

Spherical Harmonics Methods for Thermal Radiation Transport

by

Ryan G. McClarren

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Nuclear Engineering and Radiological Sciences)
in The University of Michigan
2006

Doctoral Committee:

Professor James Paul Holloway, Chair

Professor William R. Martin

Professor Kenneth Powell

Dr. Thomas A. Brunner (Sandia National Laboratories)

© Ryan McClarren 2006
All Rights Reserved

In my beginning is my end. Now the light falls
Across the open field, leaving the deep lane
Shuttered with branches, dark in the afternoon,
Where you lean against a bank while a van passes,
And the deep lane insists on the direction
Into the village, in the electric heat
Hypnotised. In a warm haze the sultry light
Is absorbed, not refracted, by grey stone.
The dahlias sleep in the empty silence.

- from T.S. Eliot's "Four Quartets"

ACKNOWLEDGEMENTS

My thanks go out to my wife Katie for being giving me loving encouragement. Without her I probably would have been working on this for another few years. I also want to thank my family, especially my parents, for their support and understanding when I tried to explain what it is I have been doing. Some of my friends: Ezequiel Berdichevsky, Mark Calaguas, Ben Heller, Arthur Farnham and Paul Litvak played a large role in my pursuing graduate school and to them I am grateful.

I also need to thank my advisor James Holloway for putting up with my incessant questions and sometimes harebrained ideas. Under his hand I have changed from a student into a scientist.

I had wonderful support (both scientifically and monetarily) from the High Energy Density Theory and Target Design group at Sandia National Laboratories. Tom Brunner has certainly provided me with much insight, and his thoughts have colored many of my ideas. Tom Mehlhorn was a great manager to have and I appreciate all the help he has given me.

TABLE OF CONTENTS

EPIGRAPH	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vii
 CHAPTER	
I. Introduction	1
1.1 Radiation Transport Methods	2
1.2 Numerical Methods for the P_n Equations	5
1.3 Analysis of the P_n Equations	5
1.4 Results	6
1.5 Outline	7
II. The Transport Equation and its Approximations	9
2.1 The Transport Equation for Thermal Photons	9
2.1.1 Isotropic Scattering	10
2.1.2 Grey Radiation Transport	11
2.1.3 Blackbody Source	11
2.1.4 Material Energy Equation	13
2.1.5 The thermal transport equation in neutron transport notation	14
2.2 Thermal Spherical Harmonics Equations	16
2.2.1 Spherical harmonics functions	16
2.2.2 Spherical harmonics moments of the grey transport equation	18
2.2.3 Truncating the Expansion: the closure issue	19
2.2.4 Two-Dimensional P_n Equations	22
2.2.5 One-Dimensional P_n Equations	23
2.3 Other Methods	23
2.3.1 Discrete ordinates methods	23
2.3.2 Flux-limited Diffusion	24
2.3.3 Implicit Monte Carlo	25

III. Upwind Finite Volume Discretization	27
3.1 Cell-Averaged P_n Equations	28
3.2 Upwinding The Equations Using A Riemann Solver	29
3.2.1 The flux between cells	29
3.2.2 Example: Two-dimensional Cartesian Grid	31
3.2.3 Example: One-dimensional Equations	33
3.2.4 Extra Dissipation	33
3.2.5 Direction Splitting Issues	33
3.3 High-Resolution Correction	34
3.4 Justification of High Resolution Method	36
IV. Time Integration	38
4.1 θ Notation for Time Integration	39
4.2 The Backward Euler Method	39
4.3 Crank-Nicolson Method	41
4.4 High-Resolution Time Integration	43
V. Implementation	46
5.1 Implicit Equations	46
5.2 Quasi-Linear High-Resolution Implementation	47
5.2.1 Analysis of the Quasi-Linear Approach	49
5.2.2 Quasi-linear time integration	52
5.2.3 Temperature coupling	53
5.3 Newton-Krylov Methods	53
5.4 GMRES	54
5.4.1 Preconditioning	55
5.4.2 Matrix Free	57
5.5 Trilinos	57
VI. Boundary Conditions	59
6.1 Mark Boundary Condition	60
6.1.1 Vacuum boundary condition	61
6.2 Reflecting Boundary Condition	61
6.2.1 Albedo Boundary Condition	66
VII. Limits of the P_n Equations	67
7.1 The Diffusion Limit	67
7.1.1 Diffusion properties of the P_n equations	69

