The Suppression of Energy Discretization Errors in Multigroup Transport Calculations

Integrated University Programs

Dr. Edward W. Larsen
University of Michigan

In collaboration with:
Oak Ridge National Laboratory

David Gerts, Technical POC

NNSA

National Nuclear Security Administration

Final Report

The Suppression of Energy Discretization Errors in Multigroup Transport Calculations

Edward W. Larsen

Department of Nuclear Engineering and Radiological Sciences
University of Michigan
Ann Arbor, Michigan 48109-2104 USA
edlarsen@umich.edu

(a) ABSTRACT

The objective of this project is to develop, implement, and test new deterministic methods to solve, as efficiently as possible, multigroup neutron transport problems having an extremely large number of groups. Our approach was to (i) use the standard CMFD method to "coarsen" the space-angle grid, yielding a multigroup diffusion equation, and (ii) use a new multigrid-in-space-and-energy technique to efficiently solve the multigroup diffusion problem. The overall strategy of (i) how to coarsen the spatial and energy grids, and (ii) how to navigate through the various grids, has the goal of minimizing the overall computational effort. This approach yields not only the fine-grid solution, but also coarse-group flux-weighted cross sections that can be used for other related problems.

(b) OBJECTIVES AND ACCOMPLISHMENTS

The most difficult and unresolved aspect of deterministic neutron transport calculations is undoubtedly the acquisition of practical multigroup cross sections. All other aspects of deterministic neutron transport problems (discretizing in space, angle, and time; solving the discrete equations) have been extensively studied during the past 60 years, and have become increasingly manageable. However, this is not the case for discretizing in energy. The reason is that the exact energy-dependent cross sections $\Sigma(E)$ are extremely "oscillatory" functions of energy E, while the multigroup approximation is based on a simple histogram function of E.
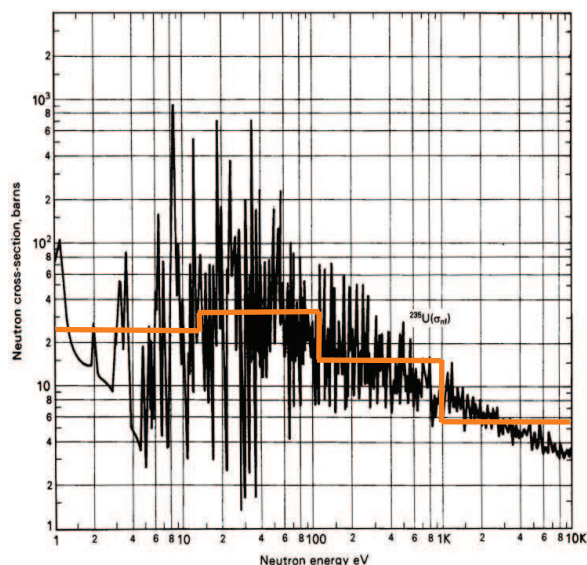

Figure 1: Continuous-Energy Cross Section (Black)
and a Hypothetical Multigroup Cross Section (Orange)

Thus, it is not possible to accurately model $\Sigma(E)$ unless an *extraordinarily* large number of energy groups is used – on the order of 10^6. However, practical problems require cross section sets with a much smaller number of energy groups – usually between 10 and 100, and sometimes less than 10.

To obtain "accurate" few-group cross sections – that yield useful results for practical problems – it is necessary for an experienced neutron data specialist to develop them. The process by which this is done is based on the experience and intuition of the specialist, and the results are variable. (Different experts will generate different multigroup cross sections.) Because the process of acquiring multigroup cross sections requires human intuition and judgment, it has not been automated. This crucial element of neutron transport simulations remains dependent on human experience and intuition – just as it did 50 years ago. To emphasize the importance of these multigroup cross sections, it is well-known that no matter how accurately the Boltzmann transport equation is discretized in angle and space, if "accurate" multigroup cross sections are not provided, the solution of the multigroup equations will not be useful.

The goal of this project was to make a major dent in this very old problem. The impetus for the project was based on the following observations.

