Chapter 5

Conclusions

Hybrid methods are and will be a realm of continued importance in radiation transport methods development. The application space and demand for hybrid methods continues to grow. With this growth, accurately and efficiently modeling the physics of increasingly complex problems is paramount for safety and security. In this dissertation, a new set of hybrid methods were proposed, implemented, and characterized. From this work, several pathways have revealed themselves for future hybrid methods work.

5.1 Assessment of the Ω-methods

The results in Chapter 4 showed that CADIS-Ω has varied performance when compared to CADIS over the problem space investigated. Depending on the geometric configuration, the material composition, and the solver options used, the method can outperform or underperform CADIS by an order of magnitude. This underscores the difficulty of developing a method that is broadly applicable to a large subset of application space. Further, it illustrates the necessity for further methods development.

Several characterization problems were formulated that contained anisotropy in the flux. The mechanisms for inducing anisotropy in the flux anisotropy were either from the source, or from physical interactions with the problem materials and geometry. The success of the Ω-methods was not directly correlated with any single physical mechanism, but both CADIS and CADIS-Ω struggled in problems primarily comprised of air.

In the single turn labyrinth, CADIS-Ω achieved lower relative errors in epithermal and fast energy groups. These groups were shown to have flux anisotropies with anisotropy distributions that were clumped around a particular anisotropy value. For the multiple turn labyrinth, CADIS achieved uniformly lower relative errors than CADIS-Ω. For both the steel beam in concrete and the u-shaped bend, CADIS-Ω achieved lower relative errors than CADIS but had runtimes 3-7x longer than those of CADIS. For the geometrically complex rebar-embedded concrete, CADIS-Ω had higher relative errors than CADIS. In high energy regions, the convergence for energy bins would take days of computational runtime to get to a relative error of less than 10%. For two heavily air-centered problems, CADIS-Ω and CADIS both had comparable relative error achievements.

In addition to checking the limitations of the Ω-methods with respect to geometry and material composition, the sensitivity of the methods to deterministic parameter selection was also studied. In particular, the effect of quadrature order and P_N order on method performance were studied. For both CADIS and CADIS-Ω, the change in quadrature order had a stronger effect on the change in the relative error and the FOM. CADIS showed stronger sensitivity to changes in both P_N order and quadrature order over CADIS-Ω. CADIS also proved to have more and higher magnitude oscillations in the relative error between different P_N and quadrature orders. Spikes in the relative error occurred in both methods, but more frequently in CADIS. Both methods showed improvement in the FOM and relative error with increasing quadrature order and P_N order. In high energies, CADIS-Ω achieved superior FOMS to CADIS for all P_N orders and quadrature orders.

Chapter 4 showed a few examples of the anisotropy metrics when they showed promising trends with I_{RE} or I_{FOM}. These metrics did provide information on the relative distribution of anisotropy in the problem, and they also showed some trends with the improvement factors. However, most problems did not have significant trends, so more work must be done to fully characterize hybrid methods using this novel analysis technique.

The Ω-methods have been characterized with their sensitivity to geometric and material configuration, as well as their sensitivity to deterministic calculation parameter choice. It is clear from the results in Sections 4.2 and 4.3 that the Ω-methods are not always the best choice for reducing the variance in problems with anisotropy. This is from a combination of many effects, but primarily the varied range of runtimes when compared to CADIS. In many problems, CADIS-Ω was able to obtain lower relative errors for tally bins than CADIS, but the runtimes were significantly longer. The generally longer runtime for CADIS-Ω negatively impacts the FOMs that it is able to achieve, thus negating its more effective transport of particles.

5.2 Suggested Future Work

While this dissertation covered the characterization of CADIS-Ω over a fairly broad spectrum of anisotropy-containing problems, there are a number of fruitful pathways by which the method could be improved or characterization expanded. Broadly, these fall into three categories: improvements to the software implementation and algorithmic design, expansion of the characterization space, and application to larger, real-life problems. The next few subsections addresses each one of these categories individually.

5.2.1 Software Improvement

This subsection addresses the improvements that could be made to the software and analysis methods to enhance understanding of the omega methods. A discussion on how improving software performance aids in future work will start this section. An explanation of how extending the breadth of analysis helps to understand the Ω-methods further will finish the section.

