HIERARCHICAL ANGULAR PRECONDITIONING FOR THE FINITE ELEMENT-SPHERICAL HARMONICS RADIATION TRANSPORT METHOD

C. R. E. de Oliveira, C. C. Pain and M. D. Eaton
Computational Physics and Geophysics Group
T. H. Huxley School of the Environment, Earth Sciences and Engineering,
Imperial College of Science, Technology and Medicine,
Prince Consort Rd, London SW7 2BP, UK
r.transport@ic.ac.uk

ABSTRACT

In this paper we describe the implementation of a hierarchical angular preconditioner for the solution of the finite element-spherical harmonics equations which is akin to the well-known multi-grid method. The theoretical details of the preconditioner are presented and its performance using the standard multi-grid cycles such as V, W and F is assessed by solving a 3D benchmark problem.

1. INTRODUCTION

The finite element-spherical harmonics (FE-P_N) method has developed into a powerful and versatile alternative to established radiation transport methods such as discrete-ordinates and Monte Carlo. A general multigroup anisotropic scattering P_N formalism[1] coupled with a high-degree of geometric flexibility enables the method to address realistic steady-state and time-dependent multi-dimensional radiation transport problems. The method has been incorporated, in its most general form, into the radiation transport code EVENT[2], which has been adopted by industry and is currently being applied to a wide range of problems. These include reactor core calculations, shielding, criticality safety assessment, radiative transfer in the atmosphere and optical tomography.

The radiation FE-P_N method is characterised by the explicit assembly and solution of of a coupled system of linear equations for the nodal angular unknowns. This system though symmetric and sparse can be very large for realistic 3D problems. To tackle this, a general and economical moment-by-moment (MBM) solution scheme has been developed based on a preconditioned conjugate gradient (PCG) iterative procedure. In this scheme neither the global matrix nor the preconditioning matrix need be explicitly calculated and stored. The scheme readily extends to

parallel processing via domain-decomposition methods[3].

The computational efficiency of the MBM PCG procedure is highly dependent on the choice of preconditioner. The MBM scheme renders itself naturally to hierarchical or multilevel preconditioning which are a powerful class of methods that are rapidly gaining widespread acceptance for the solution of large systems of linear equations. They have good convergence properties and are natural candidates for spherical harmonic discretizations.

In this paper, we describe the implementation of a hierarchical angular preconditioner which is akin to the well-known multigrid method. The central idea behind the method consists in prolongation and restriction of the P_N angular approximation so as to build coarse and fine angular "grid" from which solutions can be automatically interpolated. The P_1 (diffusion) approximation, which is the coarsest level, is solved exactly but other finer levels of angular approximation are not solved exactly; instead some partial solution is obtained - this closely corresponds to diffusion synthetic acceleration.

This approach differs from the more conventional multigrid approaches in that it uses a series of hierarchical solutions on coarse and fine grids to obtain the eventual solution, typically using some relaxation scheme on each grid level which is good at resolving the high frequency parts of the solution. Again typically, an exact solution is obtained on the coarsest level.

2. THE FINITE ELEMENT-SPHERICAL HARMONICS METHOD

The finite element-spherical harmonics method, which is the subject of this paper, is based on the second-order, even-parity form of the Boltzmann transport equation which can be written as[4]

$$-\Omega.\nabla\mathcal{G}\Omega.\nabla\psi+\mathcal{C}\psi=S^{+}-\Omega.\nabla\mathcal{G}S^{-} \tag{1}$$

where

$$\psi(\mathbf{r},\Omega)=\frac{1}{2}[\phi(\mathbf{r},\Omega)+\phi(\mathbf{r},-\Omega)]$$

$$S^{\pm}(\mathbf{r},\Omega)=\frac{1}{2}[S(\mathbf{r},\Omega)\pm[S(\mathbf{r},-\Omega)]$$

are the even-parity angular components of the flux $\phi(\mathbf{r},\Omega)$ and the parity components of the prescribed source $S(\mathbf{r},\Omega)$ respectively, and $\mathcal{C}$ and $\mathcal{G}$ are collision operators defined by