1. The fewer the number of energy groups in a multigroup problem, the more "carefully-taylored" the cross sections must be to obtain an accurate solution. (The more important is the human element in deriving the cross sections.) Conversely, as the number of groups increases, the histogram-dependence of the multigroup cross sections begins to be able to accurately model $\Sigma(E)$, and the input of the data specialist becomes less important. In the limit as the number of groups tends to infinity, $\Sigma(E)$ is modeled accurately, and no input from the data specialist is required.

2. Therefore, it will be possible to define, for each material of interest to nuclear engineers, a multigroup cross section set, with a sufficiently large number of groups, that this cross section set would be sufficiently accurate for a large range of problems of physical interest. The requisite number of groups G is likely very large: probably on the order of 100,000, or more. This number of groups is 4 to 5 orders of magnitude larger than multigroup cross sections used for routine calculations. Nonetheless, the assumption is that if problems with these multigroup cross section sets were solved, then there would be no need for data specialists to generate few-group cross section sets heuristically. The results of simulations would then (i) be no longer dependent on the skill of the data analyst, and (ii) become much more reliable and predictive.

3. This NEUP project was aimed at the mathematical problem of deriving an iterative method for solving, as efficiently as possible, multigroup transport problems with on the order of 100,000 groups. If such a method could be formulated – especially one whose cost did not increase linearly with the number of groups G – it would provide an impetus to the possibility of solving neutron transport problems in a completely computer-oriented manner, without the need for human intervention.

The work performed on this project tackled this problem with a *multigrid* approach, in which the solution of the "high-order" transport problem is iteratively obtained by using a nested sequence of "low-order" problems defined on increasingly "coarse" space-angle-energy grids. The overall strategy is depicted in Figure 2:
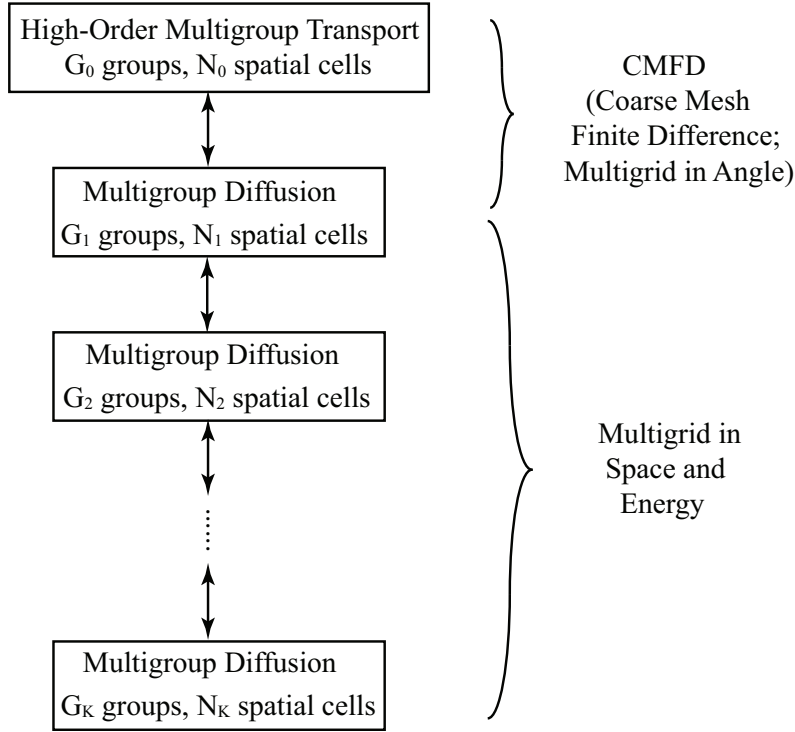

Figure 2: Multigrid Strategy in Angle, Space, and Energy

The "high-order" transport problem, with G_0 energy groups and N_0 spatial cells, is the problem that one ultimately desires to solve. This equation can be solved iteratively using conventional transport sweeps. However, for problems of typical interest, each high-order sweep would be very costly, and the number of sweeps required to converge would be very high; consequently a simple sweep-based iteration method would be unacceptably costly.

