Preconditioners for Thermal Radiation Transport

Jim E. Morel

Texas A&M University
Department of Nuclear Engineering
College Station, TX 77843-3133
morel@tamu.edu

Collaborators:

James S. Warsa
Los Alamos National Laboratory

Brian Yang
Lawrence Livermore National Laboratory

Overview

- The equations of thermal radiation transport
- Angular, energy, spatial, and temporal discretization
- Traditional solution of the transport equation
  - Source Iteration (SI)
  - Diffusion-Synthetic Acceleration (DSA)
  - Linear Multifrequency Grey Acceleration (LMFGA)
- Krylov solution of preconditioned systems from acceleration schemes
- Nested Krylov methods versus nested acceleration schemes
- LMFGA-preconditioned systems
- Computational results
- Concluding remarks

Basic Equations

- The equations of thermal radiation transport consist of a transport equation for the angular intensity $I(\overrightarrow{r}, \overrightarrow{\Omega}, E, t)$:

$$\frac{1}{c}\frac{\partial I}{\partial t} + \overrightarrow{\Omega} \cdot \overrightarrow{\nabla} I + \sigma_t I = \frac{\sigma_s}{4\pi}\phi + \sigma_a B(T)\,,$$

and an equation for the material temperature $T(\overrightarrow{r}, t)$:

$$C_v \frac{\partial T}{\partial t} = \int_0^\infty \sigma_a [\,\phi - 4\pi B(T)\,]\, dE\,.$$

- The angular intensity has units of $(energy/area - time - steradian - energy)$,

Basic Equations

- $\overrightarrow{\Omega}$ is the photon direction vector,
- $E = h\nu\ (energy)$ is the photon energy,
- $c\ (length/time)$ is the speed of light,
- $\sigma_t(\overrightarrow{r}, E, T)\ (length^{-1})$ is the total macroscopic cross section,
- $\sigma_s(\overrightarrow{r}, T)\ (length^{-1})$ is the Thompson macroscopic scattering cross section,
- $\phi(\overrightarrow{r}, E, t)\ (energy/area - time - energy)$ is angle-integrated intensity,
- $C_v(\overrightarrow{r}, T)\ (energy/volume - temperature)$ is the material heat capacity,

Basic Equations

- $B(E,T)\ (energy/area - time - steradian - energy)$ is the Planck function:

$$B(E,T) = \frac{2E^3}{h^3c^2}\left[\exp\left(\frac{E}{kT}\right) - 1\right]^{-1} \quad ,$$

- h is *Planck's constant*,
- k is *Boltzmann's constant*

Nonlinearities

- Note that the radiative transfer equations have nonlinearities arising only from the temperature dependence of the material property coefficients and the Planck function.
- The heat capacity is generally a weak function of temperature, while the absorption cross sections are strong functions of temperature.
- The radiative transfer equations are generally solved using an approximate form of Newton's method.
- The method is approximate in that the contributions to the Jacobian from the material property functions are neglected.
- Tabulations of material property data are generally not differentiable.
- Stability considerations require linearization of the Planck function but not the material property functions.

Connection with Neutronics

- The radiative transfer equations have much in common with the neutron transport equation.
- The regimes are quite different, but most of the computational technology developed for neutron transport has been applied over the last twenty years to the radiative transfer equation both in static media and with nonrelativistic material motion.
- Because radiative transfer calculations are much more demanding than neutron transport calculations, this connection has resulted in improvements in neutron transport methods that might not otherwise have occured.

Radiation-Hydrodynamics

- The radiative transfer equations are often coupled with the hydrodynamics equations to form the radiation-hydrodynamics equations.
- With non-relativistic material motion, operator splitting techniques can be applied in conjunction with standard hydrodynamics and radiative transfer solution techniques to solve the radiation-hydrodynamics equations.
- For the case of relativistic material motion, fundamentally different approaches to the solution of the radiative transfer equations are required.