7.1.2	Linear solution of the P_n equations	70
7.1.3	Asymptotic analysis of the P_n equations	71
7.1.4	Diffusion properties of the Riemann discretization	73
7.1.5	Intermediate Diffusion Limit	77
7.1.6	Modified Riemann solver in the diffusive limit	78
7.2	Free Streaming Limit and Negative Solutions	80
VIII. Verification and Test Problems		84
8.1	Analytic Solutions to the Thermal P_1 Equations	84
8.1.1	Normalized P_1 equations	85
8.1.2	Green's function	87
8.1.3	Finite source in time and space	90
8.2	Point Source	92
8.2.1	Line source	93
8.3	Linear Exact Solutions	94
8.4	Other Test Problems	95
IX. Numerical Results		99
9.1	Linear Transport Results	99
9.1.1	Implicit Riemann solver using a Newton method	100
9.2	Quasi-linear time integration results	114
9.3	Thermal Transport Results	118
9.3.1	One-dimensional problems	118
9.3.2	Two-dimensional results	126
9.4	High Resolution Time Integration	134
X. Conclusions and Future Work		141
10.1	Implicit time integration	141
10.1.1	High resolution time integration	142
10.2	Accuracy of P_n expansion	142
10.3	Future implementation issues	144
10.3.1	Diffusion limit	144
10.4	Benchmark solutions	145
10.5	Coda	146
BIBLIOGRAPHY		147

LIST OF FIGURES

Figure

2.1	A general unit vector \vec{u} defined by (θ, ϕ) . The polar angle, θ is the angle between the z axis and the vector and the azimuthal angle is the angle between the x axis and the projection of \vec{u} in the $x - y$ plane. . .	17
3.1	The initial conditions for the Riemann problem has two flat solutions discontinuously joined at $\chi = 0$. This effectively models the boundary between two cells.	30
3.2	The two-dimensional Cartesian grid. A generic cell (i, j) is shaded in grey.	31
3.3	This cartoon shows how a dimensionally split numerical method would propagate information moving skew to the grid.	34
6.1	The reflected angle equals the incident particle angle for a reflecting boundary.	61
8.1	The Green's function solution to the P_1 equations at $t = 10$	91
8.2	Analytic solutions to the P_1 and transport equations for the finite source problem.	92
8.3	The solutions shown in Fig. 8.2 on a logarithmic scale.	93
8.4	The collided energy density from the line source problem at time $ct = 5, 10$	95
8.5	The P_1 and transport solutions to the plane, pulsed source problem. . .	96
8.6	The "M" problem schematic: the maize colored regions are vacuum and the blue regions are dense material.	97
8.7	The problem of one wire heating another across a vacuum.	98

9.1	The high resolution P_1 solutions between $x = -3.5$ and 3.5 at $t = 5$	101
9.2	The high resolution P_1 solutions at $t = 10$. The analytic P_1 solution and the $\Delta t = 0.1$ solution are coincident in the middle region.	102
9.3	Comparison of high resolution and first-order solutions to the P_7 equations with different numbers of spatial grid points.	103
9.4	Steady state results from Reed's problem	104
9.5	The scalar flux for different values of Δt ; cell width is $\Delta x = 0.05$	105
9.6	P_1 results from the pulsed line source problem with $N_x = N_z = 50$ at $t = 1$. The linear color scale represents the energy density.	106
9.7	Comparison of P_1 results from high resolution and first order methods on the pulsed line source problem.	107
9.8	The effect of the time scale on P_1 results for the pulsed line source problem. All of these were computed using the high resolution spatial scheme.	108
9.9	The P_5 , steady state solution with incident beam on the left, and two regions: a strong absorber and a strong scatterer	109
9.10	The P_1 , steady-state solution to a uniform source problem with $Q = \frac{1}{2\pi}$ and $\sigma_a = 0.01, \sigma_t = 10$	110
9.11	The solution to the modified Reed's problem.	111
9.12	The P_7 solution at $t = 35$ after the initial pulse of particles	112
9.13	The P_1 solution at $t = 35$ after the initial pulse of particles	113
9.14	The P_1 scalar flux for the first test problem at $T = 10$ after the initial pulse for different Δt , solved using the quasi-linear formulation.	114
9.15	The error as a function of spatial resolution for P_1 scalar flux with $\Delta t = 0.05$ at $T = 10$ as compared with the analytic solution. Note the errors for the quasi-linear and Newton-based nonlinear method are essentially identical.	115
9.16	The CPU time as a function of error for P_1 scalar flux with $\Delta t = 0.05$ at $T = 10$ as compared with the analytic solution.	116

9.17	Results for the “M” problem at $t = 1$ after the source is turned on. . . .	119
9.18	The scalar flux for the “M” problem along $x = 0$ at $t = 1$ after the source is turned on.	120
9.19	The numerical energy density and radiation flux, F_z , at $ct = 1$ compared with the P_1 analytic solution for the Su-Olson benchmark problem.	120
9.20	Logarithm of the Error between high resolution solution and analytic solution for the Su-Olson benchmark as a function of the logarithm of the number of grid points.	122
9.21	Comparison of several P_n solutions with the transport solution for the Su-Olson benchmark.	123
9.22	Comparison of several P_n solutions with the transport solution for the Su-Olson benchmark on a logarithmic scale.	123
9.23	P_7 solutions to the variable opacity test problem with $p = 3$; the dashed lines are the material temperature and the solid lines are the radiation temperature.	125
9.24	P_7 solutions to the variable opacity test problem with $p = 5$; the dashed lines are the material temperature and the solid lines are the radiation temperature.	125
9.25	The energy density from the P_1 solution to the thermal line source problem at time $t = 5$	127
9.26	Comparison of energy density from P_1 numerical and analytic solutions at $t = 5$	127
9.27	Comparison of energy density from P_1 numerical and analytic solutions at $t = 10$	128
9.28	Radiation temperature (eV) at 2 nanoseconds for the mock wires problem using different P_n orders. All problems were run on a Cartesian mesh with $N_x = N_z$	129
9.29	Comparison of the positive and negative parts of the radiation temperature for the P_7 solution at 2 nanoseconds.	130