However, using the well-known CMFD method developed more than 20 years ago, the iterative solution of the angle-dependent transport equation can be obtained much more efficiently by alternating transport sweeps with the calculation of a "low-order" angle-independent multigroup diffusion equation, with G_1 groups and N_1 spatial cells, where $G_1 \leq G_0$ and $N_1 \leq N_0$. (This finest-grid diffusion equation can but does not have to coarsen in space and energy – but it definitely coarsens in angle, since the diffusion equation has no angle-dependence. Effectively, this step can be understood as a multigrid step in which the low-order problem retains the spatial and energy complexity, but it eliminates the angular complexity.)

Next, the finest-grid multigroup diffusion problem, with G_1 energy groups and N_1 spatial cells, is itself solved by a multigrid-in-space-and-energy method. A sequence of K increasingly coarse multigroup cross section sets, having

$$G_1 > G_2 > \cdots > G_K$$

groups is formulated. Also, a sequence of K increasingly coarse spatial grids, having

$$N_1 > N_2 > \cdots > N_K$$

cells, is formulated. As depicted in Figure 2, the estimate of the solution of the multigroup diffusion equation on the first coarse space-energy grid ($k = 1$) is accelerated using the estimate of the solution of the multigroup diffusion equation on the next coarser space-energy grid ($k = 2$), and this process continues until the coarsest $k = K$ space-energy grid, where the equation is solved exactly.

This is perhaps the place to emphasize that during the past 50 years, the multigrid method has become one of the standard iterative tools for iteratively solving elliptic partial differential equations – in particular, diffusion equations. The method was first developed in the mathematics community for spatial diffusion problems. For these, the multigrid method has no real competition – it is by far the most efficient iterative solution method known. In the 1980's, the multigrid method was applied in the nuclear engineering community for solving standard few-group diffusion problems. In these applications, the emphasis was on multigrid-in-space, not on multigrid-in-energy. (It was assumed that "accurate" few-group cross sections were provided, and the emphasis was on solving geometrically large systems with a large number of spatial cells.)

In our work, we assumed that the number of fine-grid spatial cells (N_0) and the number of fine-grid energy groups (N_0) are both very large, and we employed a general multigrid approach in which the fine-grid space-energy grids can be coarsened in arbitrary ways – individually in space-or-energy, or simultaneously in space-and-energy.

It is now possible to list and discuss the specific objectives and accomplishments of the project; we do this next.

1. GENERATION OF MULTIGROUP CROSS SECTION SETS

   The first task was to develop the ability to generate accurate multigroup cross section sets with a very large number of groups. It was not necessary that these cross sections be optimized for any specific class of problems, but we did want the cross sections to be reasonably accurate, and to be obtained in a recognizable way. Graduate student Eric Baker (who did all the implementational work on this project) used NJOY for this purpose. Eric developed a script that directed NJOY to construct simple "$1/E$"-weighted cross sections on equal-lethargy energy grids of arbitrary size (up to 10,000 groups) and format them in a convenient manner. The familiarization with NJOY and the writing of the script took about two months.

2. CONSTRUCTION OF THE 1D TEST CODE

   At the same time, Eric was formulating the blueprint for his 1-D planar geometry multigroup test code, which would read and utilize the NJOY-generated cross sections. His code was planned to utilize the standard 1D Gauss-Legendre quadrature sets up to S$_{32}$, and the standard Step Characteristic spatial differencing scheme. It allowed for very general geometric configurations, and for vacuum and reflecting boundary conditions.

   Also, Eric's code allowed the user to specify (i) the coarse grids i.e. the sequences $G_0, G_1, G_2, ..., G_K$ and $N_0, N_1, N_2, ..., N_K$, on which a specified multigrid strategy would be performed, (ii) the operations that should be done on each grid, and (iii) how the computer should navigate from one grid to the next.

   The construction of this basic code took several months. However, as we ran it, contemplated the results, and new ideas came to us, it is correct to say that this code has been continuously under construction or revision since the early stages of the project.