Discretization

- We use the S_n or discrete-ordinates angular discretization, which is basically a collocation technique.
- We use the multigroup discretization in energy, which is basically a Petrov-Galerkin method with piecewise-constant weight functions.
- We use lumped discontinous Galerkin discretizations in space because of their accuracy, robustness, and asymptotic-preserving properties.
- We generally use the Backward-Euler discretization in time, although second-order non-oscillatory schemes such as the BDF-2 and trapezoidal BDF-2 schemes can be used in conjunction with the solution techniques described here.

Solution of the Radiative Transfer Equations

- Traditional accelerated iterative solution techniques for the transport equation are closely related to multilevel or multigrid methods.
- In most instances, some type of diffusion operator is used to approximate a transport operator.
- For many years, such methods were thought to be unconditionally effective as long as the diffusion equations were differenced in a manner consistent with the spatial discretization of the transport operator.
- Unfortunately, when discontinuous Galerkin methods are used for the transport equation, the consistent diffusion discretizations are of a mixed form and can be very expensive to solve.

Solution of the Radiative Transfer Equations

- A great deal of research effort has been spent over the last 20 years or so to find ways to either use simpler diffusion discretizations or solve the full discretizations in an approximate manner without significant loss of effectiveness.
- These efforts have met with limited success.
- Furthermore, over the last five years or so, it has been recognized that traditional acceleration techniques are not uniformly effective in multidimensional calculations even when consistent diffusion discretizations are used.
- In particular, it has been found that strong material inhomogeneities can degrade effectiveness and occasionally generate instabilities.

Solution of the Radiative Transfer Equations

- It has now become clear that by recasting traditional accelerated iteration schemes as preconditioned Krylov methods, far greater lattitude in the choice of diffusion discretization is possible, the degrading effects of strong material inhomogeneities can be significantly reduced, and any associated instabilities eliminated.
- Consequently, there is currently a great deal of research within the computational transport community devoted to preconditioned Krylov methods.

Solution of the Radiative Transfer Equations

- As previously noted, the radiative transfer equations are generally solved via an approximate form of Newton's method.
- After linearization, temporal discretization (backward-Euler), and energy discretization (multigroup), we are able to eliminate the temperature from the transport equation:

$$\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} I_g + \sigma^*_{\tau,g} I = \frac{1}{4\pi} \sigma^*_{s,g} \phi_g + \frac{1}{4\pi} \nu \chi_g \sum_{k=1}^{G} \sigma^*_{a,k} \phi_k + \xi_g \quad , \quad g = 1, G,$$

- and obtain an intensity-dependent temperature equation:

$$T = T^* + \frac{\sum_{g=1}^{G} \sigma^*_{a,g} \left[\phi_g - 4\pi B^*_g \right] + \frac{C^*_v}{\Delta t^k} \left(T^n - T^* \right)}{\frac{C^*_v}{\Delta t^k} + \sum_{g=1}^{G} \sigma^*_{a,g} 4\pi \frac{\partial B^*_g}{\partial T}},$$

The Equations of Thermal Radiation Transport

- where

$$\sigma_\tau^* = \sigma_t^* + \tau\,,$$

$$\tau = \frac{1}{c\Delta t^k}\,,$$

$$\nu = \frac{\sum_{g=1}^{G} \sigma_{a,g}^* 4\pi \frac{\partial B_g^*}{\partial T}}{\frac{C_v^*}{\Delta t^k} + \sum_{g=1}^{G} \sigma_{a,g}^* 4\pi \frac{\partial B_g^*}{\partial T}}$$

$$\chi_g = \frac{\sigma_{a,g}^* \frac{\partial B_g^*}{\partial T}}{\sum_{k=1}^{G} \sigma_{a,k}^* \frac{\partial B_k^*}{\partial T}}\,,$$

$$\xi_g = \sigma_{a,g}^* B_g^* + \tau I_g^n -$$

$$\frac{1}{4\pi}\nu\chi_g \left[\sum_{k=1}^{G} \sigma_{a,k}^* 4\pi B_k^* + \frac{C_v^*}{\Delta t^k}\left(T^n - T^*\right)\right].$$