To calculate the Ω flux, a rotation of either the adjoint or forward flux matrix is required to ensure that the directional variable Ω is consistent between the forward and adjoint, or that $\Omega_{adjoint} = \Omega_{forward}$. Because quadrature sets are not always straightforward to interpolate, a rotationally symmetric quadrature set is currently required for computing the Ω-fluxes in order to perform this rotation. Should a method be developed that does have interpolatable quadrature points, it would be a good candidate to calculate Ω-fluxes for solutions that are not rotationally symmetric.

The anisotropy metrics described in Section 3.2.1 at this point have not shown significant trends with either the relative error or figure of merit improvement metrics (I_{RE} and I_{FOM}). To filter out values of each metric to regions more important to the problem solution, two filtering algorithms were proposed: one that only uses values of metrics from cells with contributon fluxes above the mean contributon flux, and the other that uses values from cells that have a flux above the median contributon flux. Using these filtering algorithms did show interesting features in the anisotropy metric distributions as well as shifts in I_{RE} and I_{FOM}. However, trends were not apparent for the majority of the metrics. A useful modification to the filtering algorithm would be to select certain percentages of high-valued contributon flux locations. For example, perhaps selecting out the cell locations containing the top 10% of contributon fluxes would reveal a trend in the improvement metrics. It is possible that too many values are being selected from the entire problem even with the existing filters, so an even stricter filtering algorithm may help.

To filter the anisotropy metrics, the contributon flux distribution was chosen as the filter base. This is an intuitively good choice because it will use values near both the forward- and adjoint- sources, and also the values between them where particles are most likely to flow. Further, the contributon flux is something that is method agnostic. That is, it can be used as a filtering algorithm for non-Ω methods and it will still reveal problem information. However, an argument could also be made to use the omega flux distribution as a filtering base, as that is the method in which we are interested. Modifying the filtering algorithms to use the Ω-flux distribution may provide trends in the method improvement metrics that are not apparent using the contributon flux.

The Ω-methods, as currently implemented in both Exnihilo and ADVANTG, are entirely serial. That is, there is no parallelization in any part of the Ω-flux calculation, or supporting code to that effect. In the results presented in Chapter 4, T_{hybrid} was calculated to remove parallelization effects so that CADIS and CADIS-Ω were comparable. While the results were adjusted accordingly, this is not the best implementation for production software or more difficult use cases. As mentioned previously, the ADVANTG software is entirely serial, so parallelization is not required for VR parameter generation. However, Exnihilo/Denovo is parallelized. The parallelization of the Ω-flux calculation in this code would significantly improve its usability. Parallelization would reduce the actual time to calculate the Ω-fluxes and anisotropy metrics.

Another algorithmic improvement to the Ω-methods is to reduce the memory requirements for both the computation of the Ω-fluxes and the anisotropy metrics. Much of this could be accomplished with parallelization. However, even the serial version of the Ω-methods could be adjusted to read in the angular flux data in "chunks" so as to not read in datasets larger than the memory available on the system. As a first order approach, the angular fluxes could be read in serially by energy group. Depending on the energy group structure, this has the potential to reduce the memory load at a particular time by 20x or 200x. At present, the Ω-methods are limited by memory requirements. Without a large computing cluster, there is no feasible way to calculate the Ω-fluxes for a problem of reasonable complexity.

Another alternative modification to the Ω-methods is to bypass writing the angular flux matrices entirely. This would reduce the I/O requirements for the method, and also not demand as much disk space. However, this is a non-trivial task, as the forward and angular fluxes for a cell must both be in memory to compute the Ω-method for that cell. To store the complete angular flux matrices in memory will present the same memory limitations that the Ω-methods currently face, so some algorithmic challenges exist should this be a path of future work.

The Ω-methods are currently implemented on a localized development version of both Exnihilo and ADVANTG. If a larger audience wishes to use or access them, they would require support beyond that of a standard software release. Depending on the continued characterization of the Ω-methods, integrating this software into future releases of Exnihilo and ADVANTG may be useful.

Each of the areas proposed in the previous paragraphs are areas in which the Ω-methods can be improved upon or areas that may improve our understanding of the Ω-methods' behavior. Expanding the filtering algorithm for the anisotropy metrics may also help us to understand more broadly how anisotropy is distributed in different problems. Expanding our understanding of the Ω-methods' strengths and deficiencies can also improve future hybrid methods.