9.30	Material temperatures (eV) at 10 nanoseconds for P_5 and P_7 . The color scale is such that any temperature $-1 eV$ or below is black.	130
9.31	Effect of time step size on the P_7 radiation energy density (J/m^3) at 8.5 ns.	132
9.32	Effect of time step size on the P_7 material temperature (eV) at 8.5 ns.	132
9.33	Radiation and material temperature (eV) at 0.9 ns for the nonlinear duct problem using different P_n orders. All problems were run on a Cartesian mesh with a time step of $\Delta t = 0.05$ ns.	135
9.34	Radiation and material temperature (eV) at 6 ns for the nonlinear duct problem from P_3 and P_7	136
9.35	Results from the 1D square pulse problem at $t = 1$ seconds and $\Delta t = 0.125$ and $CFL = 12.5$ for different methods compared with the exact solution.	137
9.36	The radiation energy density for different methods with $\Delta t = 0.25$ compared with the reference solution at $t = 1.5$ seconds. All solutions have $\Delta x = 0.033$	139
9.37	Results from the 2D square pulse problem at $t = 1.5$ seconds and $\Delta t = 0.5$ for different methods compared with the reference solution.	140
9.38	Results from the 2D square pulse problem at $t = 1.5$ seconds and $\Delta t = 0.1$ for different methods compared with the reference solution.	140

CHAPTER I

Introduction

The work in this thesis could be considered thermal radiation transport *a rebours*¹ because the spherical harmonics approach to solving time dependent transport problems has been an also ran. This is in part due to some known drawbacks of the spherical harmonics methods and the rich body of research characterizing the discrete ordinates method and its efficient implementation. Of late there has been somewhat of a resurgence of the spherical harmonics method. An impetus for this resurgence is the work of Brunner and Holloway [1, 2] who developed robust methods for solving problems of linear transport with spherical harmonics. The work in this thesis stands on the shoulders of that work to show how spherical harmonics methods could be applied to the nonlinear problem of thermal radiation transport. Specifically this thesis will illuminate the following novel research into using P_n methods

- An implicit, high resolution upwind scheme for the P_n equations
- A quasi-linear method for solving the nonlinear equations from the high resolution scheme
- The first analytic results for the P_1 thermal radiative transfer equations

¹French for against the grain and the title of a novel by J.K. Huysmans.

- A proof that the standard Riemann solver fails in the asymptotic diffusion limit
- A proof that the negatives in P_n solutions arises from the equations' linearity, hyperbolicity, and rotational invariance
- A high resolution implicit time integration scheme that gives better than first order time integration for arbitrarily large time steps
- Reflecting boundary conditions implemented using ghost cells
- Non-operator split results from coupled nonlinear radiation/material energy simulations

1.1 Radiation Transport Methods

Radiation transport methods are divided into two approaches: *deterministic* methods which attempt to solve the partial differential equations governing the physics of transport, and *stochastic* or Monte Carlo methods that simulate populations of particles (e.g. photons, neutrons, etc.) to find how radiation moves through the system. Deterministic methods are classified by how they treat the direction of flight dependence of the particles in the radiation field. This type of classification is natural because the direction of flight variable is continuous but does not appear inside a derivative. This makes the choice of how to handle these angular variables less like a choice is discretization such as whether to use finite element or finite difference and more a choice in model. For each method for treating the angular variables there is a different system of spatio-temporal partial differential equations to solve. Moreover, these models with a finite discretization of the angular variable do not limit to the same results except for problems with simple dependence on the direction of flight.

One of the best known and understood method for deterministic thermal radiation transport is the discrete ordinates (S_n) method. This relies on solving the transport equation along particular directions and using a quadrature rule to reconstruct the energy density. This method has been studied for many years [3] and there is a large body of literature on how to efficiently solve the systems of equations that arise using the S_n method. These methods have problems, however, in situations where the radiation field is highly anisotropic. For example a source radiating into a vacuum will give a solution of rays emanating out from the source rather than the correct solution of a smooth field of radiation streaming out from the source. These rays can cause problems in thermal radiation transport simulations coupled to other physical processes such as hydrodynamics. These rays can cause “hot spots” to develop in materials and wrongly influence the dynamics of the simulations. These rays cannot be totally eliminated except in the limit of an infinite number of discrete ordinates.