Source Iteration

- The traditional method for solving the transport equation is a nested source iteration.
- Denoting the iteration index by ℓ, the inner iteration can be represented as follows:

$$\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} I_g^{\ell+1} + \sigma_{\tau,g}^* I_g^{\ell+1} = \frac{1}{4\pi}\sigma_{s,g}^* \phi_g^{\ell} + \frac{1}{4\pi}\nu\chi_g \sum_{k=1}^{G} \sigma_{a,k}^* \phi_k \ + \xi_g \,,$$

- and the outer iteration can be represented as follows:

$$\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} I_g^{\ell+1} + \sigma_{\tau,g}^* I_g^{\ell+1} - \frac{1}{4\pi}\sigma_{s,g}^* \phi_g^{\ell+1} = \frac{1}{4\pi}\nu\chi_g \sum_{k=1}^{G} \sigma_{a,k}^* \phi_k^{\ell} \ + \xi_g \,.$$

Source Iteration

- The operator $\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} + \sigma^*_{\tau,g}$ involves no angular or energy coupling.
- When spatially discretized it takes on a block lower-triangular form with a block corresponding to the intensities within a single spatial cell for a single direction and energy.
- This operator is easily inverted using a “wavefront” or “sweep” algorithm.
- The attenuation of errors in ϕ_g determines the convergence rate of the inner iteration process.
- The attenuation of errors in $f = \sum_{g=1}^{G} \sigma^*_{a,g}\phi_g$ determines the convergence rate of the outer iteration process.

Source Iteration

- The inner iteration process can become arbitrarily slow to converge as $\sigma^*_{s,g} \longrightarrow \sigma^*_{\tau,g}$. This corresponds to scattering dominating absorption.
- The outer iteration can become arbitrarily slow to converge as $\nu \longrightarrow 1$ and $\tau \longrightarrow 0$. This physically corresponds to strong material-radiation coupling (small heat capacity and large absorption cross section).
- For the case of an infinite homogeneous medium, Fourier analysis can be used to demonstrate that the most slowly converging error modes for both iterations are those that slowly vary in space.
- Thus the inversion of $\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} + \sigma^*_{\tau,g}$ via sweeps is a form of relaxation: high-frequency errors are strongly attenuated, while low-frequency errors are poorly attenuated.

Diffusion-Synthetic Acceleration

- Inner source iteration with diffusion-synthetic acceleration (DSA) takes the following form:

$$\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} I_g^{\ell+\frac{1}{2}} + \sigma_{\tau,g}^* I_g^{\ell+\frac{1}{2}} = \frac{1}{4\pi}\sigma_{s,g}^* \phi_g^\ell + \frac{1}{4\pi}\nu\chi_g \sum_{k=1}^{G} \sigma_{a,k}^* \phi_k \; + \xi_g \, ,$$

$$-\overrightarrow{\nabla} \cdot \frac{1}{3\sigma_{\tau,g}^*} \overrightarrow{\nabla} \delta\phi_g + \left(\sigma_{\tau,g}^* - \sigma_{s,g}^*\right) \delta\phi_g = \sigma_{s,g}^* \left(\phi_g^{\ell+\frac{1}{2}} - \phi_g^\ell\right) \, ,$$

$$\phi_g^{\ell+1} = \phi_g^{\ell+\frac{1}{2}} + \delta\phi_g \, .$$

Diffusion-Synthetic Acceleration

- For the case of an infinite homogeneous medium, Fourier analysis can be used to demonstrate that the diffusion step completely attenuates the low-frequency error modes and grossly underestimates the high-frequency error modes, which are already strongly attenuated by the sweep.
- This is the best one can hope for in a low-rank approximate inverse.
- The scheme is unconditionally effective in 1-D and only becomes ineffective in strongly heterogeneous multidimensional problems.