$$\mathcal{C}\phi(\mathbf{r},\Omega)\equiv\sigma_{\ell}\phi(\mathbf{r},\Omega)-\int_{4\pi}\mathbf{d}\Omega'\sigma_{\mathbf{s}}^{+}(\mathbf{r},\Omega.\Omega')\phi(\mathbf{r},\Omega')$$

$$\mathcal{G}^{-1}\phi(\mathbf{r},\Omega)\equiv\sigma_{\ell}\phi(\mathbf{r},\Omega)-\int_{4\pi}d\Omega'\sigma_{s}^{-}(\mathbf{r},\Omega.\Omega')\phi(\mathbf{r},\Omega')$$

where $\sigma_{\ell}=\sigma_{t}-\sigma_{s\ell}$.

Solution of Eq. (1) subject to the bare surface (vacuum) boundary condition

$$\phi(\mathbf{r},\Omega) = 0, \quad \text{for} \quad \Omega.\mathbf{n} < 0$$

where n is outward normal to the surface of the domain, is equivalent to finding the function which minimises the quadratic functional

$$\mathcal{F}[\psi] = (\Omega\psi, \mathcal{G}\Omega.\nabla\psi) + (\psi, \mathcal{H}\psi) + <\psi,\psi> -2(\psi, S^{+}) - 2(\Omega.\nabla\psi, \mathcal{G}S^{-}) \tag{2}$$

where $(.,.) \equiv \int dV \int d\Omega$ and $<.,.> \equiv \int dS \int d\Omega |\Omega.\mathbf{n}|$. Hence, numerical solutions are possible via the well-known Ritz-Galerkin procedure,

The spatial domain is divided into E non-overlapping sub-domains or elements connected together at $\mathcal{N}$ nodal points. The trial function for the even-parity flux is given the form

$$\psi(\mathbf{r},\Omega) = \sum_{e=1}^{E} \sum_{j=1}^{\mathcal{N}_e} \sum_{k=1}^{M} B_j^e(\mathbf{r}) Q_k(\Omega) \psi_{jk}^e \tag{3}$$

or, in a more compact tensor notation:

$$\psi(\mathbf{r},\Omega) = \sum_{e=1}^{E} \mathbf{B}^{eT} \otimes \mathbf{Q}^T \psi^e \tag{4}$$

where ψ^e is a $(M\mathcal{N}_e \times 1)$ column vector of unknown nodal fluxes, $\mathbf{B}^{eT}(\mathbf{r})$ is a $(\mathcal{N}_e \times 1)$ column vector of finite element shape functions, and $\mathbf{Q}(\Omega)$ is a $(M \times 1)$ column vector of normalised spherical harmonic functions, given by :

$$Q_{k(\ell m)} = \sqrt{\frac{2\ell+1}{4\pi}(2-\delta_{m,0})\frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(\mu) \left\{ \begin{array}{c} \cos(m\chi) \\ \sin(m\chi) \end{array} \right\}$$

for ℓ even and $m = 0, 1, \ldots, \ell$. The series is truncated at $\ell = N-1$ and the approximation is denoted by P_N.

Substitution of expression (4) into (2) yields the discretized functional

$$\mathcal{F}[\psi] = \sum_{e=1}^{E} \psi^{eT} \mathbf{A}^e \psi^e - 2\psi^T \mathbf{b}^e \tag{5}$$

where

$$\begin{aligned} \mathbf{A}^e \equiv\ & (\Omega.\nabla\mathbf{B}^{eT} \otimes \mathbf{Q}^T, \mathcal{G}\Omega.\nabla\mathbf{B}^{eT} \otimes \mathbf{Q}^T) + (\mathbf{B}^{eT} \otimes \mathbf{Q}^T, \mathcal{H}\mathbf{B}^{eT} \otimes \mathbf{Q}^T) \\ & + <\mathbf{B}^{e^T} \otimes \mathbf{Q}^T, \mathbf{B}^{eT} \otimes \mathbf{Q}^T> \end{aligned}$$

$$\mathbf{b}^e \equiv (\mathbf{B}^{eT} \otimes \mathbf{Q}^T, S^{+}) + (\Omega.\nabla\mathbf{B}^{eT} \otimes \mathbf{Q}^T, \mathcal{G}\Omega.\nabla S^{-})$$

Making use of nodal compatibility, we can write the nodal moments of the flux as a $(M\mathcal{N} \times 1)$ column vector ψ, and the discretized functional takes the form:

$$\mathcal{F}[\psi] = \psi^T \mathbf{A}\psi - 2\psi^T \mathbf{b} \tag{6}$$

where

$$\mathbf{A} = \sum_{e=1}^{E} \mathbf{A}^e, \qquad \mathbf{b} = \sum_{e=1}^{E} \mathbf{b}^e$$

The summation over the elements is symbolic since the local and global quantities have different dimensions. It simply means adding the local coefficients into their proper locations in the global matrix and rhs vector.

