AN ABSTRACT OF THE THESIS OF

Nicholas T. Myers for the degree of Master of Science in Nuclear Engineering presented on May 18, 2011.
Title: Evaluating Preconditioning Options for the Jacobian-Free Newton Krylov Method Applied to Radiative Transfer With Matter Coupling

Abstract approved: _____

Todd S. Palmer

The objective of this thesis was to investigate the use of the Jacobian Free Newton Krylov Method (JFNK) with physics based preconditioning for solving coupled radiative transport and material energy balance. Modeling of these physics is important in inertial confinement fusion and other high energy density systems. In this thesis, the non-linear equations for radiative transfer and matter balance are converged within the JFNK framework. Coupling of the transport problem to the material energy balance is attempted in a manner unlike Britton Chang's Photon Free Method [5]. The effects of various physics-based preconditioners are compared. The proposed method with the various formulated preconditioners accurately solve the Su and Olson benchmark problem but difficulty is encountered when Chang's test problem is attempted. Analysis is performed to determine the limits of the proposed method. It is concluded that the most efficient physics-based preconditioning in the Su and Olson problem is found by taking the derivative with respect to temperature for the material energy balance equation. Because of the difficulties experienced when attempting to solve Chang's problem it is concluded that the proposed method in this thesis fails for tightly coupled problems. This is due to overloading the inner GMRES iterative solver for the full problem.

Evaluating Preconditioning Options for the Jacobian-Free Newton Krylov Method Applied to Radiative Transfer With Matter Coupling

by

Nicholas T. Myers

A THESIS

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Master of Science

Presented May 18, 2011
Commencement June 2011

Master of Science thesis of Nicholas T. Myers presented on May 18, 2011.

APPROVED:

Major Professor, representing Nuclear Engineering

Head of the Department of Nuclear Engineering and Radiation Health Physics

---

Dean of the Graduate School

I understand that my thesis will become part of the permanent collection of Oregon State University libraries. My signature below authorizes release of my thesis to any reader upon request.

Nicholas T. Myers, Author

ACKNOWLEDGEMENTS

I would like to acknowledge those who contributed to my work here and encouraged me throughout the process. The work that has gone into this thesis is extensive and involved many stressful nights. I would never have been able to accomplish this work without the strength and peace that the Lord Jesus Christ has given me. The scriptures say that it is the Lord who gives and takes away and He is the one who establishes a man's path. I am most thankful to God for blessing me with the mind that I have and the educational opportunities that I have had in my life.

I am also thankful for my family for encouraging and supporting me throughout this whole process. They have always been a support and encouragement through my educational career and have pushed me to give my best effort in the classroom.

I also am very grateful towards my advisor, Dr. Todd Palmer, for taking me on as his graduate student. He has been a great help in accomplishing this thesis and preparation for pursuing a doctoral degree at the University of New Mexico. I appreciate his commitment towards his students and that he is straightforward in helping every one of us to succeed at school and career.

I would also like to thank Dr. Dana Knoll for accepting me as his summer intern at Idaho National Laboratory in the summer of 2009. He taught me much about the JFNK method and multi-physics solution methods. He was a most excellent mentor at INL and has continued to help my understanding throughout the work that has gone into this thesis. I am also thankful to Dr. Jim Warsa for accepting me as his summer intern at Los Alamos National Laboratory in the summer of 2010. Dr. Warsa helped me in furthering my understanding of transport methods

and was a non-trivial reason why my summer was bearable in the midst of working on my thesis.

I have been very blessed to have such wonderful people in my life who not only were encouraging, but invested in my life to see me succeed. I am humbled to think that so many outstanding people would care about my education, career, and success. It is my hope that someday I may be as big a blessing to others as these people have been in my life.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

LIST OF FIGURES

LIST OF FIGURES (Continued)

LIST OF TABLES

LIST OF TABLES (Continued)

Evaluating Preconditioning Options for the Jacobian-Free Newton Krylov Method Applied to Radiative Transfer With Matter Coupling

1 Introduction

In the field of nuclear engineering, as well many other science and engineering disciplines, it is important to understand the distribution of radiation energy fields. Such fields can be as simple as heat coming from a fire or as complicated as the spectrum of energy emitted from a supernova. Regardless of the complexity, electromagnetic waves, or photons, are being absorbed and re-emitted continuously in all materials that we encounter.

The representation and modeling of such energy fields and interactions is important for many applications. One such application is inertial confinement fusion. It is also important in atmospheric modeling, where the interaction between solar radiation and cloud particles determines weather patterns. An understanding of the radiation distribution is also important in oxy-coal combustion and other power generation systems, where radiative heat transfer is a dominant process.

In many of these applications of interest, the problem involves multiple physical processes that are tightly coupled together. Multi-physics problems are found in the aforementioned examples as well as in coupled neutronics and thermal hydraulics in a nuclear reactor, material strain and stress coupled with thermal transfer processes in nuclear fuels, and biogeochemical transport. Any time there are competing and/or interacting physical processes, it can be described mathematically as a non-linear multi-physics problem.

Non-linearly coupled systems of partial differencial equations are difficult to solve accurately and often are mathematically intractable for calculating an ana-

lytical solution. It may be possible to make general assumptions to simplify the problem, but the solution of this simplified problem may not preserve features of the original system (for example, assuming a steady state solution will eliminate any transient solution information). These simplified problems are useful as benchmarks, but ultimately an experiment or computer simulation is needed to allow for a better understanding of a coupled model, especially when full analytical solutions are not available.

In many of these applications, particularly nuclear engineering, it is prohibitively expensive to construct experimental facilities to gather data about radiation field interactions. It is more cost efficient to perform computer simulations of the physical system in order to characterize its behavior and to perform design analysis—provided the code is verified and validated. Verification and validation of computer codes is necessary for code users, consumers of the product designed with the codes, and for regulatory agencies.

When attempting to resolve nonlinearities in a system of partial differential equations, one generally has two options available: fixed point iteration or the Newton-Raphson method. The fixed point iteration scheme (sometimes known as Picard or Richardson iteration) is very simple to implement and has a linear convergence rate with respect to the error reduction per iteration. A Newton-Raphson iteration scheme can be more challenging to implement, but allows for quadratic convergence rate with respect to the error reduction. The usual difficulty in using the Newton-Raphson method is forming the Jacobian matrix for the system. Sometimes this is very expensive to do analytically, especially for tightly coupled systems where the partial derivatives can be complex. Both methods may not converge, but Picard iteration is generally more stable than the Newton-Raphson

method.

One may instead form an inexact Newton method that approximates the action of the Jacobian on a vector rather than requiring the full formation of the matrix itself. Such inexact methods lose their guaranteed quadratic convergence, but the convergence rate can be close to quadratic with proper preconditioning. Choosing a good preconditioner is not always an easy task, and generally requires knowledge of the dominant physics of the system in question.

The focus of this thesis is the radiative transfer equation for photon transport coupled with the equation for material energy balance. The coupling of these two equations creates a highly non-linear system that is difficult to solve efficiently and accurately with numerical methods. This thesis seeks to answer whether the method formulated is effective with a variety of physics based preconditioning. This thesis also seeks to understand what is the most efficient form of physics based preconditioning for this particular problem. The physics-based preconditioners being evaluated in this thesis are: infinite medium, radiation diffusion, grey transport, S_2 transport, grey coupling, multi-frequency coupling, and the temperature derivative of the material energy balance equation. The following sections look into the current work that has been done in the fields of radiative transport theory and multi-physics modeling. After this a short description of this thesis is presented.

1.1 Literature Review

The following section is a review of the literature concerning the developement and implementation of the Jacobian Free Newton Krylov method for converging non-linear systems. A review of literature concerning solution techniques to the linear Boltzmann transport equation is also conducted.

1.1.1 Multi-physics Methods

In many practical situations where numerical modeling and simulation is required, there exist non-linearities that pose difficulties in the solution. Linear problems are generally much easier to solve and methods for these problems have been well expounded upon throughout the 20th century.

Conventional means of dealing with non-linearities generally involve an approximation to eliminate multi-physics coupling and nonlinear terms. Fixed point iteration can be used to converge non-linear terms and multi-physics coupling at each time step, but the many nested loops can result in prohibitive computation time when simulating a very large problem. The easiest way to accommodate troubling non-linearities is to "lag" the parameter at the previous step. This is dangerous since it trends a problem more towards an explicit formulation which can be unstable and potentially inaccurate for prolonged steps. Instead, a common method is to "linearize" non-linear terms in an equation and/or split the problem physics into separate problems that are solved in sequential order.

Linearization is a technique that involves performing a two term Taylor series expansion of the non-linear term about the solution at the previous step. This method is useful because it will make a system mathematically tractable via linear algebra. A study characterizing time integration errors of two numerical schemes (Linearly Implicit and Implicitly Balanced) found that a semi-implicit scheme does not always preserve the maximum principle and large time steps produce significant error in the linear implicit scheme [19].

Use of the operator split method is common in industry. Splitting multiple physics into sequentially solved single physical systems is generally straightforward and easy to accomplish. Production codes that solve single physics systems are

amenable to this method since the results from one code (i.e. a neutronics code for neutron flux profile) can be used in the second code as an input parameter to calculate the desired data (i.e. thermal profiles from a reactor thermal-hydraulics code). The danger with this approach is that the time scales for the different physics are dissimilar. This creates a system that could diverge if explicit numerical algorithms are being used. Even if the system is being solved implicitly, the proper transient behavior may be lost and the method may even converge to the wrong solution [17].

Before the JFNK method was seriously considered, fixed point methods were the norm to resolve the non-linearities. Szilard [37] attempted to treat the T^4 non-linear term by treating T^4 as the variable of interest as opposed to T alone. He then set C_v to be proportional to $T^{\frac{-3}{4}}$ such that the equations could be linearized. However, Szilard noted that this may not have been the best choice for test problems that start initially cold since it introduces a division by zero for C_v. In his study, Szilard looked at the grey radiation diffusion equation.

Another approach is to resolve the non-linearities within each time step using the Jacobian Free Newton Krylov method (JFNK). This method is similar to the well known Newton-Raphson method, but with several key differences. In JFNK, one does not form the Jacobian matrix. Instead, its action is approximated using a finite difference perturbation approximation. Because the Jacobian matrix is no longer formed explicitly, the quadratic convergence rate of the error reduction between iterations of the Newton-Raphson method is lost; however, the method is still super-linearly convergent (close to quadratic) [28]. Another feature of JFNK is that the internal linear solver is a Krylov iterative solver. Krylov solvers are non-stationary iterative methods that compute approximations to the solution in the

Krylov subspace. Some examples of these solvers are GMRES, Arnoldi Iteration, and Conjugate Gradient [38].

Research on the JFNK method has largely concerned its application in the field of neutron and photon transport with coupled physics. Most of these studies have focused on the one group (or grey) problem and are generally one dimensional in space. Knoll et al. have compared JFNK against a non-linear Picard solver for non-linearly coupled grey radiation diffusion [21]. This study highlighted the differences between linearization of the problem and using the JFNK method. Linearization was found to be a good preconditioner for the JFNK method and significant speedup in convergence was observed when using the JFNK method. Knoll and Park have applied the method in k-eigenvalue criticality calculations [22]. Knoll and Smith have also applied the method for non-linear diffusion accleration of source iteration in neutron transport problems [20].

Current research concerning the JFNK method concerns the choice and formulation of the proper preconditioner for various physical systems. Preconditioning is the key to an effective JFNK solver. Without a good preconditioner, the JFNK method cannot compete with current accelerated fixed point schemes. An effective preconditioner is one that approximates the Jacobian matrix and is computationally efficient to construct and invert. Often, the operator describing the dominant physics involved in a system is a good preconditioner. Mousseau et al. have investigated the operator-split approach as a physics-based preconditioner [28]. It was found that the operator-split method is quite easy to formulate and implement, and works effectively as a physics-based preconditioner. In tightly-coupled multiphysics systems, it is important to capture the coupling between equations in the preconditioner since this is often where the stiffness in a system comes from.

With regards to multi-frequency radiative thermal transport, Britton Chang [5] [6] has implemented the JFNK method in a unique way that treats the coupled system as having just one variable of interest–the temperature. In this "photon free" method (PFM), he showed that greater time steps can be taken with reduced time discretization error. However, it takes more work to resolve the problem for larger time steps. His work focused on comparing the PFM against other common Picard approaches such as the semi-implicit linear approximation. However, the preconditioning options for PFM are limited to the material energy balance equation. Chang only used a heat conducting preconditioner.

There is extensive work being done to implement the JFNK method in codes to specifically solve difficult tightly-coupled multi-physics problems. The multi-physics research team at INL has developed a coding framework called MOOSE [12] that provides methods to implement the JFNK algorithm in a general sense for a variety of multi-physics problems. Using this framework, researchers have successfully implemented a parallelized JFNK algorithm to solve neutronics-thermalhydraulics problems for pebble bed reactors (PRONGHORN) [31] and thermal effects coupled with material science for fuel performance (BISON) [29].

Other applications of JFNK exist in disciplines outside of the nuclear industry. For instance, Hammond et al. [14] utilized JFNK in the modeling and simulation of biogeochemical transport and found noteworthy reductions in memory and processor time required, but no wall clock time reductions. JFNK research is also ongoing in the field of weather modeling [32].

1.1.2 Radiative Transport

A challenging task in computational physics is to correctly simulate the interaction between radiation and matter. Interactions between radiation and matter are described using a discretized form of the linearized Boltzmann problem. The resulting system of equations are then solved for the desired field distribution [1]. A radiation transport problem is a function of seven independent variables–three spatial, two angular, one time, and one energy variable. Discrete solutions to the transport equation can be very challenging since the number of unknowns can be very large.

A common way to discretize the angular variable is to use the S_N or discrete-ordinates approximation. This approximation makes use of numerical quadrature to approximate integrals over angular space. The resulting transport equation is a system of differential equations–one for each discrete angle. This method has been well documented and used by researchers since the 1960's [1]. An alternative to the S_N method is to use the P_N equations from spherical harmonics.

The linear Boltzmann equation may behave differently depending on the type of physical system being represented. In an optically thick, highly scattering region, it behaves like an elliptic or parabolic partial differential equation (PDE). However, in a void, it behaves as a hyperbolic PDE. It is difficult to find a transport discretization that works well over a wide range of behavior. Advanced discretization and iteration schemes have been the subject of much research for many years [1].

The simplest iterative solver to implement in a transport system is called Source Iteration (SI). Essentially, SI is a fixed point iteration about the source term. An initial estimate of the source is chosen, the transport equation is solved for angular flux/intensity values, and the quadrature approximation of the S_N method is used

to find the new source. This process is repeated until convergence on the source is achieved to some prescribed tolerance. A danger with the SI scheme is the possibility of false convergence [1]. Source Iteration is also known to be very slow in highly scattering optically thick problems [1]. Many researchers have investigated acceleration schemes to improve the rate of convergence for these problems.

Some acceleration methods currently in use are the diffusion synthetic, quasi-diffusion, and spatial multigrid approaches [1]. Diffusion synthetic acceleration was first introduced by Kopp [15] while the quasi-diffusion acceleration was first introduced by Gol'din [39]. These methods and many others are described in detail in a review paper on fast iterative methods for neutral particle transport [1]. DSA and QD will briefly be described here.

The general idea of diffusion synthetic acceleration is to perform a standard procedure for a source iteration and then solve an approximate equation for a correction to the scalar flux. The correction is calculated using the zeroeth and first angular moment equations to form a low order diffusion approximation equation. This correction is then added to the source iteration scalar flux to get a more accurate answer.

The idea of the quasi-diffusion (QD) method is to calculate the 0th and 1st angular moments of the transport equation and combine them in a low-order equation for the scalar flux. After performing a standard transport sweep, coefficients called "Eddington factors" are calculated from the high-order angular flux values and then the low-order equation is solved for the new scalar flux value. It is important that the low-order equation is discretized in the same way as the high-order equation, otherwise the QD method will converge to an answer for a slightly different problem.

Much of the early work in radiative transport theory for photon fields was performed by Chandrasekhar [4] and has been built upon by Sobolev [34], Modest [27], and others more recently as the field has gained more prominence. Coupling to other physical processes such as material absorption and re-emission of photons as well as heat transfer quickly make the problem challenging to solve iteratively. Exact solutions to the photon transport equation are rare [27]. Modest describes four primary difficulties in the formulation of an analytical solution. These four categories are geometry, temperature field, scattering, and radiative properties.

Geometry issues arise when trying to solve the problem in anything greater than 1D. Temperature field issues arise when the transport equation is coupled to heat conduction, convection, and material energy balance. If there is scattering present in the problem, then the equation becomes an integro-differential equation. Radiative property issues are encountered when such properties are dependent on the temperature field and/or the frequency of radiation.

Thus far the work done to solve the radiation transport equation has been rather extensive. Generally, practical problems of interest have zones of great optical thickness and thus standard source iteration will always be slow to converge. Generally, coherent scattering is ignored. However, the absorption-emission coupling can be thought of as an effective scattering term.