Linear Multifrequency-Grey Acceleration

- Outer source iteration with LMFGA takes the following form:

$$\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} I_g^{\ell+\frac{1}{2}} + \sigma^*_{\tau,g} I_g^{\ell+\frac{1}{2}} - \frac{1}{4\pi}\sigma^*_{s,g}\phi_g^{\ell+\frac{1}{2}} = \frac{1}{4\pi}\nu\chi_g f^{\ell} + \xi_g \,,$$

$$-\overrightarrow{\nabla}\cdot\langle D\rangle\overrightarrow{\nabla}\delta\Phi + \left[\langle\sigma_a\rangle(1-\nu)+\tau\right]\delta\Phi = \nu\left(f^{\ell+\frac{1}{2}} - f^{\ell}\right) \,, \quad (1)$$

$$f^{\ell+1} = f^{\ell+\frac{1}{2}} + \langle\sigma_a\rangle\delta\Phi \,,$$

Linear Multifrequency-Grey Acceleration

- where

$$\langle D \rangle = \sum_{g=1}^{G} \frac{\varsigma_g}{3\sigma^*_{\tau,g}} ,$$

$$\langle \sigma_a \rangle = \sum_{g=1}^{G} \sigma^*_{a,g} \varsigma_g ,$$

$$\varsigma_g = \frac{\frac{\chi_g}{\sigma^*_{\tau,g}}}{\sum_{k=1}^{G} \frac{\chi_k}{\sigma^*_{\tau,k}}} .$$

Linear Multifrequency-Grey Acceleration

- For the case of an infinite homogeneous medium, Fourier analysis can be used to demonstrate that the diffusion step completely attenuates the low-frequency error modes and grossly underestimates the high-frequency error modes, which are already strongly attenuated by the transport solves.
- This is the best one can hope for in a low-rank approximate inverse.
- The scheme appears to be unconditionally effective in 1-D but can apparently become unstable in strongly heterogeneous multidimensional problems.
- This motivates us to develop a preconditioned Krylov method based upon the multifrequency-grey acceleration technique.

Krylov Methods

- Suppose one wants to solve a linear system of the following form:

$$\mathbf{A}\,\overrightarrow{x} = \overrightarrow{y}\,,$$

  where $\mathbf{A}$ is a matrix, $\overrightarrow{x}$ is the solution vector, and $\overrightarrow{y}$ is the source vector.

- To solve this system using a Krylov method, one must simply be able to provide the Krylov routine with the action of $\mathbf{A}$ on an arbitrary vector, $\overrightarrow{z}$, i.e., one must compute $\overrightarrow{v}$ where

$$\overrightarrow{v} = \mathbf{A}\,\overrightarrow{z}\,.$$

- For the case of radiative transfer, the matrix $\mathbf{A}$ is dense and the action of $\mathbf{A}$ must be calculated in an indirect manner.

Krylov Methods

- Characterizing the convergence of Krylov methods for general matrices is difficult, but there is one simple rule that can be followed: convergence will improve as the domain of the eigenvalues becomes smaller and as the domain moves away from the origin.
- Furthermore, all Krylov methods converge in one iteration in the limit as the matrix $\mathbf{A}$ approaches the identity.
- Preconditioning can be used to improve convergence by modifying the spectrum of the constituent operator without changing the solution of the equation.
- We use traditional acceleration schemes to provide preconditioning and we define our preconditioned equations directly from the accelerated iteration equations.

Preconditioned Systems from Iteration Schemes

- We now show how to construct a preconditioned system of equations from the accelerated iteration equations.
- We begin by expressing the accelerated iteration scheme in standard single-step form:

$$\overrightarrow{x}^{\ell+1} = \mathbf{M}x^{\ell} + \overrightarrow{q} \ ,$$

  where $\overrightarrow{x}$ is the solution vector, $\mathbf{M}$ is the iteration matrix, $\overrightarrow{q}$ is the source vector, and ℓ is the iteration index.
- Although most schemes are more easily expressed in a multi-step form, it is simply a question of algebra to obtain the single-step form.
- We next rewrite this iteration scheme in the form of Richardson iteration:

$$\overrightarrow{x}^{\ell+1} = \overrightarrow{x}^{\ell} + \overrightarrow{q} - (\mathbf{I} - \mathbf{M})\overrightarrow{x}^{\ell} \ .$$