By requiring the first variation of $\mathcal{F}$ to vanish, we obtain the matrix system for the unknown angular coefficients or moments at the nodal points:

$$\mathbf{A}\psi = \mathbf{b} \tag{7}$$

The framework above readily extends to energy and time-dependence via the multigroup approximation and time-marching schemes respectively.

3. THE MOMENT-BY-MOMENT ITERATIVE SOLUTION SCHEME

Solution of the system of equations (7) can be obtained via standard direct or iterative numerical linear solvers. For 3D problems, solution is perforce iterative due to storage demands and operations count. The FE-P_N equations render themselves quite naturally to iterative solution schemes. The Ritz-galerkin formulation above gives rise to symmetric positive-definite (SPD) matrices which, due to the local or compact support of the finite element basis functions, have a sparse structure. Iterative solution techniques such as the preconditioned conjugate gradients (PCG) method can exploit these features very effectively to keep storage demands and computational effort to a minimum.

An iterative strategy for solution of the FE-P_N equations is to iteratively sweep through the moments as follows: the matrix $\mathbf{A}$ is partitioned into $M \times M$ submatrices of dimension $\mathcal{N} \times \mathcal{N}$:

$$\mathbf{A} = \begin{pmatrix} \mathbf{A_{11}} & \cdots & \mathbf{A_{1}}_M \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{M\mathbf{1}} & \cdots & \mathbf{A}_{MM} \end{pmatrix} \tag{8}$$

and this partitioned form is used in the following iterative moment-by-moment preconditioned conjugate gradient (PCG) solution procedure:

$$
\begin{aligned}
\mathbf{r}^0 &= \mathbf{b} - \mathbf{A}\psi^0 \\
\mathcal{A}\mathbf{z}^0 &= \mathbf{r}^0 \\
\mathbf{p}^0 &= \mathbf{z}^0 \\
&\left[\begin{aligned}
\alpha^n &= \frac{\mathbf{z}^n \cdot \mathbf{r}^n}{\mathbf{p}^n \cdot \mathbf{A}\mathbf{p}^n} \\
\mathbf{r}^{n+1} &= \mathbf{r}^n - \alpha^n \mathbf{A}\mathbf{p}^n \\
\psi^{n+1} &= \psi^n + \alpha^n \mathbf{p}^n \\
\mathcal{A}\mathbf{z}^{n+1} &= \mathbf{r}^{n+1} \\
\beta^{n+1} &= \frac{\mathbf{z}^{n+1} \cdot \mathbf{r}^{n+1}}{\mathbf{z}^n \cdot \mathbf{r}^n} \\
\mathbf{p}^{n+1} &= \mathbf{z}^{n+1} + \beta^{n+1}\mathbf{p}^n
\end{aligned}\right.
\end{aligned}
$$

where **r** is the residual vector, **z** and **p** search direction vectors, and α and β are weighting and acceleration factors respectively.

In the scheme above, $\mathcal{A}$ is the preconditioning matrix, the choice of which is crucial to the efficiency of the algorithm. It should be as close approximation to **A** as possible, but simple enough so that its numerical "inversion" is computationaly cheap. It should also satisfy the SPD requirement. As an example, one may choose the preconditioner given by the expression:

$$\mathcal{A} = (\mathbf{D} + \mathbf{L})\mathbf{D}^{-1}(\mathbf{D} + \mathbf{L}^T) \approx \mathbf{A} \tag{9}$$

with

$$\mathbf{L}_{ij} = \begin{cases} \mathbf{A}_{ij} \text{ if moment } i > j; \\ 0 \quad \text{otherwise} \end{cases} \qquad \mathbf{D}_{ij} = \begin{cases} \mathbf{A}_{ij} \text{ if moment } i = j; \\ 0 \quad \text{otherwise} \end{cases} \tag{10}$$