Acceleration/preconditioning of the standard fixed point iteration has been considered by Larsen [23] and Gol'din [39]. Larsen specifically focused on the use of the grey radiative transport solution to precondition the multi-frequency transport equation system. In his study, Larsen compared a synthetic acceleration method with the so called Lambda Iteration method, which is functionally the same as a fixed point iteration. It was found that for cases where strong coupling is present,

both the fixed point iteration and the proposed acceleration method were slow to converge. Such strong coupling occurs when the time step is very large and/or when the opacity is very large. Gol'din used the QD method in a similar manner to accelerate the convergence of the fixed point tightly coupled radiative transfer system.

When discretizing the streaming term in the photon transport equation there are several options available. The simplest and easiest option is to use diamond differencing–which approximates the cell-average flux as a simple average of the two cell edge values. However, this method is known to work poorly in the thick diffusive limit. While it is incredibly easy to implement, there is also a danger of incurring negative energy densities–which are unphysical.

The step characteristic method is fairly simple to implement and essentially makes use of an integrating factor to form a characteristic closure equation. However, the emission source must be assumed constant through the cell. When this method is used, the thick diffusion limit is not maintained and the numerical scheme will be in error. The linear characteristic method does maintain the thick diffusion limit, but is more involved as it assumes that the energy source is linear over the cell.

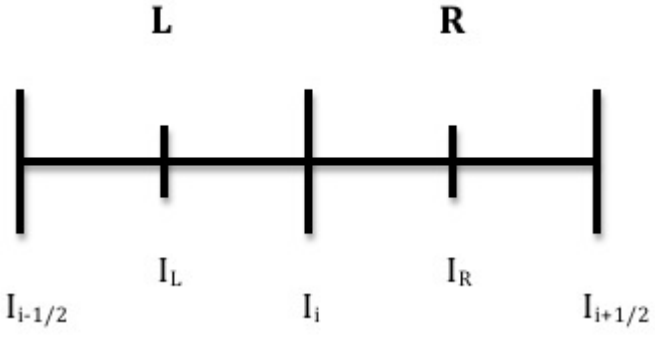

Figure 1: 1D Simple Corner Balance Node Cell

An alternative discretization suggested by Adams [25] is the Simple Corner Balance, which has been shown to work well in the thick diffusive limit for difficult problems on coarse grids. This is the scheme used in this thesis to discretize the streaming term. The transport equation is integrated over each half-cell as shown in Figure 1. The numerical approximations utilized in this method are further discussed in the Methods chapter. This method has been successfully implemented by Palmer [30] in curvilinear geometries.

As an alternative to deterministic methods, Monte Carlo transport can be used. Here, individual particles (in this case photons) are simulated in the given environment and statistical distributions are built from the use of psuedo-random number sampling. Fleck and Cummings [11] have greatly expounded upon the implicit Monte Carlo method to solve the radiative transport problem and further work in this field continues to be carried out by Gentile [13], Densmore [8], Urbatsch, Evans, Brown [26], Cleveland [7], and others.

1.2 Thesis Overview

The remainder of the thesis is organized in the following way:

II . Chapter 2 contains the underlying physical equations of the general problem. The numerical schemes used to solve the problem on the computer are then discussed. Concise descriptions of the discretization schemes used for each phase space element are given with formulations for the different test problems also presented. The JFNK method is described, followed by various preconditioner formulations. An overview of the structure of the software used to test these methods is also provided.

III . Chapter 3 describes the numerical test problems as well as the results associated with each test case. The first problem examined is the Su and Olson problem [36]. The second problem examined is the more non-linear problem of Chang [5] with heterogeneous materials. The solutions for each problem are presented in comparison with previously published solutions. Timing results for the various preconditioners are also given and remarks are made concerning the effectiveness of various preconditioners. When solutions could not be obtained, the causes for the method's breakdown are given.

IV . Chapter 4 contains the conclusions obtained from this work. A suggested optimal preconditioning strategy is given based on the various formulations used here. A more general conclusion is reached on the effectiveness of this JFNK formulation compared to that of Britton Chang's Photon Free Method.

2 Methods

2.1 Introduction

In this chapter the analytic transport equations of interest are presented and the numerical discretization of each independent variable is discussed. The Jacobian Free Newton Krylov (JFNK) method is presented and the direct implementation within the computer simulation is explained. Lastly, the physics-based preconditioners used in this study are derived.

2.2 Analytic Equations

The photon transport equation as given in equation (1) represents the transient multi-frequency problem in slab geometry. The loss terms on the left hand side include the streaming operator and the interaction of the radiation intensity with matter. On the right hand side there are three source terms: the in-scattering term, the re-emission term from interacting matter, and an arbitrary internal source $[Q]$.

$$\frac{1}{c}\frac{\partial I(x,\mu,t,\nu)}{\partial t}+\mu\frac{\partial I(x,\mu,t,\nu)}{\partial x}+\sigma_t(x,\nu,T)I(x,\mu,t,\nu)= \tag{1}$$

$$\frac{c}{2}\left[\sigma_a(x,\nu,T)B(\nu,T)+\sigma_s(x,\nu,T)E(x,t,\nu)+Q(x,\nu)\right]$$

where E is the energy density $[E]$:

$$E(x,t,\nu)=\frac{1}{c}\int_{-1}^{1} I(x,\mu,t,\nu)d\mu \tag{2}$$

and B is the Planck function:

$$B(\nu,T)=\frac{8\pi h\nu^3}{c^3}\left(e^{\frac{h\nu}{kT}}-1\right)^{-1}. \tag{3}$$

The re-emission source term is coupled via the energy density and temperature

to the material energy balance equation (4).

$$C_v(x,T)\frac{\partial T(x,t)}{\partial t} = \int_0^\infty \int_{-1}^1 c\sigma_a(x,\nu,T)\left(I(x,\nu^{'},\mu^{'}) - B(\nu^{'},T)\right) d\mu^{'} d\nu^{'}. \qquad (4)$$

In the photon transport equation the radiation intensity $[I]$ is a function of position, frequency, angle, and time. Temperature $[T]$ is a function of spatial position and time. The Planck distribution is a function of frequency and temperature. Opacity may be a function of spatial location (heterogeneous system), frequency, and temperature. In the material energy balance equation the specific heat with respect to constant volume $[C_v]$ may be dependent on spatial location and temperature. This system of equations is non-linear and non-trivial to solve numerically.

When integrated over all frequencies, the Planck function yields:

$$\int_0^\infty d\nu\ B(\nu,T) = acT^4. \qquad (5)$$

The radiation constant, a, is:

$$a = \frac{8\pi^5 k^4}{15h^3c^3}. \qquad (6)$$

The variables of interest are the energy density and the temperature. The non-linearities in these coupled equations arise from the T^4 temperature term as well as temperature and energy dependent specific heat and opacities.

2.3 Discretized Equations

In order to solve the equations described in the previous section on a computing system, one must discretize the independent variables: frequency, time, spatial position, and angle.

2.3.1 Frequency Discretization

Frequency discretization is handled with the multi-group method. A discrete set of frequency groups is chosen and a set of coupled photon transport and material energy balance equations are formed for each group. The group intensity is the integral of the frequency-dependent intensity over the frequency group:

$$I_n(z,\mu,t) = \int_{\nu_n}^{\nu_{n+h}} d\nu I(z,\nu,\mu,t). \tag{7}$$

Using this definition, and integrating equations (1) and (4) over a group, yields:

$$\frac{1}{c}\frac{\partial I_n(x,\mu,t)}{\partial t} + \mu\frac{\partial I_n(x,\mu,t)}{\partial x} + \sigma_{t_n}(x,T)I_n(x,\mu,t) = \tag{8}$$

$$\frac{c}{2}\left[\sigma_{a_n}(x,T)B_n(T) + \sigma_{s_n}(x,T)E_n(x,t) + Q_n(x)\right],$$

$$C_v(x,T)\frac{\partial T(x,t)}{\partial t} = c\sum_{n=1}^{N}\sigma_{a_n}(x,T)\int_{-1}^{1} d\mu I_n(x,\nu_n,t) - \sum_{n=1}^{N}\sigma_{a_n}(x,T)B_n(T). \tag{9}$$

The group averaged cross sections are described as follows:

$$\sigma_{z_n}(x,T) = \frac{1}{\nu_{n+\frac{1}{2}} - \nu_{n-\frac{1}{2}}}\int_{\nu_{n-\frac{1}{2}}}^{\nu_{n+\frac{1}{2}}}\sigma_z(x,\nu,T)d\nu. \tag{10}$$

The group Planckian distribution is given as follows:

$$B_n(T) = \int_{\nu_{n-\frac{1}{2}}}^{\nu_{n+\frac{1}{2}}}\frac{4\pi h\nu^3}{c^2}\frac{e^{\frac{-h\nu}{kT}}}{1-e^{\frac{-h\nu}{kT}}}d\nu. \tag{11}$$

For the Su and Olson problem [36], the frequency integrated Planckian source term is given in (12):

$$\int_0^\infty \sigma_a(x,\nu,T)B(\nu,T) = \bar{\sigma}_a(x)acT^4, \tag{12}$$

where $\bar{\sigma}_a(x)$ is the Planckian mean opacity. Chang's first problem in [5] defines the frequency integrated Planckian source term as shown in (13):

$$\int_0^\infty \sigma_a(x,\nu,T)B(\nu,T) = \sum_{n=1}^{N} \sigma_{a_n}(x,T)B_n(T). \tag{13}$$

2.3.2 Time Discretization

The time variable is discretized using a first order backward Euler formulation. This allows for the problem to be solved fully implicitly as all parameters are evaluated for the future time step. If the time step size is too large, the problem will exhibit significant error, even if it is stable. The general discretization of a first order differential for an arbitrary parameter is given in equation (14).

$$\frac{\partial F(t)}{\partial t} \approx \frac{F(t^{k+1}) - F(t^k)}{\Delta t}. \tag{14}$$

The resulting discretized system of equations are given in (15) and (16). Here the index k is used to represent a discrete point in time.

$$\frac{1}{c}\frac{I_n^{k+1}(x,\mu) - I_n^k(x,\mu)}{\Delta t} + \mu\frac{\partial I_n^{k+1}(x,\mu)}{\partial x} + \sigma_{t_n}(x,T^{k+1})I_n^{k+1}(x,\mu) =$$

$$\frac{c}{2}[\sigma_{a_n}(x,T^{k+1})B_n(T^{k+1}) + \sigma_{s_n}(x,T^{k+1})E_n^{k+1}(x) + Q_n(x)], \tag{15}$$

$$C_v(x,T^{k+1})\frac{T^{k+1}(x) - T^k(x)}{\Delta t} = \sum_{n=1}^{N}\left[\sigma_{a_n}(x,T^{k+1})E_n^{k+1}(x) - \sigma_{a_n}(x,T^{k+1})B_n(T^{k+1})\right]. \tag{16}$$

Here the time index k is written as a superscript. The error incurred in each equa-

tion from this approximation is described below from modified equation analysis:

$$\varepsilon_I = \frac{1}{2c}\Delta t \frac{\partial^2 I_n^{k+1}(x,\mu)}{\partial t^2} + O\left(\Delta t^2\right), \tag{17}$$

$$\varepsilon_T = \frac{1}{2}\Delta t \; C_v(x, T^{k+1}) \frac{\partial^2 T^{k+1}(x)}{\partial t^2} + O\left(\Delta t^2\right). \tag{18}$$

2.3.3 Spatial Discretization

Discretization in the spatial dimension is accomplished with the use of the Simple Corner Balance (SCB) method. Each node is divided into two half nodes and new angular intensity parameters are introduced (I_L and I_R). These are the left and right half cell intensity values. The Simple Corner Balance method makes use of simple numerical approximations to define the cell edge and whole-cell midpoint values. These approximations are shown in (19). In the following equations the spatial node index is i.

$$I_i = \frac{I_{L_i} + I_{R_i}}{2} \tag{19}$$

$$I_{i-\frac{1}{2}} = \begin{cases} I_{R_{i-1}} & , \; \mu > 0 \\ I_{L_i} & , \; \mu < 0 \end{cases} \tag{20}$$

$$I_{i+\frac{1}{2}} = \begin{cases} I_{R_i} & , \; \mu > 0 \\ I_{L_{i+1}} & , \; \mu < 0 \end{cases} \tag{21}$$

The cell midpoint is calculated using a simple average of the left and right half-cell averages, and the edge values are calculated using upwind differencing. Implementing this spatial discretization scheme into the coupled photon transport

and material energy balance equations gives equations (22) and (23):

$$\Delta x_i \left[I^{k+1}_{L_{i,n}}(\mu) - I^{k}_{L_{i,n}}(\mu) \right] + 2c\mu\Delta t^k \left[\frac{I^{k+1}_{L_{i,n}}(\mu) + I^{k+1}_{R_{i,n}}(\mu)}{2} - I^{k+1}_{i-\frac{1}{2},n}(\mu) \right] +$$

$$c\Delta x_i \Delta t^k \sigma_{t_{n,i}}(T^{k+1}_{L_i}) I^{k+1}_{L_{i,n}}(\mu) = \frac{c}{2}\Delta x_i \Delta t^k \Big[\sigma_{a_{n,i}}(T^{k+1}_{L_i}) B_n(T^{k+1}_{L_i}) +$$

$$\sigma_{s_{n,i}}(T^{k+1}_{L_i}) E^{k+1}_{L_{i,n}} \Big], \tag{22}$$

$$\Delta x_i \left[I^{k+1}_{R_{i,n}}(\mu) - I^{k}_{R_{i,n}}(\mu) \right] + 2c\mu\Delta t^k \left[I^{k+1}_{i+\frac{1}{2},n}(\mu) - \frac{I^{k+1}_{L_{i,n}}(\mu) + I^{k+1}_{R_{i,n}}(\mu)}{2} \right] +$$

$$c\Delta x_i \Delta t^k \sigma_{t_{n,i}}(T^{k+1}_{R_i}) I^{k+1}_{R_{i,n}}(\mu) = \frac{c}{2}\Delta x_i \Delta t^k \Big[\sigma_{a_{n,i}}(T^{k+1}_{R_i}) B_n(T^{k+1}_{R_i}) +$$

$$\sigma_{s_{n,i}}(T^{k+1}_{R_i}) E^{k+1}_{R_{i,n}} \Big]. \tag{23}$$

With these equations it becomes possible to solve for angular intensity values at discrete points given a previous guess for the energy density. This is commonly called a transport sweep. A transport sweep works by starting at a known boundary value and solving the 2x2 system of equation in each cell in the direction of μ until all angular intensities have been calculated. Figure 2 describes a transport sweep along the positive ordinate direction.

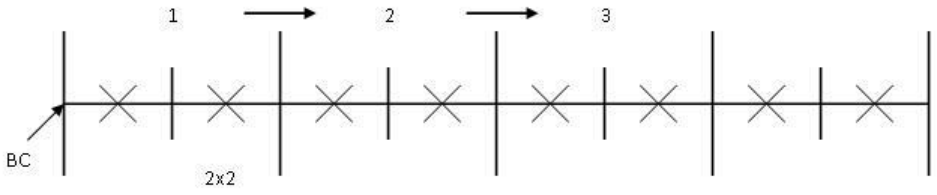

Figure 2: 1D Simple Corner Balance Transport Sweepe

2.3.4 Angular Discretization

In order to evaluate the integrals over all angles (μ), the S_N method is used. This involves using Gaussian quadrature with weights and abscissas to numerically evaluate an integral. For the S_N method, the abscissas are discrete parameters of $cos(\theta)$. Equation (24) describes how the angular integration is accomplished with this method:

$$E(x, \nu, t) \approx \frac{1}{c} \sum_{q=1}^{Q} w_q I(x, \mu_q, t, \nu). \tag{24}$$

2.4 Non-Linear Treatment

The non-linearities in the coupled radiative transfer and material energy balance equations are treated with the Jacobian Free Newton Krylov method. The variables of interest that are being converged in the scheme are the energy density, temperature, and reflecting boundary values. The converged solution to the system should satisfy equation (25),

$$F(u + \delta u) = 0, \tag{25}$$

where u represents the composite vector of all variables of interest, F represents all associated residual equations, and δu represents the step correction for the Newton iteration.

The Newton method involves evaluation of a residual, F, at the current iteration step. The Jacobian matrix of the system must also be formed. A step increment is calculated and added to the solution vector to form the new iteration solution values. The Newton iteration process is outlined in equations (26)-(28):

$$F(u^p) + \delta u^p J^p \approx 0, \tag{26}$$

$$J^p \delta u^p = -F(u^p), \tag{27}$$

$$u^{p+1} = u^p + \delta u^p. \tag{28}$$

Here p is the Newton iteration index and J is the Jacobian matrix $\left(\frac{\partial F}{\partial u}\right)$. In the JFNK method, the Jacobian matrix is not explicitly formed. Instead a finite difference scheme is used to approximate the action of the Jacobian operator on δu. This finite differencing is given in (29):

$$Jv \approx \frac{F(u + \epsilon v) - F(u)}{\epsilon}. \tag{29}$$

Different schemes exist to calculate an appropriate ϵ. The scheme used in this study is given in (30):

$$\epsilon = \frac{1}{N\|v\|_2} \sum_{i=1}^{N} (a|u_i^p| + a), \tag{30}$$

where v is the Krylov vector, N is the system dimension, u_i^p is the system parameter at the current iterate, and a is a constant with a magnitude close to the square root of machine precision. In this thesis, $a = 1.5 \times 10^{-08}$. There are other methods for the choice of ϵ that are described in Knoll and Keyes' JFNK review paper [18]. The calculation of δu is performed using GMRES—a non-stationary iterative Krylov solver. The GMRES algorithm from the PETSC (Portable Extensible Toolkit for Scientific Computing) framework produced by Argonne National Laboratory [3] was used in this research.