Another deterministic method is flux-limited diffusion which is based on an approximation of the form of the angular dependence of the radiation field. This diffusion approximation is incorrect most cases, and through the use of a flux limiter the diffusion coefficient is fixed up to compensate. The nature of this approximation is such that it turns the hyperbolic transport equation into a parabolic equation. In effect, this allows the radiation to travel at infinite speeds. Beyond this drawback there is the problem that for general cases the angular dependence of a radiation field cannot be modeled by a diffusion approximation corrected by a (generally ad hoc) flux limiter. This infinite radiation propagation velocity and poor angular resolution can lead to difficulties in simulations where the radiation is coupled to another physical process. For instance in a simulation where radiation was heating an object across a vacuum, a diffusion approach will allow light to heat the back of the object that should be dark and cold. This

could lead to an incorrect evolution of the system².

Implicit Monte Carlo (IMC) methods are the stochastic approach to solving thermal radiative transfer problems. These methods sample photons to find where radiation energy is deposited in the system and use a linearized equation to model the material temperature evolution. Since this is a stochastic method, IMC has noise in the solution. This noise is hard to squelch because the error in the solution decreases as the number of particles simulated raised to the negative one-half power. The problem with this noise is in simulations where the radiation is coupled to another physical process that is prone to instability. In this case the noise can seed an instability such as the Rayleigh-Taylor instability. Beyond this issue, IMC methods are the most computationally intensive approach to solving thermal radiation transport problems.

This leads us to spherical harmonics (P_n) methods. This deterministic approach uses a truncated expansion of the angular variables to arrive at a finite, hyperbolic system of equations. In multiple dimensions these equations give very different solutions than the S_n method. The P_n equations model radiation propagation as a series of waves whereas the S_n method moves photons in rays. This difference leads to P_n methods not having ray effects, rather it has wave effects. The main drawback to these wave effects is that the trough of a given wave can be negative, that is the amount of energy at a given point in the problem can be unphysically less than zero. This shortcoming of the P_n approach is a large reason such methods have not been thoroughly researched for large scale simulation. Yet, there may be an upside to method; in steady-state the P_n equations give positive solutions and in many applications the time scale of interest is very long compared the time scale of radiation propagation and the radiation field may

²In fact a very similar situation occurs in some radiation hydrodynamic simulations of z-pinch wire arrays. When the wires are mutually heating each other and flux-limited diffusion is used, the outside of the array can be heated and cause the wires to ablate wrongly [4].

be near steady state.

1.2 Numerical Methods for the P_n Equations

In this thesis we will see how the thermal radiation transport equations can be solved in one and two dimensions using an upwind method (based on a Riemann solver approach). This approach was first put forward by Brunner and Holloway for explicit time integration [5] and later by Eaton, Pain, and Oliveira for steady state problems [6, 7]. This thesis details the use of an implicit time integration scheme. The use of an implicit upwind method for radiation transport is the first of its kind.

These implicit methods in general require using a nonlinear solver to advance the solution to the next time step. We have developed a method for doing this using only two linear solves instead of employing a Newton based method. We call this approach the quasi-linear method. This innovation had a direct impact on the computational cost of the P_n methods and could be the goad to the use of high resolution methods for S_n methods or for other linear hyperbolic systems.

Below we will also discuss how to use the quasi-linear method to implement a high resolution time integration method [8]. This approach uses the same principles of the high resolution spatial scheme to use a second order method in smooth regions of the problem and a first order scheme where the second order method would create artificial oscillations in the solution.

1.3 Analysis of the P_n Equations

Along with the work on the numerical solution of the P_n equations we will present some analysis of the P_n equations in both discrete and continuous form. Part of this analysis shows that the P_n equations asymptotically limit to the diffusion equation in

diffusive problems but the Riemann solver based P_n equations do not limit to a discrete diffusion equation when only a diffusion length is resolved. This is a drawback because in many thermal transport problems there are regions where a mean free path of a photon is so short that resolving this distance is prohibitive.

In the other limit, the free streaming limit, we will show why the P_n equations give negative solutions in multi-dimensional problems. It turns out that this is a consequence of discretizing the transport equation in a way that is hyperbolic, rotationally invariant, and linear. We then point out how other methods forsake one of these properties to get positivity.

To test our numerical method we develop an analytic benchmark for the P_1 equations of time dependent thermal transport. Using a common form of the heat capacity we are able to obtain the Green's function of the P_1 equations and we then use this approach to build up the solution to a common one-dimensional benchmark problem and a two-dimensional problem.

Finally, to aid in the simulation of problems with an axis of symmetry we have developed the reflecting boundary conditions for the P_n equations. These boundary conditions in three-dimensions are not well disseminated. We show how to implement them with ghost cells, the most natural boundary condition type for an upwinded method.

1.4 Results

Our results show that using an implicit method allows us to obtain quality solutions with time steps larger than we could use with an explicit method. Furthermore, the use of the quasi-linear method does indeed provide a speed-up in terms of computational cost without sacrificing accuracy. In one dimension we tried our method on various benchmark problems and achieved propitious results. We were able to get good agree-

ment with an analytic transport solution to a benchmark problem using P_7 and solved a problem involving nonlinear opacities with large time steps while remaining stable.

In two dimensions the results were mixed. The energy density was negative in the solution to some problems. These negative energy densities did cause the problem to have negative material temperatures. However, despite these predictable results we were able to show that our solutions were converged in terms of time discretization error for time steps that were much larger than the radiation propagation time but smaller than the material temperature time scale.