3. CONSTRUCTION OF THE 2D TEST CODE

   Since the beginning of the project, we had planned to test our methodology in 2D as well as 1D, so Eric also made plans for a separate 2D code, using (for simplicity) a Cartesian spatial grid. This code would be based on the discrete-ordinates Step Characteristic scheme, and would employ high-order

6

angular quadrature sets. Eric's two codes were written in C++, and much of the planning of the codes was done with the idea that as much of the 1D code as possible would be used within the 2D code. As a result, the development of the 2D code has consistently lagged that of the 1D code. Also, we found that using the 1D code was simpler, and that due to memory constraints it was easier for us to run 1D problems than 2D problems. Therefore, perhaps not surprisingly, we have used the 1D test code much more than the 2D test code.

4. LINEAR VS NONLINEAR MULTIGRID

When the spatial multigrid strategy was formulated about 50 years ago, a particular *linear* form of this method was developed in which the solution of a problem posed on a coarse grid is a coarse-grid *correction* to the solution on the next finer grid. Therefore, as this multigrid process converges, all the solutions on all the coarse grids limit to zero.

However, somewhat later, a different *nonlinear* form of multigrid (CMFD) was formulated in the nuclear engineering community in which the solution of a problem posed on a coarse grid is the exact voleme-integrated solution of the problem posed on the fine grid. (CMFD was formulated as a 2-grid method, with one fine and one coarse grid, but the method can be expanded to include an arbitrary number of coarse grids, in which case it acts very much like the classic linear multigrid method.) In order to ensure that the coarse-grid solutions converge to the exact volume integrals of the fine-grid solution, the cross sections on the coarse grid must be expressed as proper flux-weighted versions of the cross sections given on the fine grid. Since these are not known at the outset, they must be calculated and iterated, along with the coarse and fine-grid solutions.

In practice, CMFD has been used as a 2-grid method for accelerating the convergence of a transport problem posed on an unstructured fine grid using a low-order diffusion problem posed on a coarser Cartesian grid. In typical light water reactors, the geometric structure of the core consists of an array of 200 to 300 fuel assemblies, in which each assembly is an array of roughly 17×17 fuel pins, and each fuel pin consists of a fuel region, surrounded by cladding, surrounded by water. Figure 3 depicts a single fuel pin (which constitutes a single Cartesian coarse cell), decomposed into five radial zones and 8 angular zones, creating 40 fine spatial cells, on which the high-order transport equation is to be solved:
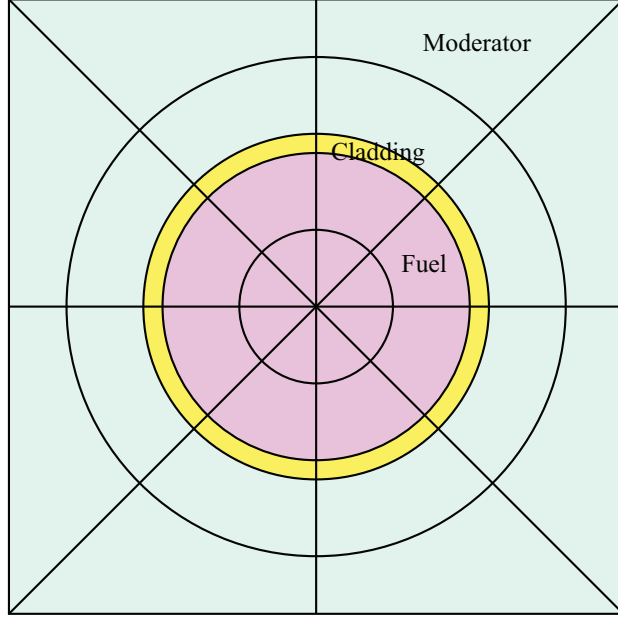
Figure 3: A Fuel Pin, with a Typical Unstructured Fine Spatial Grid

The number of fine spatial cells per coarse cell in this figure (40) is typical. In practical applications of CMFD, only the top two grids depicted in Figure 2 exist. The "low-order" diffusion equation is typically solved by an algebraic matrix-solver, which works efficiently if the number of energy groups is not too large.