For this choice of preconditioner, each iteration requires the solution of $2M$-1 matrix systems of order $\mathcal{N} \times \mathcal{N}$. Note that only the diagonal submatrices $\mathbf{D}_{ii}$ need to be explicitly assembled, but only one of these matrices is present in core at any given time. The remaining matrix operations involving the $\mathbf{L}_{ij}$ are performed in an element-by-element fashion by exploiting the tensorial properties of the global matrix **A**. This applies also to any global matrix vector multiplication involving **A**.

4. HIERARCHICAL ANGULAR PRECONDITIONING

The Spherical Harmonics functions constitute a hierarchical basis, and this allows the MBM solution scheme for the P_N equations to be formally studied in an algebraic multilevel preconditioning framework. This in turn, allows us to explore the relative merits of preconditioners in terms of

well-studied multigrid cycles. These have bcome very popular in recent years because of their convergence properties.

In a multilevel framework the matrix **A** can be written as:

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{pmatrix} = \begin{pmatrix} \mathbf{A}_{11} & \mathbf{0} \\ \mathbf{A}_{21} & \mathbf{C} \end{pmatrix} \begin{pmatrix} \mathbf{I} & {\mathbf{A}_{11}}^{-1}\mathbf{A}_{12} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \tag{11}$$

in which $\mathbf{A}_{22}$ is the matrix formed from the Spherical Harmonics basis functions of order one less than the one used to obtain the matrix **A**, and is thus the matrix formed on the coarser level (P_{N-2} approximation), and **C** is the Schur complement matrix $\mathbf{C} = \mathbf{A_{22}}$ - $\mathbf{A}_{21}\ \mathbf{A}_{11}^{-1}\ \mathbf{A}_{12}$. The matrix $\mathbf{A}_{11}$ contains the coupling between angular moments associated with the largest P_N shell. We will now attempt to obtain a preconditioner of **A** which is "optimal" (or nearly) in some sense.

The SPD preconditioner associated with the multilevel splitting of the matrix above is

$$\mathbf{M} = \begin{pmatrix} \mathbf{B} & \mathbf{0} \\ \mathbf{A}_{21} & \hat{\mathbf{A}}_{22} \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{B}^{-1}\mathbf{A}_{12} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \tag{12}$$

The matrix **B** is an approximation to $\mathbf{A}_{11}$ and is defined such that, for any vector **x**:

$$\mathbf{x}^T\mathbf{A}_{11}\mathbf{x} \leq \mathbf{x}^T\mathbf{B}\mathbf{x} \leq (1+\alpha)\mathbf{x}^T\mathbf{A}_{11}x \ \text{ for some } \ \alpha \geq 0 \tag{13}$$

which ensures optimal complexity. In other words, **B** maintains a block diagonal dominance in **M** (notice that $\mathbf{B} = \mathbf{A_{11}}$ satisfies this property). An example of **B** might be

$$\mathbf{B} = (1+w)(\mathbf{L}_{11}+\mathbf{D}_{11})^T{\mathbf{D}_{11}}^{-1}(\mathbf{L}_{11}+\mathbf{D}_{11}) \tag{14}$$

for some $w \geq 0$. In references 5 and 6 an incomplete Cholesky approximation to $\mathbf{A}_{11}$ scaled by a factor of 2 was used. These approximations to $\mathbf{A}_{11}$ when used in the preconditioner **M** under-relax the solution and thus often stabilise it. In the present context, a block diagonal or or a Backwards-Forwards-Backwards Gauss-Seidel (BFBGS) approximation would be more appropriate, for storage reasons. As pointed out in reference 6, $\mathbf{B}^{-1}$ multiplication can be viewed as a smoothing of fine correction terms and post multiplication by $\mathbf{A}_{12}$ as a prolongation operator interpolating coarse information to fine correction terms.