Approximating the action of the Jacobian on a vector requires that two system vectors be calculated—one for the solution at the current iteration and one for

a perturbed solution at the current iteration. The Newton iteration is continued until the desired norm of the residual vector is smaller than a specified tolerance. In this study the L_2 norm is used for the purposes of determining the error of the system.

The residual vector is calculated through three variable-specific residual equations for the parameters E, T, and $I_{boundary}$. The residual equations for temperature [(31),(32)] and reflecting boundary intensities (33) are straightforward and included below. The noted discretization schemes of section 2.3 are applied to the material energy balance equation and boundary value residual:

$$F^{k+1}_{T_{L_{i,n}}} = \frac{c_v(x_{L_i}, T^{k+1}_{L_{i,n}})}{\Delta t}\left(T^{k+1}_{L_{i,n}} - T^{k}_{L_{i,n}}\right) - c\sum_{n=1}^{N}\sigma_{a_{n,i}}(T^{k+1}_{L_{i,n}})\left(E^{k+1}_{L_{i,n}} - B_{L_n}(T^{k+1}_{L_{i,n}})\right), \tag{31}$$

$$F^{k+1}_{T_{R_{i,n}}} = \frac{c_v(x_{R_i}, T^{k+1}_{R_{i,n}})}{\Delta t}\left(T^{k+1}_{R_{i,n}} - T^{k}_{R_{i,n}}\right) - c\sum_{n=1}^{N}\sigma_{a_{n,i}}(T^{k+1}_{R_{i,n}})\left(E^{k+1}_{R_{i,n}} - B_{R_n}(T^{k+1}_{R_{i,n}})\right), \tag{32}$$

where $\bar{\sigma}_a$ is the Planck mean opacity.

$$F_{I_{BC}} = I(0, \mu_q, t_{k+1}, \nu_n) - I^{k+1}_{L_{1,n}}(\mu_{Q-q}) \quad , \quad \mu_q > 0. \tag{33}$$

$$F_{I_{BC}} = I(Z, \mu_q, t_{k+1}, \nu_n) - I^{k+1}_{R_{Z,n}}(\mu_{Q-q}) \quad , \quad \mu_q < 0.$$

For known boundary conditions, the boundary residual equations are neglected since the residual for such cases will always be zero.

The residual equation for the energy densities is obtained by integrating the

photon transport equation over angle:

$$F_{E_L}(x,t,\nu) = \frac{1}{c}\frac{\partial E_L(x,t,\nu)}{\partial t} + \left(\frac{\partial J(x,t,\nu)}{\partial x}\right)_L + \sigma_a(x,\nu,T_L(x,t))\left[E_L(x,t,\nu) - B_L(\nu,T_L(x,t)) - Q_L(x,t,\nu)\right] \quad (34)$$

$$F_{E_R}(x,t,\nu) = \frac{1}{c}\frac{\partial E_R(x,t,\nu)}{\partial t} + \left(\frac{\partial J(x,t,\nu)}{\partial x}\right)_R + \sigma_a(x,\nu,T_R(x,t))\left[E_R(x,t,\nu) - B_R(\nu,T_R(x,t)) - Q_R(x,t,\nu)\right] \quad (35)$$

This has the effect of reducing the overall size of the system being converged– improving performance by requiring fewer operations and less memory per Newton step. However, the amount of total work may be more since more Newton iterations may be required to converge the system to a given tolerance. Unlike the photon free method developed by Chang [5], the radiative transfer component is included. This allows for targeted preconditioning of the operators in the energy density residual.

Here, J is the photon flux. Because Eqs. (34)+ (35) are integrated over angle, the resulting residual vector will be of a smaller size than if the full set of photon transport equations were used. Applying the discretization schemes discussed in 2.3 to (34) and (35) gives us equations (36) and (37):

$$F^{k+1}_{E_{L_{i,n}}} = \frac{1}{c}\frac{E^{k+1}_{L_{i,n}} - E^{k}_{L_{i,n}}}{\Delta t} + \left(\frac{J^{k+1}_{i,n} - J^{k+1}_{i-\frac{1}{2},n}}{\frac{\Delta x}{2}}\right) + \sigma_{a_{n,i}}(T^{k+1}_{L_{i,n}})\left[E^{k+1}_{L_{i,n}} - B_{L_n}(T^{k+1}_{L_{i,n}})\right] - Q^{k+1}_{L_{i,n}}, \quad (36)$$

$$F_{E_{R_{i,n}}}^{k+1} = \frac{1}{c}\frac{E_{R_{i,n}}^{k+1} - E_{R_{i,n}}^{k}}{\Delta t} + \left(\frac{J_{i+\frac{1}{2},n}^{k+1} - J_{i,n}^{k+1}}{\frac{\Delta x}{2}}\right) + \sigma_{a_{n,i}}(T_{R_{i,n}}^{k+1})\left[E_{R_{i,n}}^{k+1} - B_{R_n}(T_{R_{i,n}}^{k+1})\right] - Q_{R_{i,n}}^{k+1}. \tag{37}$$

The photon flux values are calculated by performing Gaussian quadrature integration using known angular intensities. These angular intensities are calculated in a separate transport sweep of the original photon transport equation. The currents are built using equations (38), (39), and (40):

$$J_{i,n}^{k+1} \approx \frac{1}{2}\sum_{q=1}^{Q} \omega_q \mu_q \left[I_{L_{i,n}}^{k+1}(\mu_q) + I_{R_{i,n}}^{k+1}(\mu_q)\right], \tag{38}$$

$$J_{i-\frac{1}{2},n}^{k+1} = \sum_{q=1}^{Q} \omega_q \mu_q I_{i-\frac{1}{2},n}^{k+1}(\mu_q), \tag{39}$$

$$J_{i+\frac{1}{2},n}^{k+1} = \sum_{q=1}^{Q} \omega_q \mu_q I_{i+\frac{1}{2},n}^{k+1}(\mu_q), \tag{40}$$

where

$$I_{i-\frac{1}{2},n}^{k+1}(\mu_q) = \begin{cases} I_{n,q}^{k+1}(0) & , i = 1,\ \mu > 0, \\ I_{R,i-1,n}^{k+1} & , i \neq 1,\ \mu > 0, \\ I_{L,i,n}^{k+1} & , \mu < 0, \end{cases} \tag{41}$$

$$I_{i+\frac{1}{2},n}^{k+1}(\mu_q) = \begin{cases} I_{n,q}^{k+1}(Z) & , i = NZ,\ \mu < 0, \\ I_{L+1,n,q}^{k+1} & , i \neq NZ,\ \mu < 0, \\ I_{R,n,q}^{k+1} & , \mu > 0. \end{cases} \tag{42}$$

Here Z is the length of the problem and NZ is the maximum number of cells. Once all the entries of the system residual vector are calculated then GMRES can be called to solve for δu at the current Newton iteration step.

2.5 Preconditioning

When using the JFNK method, good choices of preconditioning are required to realize improvements in the convergence rate. Without preconditioning, GMRES may be very slow to converge during each Newton iteration. The goal of a good preconditioner is to cluster eigenvalues in the linear system, minimizing Krylov iterations per Newton iteration [20]. Right preconditioning of the linear Newton problem is expressed in equation (43):

$$J^p M^{-1} M v = -F(u^p). \tag{43}$$

The matrix M is the preconditioning operator and its inverse is the preconditioning process. The matrix M should be close to J, but not exactly J. This is implemented in the matrix-vector multiply formulation used to approximate the action of the Jacobian on a vector (44):

$$J^p M^{-1} y \approx \frac{F(u^p + \epsilon M^{-1} y) - F(u^p)}{\epsilon}. \tag{44}$$

This is generally a two step process involving the calculation of $M^{-1}y$ followed by the matrix-vector multiply. The calculation of $M^{-1}y$ may be done with direct or iterative solver methods. This is an added cost to the simulation.

In this work, several different physics-based preconditioners are examined: multi-frequency radiation diffusion, infinite medium, linearized temperature derivative of the material energy balance, grey radiation transport, and S_2 multi-frequency radiation transport. Each of these preconditioners neglects the non-linear coupling gradients in the Jacobian. Two more preconditioners are explored: infinite medium with grey coupling and infinite medium with full frequency group coupling. These coupling preconditioners will show the importance of including the non-linear cou-

pling physics for preconditioning the system. All preconditioners are derived using the "delta" method described by Knoll and Keyes [18].

2.5.1 Radiation Diffusion

The radiation diffusion preconditioner is formed by taking the 1st angular moment of the photon transport equation and using Fick's Law (45) to replace the gradient of the flux with a diffusion operator:

$$\frac{\partial J}{\partial x} \approx -D\frac{\partial^2 E}{\partial x^2}, \tag{45}$$

$$D = \frac{1}{3\sigma_t}. \tag{46}$$

Diffusion allows for direct preconditioning of the energy density variable without performing any transport sweeps. This saves memory and the time spent in one call to the preconditioner.

The multi-frequency form of the transient radiation diffusion equation is shown in equation (47):

$$\frac{1}{c}\frac{\partial \delta E(x,t,\nu)}{\partial t} - \frac{1}{3\sigma_t(x,\nu,T)}\frac{\partial^2 \delta E(x,t,\nu)}{\partial x^2} + c\sigma_t(x,\nu,T)\delta E(x,t,\nu) = -F_E(x,t,\nu). \tag{47}$$

The diffusion operator is solved using a Gauss-Seidel iteration. Since the preconditioner need not be perfect in inversion, the number of Gauss-Seidel iterations can be very few and still give quality preconditioning. The discrete radiation equation for the left half and right cells are:

$$\frac{1}{c}\frac{\delta E^{k+1}_{L_{i,n}} - \delta E^{k}_{L_{i,n}}}{\Delta t} - \frac{1}{3\sigma_t(x_{L_i},\nu_n,T^{k+1}_{L_i})}\frac{\delta E^{k+1}_{R_{i-1,n}} - 2\delta E^{k+1}_{L_{i,n}} + \delta E^{k+1}_{R_{i+1,n}}}{\Delta x^2} +$$

$$c\sigma_t(x_{L_i},\nu_n,T^{k+1}_{L_i})\delta E^{k+1}_{L_{i,n}} = -F^{k+1}_{E_{L_{i,n}}}, \tag{48}$$

$$\frac{1}{c}\frac{\delta E_{R_{i,n}}^{k+1} - \delta E_{R_{i,n}}^{k}}{\Delta t} - \frac{1}{3\sigma_t(x_{R_i}, \nu_n, T_{R_i}^{k+1})}\frac{\delta E_{L_{i,n}}^{k+1} - 2\delta E_{R_{i,n}}^{k+1} + \delta E_{L_{i+1,n}}^{k+1}}{\Delta x^2} +$$

$$c\sigma_t(x_{R_i}, \nu_n, T_{R_i}^{k+1})\delta E_{R_{i,n}}^{k+1} = -F_{E_{R_{i,n}}}^{k+1}. \tag{49}$$

2.5.2 Grey Radiation Transport

Preconditioning based on the grey radiation transport operator is relatively straightforward. The energy density residual is passed into the transport sweep function as the right hand side term (50):

$$L\delta I = -\frac{F_E}{2}, \tag{50}$$

where $L = \left[\frac{1}{c}\frac{\partial}{\partial t} + \mu\frac{\partial}{\partial x} + \bar{\sigma}_t(x, T)\right]$.

The transport sweep action is essentially the inverse of the transport operator. The transport object that is used to perform the sweep action is constructed with one energy group. The Planckian mean is used for all opacity values in the grey problem. Boundary values for the angular intensity, I, are calculated by taking the group weighted average of the multi-frequency boundary conditions. Since each group is given equal weight, the evaluation is as shown in (51):

$$\bar{I}(x, \mu, t) = \frac{\int_\nu d\nu I(x, \mu, t, \nu)}{\int_\nu d\nu}. \tag{51}$$

The solution for δE is constructed numerically via the S_N quadrature approximation shown in (52):

$$\delta E(x, t) \approx \frac{1}{c}\sum_{q=1}^{Q} w_q \delta I(x, \mu_q, t). \tag{52}$$

The transport operator, L, is for the grey problem. For this preconditioner with the transport operator of the transport equation, the non-linear coupling within a

time step is left untouched. For any kind of significant non-linear coupling, this preconditioner alone is unlikely to be of much use. The preconditioning process is accomplished by performing a group collapse of the necessary opacities, boundary values, residual, and initial guess data. A transport object is created with the appropriate parameters and the current GMRES residual is passed into the transport sweep function. The calculated angular intensity values are numerically integrated using gaussian quadrature and the result is stored in the preconditioning vector.

This preconditioner should be most effective when dealing with a problem where there are heterogeneous materials and many frequency groups. Heterogeneous materials are poorly modeled by radiation diffusion and a group collapse of the opacities into a grey model may provides a consistent approximation for efficient preconditioning.

2.5.3 S_2 Multi-frequency Radiation Transport

The S_2 preconditioner approximately inverts the transport operator for a system of equations with two discrete angles. Equation (50) is used again but with the transport operator, L, for the multi-frequency problem. The numerical computation of δE is performed with a two point Gaussian Quadrature. Such a preconditioner should be useful for speeding up problems when a very large angular quadrature is present. The boundary conditions for the S_2 problem are computed by performing a weighted average with respect to the previous Gaussian quadrature weights:

$$\bar{I}^{BC} = \frac{\sum_{q=1}^{Q} \omega_q I_q^{BC}}{\sum_{q=1}^{Q} \omega_q}. \tag{53}$$

Frequency dependence is maintained since only quadrature reduction is performed.

2.5.4 Infinite Medium

The infinite medium preconditioner is one way of preconditioning the spatial dependence of the energy density residuals. It is easily constructed since the streaming operator in the transport equation is ignored. A cell by cell calculation of the preconditioning effect may be accomplished. No iterative solvers are required and no transport sweeps are needed. This preconditioner option is very cheap, but not a good approximation when spatial effects are dominant. The infinite medium residual formulation is given in equation (54):

$$\frac{1}{c}\frac{\partial E(x,\nu,t)}{\partial t} + \sigma_a(x,\nu,T)E(x,\nu,t) = -F_E(x,\nu,t). \tag{54}$$

Using the delta form of preconditioning, equation (55) is used to precondition the energy density portion of the Jacobian action:

$$\delta E^{k+1}_{L_{i,n}} = \frac{-F_{E_{L_{i,n}}}}{\left[\frac{\Delta x}{2c} + \frac{1}{2}\Delta x \Delta t \sigma_a(x_i,\nu_n,T^{k+1}_{L_{i,n}})\right]} \tag{55}$$

This method does not take into account the the non-linear coupling between the transport and material energy balance equations within the time step.