Preconditioned Systems from Iteration Equations

- Note that Richardson iteration is being performed on the following system:

$$(\mathbf{I} - \mathbf{M})\overrightarrow{x} = \overrightarrow{q} \ .$$

- This the preconditioned system solved with the Krylov method.
- Note that the generating the action of this operator almost identical to perfoming an iteration.
- If the accelerated iteration scheme is highly effective, the iteration matrix $\mathbf{M}$ must be ”close” to the zero matrix, and if this is so, then the matrix $(\mathbf{I} - \mathbf{M})$ must be “close” to the identity matrix.
- Thus it follows that a Krylov method should be highly effective for solving the preconditioned system.

Preconditioned Systems from Iteration Equations

- It would seem from this result that there is only one way to formulate the preconditioned system, given an accelerated iteration scheme, but the same iteration scheme can sometimes be expressed in terms of different unknowns.
- One generally wants the preconditioned system to have the unknown of lowest possible rank.
- This rank is given by the rank of the operator that is being iterated upon.

Nested traditional iteration versus nested Krylov

- There is one disadvantage of nested Krylov methods relative to nested acceleration schemes that is not immediately obvious, but easily understood.
- We define a nested scheme to be one in which one or more equations must be iteratively solved to obtain the action of the consituent operator.
- For an acceleration scheme, the solution for each of the nested systems converges as the outer iteration proceeds. Thus one has an increasingly better initial guess for the nested iterations.
- This is not the case for a nested Krylov vector because the Krylov vectors do not converge to the solution.
- Thus the nested iterations associated with Krylov methods can take much longer on the average than those associated with acceleration schemes.

Nested traditional iteration versus nested Krylov

- There is some theory regarding techniques for maximizing the nested convergence tolerance in Krylov iterations.
  - A. Bouras and V. Frayssé, “A Relaxation Strategy for Inexact Matrix-Vector Products for Krylov Methods,” European Centre for Research and Advanced Training in Scientific Computation, Toulouse, France, CERFACS TR/PA/00/15 (2000), Submitted to *SIAM J. Matrix Anal. Appl.*, http://www.cerfacs.fr/algor/reports/2000/TR_PA_00_15.ps.gz.
  - V. Simoncini and D. Szyld, “Theory of Inexact Krylov Subspace Methods and Applications to Scientific Computing,” *SIAM J. Sci. Comput.*, **25**, 454–477 (2003).
- Krylov methods may have concrete advantages over acceleration schemes for maximizing the nested convergence tolerance.

LMFG-Preconditioned Systems

- We consider two LMFG-accelerated iteration schemes and associated preconditioned systems.
- The first system corresponds to the standard LMFG iteration scheme, but its unknown is the absorption rate f rather than the intensity.
- A significant advantage of this system is that the f-vector is low-rank relative to the I-vector.
- A disadvantage of this system is that it is nested with G independent transport equations and one grey diffusion equation that must be solved to obtain the action of the associated operator.
- Once the absorption rate has been obtained, the intensities are obtained by solving G independent transport equations:

$$\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} I_g + \sigma^*_{\tau,g} I_g - \frac{1}{4\pi} \sigma^*_{s,g} \phi_g = \frac{1}{4\pi} \nu \chi_g f \ + \xi_g \, .$$

First Preconditioned System

$$\mathbf{CB}f = \mathbf{C}\sum_{g=1}^{G}\sigma_{a,g}^{*}\mathbf{PA}_{g}^{-1}\xi_{g}\,,$$

where

$$\mathbf{P}\langle\cdot\rangle = \int_{4\pi}\langle\cdot\rangle\,d\Omega\,,$$

$$\mathbf{A}_{g} \equiv \overrightarrow{\Omega}\cdot\overrightarrow{\nabla} + \sigma_{\tau,g}^{*} - \frac{1}{4\pi}\sigma_{s,g}\mathbf{P}\,,$$