The approximation of **B** in the above algorithm can be improved when **B** is block diagonal. Suppose the matrix preconditioner is

$$\left.\begin{aligned} \mathbf{M} &= \begin{pmatrix} \mathbf{D_{11}}+\mathbf{L_{11}} & \mathbf{0} \\ \mathbf{A}_{21} & \hat{\mathbf{A}}_{22} \end{pmatrix} \begin{pmatrix} \mathbf{D_{11}^{-1}} & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{A}}_{22}^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{D_{11}}+{\mathbf{L_{11}}}^T & \mathbf{A}_{12} \\ \mathbf{0} & \hat{\mathbf{A}}_{22} \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{D_{11}}+\mathbf{L_{11}} & \mathbf{0} \\ \mathbf{A}_{21} & \hat{\mathbf{A}}_{22} \end{pmatrix} \begin{pmatrix} \mathbf{I}+{\mathbf{D_{11}}}^{-1}{\mathbf{L_{11}}}^T & {\mathbf{D_{11}}}^{-1}\mathbf{A}_{12} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \end{aligned}\right\} \tag{15}$$

which is SPD provided that $\hat{\mathbf{A}}_{22}$ is SPD. An approximation to $\hat{\mathbf{A}}_{22}$ is

$$\hat{\mathbf{A}}_{22}^{-1} = [\alpha_0\mathbf{I}+\alpha_1\mathbf{G}+\ldots+\alpha_{m-1}\mathbf{G}^{m-1}]\mathcal{M}^{-1} \tag{16}$$

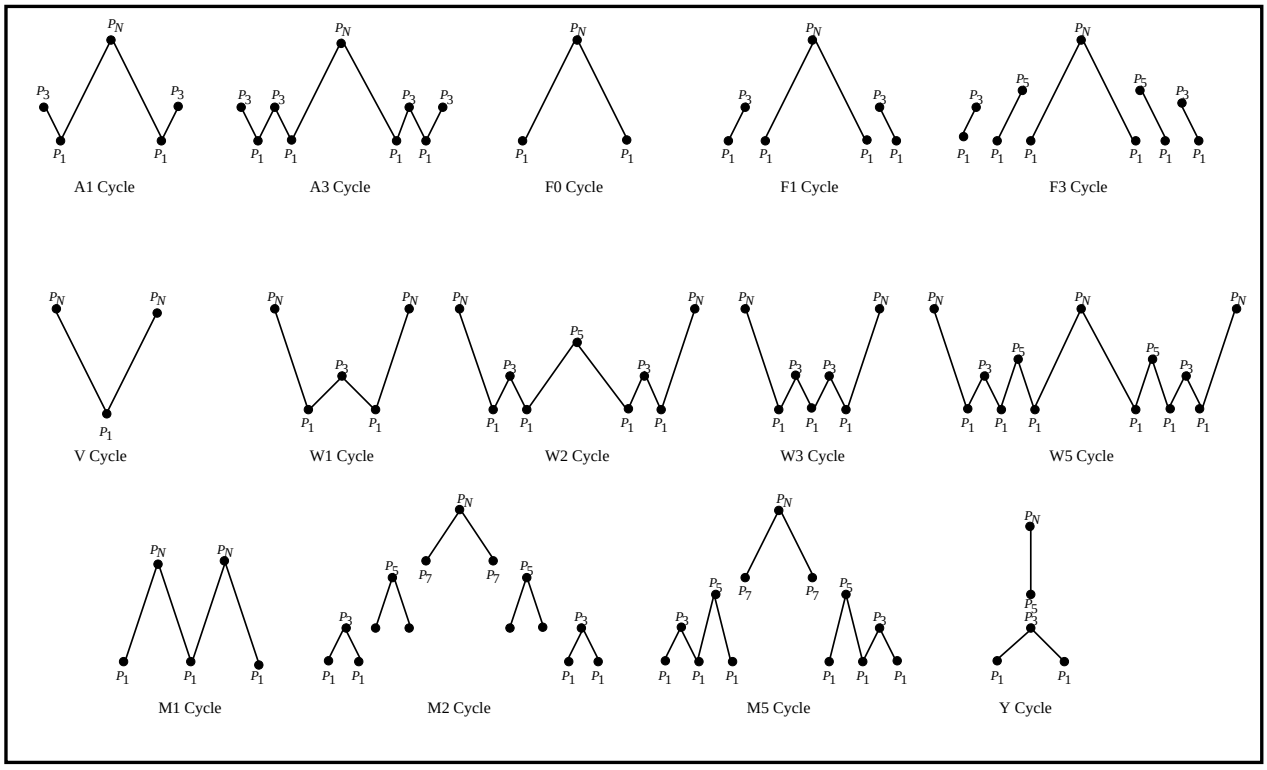

Figure 1: Multi-level P_N cycles

where α_i are chosen to make this approximation as close as possible to $\mathbf{A}_{22}$, and the iteration matrix $\mathbf{G}$ is defined by