5.2.2 Characterization Problem Extension

Broadening the scope of the characterization problem study is another fruitful avenue for exploration. In this vein, there exits a two-pronged approach: first extending the types of problems (more diverse materials, less air in problems, more diverse geometries) will enhance knowledge of the methods. Next, extending the scope of the parametric studies will help to inform how resilient the Ω-methods might be to changes in the solutions space that indirectly impacts angle. In this realm, the deterministic calculation specifics, like quadrature type will be addressed.

The characterization problems studied covered a broad range of anisotropy-inducing physics. The geometries chosen were fairly simple, with very few materials. The majority of the problems used air in some portion of their geometry to have streaming-induced anisotropy of the flux. Depending on their geometries, this caused sampling issues and slowdown of the CADIS-Ω method. For example, the problem variants of the steel beam embedded in concrete geometry illustrated the CADIS-Ω method's susceptibility to air. In the air-filled beam variant of the problem, the Ω-method had the lowest improvement margin when compared to CADIS of the three matierial variants. A beneficial extension of the characterization problem study would be to replace the air in this geometry with a high atomic mass material that maintains scattering anisotropy but includes more sampling interaction points. Using problems with greater material diversity and more problems with preferential flowpaths (that are not air), would be an interesting extension to the characterization problem materials.

While the characterization problems are fairly simple geometrically, it may be advantageous to investigate simpler problem geometries with even less geometric complexity. In comparing the single- and multiple-turn labyrinths, we observed that with too little anisotropy in the problem, the Ω method's performance suffers. However, a simpler geometry of the labyrinth (perhaps an elbow bend), or a hallway in concrete with no air rooms, can show if there is a turnover in labyrinth anisotropy in which the Ω methods perform the best.

The results presented in Section 4.3 showed that CADIS-Ω is generally more resilient than CADIS to changes in quadrature and P_N discretization. As a result, CADIS-Ω can use a coarser problem discretization to obtain variance reduction parameters, saving computational cost in terms of both runtime and memory. The results in Section 4.3 also showed that CADIS-Ω was less susceptible to large fluctuations in the relative errors in intermediate energy energy bins.

Beyond sensitivity to quadrature order and P_N order, it may be worth investigating the sensitivity of each method to other deterministic calculation parameters. If, like quadrature order and P_N order, CADIS-Ω generates better importance maps with lower-fidelity solutions in other deterministic parameters, then even more computational time could be saved. For something like mesh refinement, the number of mesh cells can significantly alter the speed at which the deterministic solution converges.

Investigating the impact of quadrature type may also be an area of future work. In Section 4.3, it was observed that both CADIS and CADIS-Ω showed greater sensitivity to changes in quadrature order than P_N order. CADIS showed a greater sensitivity to changes in quadrature order than CADIS-Ω. We expect that the behavior of other quadrature sets will be similar, but this may be worth verifying in future use cases. It is possible that the different properties of different quadrature sets may more strongly affect the Ω-methods' performance.

In addition to characterizing the performance of CADIS-Ω, it will be important to characterize FW-CADIS-Ω. In Chapter 3, the Ω-method theory for both CADIS and FW-CADIS were presented. Indeed, FW-CADIS-Ω has also been implemented in Exnihilo and ADVANTG. The scope of this project did not extend to the characterization of FW-CADIS-Ω, though it could prove useful to characterize for large, global calculations. A similar set of characterization problems can be designed for FW-CADIS-Ω, but with global mesh tallies rather than small detectors.

It would also be beneficial to perform a thorough investigation into the Ω-methods' mitigation or multiplication of ray effects. Both the forward and the adjoint angular fluxes will have ray effects in problems with long mean free paths. As discussed in Section 4.2, the ray effects may also be multiplied depending on the geometric configuration of the problem. The degree to which the Ω-flux exacerbates or minimizes ray effects as a function of these locations would be an interesting study and may help in further specifying to which problems the Ω-methods is suited. Further, the difference in the construction of the adjoint between CADIS and FW-CADIS means that CADIS-Ω and FW-CADIS-Ω will have different sensitivities to ray effects.