2.5.5 Grey Non-Linear Coupling

The grey non-linear coupling preconditioner captures the effect of the non-linear coupling between the energy density and temperature within a time step. The streaming operator is not accounted for in this preconditioner, allowing for a cell by cell evaluation of the coupling effect, and may be viewed as a more complex infinite medium approach. For the grey case, a 2x2 matrix must be inverted after using the delta method discussed by Knoll and Keyes [18]. One applies the numerical

discretizations discussed earlier in Eqs. (36) (37) (31) (32). The streaming term is neglected and the following substitutions are made:

$$E^{k+1} = E^* + \delta E, \tag{56}$$

$$T^{k+1} = T^* + \delta T. \tag{57}$$

One more approximation must be made for the T^4 term found in the frequency integrated Planckian. Because this is a preconditioner and it is only required that an approximate answer be calculated for preconditioning purposes, a simple approximation of T^4 will be used:

$$\left(T^{k+1}\right)^4 = \left(T^k\right)^3 T^{k+1}. \tag{58}$$

Inserting (56)-(58) into the discretized transport equation without the streaming term yields:

$$\frac{1}{c}\frac{(E_i^* + \delta E_i) - E_i^k}{\Delta t} + \bar{\sigma}_{a_i}(T_i^k)(E_i^* + \delta E_i) = \bar{\sigma}_{a_i}(T_i^k)ca(T_i^k)^3(T_i^* + \delta T_i) + Q_i^{k+1}, \tag{59}$$

$$c_v(x_i, T_i^{k+1})\frac{(T_i^* + \delta T_i) - T^k}{\Delta t} = \bar{\sigma}_{a_i}(T_i^k)(E_i^* + \delta E_i) - \bar{\sigma}_{a_i}(T_i^k)ca(T_i^k)^3(T_i^* + \delta T_i). \tag{60}$$

These equations can be manipulated into the form:

$$\frac{1}{c}\frac{\delta E_i}{\Delta t} + \bar{\sigma}_{a_i}(T_i^k)\delta E_i - \bar{\sigma}_{a_i}(T_i^k)ca(T_i^k)^3\delta T_i =$$

$$-\left(\frac{1}{c}\frac{E_i^* - E_i^k}{\Delta t} + \bar{\sigma}_{a_i}(T_i^k)E_i^* - \bar{\sigma}_{a_i}(T_i^k)ca(T_i^k)^3T_i^* + Q_i^{k+1}\right), \tag{61}$$

$$c_v(x_i, T_i^{k+1})\frac{\delta T_i}{\Delta t} - \bar{\sigma}_{a_i}(T_i^k)\delta E_i - \bar{\sigma}_{a_i}(T_i^k)ca(T_i^k)^3\delta T_i =$$

$$-\left(c_v(x_i, T_i^{k+1})\frac{T_i^* - T_i^k}{\Delta t} - \bar{\sigma}_{a_i}(T_i^k)E_i^* - \bar{\sigma}_{a_i}(T_i^k)ca\left(T_i^k\right)^3T_i^*\right), \tag{62}$$

which is very similar to the original form of the equations:

$$\frac{1}{c}\frac{\delta E_i}{\Delta t} + \bar{\sigma}_{a_i}(T_i^k)\delta E_i - \bar{\sigma}_{a_i}(T_i^k)ca(T_i^k)^3\delta T_i \approx \frac{-F_{E_i}^{k+1}}{\Delta t}, \tag{63}$$

$$C_v(x_i, T_i^{k+1})\frac{\delta T_i}{\Delta t} - \bar{\sigma}_{a_i}(T_i^k)\delta E_i - \bar{\sigma}_{a_i}(T_i^k)ca(T_i^k)^3\delta T_i \approx \frac{-F_{T_i}^{k+1}}{\Delta t}. \tag{64}$$

Equations (61) and (63) can be written as a linear system where F represents the residual of the system: $M\delta u^p = -F_{u^p}$

The solution vector is:

$$\delta u^p = \begin{bmatrix}\delta E \\ \delta T\end{bmatrix}^p. \tag{65}$$

The source vector is:

$$F_u^p = \begin{bmatrix}F_E \\ F_T\end{bmatrix}^p. \tag{66}$$

The preconditioning matrix M is a 2x2 matrix in each spatial cell, and is defined as:

$$M = \begin{bmatrix}A & B \\ C & D\end{bmatrix}_i^{k+1}, \tag{67}$$

where

$$A_i^{k+1} = \frac{1}{c} + \Delta t\bar{\sigma}_a(x_i), \tag{68}$$

$$B_i^{k+1} = -\Delta t\bar{\sigma}_{a_i}(T_i^k)ac\left(T_i^k\right)^3, \tag{69}$$

$$C_i^{k+1} = -\Delta t\bar{\sigma}_{a_i}(T_i^k), \tag{70}$$

$$D_i^{k+1} = c_v(x_i, T_i^{k+1}) + \Delta t\bar{\sigma}_{a_i}(T_i^k)\left(T_i^k\right)^3. \tag{71}$$

This 2x2 system of equations is directly solved in the grey coupling preconditioner algorithm.

2.5.6 Multi-Frequency Non-Linear Coupling

It is generally not sufficient to simply precondition based on the dominant physics associated with each residual equation. For tightly coupled systems, it may be important to capture the coupling in order to maintain accuracy. The following derivation constructs the multi-frequency preconditioning matrix for the coupling physics. The streaming term is neglected so the linear system of equations can be solved cell-by-cell. Opacities may be lagged at the previous time step since preconditioning need not be exact.

In this preconditioner, we perform a first order Taylor series expansion of the Planckian distribution about the old time step value for temperature. This linearizes the Planckian emission term. This is acceptable in the preconditioner since strict accuracy is not generally required.

$$B_n\left(T_i^{k+1}\right) \approx B_n\left(T_i^k\right) + \left(T_i^{k+1} - T_i^k\right) \frac{\partial B_n(T_i^k)}{\partial T_i^k} \tag{72}$$

Using equations (59), (58), and (72) we find the residual equations become:

$$\frac{1}{c}\frac{(E_{i,n}^* + \delta E_{i,n}) - E_{i,n}^k}{\Delta t} + \sigma_{a_{i,n}}(T_i^k)(E_{i,n}^* + \delta E_{i,n}) =$$

$$\sigma_{a_{i,n}}(T_i^k)\left[B_n\left(T_i^k\right) + \left(T_i^* + \delta T_i - T_i^k\right)\frac{\partial B_n(T_i^k)}{\partial T_i^k}\right] + Q_{i,n}^{k+1}, \tag{73}$$

$$C_v(x_i, T_i^{k+1})\frac{(T_i^* + \delta T_i) - T_i^k}{\Delta t} = \sum_{n=1}^{N} \sigma_{a_{i,n}}(T_i^k)(E_{i,n}^* + \delta E_{i,n})$$

$$- \sum_{n=1}^{N}\left[B_n\left(T_i^k\right) + \left(T_i^* + \delta T_i - T_i^k\right)\frac{\partial B_n(T_i^k)}{\partial T_i^k}\right]. \tag{74}$$

After minor rearrangement, the equations become

$$\frac{1}{c}\frac{\delta E_{i,n}}{\Delta t} + \sigma_{a_{i,n}}(T_i^k)\delta E_{i,n} - \sigma_{a_{i,n}}(T_i^k)\frac{\partial B_n(T_i^k)}{\partial T_i^k}\delta T_i \approx \frac{-F_{E_{i,n}}^{k+1}}{\Delta t}, \tag{75}$$

$$c_v(x_i, T_i^{k+1})\frac{\delta T_i}{\Delta t} - \sum_{n=1}^{N}\sigma_{a_{i,n}}(T_i^k)\left[\delta E_{i,n} - \frac{\partial B_n(T_i^k)}{\partial T_i^k}\delta T_i\right] \approx \frac{-F_{T_{i,n}}^{k+1}}{\Delta t}. \tag{76}$$

A linear system is constructed as in the grey coupling case. The preconditioning matrix M is given in (77) and is of size $(N+1)\times(N+1)$. Use of the SCB discretization doubles the row and column size. Solving this linear system gives the updated approximate step values for energy density and temperature in the general preconditioning procedure. The matrix M and vectors δu and $-F_u$ are shown below.

$$M = \begin{bmatrix} \frac{1}{c} + \Delta t \sigma_{a_{i,1}}(T_i^k) & 0 & \cdots & 0 & -\Delta t \sigma_{a_{i,1}}(T_i^k) \frac{\partial B_1(T_i^k)}{\partial T_i^k} \\ 0 & \frac{1}{c} + \Delta t \sigma_{a_{i,2}}(T_i^k) & 0 & \cdots & -\Delta t \sigma_{a_{i,2}}(T_i^k) \frac{\partial B_2(T_i^k)}{\partial T_i^k} \\ \vdots & & \ddots & & \vdots \\ 0 & \cdots & 0 & \frac{1}{c} + \Delta t \sigma_{a_{i,n}}(T_i^k) & -\Delta t \sigma_{a_{i,n}}(T_i^k) \frac{\partial B_n(T_i^k)}{\partial T_i^k} \\ -\Delta t \sigma_{a_{i,1}}(T_i^k) & -\Delta t \sigma_{a_{i,2}}(T_i^k) & \cdots & -\Delta t \sigma_{a_{i,n}}(T_i^k) & c_v(x_i, T_i^{k+1}) + \Delta t \sum_{n=1}^{N} \sigma_{a_{i,n}}(T_i^k) \frac{\partial B_1(T_i^k)}{\partial T_i^k} \end{bmatrix}, \tag{77}$$

$$\delta \vec{u} = \left[\delta E_1, \delta E_2, \cdots \delta E_n, \delta T\right]^T, \tag{78}$$

$$F(\vec{u}) = \left[F_{E_1}, F_{E_2}, \cdots F_{E_n}, F_T\right]^T. \tag{79}$$

This matrix has a peculiar structure that can be taken advantage of when inverting. There is a single diagonal with the end row and column terms present. In fact, given our system of equations, there are essentially four term structures to account for: the diagonal terms, the end row terms, the end column terms, and the bottom right corner term. Since these coefficient structures are known, one only needs to worry about the correct frequency group data. The procedure follows thus: put all step values of the energy density at each group in terms of the step value of temperature, solve the last row equation for the step temperature value, then loop back and solve for all step energy density values.

A simple symbolic representation of the solution method is presented below. First the matrix is formed given known coefficient forms a, b , c, and d.

$$\begin{bmatrix} a_1 & 0 & \cdots & 0 & b_1 \\ 0 & a_2 & 0 & \cdots & b_2 \\ \vdots & & \ddots & & \vdots \\ 0 & \cdots & & a_g & b_g \\ c_1 & c_2 & \cdots & c_g & d \end{bmatrix} \tag{80}$$

Then the system of equations is formulated and all step values of the energy density are put in terms of the step value of temperature–δT.

$$a_1\delta E_1 + b_1\delta T = -F_{E_1}$$

$$a_2\delta E_2 + b_2\delta T = -F_{E_2}$$

$$\vdots$$

$$a_g\delta E_g + b_g\delta T = -F_{E_g} \tag{81}$$

$$\sum_{i=1}^{g}(c_i\delta E_i) + d\delta T = -F_T \tag{82}$$

$$\delta E_i = \frac{(-F_{E_i} - b_i\delta T)}{a_i} \tag{83}$$

$$\sum_{i=1}^{g}\left(c_i\frac{(-F_{E_i} - b_i\delta T)}{a_i}\right) + d\delta T = -F_T \tag{84}$$

With every δE_i in terms of δT, the last equation can be solved for the value of δT at the current node. Once δT has been solved for on the current node, one can substitute the answer back into the system to calculate δE_i for each energy group.

$$\delta T = \left[\frac{-F_T + \sum_{i=1}^{g}\left(F_{E_i}\frac{c_i}{a_i}\right)}{d - \sum_{i=1}^{g}\left(\frac{b_i c_i}{a_i}\right)}\right] \tag{85}$$

2.6 Code Execution Algorithm

First the initial environment data needs to be established (such as cross sections, global constants, initial conditions, boundary conditions, etc.). Then the first time step solution is undertaken. Inside the time loop, the Newton loop is formed to converge the non-linearities within the time step. Inside the Newton loop, the residual of the system is formed. Then the L_2 norm of this residual is calculated as a means of measuring the error of the solution. If the error is greater than the prescribed tolerance, the Newton loop continues. Using the already formed

Figure 3: Basic Code Execution Algorithm

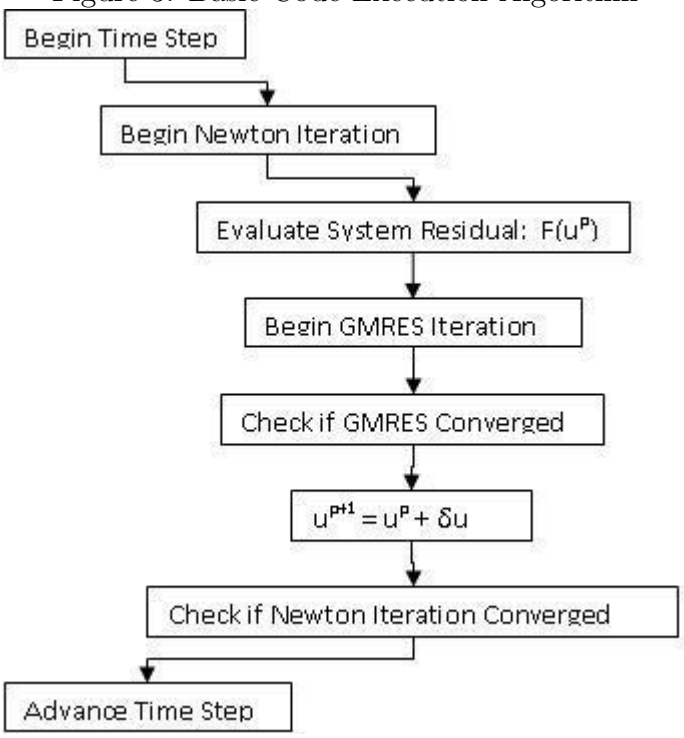

residual, the Krylov solver–GMRES–is called with a pointer to a matrix free Jacobian evaluation function. Matrix free preconditioner functions are applied to δu as GMRES attempts to solve the internal linear system for the step size to add to

the solution vector. After GMRES finishes, δu is added to the previous solution vector and a new solution vector is determined.

The Newton loop continues onwards until the residual error is small enough. Then the Newton loop exits and the system advances forwards a time step until the maximum time is reached for the problem. After determining the final solution vector, the results are output into a data file. The program then deletes/frees all assigned dynamic memory and quits.

3 Results

There are two test problems used to examine the effectiveness of the various preconditioners in the JFNK code. The first problem examined is the Su and Olson benchmark for multi-frequency radiative transport [36]. Su and Olson have provided an analytic equation to evaluate the answer for given spatial and temporal coordinates. The second problem is defined in Chang's paper concerning the Photon Free Method [5]. The specifications and results for these problems are presented in this section. Table 2.6 contains the acronyms used for the various precondition options tested.

Table 1: Preconditioner Options

Acronym	Preconditioner
NULL PC	No preconditioning is used
DT	Derivative with respect to Temperature on Material Energy Balance
RD	Multi-frequency Radiation Diffusion
RDDT	Rad-Diffusion + DT
GT	Grey Transport
GTDT	Grey Transport + DT
S2	S2 Quadrature Multi-frequency Transport
S2DT	S2 + DT
INF	Infinite Medium
MC	Multi-frequency Coupling
INFMC	Infinite Medium + MC

3.1 Su and Olson Multi-Frequency Problem

In this problem, the opacities for the different frequencies are given by the picket fence model. The non-linearities in the system are found in the T^4 and $C_v(T)$ terms. Su and Olson have defined the specific heat with respect to volume as $C_v(T) = 4aT^3$ such that the equation is linearized by solving for T^4. In this code, T is the parameter updated from one Newton iteration to another. This preserves the non-linear treatment of the specific heat. All tabulated results are compared against the analytic results generated from the functional evaluation derived by Su and Olson [36].

This transport problem is essentially a two frequency group problem. More groups may be defined, but the result is the same as source distribution is evenly divided amongst groups. This suggests that the grey transport preconditioner may not be very effective with only two distinct frequency groups present.

3.1.1 Problem Specifications

This one-dimensional slab geometry problem has a reflecting boundary condition on the left and a vacuum boundary on the right. The slab is ten length units longer than the reported significant analytic results in Su and Olson's paper [36]. The slab is initially cold, but an internal source is present from $0 \leq x \leq x_0$, from time $t = 0$ to $t = \tau_0$. The data for this problem is included in Table 3.1.1. Implementation of the picket fence model in addition to the definition of $C_v(T)$ we arrive at equations defined by Su and Olson in their benchmark paper (86) and (87). The picket fence model allows for a full integration over frequency space. The opacities are homogeneous within their respective frequency group and can be treated as constants when dealing with frequency discrete equations.

The radiative transport and material energy balance equations for this problem are:

$$\frac{1}{c}\frac{\partial I_n(x,\mu,t)}{\partial t} + \mu\frac{\partial I_n(x,\mu,t)}{\partial x} =$$

$$\sigma_{a_n}\left[p_n\frac{acT(x,t)^4}{4\pi} - I_n(x,\mu,t)\right] + p_n\frac{Q_0}{4\pi} \quad , \quad n = 1,2, \tag{86}$$

$$4ac\left(T(x,t)\right)^3\frac{\partial T(x,t)}{\partial t} =$$

$$2\pi\sum_{n=1}^{2}\sigma_{a_n}\int_{-1}^{1}d\mu\ I_n(x,\mu,t) - \bar{\sigma}_a ac\left(T(x,t)\right)^4, \tag{87}$$

where p_n is the fractional contribution of a source associated with the frequency group. In a dimensionless form described by Su and Olson, they are

$$\varepsilon\frac{\partial U_n(x,\tau)}{\partial\tau} + \mu\frac{\partial U_n(x,\tau)}{\partial x} = w_n\left[\frac{p_n V(x,\tau)}{2} - U_n(x,\tau)\right] + p_n\frac{\tilde{Q}(x,\tau)}{2} \quad , \quad n = 1,2 \tag{88}$$

$$\frac{\partial V(x,\tau)}{\partial\tau} = \sum_{n=1}^{2} w_n\int_{-1}^{1}d\mu\ u_n(x,\tau) - V(x,\tau), \tag{89}$$

where n is the group-opacity index for the picket fence model. Note that the JFNK algorithm solves for $T(x,\tau)$ instead of $V(x,\tau)$. The analytic solution for this problem solves for $V(x,\tau)$, but with a reference temperature, T_0, of 1.0 it is trivial to convert between the two.