$$\mathbf{G} = \mathcal{M}^{-1}(\mathcal{M} - \mathbf{A}_{22}) \tag{17}$$

The matrix $\mathcal{M}$ is obtained by a repeated application of Eq. (12) or when the coarsest level is reached, $\mathcal{M} = \mathbf{A}_{22}$.

Multiplication by $\hat{\mathbf{A}}_{22}^{-1}$ is effected by applying m parameterised BFBGS iterations starting from an initial guess of zero. Then recursively feeding new $\mathbf{x}^n$ into this equation up to the m^{th} iteration gives a solution:

$$\mathbf{x}^m = \hat{\mathbf{A}}_{22}^{-1} b \tag{18}$$

In practice, $\mathbf{x} = \hat{\mathbf{A}}_{22}^{-1}\mathbf{b}$ is obtained by Horners method[5] initialising with $\mathbf{x}^0 = 0$ for n=1 to m solve:

$$\mathcal{M}x^N = \mathbf{A}_{22}x^{n-1} + \alpha_{m-n}\mathbf{b} \tag{19}$$

The constants $\alpha_0, \ldots, \alpha_{m-1}$ are chosen so that the eigenvalues of $\hat{\mathbf{A}}_{22}^{-1}\mathbf{A}_{22}$ are positive and between the max and min eigenvalues of $\mathcal{M}^{-1}\mathbf{A}_{22}$, inclusive, and as close to unity, in a least squares sense, as possible.

W-cycle

The above algorithm as described is a subset of W-cycles (see Figure 1) and Eq. 15. For example the W1-cycle obtains an approximation $\hat{\mathbf{A}}_{\mathbf{11}}$ to $\mathbf{A}_{11}$ using two block forward backward Gauss Siedel iteration. A W-cycle can be viewed as an extension of Block Gauss Seidel iteration. For example, suppose we have a P_5 approximation then we can solve each of the higher moments in turn of the P_5 shell sweeping down through the moments using the information we have gained in a Gauss Seidel fashion. When we get to the P_3 solution we put the residual on the right hand side and then form an additional approximation to P_1 to P_3 shells. Once this approximation is formed then the result is placed on the right hand side and a Gauss Seidel sweep of the lower moments to the highest moments of the P_5 shell is performed. If a symmetric Gauss Seidel sweep was applied to the P_1 and P_3 shells and a forward sweep through this P_5 shell was performed then this would result in the V-cycle. However, it may be desirable to achieve a greater accuracy for the P_1 and P_3 shells, say, by using two FBGS iterations - this would result in a W-cycle. The particular W-cycle we have described is given in Figure 1 and denoted by W1-cycle but many other possibilities exist as shown in the figure. Often one would nest these cycles with say two iterations to P_5 and two FBGS iterations to P_3 see Figure 1.

Since more attention is paid in the W-cycle to the lower moments the method converges faster if the higher moments or fine correction terms of the spherical harmonics hierarchical basis are small which is the case for smooth angular functions. This method is guaranteed to produce a symmetric positive definite preconditioner from Eq. 15 as well as ensuring that both matrices D_{11} and $\hat{A}_{22}$ are symmetric-positive-definite. This does not preclude the possibility that under relaxation such as given by Eq. 14 may be beneficial e.g. multiplying the block diagonal approximation by a quantity greater than unity. However, this is not done here as we are not aware of appropriate theory to guide us. These methods are made more efficient by taking into account that the spherical shells are coupled in the global matrix by block tri-diagonal matrix coefficients.

