\text{FOM}_{MC,avg}$, $\text{FOM}_{MC,max}$, $\text{FOM}_{hybrid,avg}$, $\text{FOM}_{hybrid,avg}$. Together, they provide an overall picture of the performance of the Ω-method's performance with respect to relative error and time, rather than of a single criteria. Because anisotropy has the ability to affect energy groups differently, resulting in different relative errors achieved in different energy bins, separating out different FOMs helps to isolate interesting behavior in the methods.

The Ω-methods are designed to work in problems with strong anisotropies in the flux. As a result, several anisotropy metrics with which to investigate flux anisotropy were proposed.

Using these metrics and comparing them to the relative errors or FOMs in each tally region, we can try to understand the effect that anisotropy has on the Ω-method performance. Each metric quantifies the anisotropy in cells differently, so each has the potential to capture different information. Denovo was modified to output angular fluxes to generate the Ω-flux for the Ω-methods. As a result, the anisotropy metrics use data generated from the existing Ω-method calculation.

Using the methodology described in this chapter, the Ω-methods' performance can be fully characterized. Further, the characterization presented in this chapter has been extended from standard FOM performance metrics to include anisotropy quantification. By implementing the Ω-methods into production-level software, it is accessible to any user beyond the author. The generation of the anisotropy metrics is also incorporated into the codebase, meaning that any user could feasibly perform an investigation of the Ω-performance consistent with what is proposed herein. The use of the various FOMs and of the aniostropy metrics helps the understanding of the Ω-method performance as a function of time, error, and anisotropy.