$$\mathbf{B} = \left[\mathbf{I} - \sum_{g=1}^{G}\sigma_{a,g}^{*}\mathbf{PA}_{g}^{-1}\frac{1}{4\pi}\nu\chi_{g}\right]\,,$$

$$\mathbf{C} \equiv \left(\mathbf{I} + \langle\sigma_{a}\rangle\mathbf{H}^{-1}\nu\right)\,,$$

$$\mathbf{H} \equiv -\overrightarrow{\nabla}\cdot\langle D\rangle\overrightarrow{\nabla} + \left[\langle\sigma_{a}\rangle(1-\nu) + \tau\right]\,.$$

LMFG-Accelerated Preconditioned Systems

- The second system corresponds tc a “modified” LMFG iteration scheme, but its unknown is the group-dependent angle-integrated intensities $\{\phi_g\}_{g=1}^{G}$.
- In particular, the modified iteration scheme corresponds to the substitution of one sweep for each of the G independent transport solves:

$$\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} I_g^{\ell+\frac{1}{2}} + \sigma_{\tau,g}^* I_g^{\ell+\frac{1}{2}} = \frac{1}{4\pi}\sigma_{s,g}^* \phi_g^{\ell} + \frac{1}{4\pi}\nu\chi_g f^{\ell} + \xi_g \,,$$

$$-\overrightarrow{\nabla}\cdot\langle D\rangle\overrightarrow{\nabla}\,\delta\Phi + \left[\langle\sigma_a\rangle(1-\nu)+\tau\right]\delta\Phi = \nu\left(f^{\ell+\frac{1}{2}} - f^{\ell}\right) \,,$$

$$f^{\ell+1} = f^{\ell+\frac{1}{2}} + \langle\sigma_a\rangle\delta\Phi \,.$$

LMFG-Accelerated Preconditioned Systems

- An advantage of this system is that the ϕ-vector is low-rank relative to the I-vector, but a disadvantage is that the ϕ-vector is high-dimensional relative to the f-vector.
- An advantage of this system is that only G independent "sweeps" and one grey diffusion solve must be performed to obtain the action of the associated operator.
- A disadvantage of this system is that it corresponds to less effective version of the LMFG method, however it is only intended to be applied in problems with weak scattering.
- Once the angle-integrated intensities have been obtained, the intensities are obtained by performing G independent sweeps:

$$\overrightarrow{\Omega} \cdot \overrightarrow{\nabla} I_g + \sigma^*_{\tau,g} I_g = \frac{1}{4\pi} \sigma^*_{s,g} \phi_g + \frac{1}{4\pi} \nu \chi_g f \ + \xi_g \, .$$

Computational Results

- We do not yet have computational results for two LMFG-preconditioned systems.
- We do have results that:
  - Compare a traditional acceleration scheme and its preconditioned Krylov counterpart.
  - Demonstate the disadvantage of Krylov methods with respect to nested iterations.

A Multigroup Diffusion Example

- We have tested a LMFG-preconditioned Krylov method that solves the following equation:

$$-\overrightarrow{\nabla}\cdot D_g^*\overrightarrow{\nabla}\phi_g+\sigma_{\tau,g}^*\phi_g-\nu\chi_g\sum_{k=1}^{G}\sigma_{a,k}^*\phi_k=\xi_g\,,\quad g=1,G,$$

- The method is analogous to our transport method.
- The only significant difference is that

$$\mathbf{A}_g\equiv-\overrightarrow{\nabla}\cdot D_g\overrightarrow{\nabla}+\sigma_{\tau,g}^*\,.$$

- Thus we solve a diffusion equation rather than a transport equation for each group.

A Multigroup Diffusion Example

- There are really only two material regions, the middle region facilitates non-uniform zoning.
- There are two problems: one with a uniform density everywhere and the other with a density jump of 1000 between the inner and outer regions - both were performed with the same 10-group cross sections.
- The radiation source is turned on at $t = 0$ and radiation propagates through the system.

Cylindrical Geometry for Diffusion Test Problems.