The non-dimensionalized parameters are defined by:

$$\bar{\kappa} = \kappa_a + \kappa_s, \tag{90}$$

$$w_n = \frac{\kappa_n}{\bar{\kappa}}, \tag{91}$$

$$\epsilon = \frac{4a}{\alpha}, \tag{92}$$

$$x = \bar{\kappa} z, \tag{93}$$

$$\tau = \epsilon c \bar{\kappa} t, \tag{94}$$

$$1 = p_1 w_1 + p_2 w_2, \tag{95}$$

$$U_n(x, \mu, \tau) = \frac{I_n(x, \mu, \tau)}{a T_0^4}, \tag{96}$$

$$V(x, \tau) = \left[\frac{T(x, \tau)}{T_0} \right]^4, \tag{97}$$

$$\bar{Q}(x, \tau) = \frac{Q_0(x, \tau)}{ac\bar{\sigma}(x) T_0^4}, \tag{98}$$

$$W(x, \tau) = \int_{-1}^{1} d\mu U(x, \mu, \tau). \tag{99}$$

Table 2: Opacity Specification for Su and Olson Test Cases

Parameter	Case A	Case B	Case C
w_1	1	2/11	2/101
w_2	-	20/11	200/101
w_2/w_1	-	10	100

Case A is the grey case and of little interest in this study of preconditioner behaviour. Case A was used to verify that the code was working properly before

Table 3: Su-Olson Problem Data

Parameter	Value
$\Delta\tau$	0.01
τ_0	10.0
Δx	0.025
NZ	1200
Length	30
S_N Quadrature	12
x_0	5
Q	10
GMRES Tolerance	1e-6
Newton Tolerance	1e-3

Table 4: Table of Physical Constants used for Simulation

Constant	*Value*	*Units*
h	6.626E-34	Joule-sec
k	1.381E-23	Joule/K
c	299.8	Mm/sec

simulating Cases B and C. The more interesting problems are the multi-frequency cases: Case B and Case C. The only difference between Case B and Case C is the ratio $\frac{w_2}{w_1}$.

Initially the slab is cold and the energy density, angular intensity, and temperature are all specified at 0.0. For the first Newton iteration we choose $E = 2.0$ and the angular intensity as constant , isotropic, and equal to 1.0 for the anticipated solution. The initial guess for the temperature distribution in the slab is chosen as 1.0. These chosen anticipated solutions are sufficient to start the Newton iteration converging a solution. Every subsequent time step uses the old time step parameter values as the anticipated solution to start the Newton iteration.

Total CPU time required to reach the simulation end time is tabulated for each preconditioner. The numerical solution is reported graphically and compared

against the analytic solution provided by Su and Olson [36]. The relative error is reported in the tables for the various spatial locations. The analytic benchmark requires a numerical double integration that oscillates from positive to negative depending on the integration boundaries. The analytic integration bounds are from zero to infinity, but discrete boundaries must be chosen for numerical evaluation. An upper limit of 2000π ensures three stable significant figures in the numerical integration.

3.1.2 Results

The total CPU time (seconds) required for each simulation is reported in Tables 5 and 6. The time taken for the NULL PC option–where no preconditioning exists–is prohibitive. Beyond $\tau = 1.0$ it becomes obvious that running without preconditioning is pointless with the JFNK method. However, where a decent preconditioner exists, the JFNK method converges much more quickly.

It is clear that the simple derivative with respect to temperature of the material energy balance equation (DT) is the best preconditioner for this problem. The most effective preconditioners for this problem incorporate this treatment of the material energy balance terms. The dominant physics appears to be the Planckian absorption in the material energy balance equation.

The radiation diffusion (RD) preconditioner is only effective when the DT portion is included. Even the infinite medium preconditioner fails without preconditioning of the material energy balance equation. The time scale of the energy density equation must be well within the time step taken in the simulation or the energy density portion of the problem would require significant preconditioning to run quickly. It is interesting that the increasing number of Gauss-Seidel iter-

Table 5: Case B - Timing Results for Different Preconditioners

Preconditioner	τ=0.1	τ=1.0	τ=10.0	τ=30.0
NULL PC	74.62	210860	-	-
RD1	86.5	204363	-	-
INF	86.0	205922	-	-
DT	6.18	31.7	223.7	547.1
RDDT-1	6.66	34.0	238.8	583.0
RDDT-2	6.76	34.5	242.5	591.1
RDDT-4	7.61	38.8	271.1	660.1
INFGC	8.69	38.4	298.4	778.5
INFMC	6.31	32.3	228.6	558.8
GTDT	8.14	95.8	3658.1	20342
S2DT	3.68	41.3	847.6	3875.8

ations yields no net gain. This could be because the Planckian re-emission term causes a loss of diagonal dominance, or that the cost of the Gauss-Seidel iterations outweighs the benefits from reduced GMRES iterations.

An interesting comparison is the performance of the diffusion and the grey transport preconditioners. Because the S_2 equations in 1D are equivalent to the diffusion approximation, the S_2 transport preconditioner can be compared against the grey transport preconditioner since the implementation for each is similar. However, the S2DT transport preconditioner failed for the case C problem. The solution diverged on the second time step in the GMRES iterations. It is believed that this is due to ray effect phenomena causing a divergence in the GMRES iteration when S2DT preconditioning is applied. In Case B, the S2DT preconditioner was faster than the GTDT preconditioner, especially as the simulation progressed and the shock front slowes and the problem is more diffuse. Because the material is homogeneous, there is not much value in using the transport preconditioners. If there were many material interfaces to deal with, these transport preconditioners would likely help capture this physics.

Table 6: Case C - Timing Results for Different Preconditioners

Preconditioner	$\tau = 0.1$	$\tau = 1.0$	$\tau = 10.0$	$\tau = 30.0$
NULL PC	60.4	205,264	-	-
RD1	86.4	156813	-	-
INF	170.4	189866	-	-
DT	6.23	31.4	218.3	507.8
RDDT-1	6.61	33.2	229.8	537.1
RDDT-2	6.91	35.0	240.8	560.9
RDDT-4	7.39	37.4	258.4	600.1
INFGC	5.15	35.4	298.9	1038.9
INFMC	6.78	36.1	265.9	619.7
GTDT	7.94	94.0	3513.9	19575.8
S2DT	-	-	-	-

The numerical results for Case B of the Su and Olson benchmark that are produced by the JFNK code are given in the many tables and plots that follow. There is appreciable error in the temperature at early times, but this is due to the time step size chosen. The error in the temperature is reduced as the simulation progresses.

There is also a concern at the point when the internal source is turned off. At $\tau = 10.0$ the solution in the first frequency group, $U1$, is less smooth than the analytic solution. The opacity associated with this frequency group is: $\omega_1 = 2/11$. When the internal source is turned off there is nothing driving the transport except the re-emission source term and the previous time step values.

It is possible that with a small re-emission source this degradation in smoothness is caused by ray effects. Aside from these mentioned problem areas, the results are usually within a couple percent of the analytic solution. There was no failure of any of the preconditioner options in this problem. Effectiveness of each preconditioner option is measured by the CPU time required to perform the simulation. All tabulated and plotted results are produced using the infinite medium

multi-coupling (INFMC) preconditioner.

Case B Results for τ=0.1: The tabulated results for the Su/Olson case B problem at $\tau = 0.1$ are listed here along with a plot of the data. The only data point of concern is at $x = 0.5375$, the leading front of the propagating wave.

Table 7: Case B - Relative Error for U1 at $\tau = 0.1$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.04956	0.050	0.084
0.4625	0.04322	0.043	1.508
0.5125	0.01521	0.016	1.669
0.5375	0.00634	0.007	9.621

The non-dimensionalized energy densities are within 10% of the predicted analytical solution. The temperature field is 30% in error compared with the analytic result. This error is associated with the temporal discretization and relatively large timestep being used. As time progresses and the time step is not such a large fraction of the total time taken, this error diminishes.

Table 8: Case B - Relative Error for U2 at $\tau = 0.1$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.04585	0.045	0.799
0.4625	0.04024	0.039	2.053
0.5125	0.01378	0.014	0.777
0.5375	0.00561	0.006	8.19

Table 9: Case B - Relative Error for V at $\tau = 0.1$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.00458	0.004282899	6.487
0.4625	0.00431	0.003899737	9.52
0.5125	0.00105	0.001020441	2.815
0.5375	0.00027	0.000352789	30.66

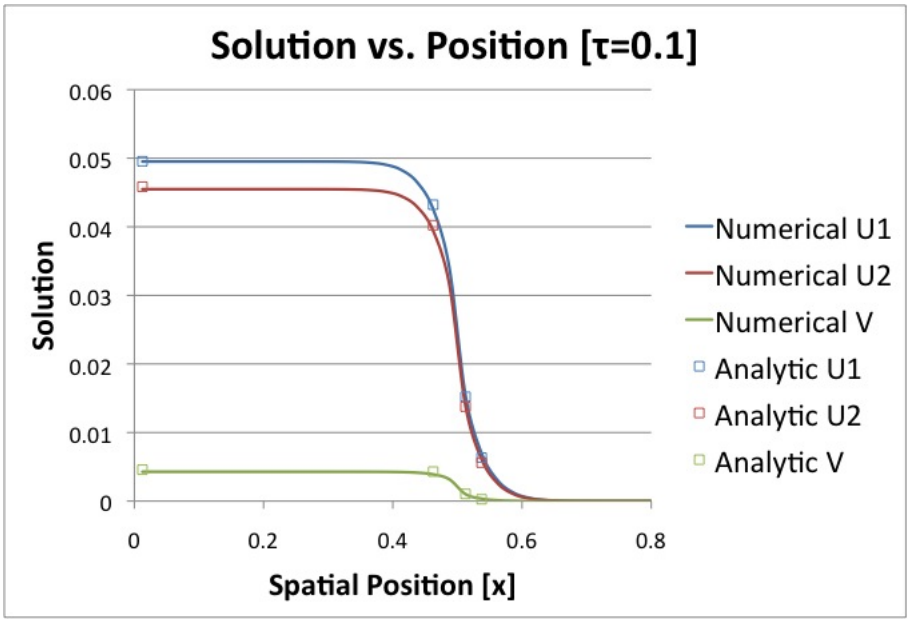

Figure 4: Case B at $\tau = 0.1$ solution

Case B Results for τ=1.0: The tabulated results for the Su/Olson case B problem at $\tau = 1.0$ are listed here along with a plot of the data. The relative error of each data point is small as the wavefront propagates further in the simulation.

Table 10: Case B - Relative Error for U1 at $\tau = 1.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.3988	0.397	0.403
0.3125	0.3492	0.348	0.269
0.5125	0.2163	0.221	2.41
0.6625	0.1226	0.121	1.57
0.9375	0.0439	0.044	0.986

Table 11: Case B - Relative Error for U2 at $\tau = 1.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.26669	0.265	0.477
0.3125	0.23551	0.234	0.486
0.5125	0.12879	0.131	1.77
0.6625	0.05723	0.057	0.986
0.9375	0.01439	0.014	0.195

Table 12: Case B - Relative Error for V at $\tau = 1.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.23881	0.236	1.174
0.3125	0.21399	0.211	1.22
0.5125	0.10731	0.109	1.185
0.6625	0.03873	0.039	0.564
0.9375	0.00673	0.007	2.13

Case B Results for τ=10.0: The tabulated results for the Su/Olson case B problem at $\tau = 10.0$ are listed here along with a plot of the data. The relative error in the $U1$ solution has increased significantly. This could be attributed to the internal source turning off, but data from the next future time point suggests that this is likely due to ray effects.

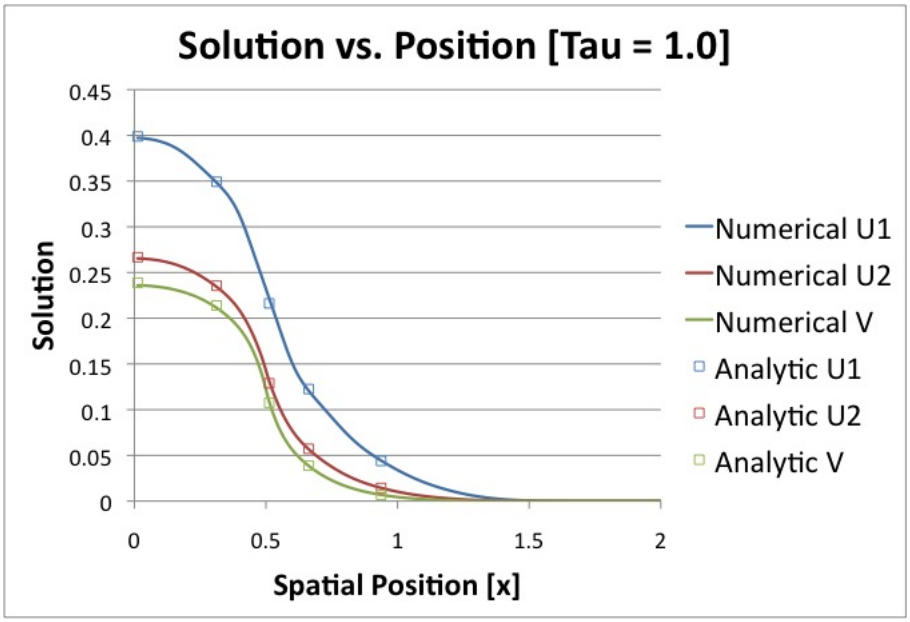

Figure 5: Case B at $\tau = 1.0$ solution

Table 13: Case B - Relative Error for U1 at $\tau = 10.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.9918	0.950	4.17
0.3625	0.9056	0.873	3.56
0.7625	0.6089	0.659	8.18
1.8125	0.3118	0.292	6.23
3.2875	0.1455	0.152	4.32

Table 14: Case B - Relative Error for U2 at $\tau = 10.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	1.0932	1.095	0.333
0.3625	0.9607	0.958	0.25
0.7625	0.5425	0.544	0.357
1.8125	0.1486	0.147	0.864
3.2875	0.0346	0.035	0.067

Table 15: Case B - Relative Error for V at $\tau = 10.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	2.0400	2.0282881	0.572
0.3625	1.7902	1.7811676	0.504
0.7625	0.9951	1.0051479	1.01
1.8125	0.2782	0.27425488	1.433
3.2875	0.0722	0.072286217	0.085

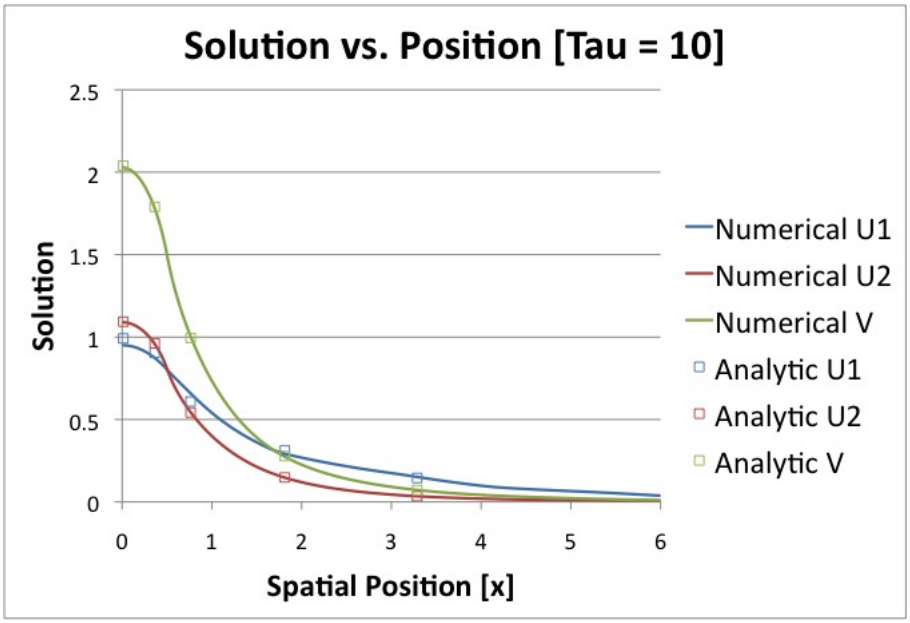

Figure 6: Case B at $\tau = 10.0$ solution

Case B Results for τ=30.0: The tabulated results for the Su/Olson case B problem at $\tau = 30.0$ are listed here along with a plot of the data. The wavefront for U1 seems to have slowed while U2 has decreased since the internal source was turned off. The relative errors associated with $U1$ are still large compared to the analytic solution. This large error is not observed with the $U2$ parameter. Case C results shed more light on the issue at hand.