The basic point is that a multigrid strategy can be formulated either linearly, or nonlinearly. If linear, the solutions on a coarse grid are corrections to the solution on the next finer grid. If nonlinear, the solutions on a coarse grid are volume-integrals (coarse-grid representations) of the solution on the next finer grid. In the early stages of our work, we implemented *both* methods, to see if either method was advantageous. The linear method has the advantage that the cross sections on the coarser grids are simple volume averages of the cross sections on the finer grids. Thus, the algebra to implement the linear method is less. However, the nonlinear method uses flux-weighted cross sections, which although are more costly to calculate and must be updated as the iterations progress, nonetheless preserve the physics much more precisely between the various grids.

After considerable testing in 1D, we found that the nonlinear method was consistently advantageous. Although the amount of algebra per iteration required by the nonlinear method was higher, the total number of iterations was sufficiently reduced that this method was unambiguously beneficial. Other

benefits to this approach became apparent later, and are discussed below.

5. 1D IMPLEMENTATION STRATEGIES (V-CYCLES)

Although optimized strategies of how to perform multigrid in space and energy are still unsettled, Eric has shown that the following multigrid strategy works well for a wide variety of problems. To describe some typical results, consider:

(a) The high-order problem consists of the discrete-ordinates transport problem, formulated with G_0 energy groups and N_0 spatial cells.

(b) The first coarse grid problem is the CMFD multigroup diffusion problem, formulated with $G_1 = G_0$ energy groups and $N_1 = N_0/8$ spatial cells. (The problem on this grid consists of a collapse in the angular variables and space, but not in energy.)

(c) Next, a sequence of increasingly coarse grids is defined, first with $G_2 = G_0/2$ groups and $N_2 = N_0/8$ cells, then with $G_3 = G_0/4$ groups and $N_3 = N_0/8$ cells, then with $G_4 = G_0/8$ groups and $N_4 = N_0/8$ cells, and finally 2 groups and $N_0/8$ cells. In this process, the energy groups are coarsened by a factor of 2 and the spatial grids are held constant. (The coarsening is modified when the number of groups is not divisible by 2.)

(d) Next, a sequence of increasingly coarse grids is considered, first with 2 groups and $N_0/16$ cells, then with 2 groups and $N_0/32$ cells, and finally 2 groups and 2 cells. In this process, the energy groups are held constant at 2 and the spatial grids are coarsened by a factor of 2. (Again, the coarsening is modified when the number of spatial grids is not divisible by 2.)

(e) The coarsest group problem is solved explicitly.

Given this setup, we define a single *V-cycle* to consist of performing one transport sweep (a), followed by Gauss Seidel calculations on each of the coarser grids (b), (c), (d), (e), and then followed by a prolongation from the coarser to the finer grids. Related strategies, e.g. coarsening or refining by other than a factor of 2 are possible, but have not been shown to be advantageous. However, replacing steps (c) and (d) above by the following can be advantageous, because it reduces the number of unknowns in the coarser grids more rapidly, and it also reduces the overall number of coarser grids:

(c', d') Next, a sequence of increasingly coarse grids is considered, first with G/2 groups and N/16 cells, then with G/8 groups and N/32 cells, then with

G/16 groups and N/62 cells, and finally with 2 groups and 2 cells. In this process, the energy groups and the spatial cells are simultaneously coarsened by a factor of 2. (Again, the coarsening must be altered when the number of groups or cells is not divisible by 2.)

We call the multigrid method in which a V-cycle consists of steps (a)-(e) the M_1 method. The multigrid method in which a V-cycle consists of steps (a), (b), (c', d'), (e) is denoted as the M_2 method.

We have generated results for a large 1-D problem representing a reactor core of width 306 cm, having 17 assemblies (15 fuel + 2 reflectors), 16 pin cells per assembly, 24 spatial cells per pin cell, a total of 6,528 spatial cells, and different fuel and moderator regions containing water, U02, Mox fuel, Gd fuel, and other materials. Results were obtained for:

(A) A standard S_N sweep accelerated by a multigroup CMFD diffusion calculation solved by line relaxation in space.