We also presented some results from the high resolution time integration scheme. Therein we show why the Crank-Nicolson method, although stable, can give oscillatory solutions whereas the backward Euler method is too dissipative. The high resolution time integration method avoids the oscillations of the Crank-Nicolson method while minimizing the smoothing of the backward Euler method.

1.5 Outline

This thesis is structured to begin with an introduction to thermal radiation transport and the P_n equations in Chapter II. Thermal transport is introduced in a way that is translatable to the notation and terminology of nuclear engineering – something I hope will benefit students from this discipline who are drawn to thermal transport. Chapters III and IV present the spatial and temporal discretizations of the P_n equations. The implementation of these methods using a quasi-linear approach is discussed in Chapter V which is followed by the development of boundary conditions in Chapter VI. In Chapter VII the P_n equations in their continuous and discrete forms are analyzed to present important behavior in both the diffusion and free-streaming limit. The P_1 thermal transport equations are solved exactly in Chapter VIII and some other test problems

are detailed. The results from numerical calculations are shown in Chapter IX which is followed by conclusions and directions for future work in Chapter X.

CHAPTER II

The Transport Equation and its Approximations

2.1 The Transport Equation for Thermal Photons

The linear Boltzmann transport equation is a special case of the Boltzmann equation [9] that describes the flow of point particles moving in a straight line (i.e. linearly) and potentially colliding with a background medium. When the particles are photons, they are often referred to as radiation and the transport equation is written in terms of the specific intensity of the radiation. The specific intensity, $I(\vec{r}, \nu, \hat{\Omega}, t)$ is a function of position, \vec{r} , frequency of the photons, ν , direction of flight, $\hat{\Omega}$ (the hat denotes that this is a unit vector), and time, t . To be more precise $I(\vec{r}, \nu, \hat{\Omega}, t) d\vec{r} d\hat{\Omega} d\nu =$ the expected amount of energy in the volume element $d\vec{r}$ about \vec{r} with direction of flight in $d\hat{\Omega}$ about $\hat{\Omega}$ with frequency in $d\nu$ about ν . The unit vector $\hat{\Omega}$ is a three-dimensional vector that describes the direction of flight of a photon in terms of the cosine of the polar angle, θ , and the azimuthal angle, φ . Note that the specific intensity is a function that maps seven dimensions to a scalar: three dimensions for space, one for time, two for direction of flight (θ, φ), and one for frequency. The units of the specific intensity are energy per unit area per unit time per unit solid angle per unit frequency.

The transport equation that describes the evolution of a radiation field is given by

$$\frac{1}{c} \frac{\partial I}{\partial t} + \hat{\Omega} \cdot \nabla I + \sigma_t I = \int_0^\infty d\nu' \int_{4\pi} d\hat{\Omega}' \frac{\nu}{\nu'} \sigma_s(\nu' \rightarrow \nu, \hat{\Omega}' \cdot \hat{\Omega}) I(\cdot, \nu', \hat{\Omega}') + S. \quad (2.1)$$

In this equation, the photons are moving through a background material and we have written the differential scattering cross-section as $\sigma_s(\nu' \rightarrow \nu, \Omega' \cdot \Omega)$. Also, we will refer to the scattering cross-section as $\sigma_s = \int d\nu' \int d\Omega' \sigma_s(\nu' \rightarrow \nu, \Omega' \cdot \Omega)$; the total cross-section, σ_t is the sum of σ_s and the absorption cross-section, σ_a ($\sigma_t = \sigma_s + \sigma_a$). . These quantities have units of inverse distance and represent the inverse of an interaction mean free path for a photon moving through the background material. All of these cross-sections can be a function of space, time, frequency, and direction of flight. These cross-sections are often called opacities in the context of thermal transport. In Eq. (2.1) we have also denoted the speed of light as c and a prescribed source as S .

The transport equation is an integro-differential equation and has been derived by many authors. In fact one can find a derivation in almost any flavor desired in the standard transport monographs: Refs. 3, 10–15 carry out a fairly standard derivation by taking a general piece of phase space and performing an exercise in bookkeeping by accounting for how particles could enter or leave that phase space. Pomraning [16] derives the equation in this fashion but also uses a Lagrangian and variational approach to derive Eq. (2.1) .

2.1.1 Isotropic Scattering

For our purposes we will only consider isotropic scattering. This means that any outgoing direction of flight is equally likely for the photon leaving a scattering event. This is manifest in the differential scattering cross-section: in isotropic scattering

$$\int_0^\infty d\nu' \int_{4\pi} d\hat{\Omega}' \frac{\nu}{\nu'} \sigma_s(\nu' \rightarrow \nu, \hat{\Omega}' \cdot \hat{\Omega}) I(\cdot, \nu', \hat{\Omega}') = \frac{1}{4\pi} \int_0^\infty d\nu' \int_{4\pi} d\hat{\Omega}' \frac{\nu}{\nu'} \sigma_s(\nu' \rightarrow \nu) I(\cdot, \nu', \hat{\Omega}'). \quad (2.2)$$

In general isotropic scattering is not a valid assumption, however, it can be relaxed without much difficulty (for instance the neutron transport community this is often

handled by expanding the direction of flight dependence of the differential scattering cross-section in spherical harmonics [13]).