Table 16: Case B - Relative Error for U1 at $\tau = 30.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.9918	0.950	4.17
0.3625	0.9056	0.873	3.56
0.7625	0.6089	0.659	8.18
1.8125	0.3118	0.292	6.23
3.2875	0.1455	0.152	4.32

Table 17: Case B - Relative Error for U2 at $\tau = 30.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	1.0932	1.090	0.333
0.3625	0.9607	0.958	0.25
0.7625	0.5425	0.544	0.357
1.8125	0.1486	0.147	0.864
3.2875	0.0346	0.035	0.067

Table 18: Case B - Relative Error for V at $\tau = 30.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	2.0400	2.028	0.572
0.3625	1.7902	1.781	0.504
0.7625	0.9951	1.005	1.01
1.8125	0.2782	0.274	1.433
3.2875	0.0722	0.072	0.085

Time Evolution of Wavefronts for Case B: In this sub-section three plots are given that graphically demonstrate the evolution of the variables $U1$, $U2$, and V

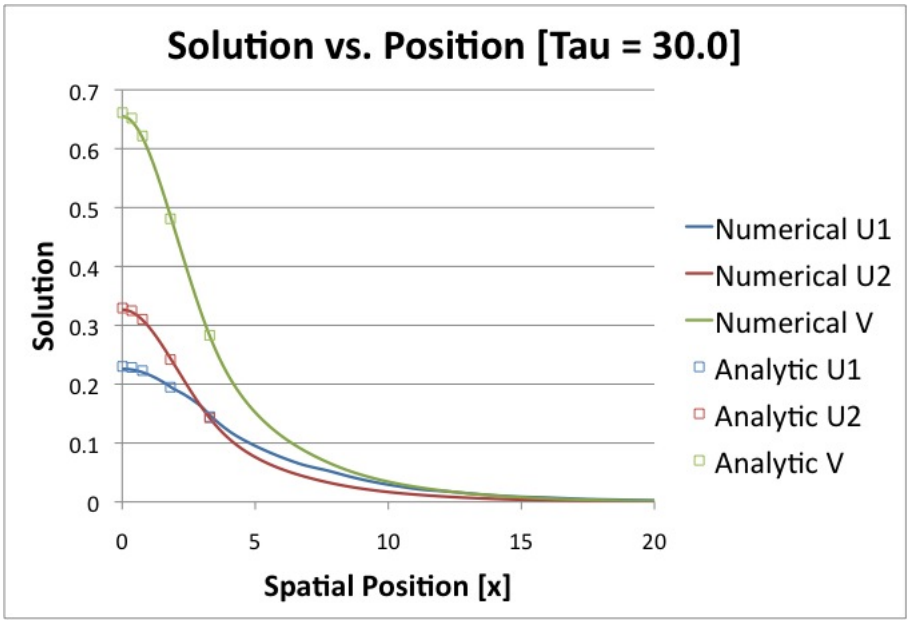

Figure 7: Case B at $\tau = 30.0$ solution

as a function of time. The solution for $U1$ demonstrates some type of stratification as time passes (most easily seen in $\tau = 10.0$ case) where the shape of the wavefront no longer appears to be differentiable. The solution in the $U1$ parameter begins to look piecewise continuous. This is not readily observable to the naked eye in $U2$ and V results. Since $U1$ is the frequency group with the smallest opacity, a greater effect should be observed in Case C, where the opacity is even smaller.

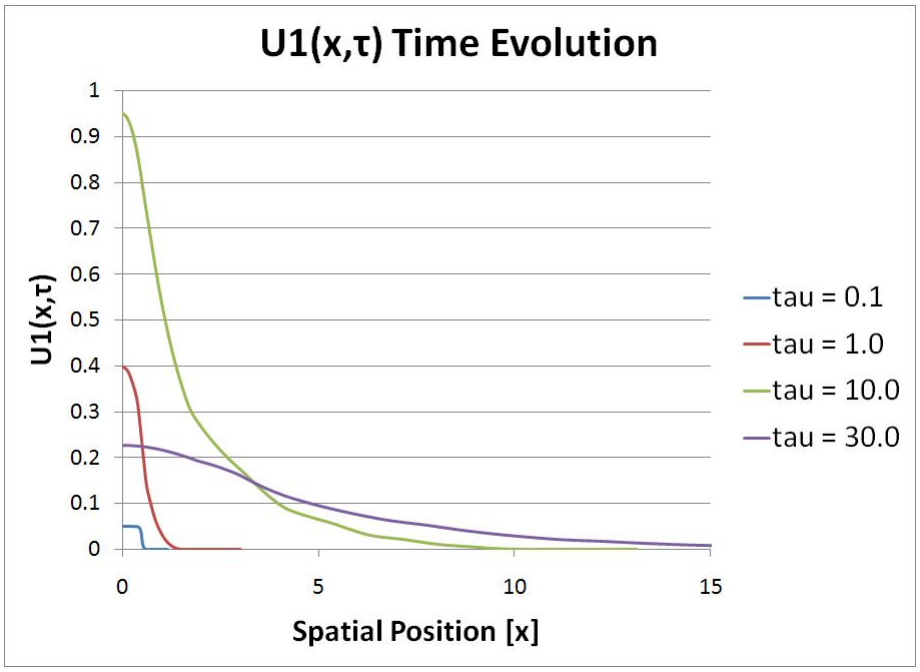

Figure 8: Case B - Time Evolution of U1(x,τ)

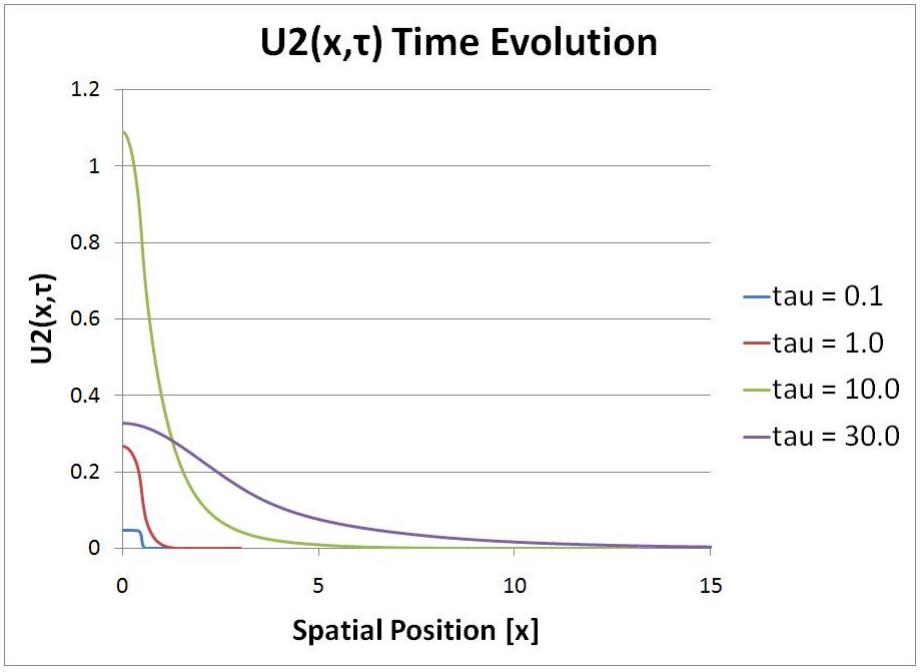

Figure 9: Case B - Time Evolution of U2(x,τ)

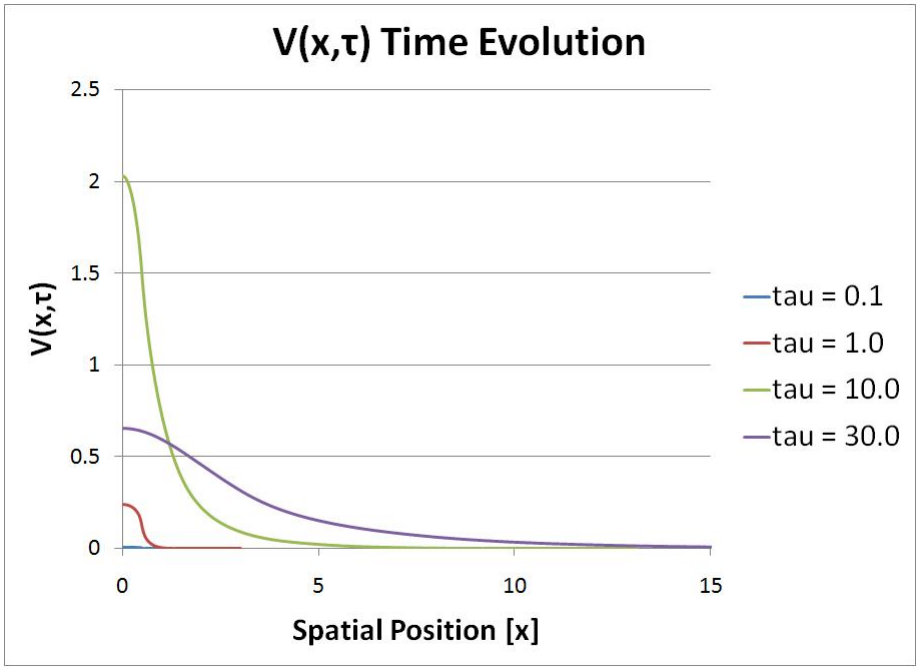

Figure 10: Case B - Time Evolution of V(x,τ)

Case C Results for τ=0.1: The tabulated results for the Su/Olson case C problem at $\tau = 0.1$ are listed here along with a plot of the data. The only data point of concern is once again when $x = 0.5375$, at the leading front of the propagating wave. The non-dimensionalized energy densities are within 10% of the predicted analytical solution. The temperature field is once again 30% in error compared with the analytic result.

Table 19: Case C - Relative Error for U1 at $\tau = 0.1$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.04993	0.050	0.041
0.4625	0.04351	0.043	1.397
0.5125	0.01536	0.016	1.788
0.5375	0.00642	0.007	9.735

Table 20: Case C - Relative Error for U2 at $\tau = 0.1$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.04548	0.045	0.816
0.4625	0.03994	0.039	2.042
0.5125	0.01365	0.014	0.678
0.5375	0.00555	0.006	8.014

Table 21: Case C - Relative Error for V at $\tau = 0.1$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.00454	0.0042	6.74
0.4625	0.00427	0.0039	9.67
0.5125	0.00103	0.0010	2.70
0.5375	0.00027	0.00035	28.9

Case C Results for τ=1.0: The tabulated results for the Su/Olson case C problem at $\tau = 1.0$ are listed here along with a plot of the data. All data points agree with the analytic solution to within 2%.

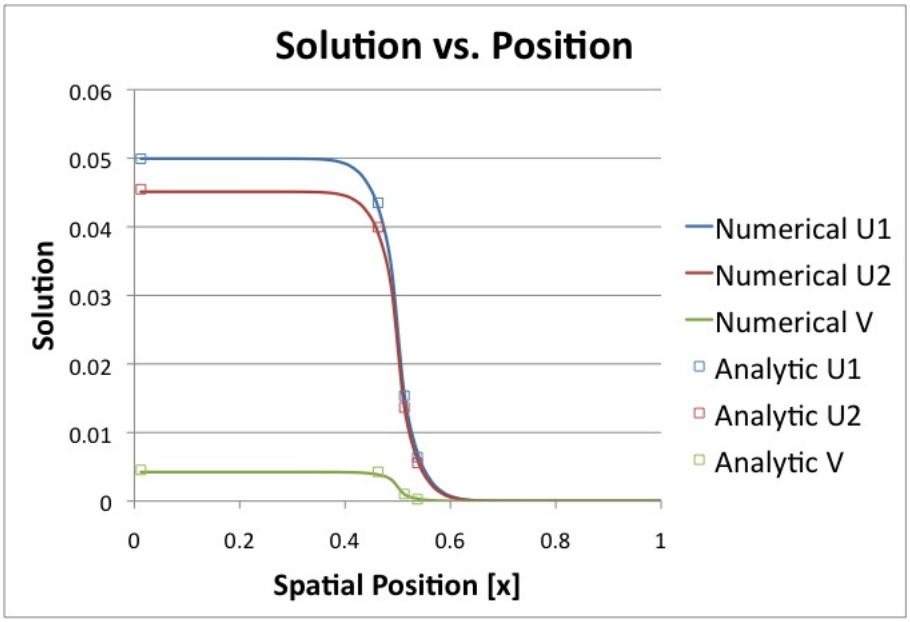

Figure 11: Case C at $\tau = 0.1$ solution

Table 22: Case C - Relative Error for U1 at $\tau = 1.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.42034	0.419	0.385
0.3125	0.36813	0.367	0.217
0.5125	0.23121	0.237	2.49
0.6625	0.13404	0.132	1.59
0.9375	0.04947	0.050	1.16

Table 23: Case C - Relative Error for U2 at $\tau = 1.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.25588	0.255	0.492
0.3125	0.22651	0.225	0.508
0.5125	0.12236	0.124	1.71
0.6625	0.05285	0.052	0.977
0.9375	0.01279	0.0128	0.088

Table 24: Case C - Relative Error for V at $\tau = 1.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.22352	0.221	1.23
0.3125	0.20139	0.199	1.28
0.5125	0.09899	0.100	0.946
0.6625	0.03367	0.0334	0.684
0.9375	0.00531	0.00541	1.82

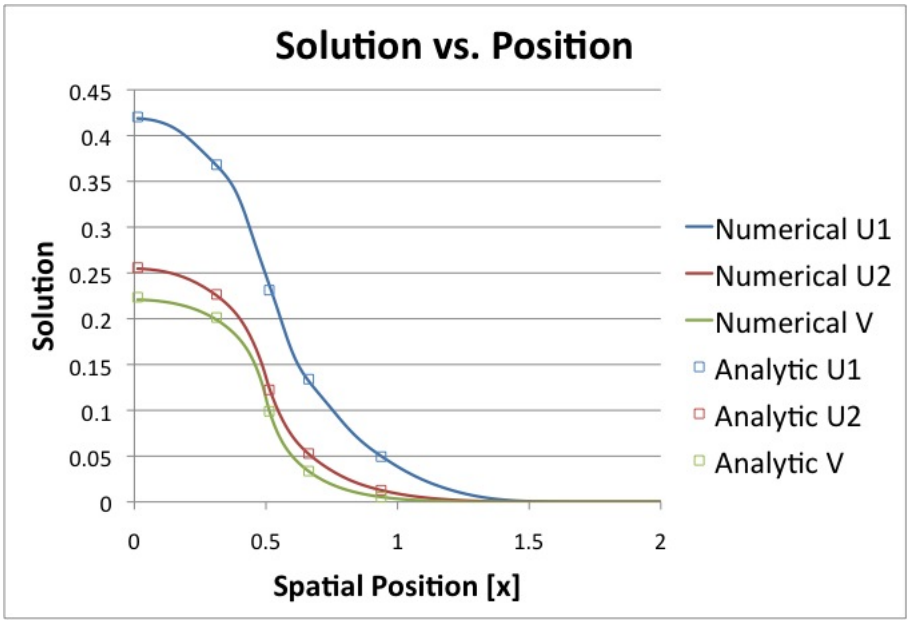

Figure 12: Case C at $\tau = 1.0$ solution

Case C Results for τ=10.0: The tabulated results for the Su/Olson case C problem at $\tau = 10.0$ are listed here along with a plot of the data. The relative error of points for $U1$ are again larger than at early times. This could be attributed to the internal source turning off, but data from the next future time point suggests that this is likely due to ray effects.

Table 25: Case C - Relative Error for U1 at $\tau = 10.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.98996	0.849	14.2
0.3625	0.91519	0.837	8.58
0.7625	0.64933	0.787	21.2
1.8125	0.40745	0.333	18.3
3.2875	0.25501	0.288	12.8

Table 26: Case C - Relative Error for U2 at $\tau = 10.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	1.04982	1.0488	0.197
0.3625	0.90944	0.908	0.165
0.7625	0.47289	0.473	0.017
1.8125	0.08680	0.0862	0.721
3.2875	0.01029	0.0103	0.387

Table 27: Case C - Relative Error for V at $\tau = 10.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	1.96445	1.958	0.306
0.3625	1.69481	1.691	0.242
0.7625	0.85229	0.854	0.168
1.8125	0.14889	0.147	1.27
3.2875	0.01994	0.0201	0.964

Case C Results for τ=30.0: The tabulated results for the Su/Olson case C problem at $\tau = 30.0$ are listed here along with a plot of the data. The relative

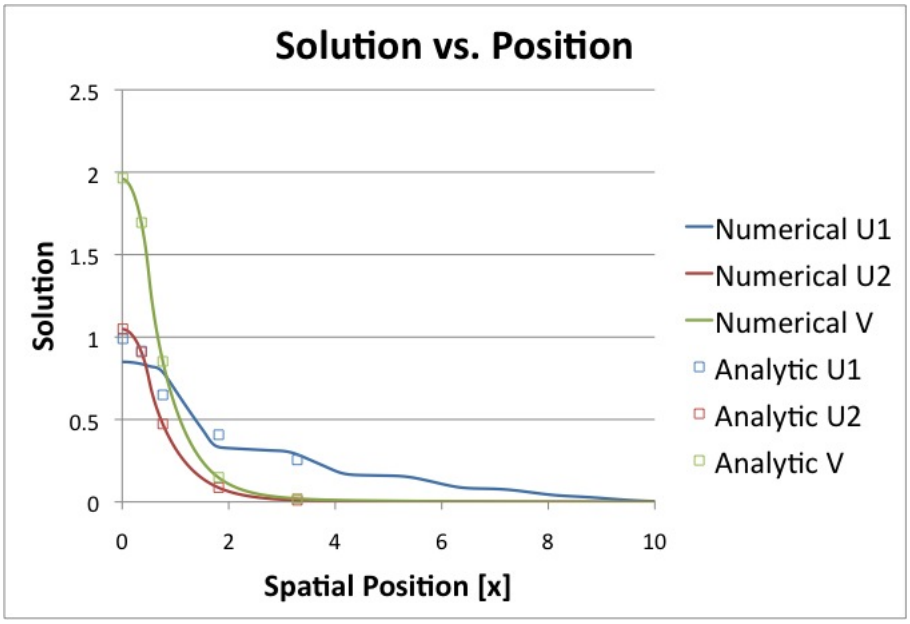

Figure 13: Case C at $\tau = 10.0$ solution

error in the $U1$ solution is very large. The bears no resemblence to the analytic solution. However, $U2$ and V both appear to be correct. $U1$ contributes trivially to V since the opacity for group one is very small. The opacity in group 2 appears to be large enough to be computed accurately.