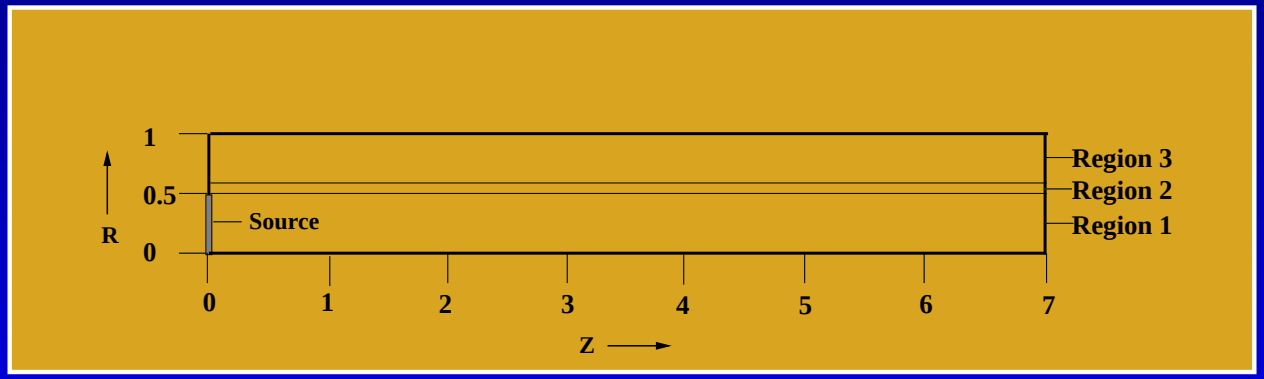

A Multigroup Diffusion Example

Comparisons for the variable–density calculations.

$\Delta t\,(sh)$	Method	Time (s)	No. of "Outer" Iterations	No. of "Inner" Iterations
0.1	LMFGK	1.553	7.98	1424.42
	LMFGA	2.992	25.73	2929.14
0.01	LMFGK	0.848	3.55	668.40
	LMFGA	0.986	8.07	785.00
0.001	LMFGK	0.450	1.26	260.32
	LMFGA	0.417	1.56	190.12

A Multigroup Diffusion Example

Comparisons for the uniform–density calculations.

$\Delta t\,(sh)$	Method	Time (s)	No. of "Outer" Iterations	No. of "inner" Iterations
0.2	LMFGK	0.601	3.41	455.71
	LMFGA	0.646	6.14	472.31
0.1	LMFGK	0.561	2.83	403.83
	LMFGA	0.561	4.57	376.99
0.01	LMFGK	0.425	1.228	249.49
	LMFGA	0.388	1.375	180.97

A Multigroup Diffusion Example

Outer iteration counts as a function of time-step number for the variable–density problem with a 0.1 sh time step.