F-cycle

In order to give the lower moments even more attention we can start the Gauss Seidel style iteration (or approximation) at the lower moments and finish on the lower moments. The result are cycles of the form seen in figure 1 and denoted by F1 and F3. F-cycle multi-grid methods tend to be amongst the most efficient multi-grid cycles. We have symmetrised the cycle for the benefit of the conjugate gradient iteration but this is usually not necessary and an F-cycle would finish once it has reached the highest level or finest correction terms.

Discussion

Once the basis is hierarchical, it is relatively easy, using the method described here, to use different cycles as the prolongation and restriction operators are already contained within the matrix of the equation system to be solved. If the basis is not hierarchical, for example if one is using a finite volume or finite element method, then one would convert the system, say FE system, from the standard nodal basis to a hierarchical basis[6]. Others have accelerated S_N calculations using multi-grid methods[7] (not in hierarchical basis) and have taken the usual approach of using an increasingly lower-order transport sweep, but in addition apply a diffusion synthetic acceration as the final step.

Finally, notice that the M-cycles (Figure 1) take their inspiration from Eq. 12 with a block FBGS approximation to matrix **B**.

Figure 1 depicts the various families of multi-level P_N cycles. The derivation of the W- and F-cycles has been discussed above. The M-cycle family arises from Eq. 12. The remainder of the cycles are hybrid schemes.

4. NUMERICAL RESULTS

In this section we compare the performance of the different multi-grid cycles for one of the Kobayashi 3-D deterministic transport benchmark problems[8]. The geometry of the problem is shown in Figure 2. Reflective boundary conditions apply on planes $x = 0, y = 0, z = 0$ and vacuum on all other boundaries. The media is homogeneous with total cross section of $0.1cm^{-1}$, and a distributed isotropic source with unit strength is located in the $10 \times 10 \times 10$cm small cube. The problem was analysed for two cases: purely absorbing medium and purely scattering medium. Both these cases were solved using EVENT with P_5 , P_7 and P_9 angular approximations and the MBM preconditioned conjugate gradient method using the cases shown in Figure 1. The number of outer conjugate gradient iterations and the number of inner iterations required in the solution of the two cases are presented in Tables 1 and 2, respectively.

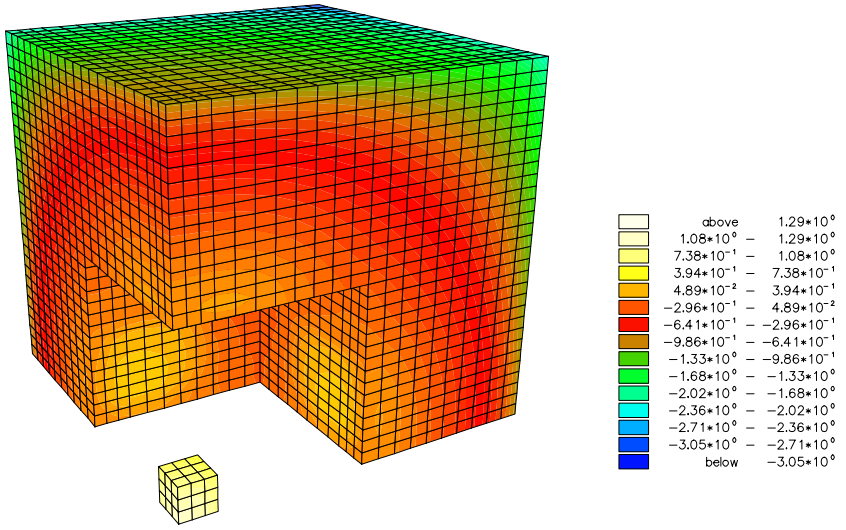

Figure 2: The domain and finite element mesh used to solve the Kobayashi problem. The scalar flux for the purely absorbing case is superimposed onto the mesh.

We see that the best cycle in terms of reducing the number of global PCG iterations seems to be the F-cycles followed by M- and W-cycles. However, the solution times are mainly determined by the number of inner iterations - solution for a single moment - and therefore, the W-cycles show a greater potential in this respect. The efficiency of these particular cycles will be improved with parameterisation and relaxation (see Eq. 14) - not included here for a fair comparison with other

methods. In addition, we can dispense the with the initial PCG (which is preconditioned with one FBGS iterations) and instead replace it with a number of FBGS iterations (probably parameterised) to provide a symmetric-positive-definite preconditioner which is efficient. Our initial experiments with this suggest relaxation is crucial to obtain a good preconditioner.