Table 28: Case C - Relative Error for U1 at $\tau = 30.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.11346	0.0455	59.9
0.3625	0.11308	0.0456	59.7
0.7625	0.11182	0.0458	59.1
1.8125	0.10568	0.0436	58.7
3.2875	0.09598	0.3174	230.6

Table 29: Case C - Relative Error for U2 at $\tau = 30.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.39036	0.385	1.435
0.3625	0.36106	0.376	4.112
0.7625	0.35123	0.348	1.054
1.8125	0.22255	0.226	1.329
3.2875	0.08083	0.0836	3.369

Table 30: Case C - Relative Error for V at $\tau = 30.0$

x	Analytic Result	Numerical Result	Relative Error (%)
0.0125	0.79475	0.782	1.559
0.3625	0.77484	0.763	1.477
0.7625	0.71124	0.703	1.176
1.8125	0.44191	0.448	1.307
3.2875	0.15656	0.163	4.359

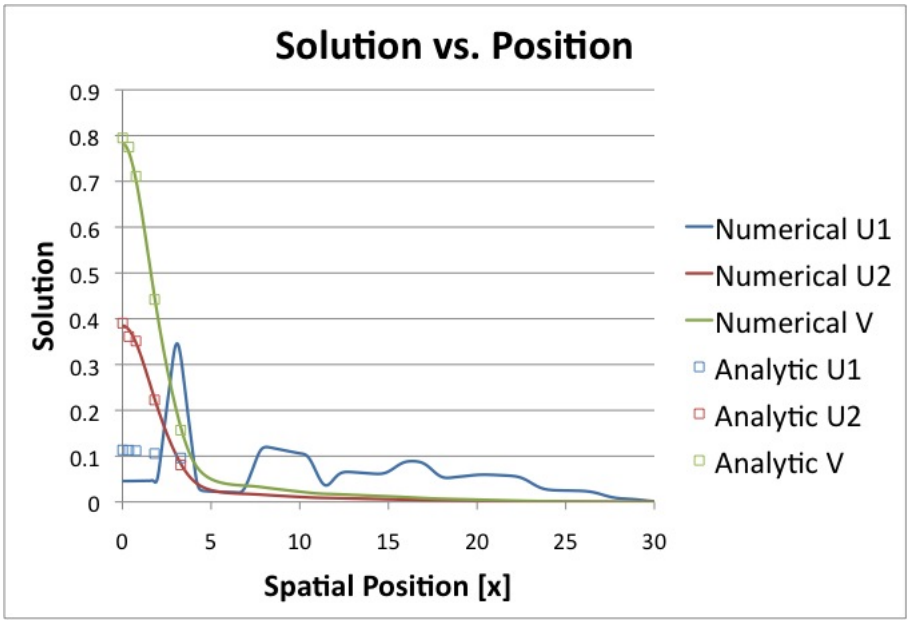

Figure 14: Case C at $\tau = 30.0$ solution

Time Evolution of Wavefronts for Case C: In this sub-section three plots demonstrate the evolution of the variables $U1$, $U2$, and V over simulation time. It is easily observed that the solution for $U1$ demonstrates some type of stratification as time passes. The problem is most evident when the internal source has been turned off. Turning off the source itself likely did not cause the issue, but was merely masking the full extent of the problem. The issue is not readily observable in the $U2$ and V results, probably because the opacity in frequency group one is again very small and the opacity for group two is large enough to be computed accurately.

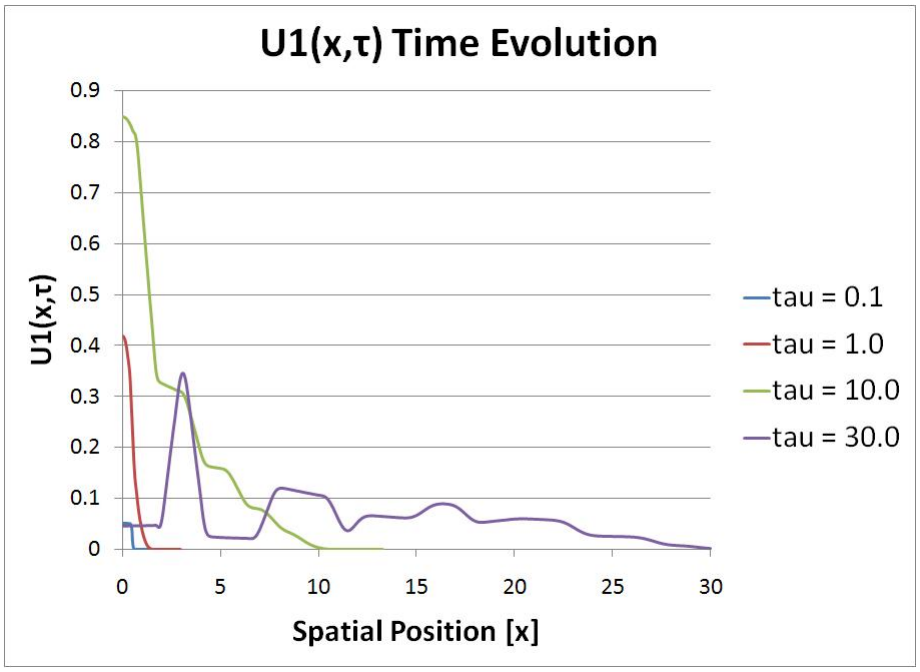

Figure 15: Case C - Time Evolution of U1(x,τ)

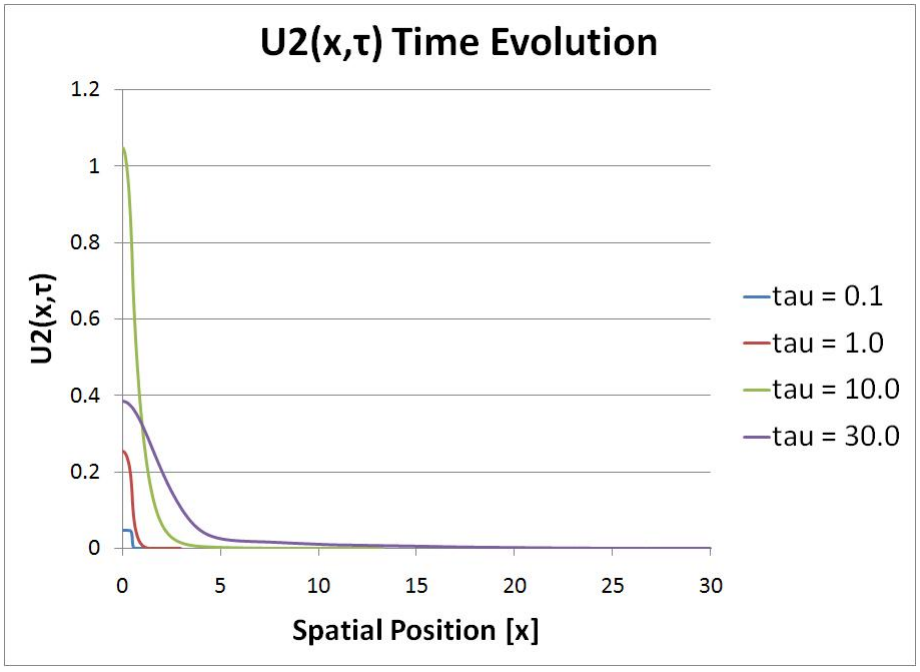

Figure 16: Case C - Time Evolution of U2(x,τ)

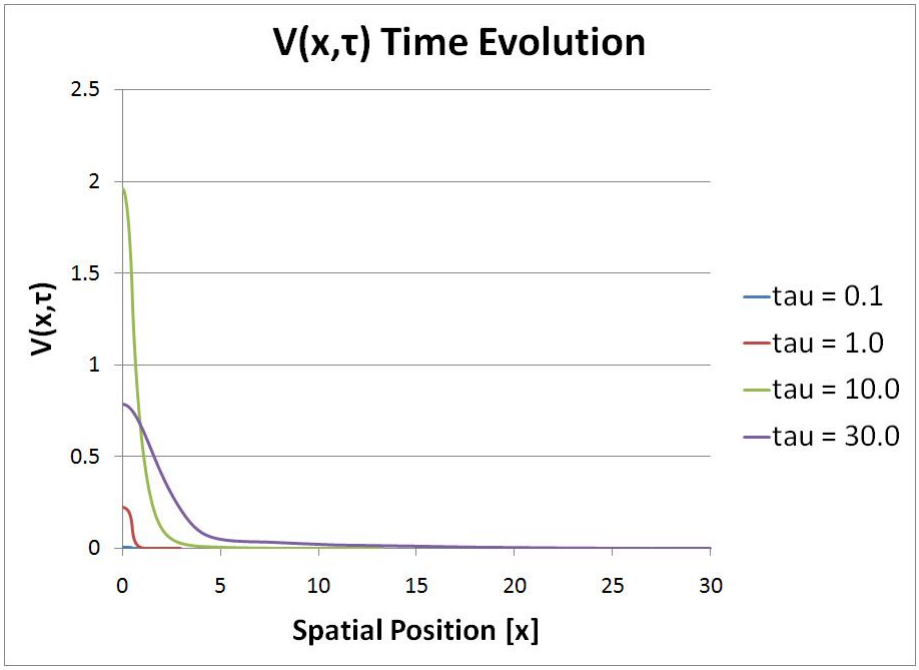

Figure 17: Case C - Time Evolution of V(x,τ)

Remarks on Oscillations: The space, time, and angular variables were refined to further investigate the character of the solution to this problem. Refinement of the spatial mesh did not improve the solution. Refinement of the temporal treatement with much smaller timesteps also did not improve the solution. However, refinement of the angular quadrature did produce a distinct improvement. Figure 18 shows a set of results for various angular quadrature sets up to S_{32}.

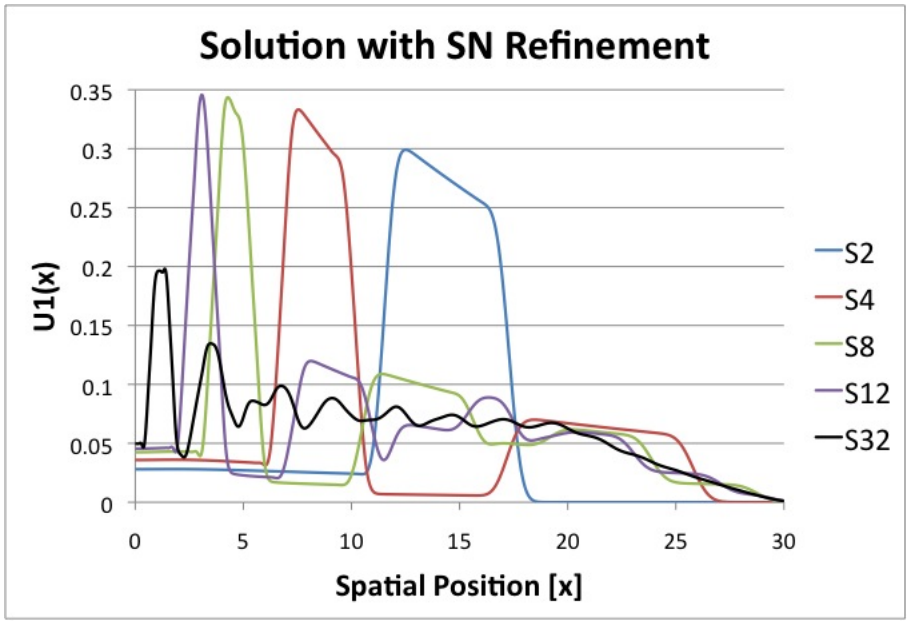

Figure 18: Solution with respect to angular refinement

The research by Chai et.al. [16] discusses a known issue called ray effects when using the S_N method. This phenomenon occurs because of the approximation that a continuous angular distribution can be represented as a set of discrete ordinate equations. Ray effects are a function of the quadrature approximation and are independent of the spatial discretization. The $U1$ parameter is associated with the smallest opacity. In the Case C problem, the opacity is small enough that the phase space is essentially a vacuum with negligible absorption and re-emission. The main driving sources are from the previous timestep and the internal source.

While the isotropic distribution of the internal source helps dampen the observed effects, the previous time step source causes the ray effect error to persist as the simulation progresses through each time step.

3.2 Chang's Model Problem

The second model problem studied is found in a journal article by Britton Chang [5]. The problem is a derivative of one used by Fleck and Cummings to test their Implicit Monte Carlo method [11] and was later reformulated by Larsen for multi-group deterministic transport and with added heterogeneous opacities [23]. The problem is non-linear in the opacities and the coupling between the radiation transport and material energy balance through the absorption/re-emission term. The increased non-linearity provides a more significant challenge for the physics-based preconditioners. The heterogenous materials will stress most diffusion based preconditioners.

This problem has been successfully solved by Larsen using the Grey Transport Acceleration [23] technique for fixed point iteration. It has also been solved by Fleck and Cummings using the Implicit Monte Carlo Method [11]. Chang solved it using the Photon Free Method (PMT) [6] in which the transport sweep is hidden as a functional evaluation to create the energy flux values through Gaussian quadrature.

3.2.1 Problem Specification

This model problem has a blackbody Planckian source at a temperature of 1 keV isotropically distributed on the left side of a 1D slab 4 cm thick. A vacuum boundary condition is imposed on the right hand side. There is no internal source other than the re-emission coupling with the material energy balance. The numerical parameters for this simulation are given in Table 31. The various pre-

Table 31: Numerical Specification for Chang's Model Problem

Parameter	Value	Units
Δt	30	ps
Δx	0.2	cm
NX	20	-
Length	4	cm
S_N Quadrature	32	-
GMRES Absolute Tolerance	1e-6	-
GMRES Relative Tolerance	1e-6	-
Newton Tolerance	1e-3	-

Table 32: Physical Constants for Chang's Model Problem

Name	*Parameter*	*Value*	*Units*
Planck's Constant	h	4.13567E-6	KeV-ps
Boltzmann Constant	k	8.61734E-08	KeV/K
Speed of Light	c	0.02998	cm/ps

conditioners described in section 2.5 are also used in this problem. The physical constants used in this problem are given in Table 32.

The frequency and temperature dependent opacity is defined by:

$$\sigma_a(x,\nu,T) = \frac{\gamma(x)}{(h\nu)^3}\left(1 - e^{-\frac{h\nu}{kT}}\right), \tag{100}$$

and the group averaged data are computed using (10). Numerical integration of the opacity and Planckian distribution over a group is performed using Simpson's Rule. The coefficient,$\gamma(x)$, is defined in the following way:

$$\gamma(x) = \begin{cases} 1.0 & 0 \le x \le 2 \\ 1000.0 & 2 \le x \le 3 \\ 1.0 & 3 \le x \le 4 \end{cases} \tag{101}$$

The frequency domain is discretized into fifty groups logarithmically spaced from $0.00001[keV]$ to $10.0[keV]$. The boundary and initial conditions are given in Eqs. (102) - (106) respectively:

$$I_n(0,\mu,t) = 0.0\left[\frac{keV}{cm^3 - str}\right], \tag{102}$$

$$I_n(4,\mu,t) = B_n(1.0[KeV]), \tag{103}$$

$$I_n(x,\mu,0) = \frac{1}{2}B_n(0.001[keV]), \tag{104}$$

$$E_n(x,0) = B_n(0.001[keV]), \tag{105}$$

$$T(x,0) = 0.001[keV]. \tag{106}$$

3.2.2 Chang's Problem Results

After extensive testing of the method, no results could be converged for the Chang non-linear problem. No preconditioning method worked. No refinement of phase space helped. A modification of the code to run the Photon Free Method (PFM) was performed and results were obtained that matched those in Chang's paper. His results are plotted along with the results generated from the PFM formulation of the code developed for this thesis in Figure 19. Several variants of this test problem were investigated to determine the limits of the method proposed in this thesis. These tests were: lagging opacities, use of the Photon Free Method, altering the Newton Centering, and running fewer frequency problems.

If the opacities are held fixed at the previous timestep values, a non-linearity is removed from the problem. However, the JFNK solver does not converge with any choice of preconditioner.

The results generated from implementation of the Photon Free Method using the developed code for this thesis are plotted as orange dots on the plot from in

Chang's paper [5].