(B) A standard S_N sweep accelerated by a multigroup CMFD diffusion calculation, which in turn is accelerated by a 2-group CMFD diffusion calculation.

(C) The multigrid M_1 method defined above. (In the CMFD calculation, energy is coarsened before space.)

(D) The multigrid M_2 method defined above. (In the CMFD calculation, energy and space are coarsened simultaneously.)

For each method, the problem was solved with $G = 50$, 187, and 640 groups. The number of seconds required to converge each problem is given below:

G	Method A	Method B	Method C	Method D
50	81	93	75	86
187	175	156	137	132
640	1825	1434	1370	1021

Table 1: Timings (Sec) for Four Iteration Strategies

Unfortunately, the gains seen by the multigrid methods are modest. Investigation revealed that this is because for the given problem, the time required to calculate the solution is not dominated by the time required to solve the multigroup CMFD diffusion calculation. The multigrid-in-space-and-energy strategy does efficiently accelerate the convergence of the CMFD equation. However,

since the cost of solving this equation does not dominate the solution process, the overall gain in run times for large G are not major.

This result informed us that the cost of performing the high-order transport sweeps was dominant, and that a more effective iteration strategy would be one that reduced the number of high-order transport sweeps. This led to the ideas described next.

6. CMFD ACCELERATION OF WITHIN-GROUP SCATTERING

In an attempt to reduce the cost of the high-order transport sweeps, we included – at the conclusion of each sweep in each energy group – a monoenergetic CMFD calculation to accelerate the convergence of the within-group scattering source. Unfortunately, this had little benefit. The reason is that within-group scattering is significant only for relatively course groups; it becomes increasingly less important as the number of groups increases. Therefore, accelerating the convergence of this term has little effect on problems with large numbers of groups; the cost of implementing this technique is not outweighed by a reduction in the overall number of iterations.

7. A "BOTTOM-UP" APPROACH

This has sometimes been called the "full multigrid" approach. The idea is to minimize the total number of high-order transport sweeps by first calculating the best possible initial guess for the high-order sweep, using information from the coarser grids. For example, the best possible initial guess for the high-order transport calculation on the finest (G_0, N_0) grid would arguably be the solution of the transport calculation on the first coarse (G_1, N_1) grid; the best initial guess for that solution would be the solution of the transport calculation on the next coarse (G_2, N_2) grid; and so on.

We spent considerable time implementing this approach, but unfortunately, it did not perform well. The apparent reason for this is that when the calculation starts on a coarse grid, it is not possible to use the corect flux-weighted cross sections from finer grids, because calculations on these grids have not been done yet. Therefore, the coarse grid solutions are less accurate than they would be otherwise. These solutions do become more accurate as the energy grid is refined, because as the number of energy groups increases, the flux-weighting becomes less important. In the end, this process does produce a more accurate initial guess for the high-order transport calculation. Unfortunately, the

extra expense associated with this "initializing" calculation does not lead to a correspondingly large reduction in the number of required high-order transport iterations.

8. HIGH-ORDER MULTIGROUP DIFFUSION INITIAL GUESS

An alternate approach is instead, to initialize the high-order (G_0, N_0) transport calculation with the solution of the corresponding high-order (G_0, N_0) diffusion calculation. In this case, the exact high-order cross sections are used, and the issue of not having flux-weighted cross sections on the coarser grids is not present. Nonetheless, this method also did not give meaningfully benficial results.

(c) PUBLICATIONS

No publications have yet been generated by this work. The only document is a work-in-progress, the PhD thesis of student Eric Baker. We anticipate that this thesis will be written during the coming summer and fall, and that Eric will defend sometime during the fall 2013 semester.

(d) OTHER RELATED INFORMATION

When the NEUP funding for this project ended in December 2012, CASL agreed to provide funding until the end of August 31, 2013. CASL was interested in the project because of its goal of reducing or eliminating the "Achilles Heel" of deterministic transport calculations – the determination of multigroup cross section sets. Unfortunately, our results to date are not particularly exciting. Nonetheless, certain elements of our work, which were not part of our original research goal, are definitely of interest. In particular:

1. In practical multigroup problems run today, an attempt is often made to accelerate convergence by using a two-stage CMFD (essentially, a multigrid) approach in which:

    - The convergence of the high-order multigroup S_N problem is accelerated by solving a multigrid diffusion problem, and:
    - The convergence of the multigrid diffusion problem is accelerated by solving a 2-group diffusion problem. (The two groups correspond to fast neutrons, and slow neutrons.)