2.1.2 Grey Radiation Transport

The specific radiation intensity integrated over all frequencies,

$$\bar{I}(\vec{r}, \hat{\Omega}, t) = \int_0^\infty I(\vec{r}, \nu, \hat{\Omega}, t) d\nu, \quad (2.3)$$

is called the radiation intensity. The radiation intensity has units of energy per unit area per unit time per unit solid angle. If we define similarly averaged cross-sections and assume isotropic scattering

$$\bar{\sigma}_t = \frac{\int_0^\infty \sigma_t I(\vec{r}, \nu, \hat{\Omega}, t) d\nu}{\int_0^\infty I(\vec{r}, \nu, \hat{\Omega}, t) d\nu}, \quad (2.4)$$

$$\bar{\sigma}_s = \frac{\int_0^\infty d\nu \int_0^\infty d\nu' \frac{\nu}{\nu'} \sigma_s(\nu' \rightarrow \nu) I(\cdot, \nu', \hat{\Omega}')}{\int_0^\infty I(\vec{r}, \nu, \hat{\Omega}', t) d\nu}, \quad (2.5)$$

we can define the transport equation for grey radiation as

$$\frac{1}{c} \frac{\partial \bar{I}}{\partial t} + \hat{\Omega} \cdot \nabla \bar{I} + \bar{\sigma}_t \bar{I} = \frac{1}{4\pi} \int_{4\pi} d\hat{\Omega}' \bar{\sigma}_s \bar{I}(\cdot, \hat{\Omega}') + S. \quad (2.6)$$

From this point on we will drop the bar from the radiation intensity and the average over frequency will be understood.

2.1.3 Blackbody Source

In thermal radiation transport the source term in Eq. (2.1) is related to the thermal emission of photons by the background media. In most cases this source is given using an approximation known as local thermodynamic equilibrium (LTE). LTE basically means that we can define a temperature for the background media and that we can treat the media as a blackbody source at this temperature, even though the photons are not in

equilibrium with the material. This requires a physical mechanism to keep the material radiating in a blackbody spectrum. In high-energy density physics this mechanism is collisions between the constituent particles of the material [17], or a dense, optically thick material where photons are absorbed in a region with a similar temperature to that where they were emitted [18].

A blackbody source was properly described by Planck [16, 19] as

$$B(\nu, T) = \frac{2h\nu^3}{c^2} \frac{1}{e^{h\nu/kT} - 1}, \quad (2.7)$$

where k is the Boltzmann constant, h is Planck's constant, and T is the temperature of the blackbody. If we integrate the blackbody source over all directions of flight and frequencies, we arrive at the mean radiation intensity from a blackbody at temperature T :

$$\int_0^\infty d\nu \int_{4\pi} d\hat{\Omega} \frac{2h\nu^3}{c^2} \frac{1}{e^{h\nu/kT} - 1} = \frac{8\pi^5 k^4}{15h^3 c^2} T^4 \equiv acT^4, \quad (2.8)$$

where

$$a = \frac{8\pi^5 k^4}{15h^3 c^3}. \quad (2.9)$$

The source in Eq. (2.6) is given by Kirchoffs law [20], which states that the emission and absorption of radiation must be equal in equilibrium. The intensity of a blackbody, Eq. (2.8), gives the amount of radiation emitted into 4π directions. Hence, the specific intensity $acT^4/4\pi$. This intensity times the absorption opacity is the rate of absorption, and hence also the rate of emission, giving the emission source source,

$$S = \frac{\sigma_a acT^4}{4\pi} \quad (2.10)$$

2.1.4 Material Energy Equation

The temperature in the background material changes based on the first law of thermodynamics. This law is a simple energy balance; the change in internal energy is the energy added to the material minus the energy lost by the material. The energy gained by the material is the amount of photon energy absorbed by the material. The energy lost by the material is the blackbody source. There could possibly other sources (or sinks) of energy if the background material is moving, but we do not consider this here. The material internal energy is defined by

$$U(\vec{r}, t) = \int_0^T C_v(T') dT', \quad (2.11)$$

where C_v is the heat capacity, which is related to material density, ρ , and the specific heat, c_v , by $C_v = \rho c_v$. By Eq. (2.11) the time rate of change of the material energy is

$$\frac{\partial U}{\partial t} = \frac{\partial U}{\partial T} \frac{\partial T}{\partial t} = C_v \frac{\partial T}{\partial t}. \quad (2.12)$$

To calculate the amount of energy absorbed in the material we will define the radiation energy density, E , as

$$E(\vec{r}, t) = \frac{1}{c} \int_{4\pi} I(\vec{r}, \hat{\Omega}, t) d\Omega. \quad (2.13)$$