Further characterization of the Ω methods' performance with different problem geometry and material configurations will deepen our understanding for which applications the methods may be best suited. For large scale, high-impact, high-complexity problems, issues observed in the characterization problem studies may be exacerbated. Before applying this method to application problems, it will be important to have confidence that the methods will achieve better results than other methods.

5.2.3 Application Problems

Based on the data presented in Sections 4.2 and 4.3 we believe that the CADIS-Ω-method has the potential to be applied to a number of application problems. These problems include, but are not limited to: detectors near dry cask nuclear waste storage, dry cask storage beds, nuclear containment buildings, and nuclear spent fuel cooling pools.

The dry casks are a promising use case for the CADIS-Ω-method because they have small air channels for ventilation, but their body is primarily metal tubes containing nuclear fuel surrounded by concrete. These rods are pointed towards the ventilation ducts, and so the results from the steel bar embedded in concrete suggest that this may be a more complex application of the physics it represents.

Further, a bed of dry cask storage containers will have several spaces through which particles may travel. A use case of this may be to calculate the dose rate standing at the boundary of such a facility, or to consider if the cask loading matches the owner-provided loading list. Because this problem has so much air, it may be more difficult for the Ω-methods. However, with the thick soil boundary in the $z-$plane the Ω-methods may still perform well.

Nuclear spent cooling pools have used fuel rods clustered in assemblies arranged in rows submerged in water. These rods emit a range of highly energetic particles. Spent fuel cooling pools will be an interesting extension of the steel beam in concrete, as water is a highly moderating material not dissimilar to the concrete from the characterization problem. The fuel rods act as both a source and a preferential flowpath, so the differing source distribution in this problem may yield interesting results.

Each of these application problems uses the physics modeled in the characterization problems but applies them to a more geometrically and materially complex problem. In extending the Ω-methods to these problems and comparing them to CADIS and FW-CADIS, we can also understand how sensitive the Ω-methods are to more difficult problems. If, as noted in the characterization problems subsection, the Ω-methods are more resilient to deterministic problem solution fidelity in larger more complex problems, these problems will benefit significantly from the decreased deterministic solution time and the lower computational burden demanded by the Ω-methods.

5.3 Concluding Remarks

In this dissertation, a new group of hybrid methods called the Ω-methods were proposed. The Ω-methods are built on the foundational work of CADIS and FW-CADIS to generate angle-informed variance reduction parameters. The two new methods proposed were CADIS-Ω and FW-CADIS-Ω. Both methods use the Ω-flux, a form of the adjoint scalar flux calculated by weighting the adjoint angular flux with the forward angular flux, to generate source biasing and weight window values. By using the forward angular flux normalization, the importance map generated for the Ω-methods is adjusted to include the directionality of the forward and the adjoint particles without explicitly including angle in the source biasing or weight window values.

The Ω-methods were implemented in two software packages developed at Oak Ridge National Laboratory: Exnihilo and ADVANTG. The functionality to generate the Ω-fluxes were implemented in Exnihilo, which contains the deterministic transport solver Denovo. The infrastructure to generate variance reduction parameters consistent with CADIS and FW-CADIS was implemented in ADVANTG. The development of these methods now allows for any user to use the Ω-methods, should they have access to the software.

In addition to the Ω-methods method proposal and implementation, CADIS-Ω has been characterized on a wide variety of problems with flux anisotropies. The problems were designed to understand the method's limitations and in what parameter space the method can and should be used. To more fully understand the method's' behavior and how flux anisotropy affected its ability to perform, a number of anisotropy metrics were proposed. These metrics were then used to investigate if performance improvement could be correlated with anisotropy in any way.

The anisotropy metrics did not show significant trends with the FOM or the solution relative error, but their distributions did help reveal more about the distribution of anisotropy in the problems. In particular, it was easily observable how the distribution of anisotropy changed between energy groups for a particular problem. Future use of these metrics may also aid us in more fully understanding other hybrid methods' performance.

CADIS-Ω is a promising hybrid method. If used with a well-suited problem, it has the potential to improve the FOM over traditional methods by an order of magnitude. This offers significant time and energy savings. However, the Ω-methods are not without their drawbacks. If used in a poorly-suited problem they can take substantially more time to transport particles in Monte Carlo. The Ω-methods' characterization and performance study presented in this dissertation have contributed a broader understanding of these types of hybrid methods, and have created ample pathways forward for future hybrid methods analysis.