In effect this method corresponds to our Figure 2, in which there are two course grids, the coarsest one having 2 energy groups.

Much of the time this strategy works, and the solution is obtained more efficiently. However, it often happens that this procedure is unstable (it does not converge). In this case the only recourse is to solve the problem using the original CMFD method, which has only the top two grids in Figure 2 and lacks the 2-group calculation.

The work done in this thesis shows, at least experimentally, that problems of this latter type can always be accelerated by using a sequence of energy grids that collapses *more slowly* from G_0 to 2. Experimentally, we have found that collapsing by a factor of 4 works nicely. For example, if $G_0 = 32$, then instead of collapsing from 32 groups directly to 2, which is sometimes unstable, we recommend collapsing from 32 to 8, and then from 8 to 2. The idea is simple, but it works.

2. Another important issue is that the nonlinear multigrid process employed in our work produces not only the high-order transport solution (albeit, more slowly than we had hoped), but it also automatically produces coarse grid flux-weighted cross sections on each of the coarse space-energy grids. This, perhaps, is more important than obtaining the high-order transport solution (the stated goal of this project). This leads to our final observations:

3. If it turns out to be impossible to efficiently solve transport problems with 100,000 energy groups, the problem of deriving "accurate" few-group cross sections will remain – because these kinds of cross sections are inevitably the ones that are now, and will be in the future, used for most practical applications. It may well be that a redirection of this project – from obtaining the high-order transport solution as efficiently as possible, to obtaining "accurate" few-group cross sections as efficiently as possible – would lead to a higher chance of success. The reason: if one has the high-order transport solution, then one can calculate any type of coarse-grid flux-weighted cross section using this exact solution. However, it may be that to calculate accurate few-group cross sections, it is not necessary to calculate the high-order transport solution.

For example, in the "bottom-up" approach described above, one deliberately begins a calculation on a coarse grid, using coarse multigroup cross sections that are not flux-weighted. The diffusion solution on this grid is obtained, and is used to initiate the iteration process on the next finer grid. After the solution

on this second grid has been obtained, this solution is used to initiate the iteration process on the third grid; and so on. Suppose that at the conclusion of the calculation on the second grid, and on the third grid, etc, one has the S_N solution on that grid. That solution could be used to generate flux-weighted cross sections for the coarsest grid. Therefore, at the end of the calculation on each increasingly fine grid, one has new flux-weighted cross sections on the coarsest grid. Eventually, these coarsest-grid cross sections will converge. When they do, it will not be necessary to perform calculations on any of the finer space-energy grids. Moreover, if the convergence criterion on the coarse grid cross sections is not too extreme, it is likely that the coarsest-grid cross sections will have converged before the finest grid reaches 100,000 groups.

This would constitute a reconfiguration of the stated primary goal of the project – but in a way that directly addresses the long-term problem of calculating few-group cross section sets. Even if this process were to work effectively, other details would have to be considered. For instance, in any problem defined on a spatial grid, the cross sections that are generated will vary slightly from one spatial cell to another within each material region. How should the resulting multigroup cross sections be "homogenized" across each specified region (where a single set of flux-weighted cross sections are desired)? Although this problem may not be simple, it seem much more tractable than the current problem of generating multigroup cross sections by hand.

In conclusion, it must be admitted that the results of this project have thus far been disappointing. However, the problem that we set out to solve is very difficult. Part of this difficulty is reflected in the fact that multigroup cross sections are obtained today in the same way they were obtained 50 years ago. Nonetheless, we have learned a great deal about the problem, and although the funding and the time for this project have ended, there still are ideas from this work that may prove fruitful in the future. If possible, we will investigate at least some of them before Eric defends his PhD thesis.