The units of the energy density are what one should expect (energy per unit volume). Using this definition we discover that the energy density of a blackbody source is aT^4 . The time scale of radiation absorption and emission is the mean free time, $1/c\sigma_a$. This makes the net amount of radiation energy absorbed per second

$$\frac{\text{radiation energy absorbed}}{\text{unit time}} = c\sigma_a(E - aT^4). \quad (2.14)$$

If we assume this is the only important mechanism for energy transfer in the material, then the amount of radiation absorbed, Eq. (2.14), should balance the time rate of change

in material energy, Eq. (2.12), giving the material energy equation:

$$C_v(T) \frac{\partial T}{\partial t} = c\sigma_a(E(\vec{r}, t) - aT^4). \quad (2.15)$$

This equation will govern the blackbody source in the transport equation. When the radiation field and the blackbody source are out of equilibrium this equation drives them toward equilibrium. It is also apparent that this equation is nonlinear in the temperature both in the fourth power found in the blackbody source and possibly even in the temperature dependence of the heat capacity and opacities.

Eqs. 2.1 and 2.15 have two different time scales. The transport equation has is the photon mean free time as an important time scale, $1/c\sigma_a$. In vacuum regions, the time scale would be some reference distance divided by the speed of light. For the material energy equation, Eq. (2.15) the important time scale is the e -folding time for T^4 . This time scale can be found by writing Eq. (2.15) in terms of T^4 ,

$$\frac{C_v}{4ac\sigma_a T^3} \frac{\partial T^4}{\partial t} = \frac{E}{a} - T^4. \quad (2.16)$$

Therefore the material energy time scale is

$$\tau_{\text{material}} = \frac{C_v}{4ac\sigma_a T^3}. \quad (2.17)$$

This quantity is a positive constant times the mean free time for radiation. For most values of C_v and T of importance in thermal radiation transport the value of Eq. (2.17) is much longer than the mean free time for radiation.

2.1.5 The thermal transport equation in neutron transport notation

The notation used above is the standard notation in the astrophysical and high-energy density physics community. But, the equations above can also be cast in a notation similar to that used in neutron transport. It is my hope that by bridging the nota-

tion gap, nuclear engineering students after me can avoid feeling like Kaspar Hauser in the notation of thermal transport¹

The thermal transport equation system (including the ODE describing the evolution of the temperature) is written as

$$\frac{1}{c} \frac{\partial I}{\partial t} + \hat{\Omega} \cdot \nabla I + \sigma_s I = \frac{1}{4\pi} \int_{4\pi} d\hat{\Omega}' \sigma_s I(\cdot, \hat{\Omega}') + \sigma_a \left(\frac{acT^4}{4\pi} - I \right) \quad (2.18)$$

$$C_v(T) \frac{\partial T}{\partial t} = c\sigma_a (E(\vec{r}, t) - aT^4). \quad (2.19)$$

Commonly, in applications of linear transport to nuclear engineering the variable used is the angular flux $\psi(\vec{r}, \hat{\Omega}, E, t)$. The angular flux is the product of the particle speed and the density of particles in a piece of phase space. The units of the specific intensity, I , are similar; the intensity is the particle speed (the speed of light) times the energy per particle ($h\nu$) times the phase space density of photons. More simply put the specific intensity is the angular flux times the energy per particle ($\psi = I/h\nu$).

Due to this simple relation some authors use the angular flux notation instead of the specific intensity notation [23]. This is done simply by swapping $I \rightarrow \psi$ in the transport equation. However, confusion can arise because by doing this change in notation, ψ loses its common meaning. The upside to this switch is that some other nice definitions “pop out”. The scalar flux, φ , often used in nuclear engineering is the integral of the angular flux over all directions of flight. After swapping variables the “scalar flux” becomes related to the energy density

$$cE = \int_{4\pi} I d\hat{\Omega} \quad \rightarrow \quad \varphi = \int_{4\pi} \psi d\hat{\Omega} = cE. \quad (2.20)$$

¹Kaspar Hauser was a foundling who appeared in Nuremberg in the 19th century. He had previously been kept in solitary confinement for his entire life, and he was bewildered by the sites of a bustling city. His life and appearance was one of the biggest mysteries of that century and the enigma of his appearance has never been solved [21,22].

Using the nuclear engineering notation the thermal transport system becomes

$$\frac{1}{c} \frac{\partial \psi}{\partial t} + \hat{\Omega} \cdot \nabla \psi + \sigma_s \psi = \frac{1}{4\pi} \int_{4\pi} d\hat{\Omega}' \sigma_s \psi(\cdot, \hat{\Omega}') + \sigma_a \left(\frac{acT^4}{4\pi} - \varphi \right) \quad (2.21)$$

$$C_v(T) \frac{\partial T}{\partial t} = \sigma_a (\varphi - acT^4). \quad (2.22)$$

However, it must be remembered that when using this notation the meaning and units of ψ and φ are not those of neutron transport. The silver lining on this whole discussion is that a nuclear engineer, upon seeing the thermal transport equations, can pell-mell replace I with ψ and E with φ/c as long as the change in meaning is understood.