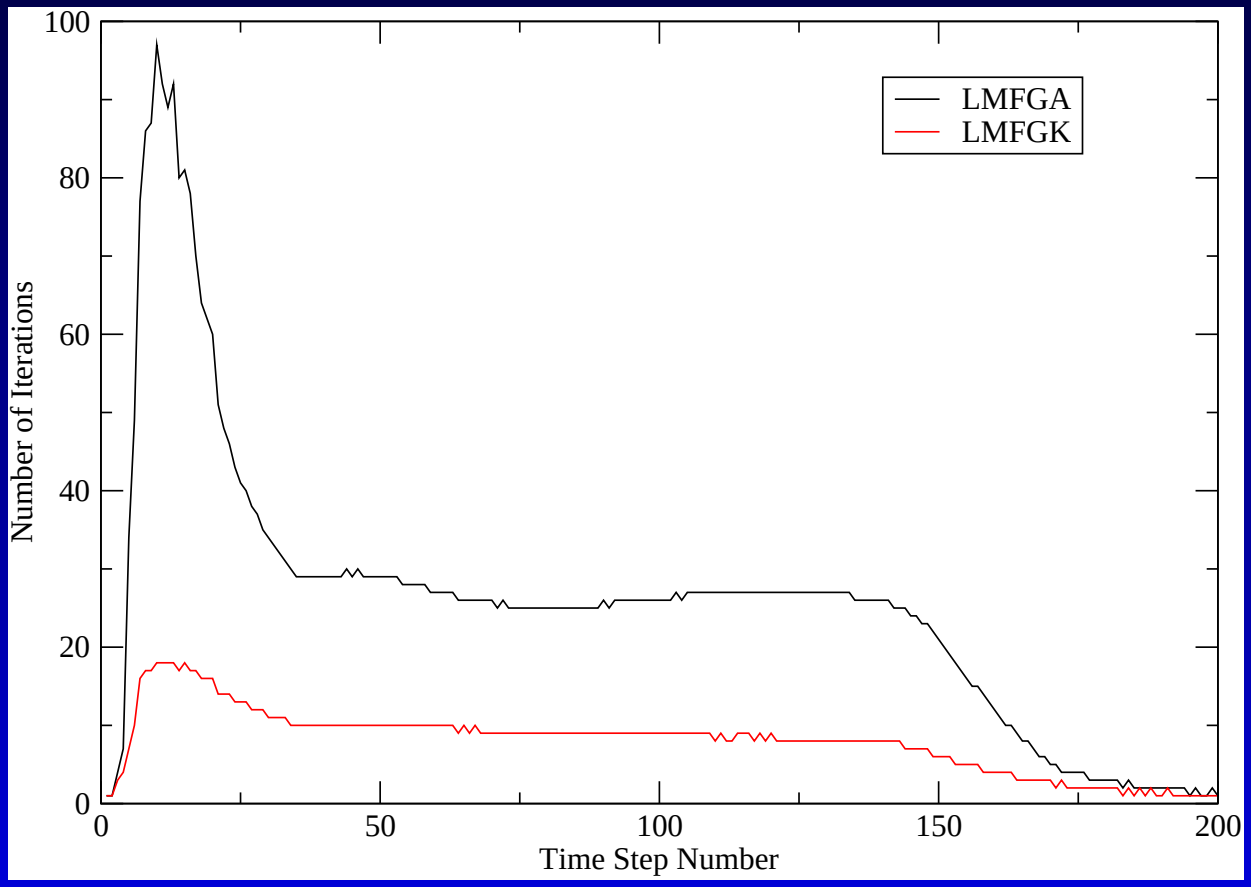

A Multigroup Diffusion Example

- The Krylov method always takes fewer outer iterations than the acceleration method.
- The Krylov method takes significantly fewer outer iterations with a large discontinuity in density.
- The relative CPU times are not proportional to the number of outer iterations because the Krylov method requires more inner iterations per outer iteration than the acceleration method.
- This is the effect that was previously explained.
- The inner solves for the acceleration technique have increasingly better initial guesses as the solution converges - the guesses are the solutions from the previous outer iteration.
- The inner iteration solutions for the Krylov method are independent for each outer iteration, so an initial guess of zero is always used.

A Multigroup Diffusion Example

- This effect causes the Krylov method to be about 10 percent less efficient than the acceleration scheme in the worst case.
- Nonetheless, the Krylov method is about two times more efficient than the acceleration scheme in the best case.
- The ratio of iterations between the two methods varies significantly between time steps.
- The acceleration method has been observed to go unstable in complex and highly heterogeneous calculations.
- The Krylov method has so far remained effective in such calculations.
- Heuristic methods have been developed to deal with the instabilities of the acceleration technique, but a robust method is always preferable.

Concluding Remarks

- The robustness and versatility of Krylov methods eliminates many of the problems associated with traditional accelerated iteration schemes.
- There may be many simplified variants of traditional iteration schemes that are themselves unstable but still may provide effective preconditioned systems for solution via Krylov methods.
- Ongoing work:
  - Computationally compare the two LMFG-preconditioned systems.
  - Investigate nested tolerance maximization algorithms for both Krylov solvers and traditional acceleration schemes.