	No of external iterations			No of internal F/B G-S iterations		
cycle	P_5	P_7	P_9	P_5	P_7	P_9
A1	44	46	54	1716	2990	5346
A3	42	44	53	2730	4044	6625
Y	75	96	120	1500	3168	6000
M1	46	49	61	2622	5341	10797
M2	36	43	51	2808	4429	6987
M5	35	41	50	2730	4223	6850
V	55	54	64	1595	2970	5696
W1	44	49	52	1716	3185	5148
W2	35	44	55	2695	4532	6850
W3	44	49	51	2156	3675	5559
W5	35	40	47	3185	5200	8507
F0	50	51	58	1450	2805	5162
F1	42	44	53	1722	3948	5353
F3	34	38	49	2414	3686	6419

Table 1: Table showing the total number of internal and the number of external PCG iterations required to solve the Kobayashi problem with pure absorption.

	No of external iterations			No of internal F/B G-S iterations		
cycle	P_5	P_7	P_9	P_5	P_7	P_9
V	138	198	262	4002	10890	23318
W1	117	162	244	4563	10530	24156
F0	123	170	251	3567	9350	22339
F1	109	150	226	4469	10050	22826

Table 2: Table showing the total number of internal and the number of external PCG iterations required to solve the pure scattering problem.

5. CONCLUSIONS

We conclude that the F-cycles described in the paper are the most efficient hierarchical (multi-grid) preconditioners for the problems solved here closely followed by the M-cycles and W-cycles

methods. The F-cycles put more effort in solving the lower moments than the other cycles, which probably contributes to their success. However, for both the purely absorbing and scattering problems solved here, it was found that the simple cycles performed the best. This can be attributed to the feature of the conjugate gradient iteration which iterates only marginally more when used with the simple preconditioners. The extra work in forming the more elaborate preconditions does not offset the extra global or outer conjugate gradient iterations. However, appropriate parameterisations and relaxation of the W-cycles may, with tuning, make them more competitive.

REFERENCES

1. C. R. E. de Oliveira, "An Arbitrary Geometry Finite Element Method for Multigroup Neutron Transport with Anisotropic Scattering", *Prog. in Nucl. Energy* , **18**, pp. 227-236 (1986).

2. C. R. E. de Oliveira and A. J. H. Goddard, "EVENT - A Multidimensional Finite Element-Spherical Radiation Transport Code", OECD Proceedings of 3-D Deterministic Radiation Transport Computer Programs, pp. 81-90, Paris, France, (2-3 December 1996).

3. C. R. E. de Oliveira, A. J. H. Goddard, C. C. Pain and A. P. Umpleby, "Parallel Domain Decomposition Methods for Large-Scale Finite Element Transport modelling", Proceedings of the ANS International Conference on Mathematics aand Computations, Reactor Physics and Environmental Analyses, Portland, Oregon, (20 April-4 May 1995).

4. E. E. Lewis and W. Miller Jr., *Computational Methods of Neutron Transport*, John Wiley and Sons, New York (1984).

5. O. Axelsson and I. Gustafsson, "Preconditioning and two-level multi-grid methods of arbitrary degree of approximation", *Math. Comp.*, **40**, pp. 219-242, (1993).

6. O. Axelsson and P.S. Vassilevski, "Algebraic multi-level preconditioning methods, II", *SIAM J. Num. Anal.*, **27**, pp. 1569-1590, (1990).

7. S. D. Pautz, J. E. Morel and M. L. Adams, "An Angular Multi-grid Acceleration Method for S_N Equations with Highly Forward-peaked Scattering", Mathematics and computation, Reactor Physics and Environmental Analysis in Nuclear Applications, pp. 647-656, Madrid, Spain, (26-30 September, 1999).

8. K. Kobayashi, "A Proposal for 3-D Radiation Transport Benchmarks for Simple Geometries with Void Region", OECD Proceedings of 3-D Deterministic Radiation Transport Computer Programs, pp. 403-413, Paris, France, (2-3 December 1996).