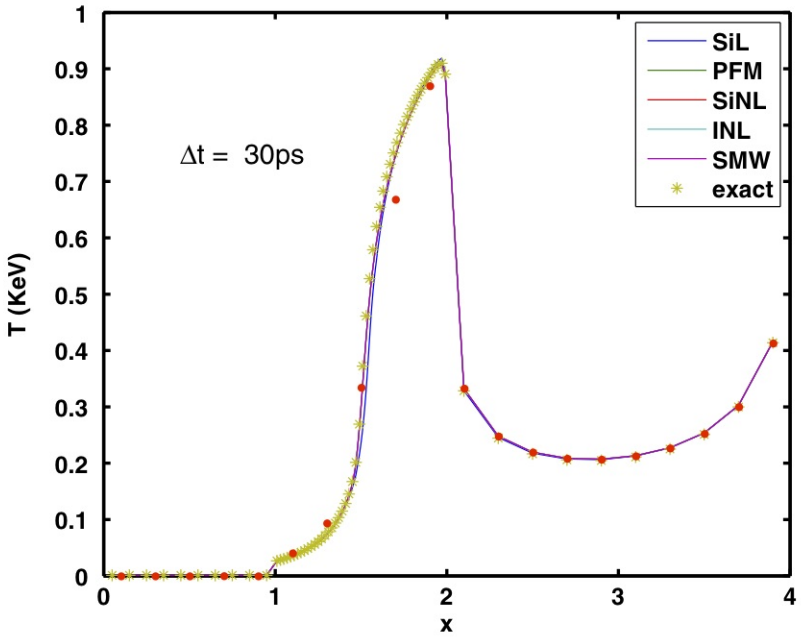

Figure 19: PFM Comparison with Chang's Results for Δt=30ps

Another factor possibly contributing to the code instability is the energy density of values within the Newton scheme. With the PFM scheme, the energy density is calculated by performing a transport sweep within the Newton loop. This value of E lies at the iteration index $k + \frac{1}{2}$ for the Newton loop index k.

In the method used in this thesis, only the spatial gradient of flux, $\frac{\partial J}{\partial x}$ is evaluated at $k + \frac{1}{2}$. The energy density remains at the current Newton iterate, k, for the purpose of converging quadratically within the JFNK framework. This also allows for direct preconditioning based on the dominant physics associated with the energy density. This centering issue is an important difference between the

methods.

In an attempt to converge an answer, the centering for the energy density was altered in the following ways shown in Figures 20 and 21. Figure 20 with the right diagram in Figure 21 is essentially the PFM formulation but with the residual equation for the energy density solved within the JFNK framework. Figure 20 with the left diagram in Figure 21 is a hybrid where the energy density residual equation is evaluated at $p + \frac{1}{2}$ and the material energy balance residual for temperature is evaluated at p, where p is here defined as the Newton iteration index.

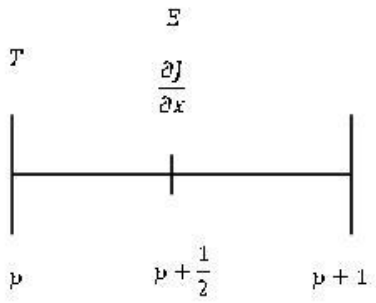

Figure 20: Newton Centering for F_E

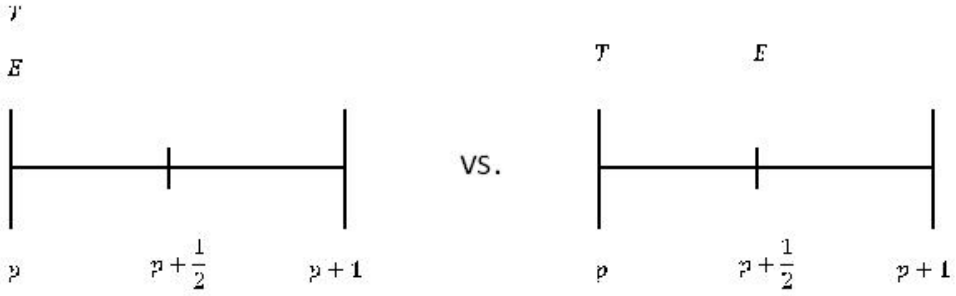

Figure 21: Newton Centering Options for F_T

The centering method expressed in the right diagram of Figure 21 did work, but only when the residual values had been so scaled as to be essentially zero– recovering the PFM scheme that Britton Chang used. The other centering method proved unstable.

When trying to run this problem with fewer frequency groups–or even the one group case–the problem is able to converge at the GMRES inner iteration level for calculating the non-linear step. However, the problem does not converge in the Newton iteration on the non-linearities.

3.2.3 Remarks

It is interesting that the PFM formulation of the problem works quite well, but not the formulation that converges both E and T non-linearly. Perhaps ray effects seen in the previous Su/Olson model problem are causing issues in the convergence of the energy densities in the JFNK framework. It may also be possible that the equations for the energy density residuals are simply hard to get into the zone of convergence for Newton type methods. If this is the case, then the embedding of the transport sweep in the material energy balance equation would require only that the JFNK application on the material energy balance equation be within the zone of convergence. However, the difficulty in the internal GMRES convergence points to some other cause for this behaviour.

It may be that widely varying magnitudes of E could be causing difficulty within the JFNK framework in approximating the action of the Jacobian. GMRES struggles when eigenvalues are widely varying, so this would seem to make sense. Two methods were used to try and address this. First, conversion factors were used to change the units of E such that the maximum values were comparable to the temperature. This approach did not address the widely varying magnitudes. The second method was to allow for the value of ϵ used in the finite difference approximation of the Jacobian action (29) to vary for each element of the residual

vector. This different method is shown in equation (107):

$$\epsilon_z = \left(1.5 * 10^{-8}\right)(u_z + 1). \tag{107}$$

Both of these methods failed to provide for a stable convergence in the GMRES iterative solver within the JFNK framework.

The possibility that a bug could exist in the code must still be mentioned, but the success of the developed code in replicating the Photon Free Method suggests that the transport solver works, the JFNK algorithm works, and the time stepping is on target. The only possible errors would have to occur in the energy density residual formulation, a formulation that worked in the Su and Olson problem. It would seem that the developed method simply fails in the presence of extreme non-linearity.

4 Conclusions

4.1 Introduction

In this chapter, we discuss the effectiveness of the physics-based preconditioners for the model problems considered. We also present conclusions on the practical value of the methods used in this thesis and reasons for the lack of convergence in the tightly coupled case. We also discuss ideas for further investigating phyisics-based preconditioning of JFNK for non-linear radiative transfer.

4.2 Preconditioning Conclusions

Based on the results of the Su/Olson problem, the most effective preconditioning was performed just on the material energy balance residual for the temperature dependent Planckian absorption term. The absorption/re-emission of photons appears to be the dominant physics present and the very simple DT preconditioner sufficient to produce a rapidly convergent solution. While the INFMC preconditioner includes the coupling terms between equations, the added matrix inversion cost seems to outweigh the benefits when used with problems where the coupling is not very strong.

Added time to converge (increased Gauss-Seidel sweeps) with the radiation diffusion preconditioner suggests that the system may not be diagonally dominant and care should be taken to check for this if an iterative inversion scheme is used. Grey preconditioners do not seem to be worthwile as the simulation time was greater than other alternatives. Specifically the INFMC preconditioner performed generally better than INFGC. The transport-based preconditioners were unreliable, possibly due to the ray effect issues present in the problem. They are also costly when compared to the simpler direct Planckian preconditioners.

Preconditioning the streaming term alone was about as effective as no preconditioning at all. This is evident from the results in Table 5. The reductions in simulation time were negligible. All of these results are generated with relatively large time steps for the simulation. When time steps are smaller, simulations without preconditioning become more efficient due to the zero added computation cost.

4.3 Method Conclusions

The Chang model problem is truly tightly coupled and significantly non-linear. No comparison of preconditioners could be achieved because no solution could ever be converged with the approach utilized in this thesis. Much time was spent attempting to ascertain the cause of the failure. Chang's Photon Free Method was successfully used within the written code for this thesis by altering variable usage. While it is impossible to prove the code developed is free from mistakes, these results indicate that the code failure is not due to human coding error. Rather, the set of energy density balance residuals being converged simultaneously with the material energy balance equation overstresses the inner GMRES solver. The variables do have dimensions and the varying magnitudes of the energy density within the many frequency groups might spread the eigenvalues too far in the GMRES operator. None of the preconditioners constructed in this thesis could overcome this problem.

When the problem was coarsened in the frequency domain, the problem would fail to converge at the Newton iteration level. This problem is sensitive to the region of convergence associated with the Newton method. The energy density residual equation seems to specifically be the culprit here. Since the PFM scheme is convergent for this problem, it can be said that the material energy balance

residual has a more stable zone of convergence.

Choice of iteration order and convergence parameters (i.e. energy density and temperature) in the Newton method is vital and can drastically change the stability of the problem. While there is greater flexibility with preconditioning using the approach in this thesis, the added problems of instability prevent this method from being more effective than Chang's PFM scheme. Perhaps a better formulation for operating residual equations can be found than those used in this body of work. A different treatment of the angular discretization such as the P_n equations [30] may help overcome the influence of ray effects.

4.4 Overall Conclusions and Future Work

Use of the next angular moment of the photon transport equation in the operating residual norm may improve convergence. The Eddington tensors present in this formulation should be slowly varying and provide more stability within the JFNK framework for converging non-linearities. There are many other preconditioner options that could be tried. Several mathematically based options may improve GMRES convergence. Continuation methods might provide a better initial guess to start the problem. Higher order finite differencing could be used as well for approximating the action of the Jacobian on a vector. Initial efforts in this direction do not appear to yield substantial performance gains. Other choices for the parameter ϵ, which impacts approximated action of the Jacobian on a vector, could also be explored.

The objective of this thesis was to compare the effectiveness of various physics based preconditioners for the solution of a tightly coupled non-linear radiative transfer problems. This goal was only partially achieved. The comparison was

possible for the Su/Olson problem, but this problem has been formulated such that it is linearized around the integrated Planckian distribution source. Though the code written for this thesis doesn't specificially rely on this property, it is difficult to claim that this is a tightly coupled non-linear problem. This is especially true since the opacities alternate between two values across the frequency range. Many questions were raised through the research conducted in this thesis and much work remains to be accomplished in this field.

BIBLIOGRAPHY

[1] Marvin L. Adams and Edward W. Larsen. Fast Iterative Methods For Discrete-Ordinates Particle Transport Calculations. *Progress in Nuclear Energy*, 40(1):3–159, 2002.

[2] Marvin L. Adams and Paul F. Nowak. Asymptotic Analysis of a Computational Method for Time - and Frequency - Dependent Radiative Transfer. *Journal of Computational Physics*, 146:366–403, 1998.

[3] Satish Balay, Jed Brown, Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, and Lois Curfman. PETSc Users Manual. Technical report, Argonne National Laboratory, 2010.

[4] S. Chandrasekhar. *Radiative Transfer*. Dover Publications Inc., 1960.

[5] Britton Chang. A Deterministic Photon Free Method to Solve Radiation Transfer Equations. *Journal of Computational Physics*, 222:71–85, 2007.

[6] Britton Chang. The Incorporation of the Semi-Implicit Linear Equations into Newton's Method to Solve Radiation Transfer Equations. *Journal of Computational Physics*, 226:852–878, 2007.

[7] Mathew Cleveland. Extending the Applicability of Implicit Monte Carlo Diffusion: Frequency Dependence and Variance Reduction Using the Difference Formulation. Master of Science, Oregon State University, July 2008.

[8] Jeffery D. Densmore, Todd J. Urbatsch, Thomas M. Evans, and Michael W. Buksas. A Hybrid Transport-Diffusion Method for Monte Carlo Radiative-Transfer Simulation. *Journal of Computational Physics*, 222(2):485–503, March 2007.

[9] eFunda Inc. Abscissas and Weights of Gauss-Legendre Integration. Webpage, 2010.

[10] R. Elaloufi, R. Carminati, and Jean-Jacques Greffet. Time-Dependent Transport Through Scattering Media: From Radiative Transport to Diffusion. *Journal of Optics A: Pure and Applied Optics*, 4:S103–S108, 2002.

[11] J.A. Fleck and J.R. Cummings. An Implicit Monte Carlo Scheme for Calculating Time and Frequency Dependent Nonlinear Radiation Transport. *Journal of Computational Physics*, 8:313–342, 1971.

[12] D. Gaston, C. Newman, G. Hansen, and D. Lebrun-Grandie. MOOSE: A parallel computational framework for coupled systems of nonlinear equations. *Journal of Nuclear Engineering and Design*, 239(10):1768–1778, October 2009.

[13] N. A. Gentile. Implicit Monte Carlo Diffusion–An Acceleration Method for Monte Carlo Time Dependent Radiative Transfer Simulations. In *Nuclear Explosives Code Developers Conference*. Lawrence Livermore National Laboratory, October 2000.

[14] G.E. Hammond, A.J. Valocchi, and P.C. Lichtner. Application of the Jacobian-Free Newton-Krylov with Physics-Based Preconditioning to Biogeochemical Transport. *Advances in Water Resources*, 28:359–376, 2005.

[15] Kopp H.J. Synthetic Method Solution of the Transport Equation. *Nuclear Science and Engineering*, 17(65), 1963.

[16] Chai J., Lee H., and Patankar S. Ray Effect and False Scattering in the Discrete Ordinates Method. *University of Minnesota Supercomputer Institute Research Report (UMSI)*, 106, 1993.

[17] D.A. Knoll, L. Chacon, L.G. Margolin, and V.A. Mousseau. On Balanced Approximations for Time Integration of Multiple Time Scale Systems. *Journal of Computational Physics*, 185:583–611, 2003.

[18] D.A. Knoll and D.E. Keyes. Jacobian-free Newton–Krylov methods: a survey of approaches and applications. *Journal of Computational Physics*, 193:357–397, 2004.

[19] D.A. Knoll, R.B. Lowrie, and J.E. Morel. Numerical Analysis of Time Integration Errors for Nonequilibrium Radiation Diffusion. *Journal of Computational Physics*, 226:1332–1347, 2007.

[20] D.A. Knoll, H. Park, and Kord Smith. Application of the Jacobian-Free Newton-Krylov Method in Computational Reactor Physics. In *American Nuclear Society 2009 International Conference on Advances in Mathematics, Computational Methods, and Reactor Physics*, 2009.

[21] D.A. Knoll, W.J. Rider, and G.L. Olson. An Efficient Nonlinear Solution Method for Non-equlibrium Radiation Diffusion. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 63:15–29, 1999.

[22] Dana Knoll, HyeongKae Park, and Chris Newman. Acceleration of k-Eigenvalue / Criticality Calculations using the Jacobian-Free Newton-Krylov Method. *Nuclear Science and Engineering*, 167(2), 2011.

[23] Edward Larsen. A Grey Transport Acceleration Method for Time-Dependent Radiative Transfer Problems. *Journal of Computational Physics*, 78:459–480, 1988.

[24] Edward W. Larsen, Guido Thoemmes, Axel Klar, Mohammed Seaied, and Thomas Goetz. Simplified Pn Approximations to the Equations of Radiative Heat Transfer and Applications. *Journal of Computational Physics*, 183:652–675, 2002.

[25] Adams M., editor. *A New Transport Discretization Scheme for Arbitrary Spatial Meshes in XY Geometry*, Pittsburgh, PA, April 28 - May 1 1991. Lawrence Livermore National Laboratory, American Nuclear Society International Topical Meeting.

[26] W. Martin and F. Brown. Error modes in implicit monte carlo. In *American Nuclear Society 2001 Winter Meeting*. Los Alamos, November 2001.

[27] Michael F. Modest. *Radiative Heat Transfer*. Academic Press, second edition, 2003.

[28] V. A. Mousseau, D. A. Knoll, and W. J. Rider. Physics-Based Preconditioning and the Newton-Krylov Method for Non-equilibrium Radiation Diffusion. *Journal of Computational Physics*, 160:743–765, 2000.

[29] C. Newman, G. Hansen, and D. Gaston. Three Dimensional Coupled Simulation of Thermomechanics, Heat, and Oxygen Diffusion in UO2 Nuclear Fuel Rods. *Journal of Nuclear Materials*, 392:6–15, 2009.

[30] Todd Palmer. Curvilinear Geometry Transport Discretizations in Thick Diffusive Regions. Master of science, University of California - Lawrence Livermore National Laboratory, June 1993.

[31] H. Park, D. Gaston, S. Kadioglu, D. Knoll, R. Martineau, and W. Taitano. Tightly Coupled Multiphysics Simulations For Pebble Bed Reactors. In *American Nuclear Society 2009 International Conference on Advances in Mathematics, Computational Methods, and Reactor Physics*, Saratoga Springs, New York, May 3-7 2009.

[32] J.M. Reisner, V.A. Mousseau, A.A. Wyszogrodzki, and D.A. Knoll. An Implicitly Balanced Hurricane Model with Physics-Based Preconditioning. *Monthly Weather Review*, 133:1003–1022, April 2005.

[33] C.E. Siewert. A Discrete-Ordinates Solution for Multigroup Transport Theory with Upscattering. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 64:255–273, 2000.

[34] V. Sobolev. A Treatise on Radiative Transfer. *Princeton, NJ: Van Nostrand*, 1963.

[35] Bingjing Su and Gordon L. Olson. An Analytical Benchmark For Non-Equilibrium Radiative Transfer In An Isotropically Scattering Medium. *Annals of Nuclear Energy*, 24(13):1035–1055, 1997.

[36] Bingjing Su and Gordon L. Olson. Non-Grey Benchmark Results for Two Temperature Non-Equilibrium Radiative Transfer. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 62:279–302, 1999.

[37] R.H. Szilard and G.C. Pomraning. Numerical Transport and Diffusion Methods in Radiative Transfer. *Nuclear Science and Engineering*, 112:256–269, 1992.

[38] L. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, 1997.

[39] Gol'din V.Ya. A Quasi-Diffusion Method for Solving the Kinetic Equation. *USSR Comp. Math. and Math. Phys.*, 4(136), 1967.