2.2 Thermal Spherical Harmonics Equations

Methods of numerically solving the transport equation are generally grouped in terms of how they treat the angular variable $\hat{\Omega}$. The approach that we will be pursuing in this thesis is the spherical harmonics method (also known as the P_n method²). In this approach the dependence of I on the direction of flight is expanded in a finite series of spherical harmonics. The derivation of the three-dimensional spherical harmonics equations has been carried out in detail by Brunner [1] for neutron transport and here we will mostly just sketch the derivation and point to other sources for any metutials.

2.2.1 Spherical harmonics functions

The spherical harmonics are a set of complex-valued transcendental functions that take as input a point on the unit sphere. A point on the unit sphere can be described by a polar angle θ and an azimuthal angle φ , which take values $\theta \in [0, \pi]$, $\varphi \in [0, 2\pi]$, as diagrammed in Fig. 2.1. Using these angles the spherical harmonics have the form

$$Y_l^m(\theta, \varphi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) e^{im\varphi}. \quad (2.23)$$

²This has also been called the P_l method [10]




Figure 2.1: A general unit vector \vec{u} defined by (θ, ϕ) . The polar angle, θ is the angle between the z axis and the vector and the azimuthal angle is the angle between the x axis and the projection of \vec{u} in the $x - y$ plane.

In this equation l, m are integers that specify the order of the spherical harmonic. The associated Legendre functions are denoted by P_l^m , these are defined as [13]

$$P_l^m(x) = (-1)^m (1 - x^2)^{m/2} \frac{\partial^m P_l(x)}{\partial x^m} \quad m \in [0, l] \quad (2.24)$$

$$P_l^m(x) = (-1)^m \frac{(l - |m|)!}{(l + |m|)!} P_l^{|m|} \quad m \in [-l, 0), \quad (2.25)$$

which uses the definition of the Legendre polynomials

$$P_0 = 1 \quad (2.26)$$

$$P_l = \frac{1}{2^l l!} \frac{\partial^l}{\partial x^l} (x^2 - 1)^l. \quad (2.27)$$

It is often beneficial from a notation standpoint to write the spherical harmonics in terms of $\mu = \cos \theta$,

$$Y_l^m(\mu, \varphi) = \sqrt{\frac{2l + 1}{4\pi} \frac{(l - m)!}{(l + m)!}} P_l^m(\mu) e^{im\varphi}. \quad (2.28)$$

The spherical harmonics form a complete set of orthonormal basis functions [24] so

$$\int_{-1}^1 d\mu \int_0^{2\pi} d\varphi Y_l^m Y_{l'}^{*m'} = \delta_{mm'} \delta_{ll'}. \quad (2.29)$$

Hence, the intensity can be expanded in terms of spherical harmonics

$$I(\vec{r}, \hat{\Omega}, t) = \sum_{l=0}^{\infty} \sum_{m=-l}^l I_l^m(\vec{r}, t) Y_m^l(\mu, \varphi), \quad (2.30)$$

with expansion coefficients

$$I_l^m(\vec{r}, t) = \int_{-1}^1 d\mu \int_0^{2\pi} d\varphi I(\vec{r}, \mu, \varphi, t) Y_l^{*m}(\mu, \varphi). \quad (2.31)$$

At this point we note that the I_0^0 coefficient is related to the energy density of the radiation field by

$$E = \frac{I_0^0}{2c\sqrt{\pi}}. \quad (2.32)$$

The $2\sqrt{\pi}$ factor arises from the normalization constant in the spherical harmonics definition.

2.2.2 Spherical harmonics moments of the grey transport equation

The P_n equations are found by taking spherical harmonics moments of the transport equation, Eq. (2.6). Many of the resulting terms are easy to integrate

$$\int_{-1}^1 d\mu \int_0^{2\pi} d\varphi Y_l^{*m}(\mu, \varphi) \frac{\partial I}{\partial t} = \frac{\partial I_l^m}{\partial t} \quad (2.33)$$

$$\int_{-1}^1 d\mu \int_0^{2\pi} d\varphi Y_l^{*m}(\mu, \varphi) \sigma_a I = \sigma_a I_l^m \quad (2.34)$$

$$\int_{-1}^1 d\mu \int_0^{2\pi} d\varphi Y_l^{*m}(\mu, \varphi) \frac{acT^4}{4\pi} = \delta_{l0} \delta_{m0} \frac{acT^4}{2\sqrt{\pi}} \quad (2.35)$$

$$\int_{-1}^1 d\mu \int_0^{2\pi} d\varphi Y_l^{*m}(\mu, \varphi) \int_{-1}^1 d\mu' \int_0^{2\pi} d\varphi' \sigma_s I(\cdot, \mu', \varphi') = \sigma_s I_0^0 \delta_{l0} \delta_{m0}. \quad (2.36)$$

The term that is difficult to integrate is the streaming term, $\hat{\Omega} \cdot \nabla I$, and will require some cagey use of the spherical harmonics.

There are many nice properties of the spherical harmonics functions, including the addition theorem and a host of useful recursion relationships [24, 25]. The properties

that we will need are the following recursion relations

$$\mu Y_l^m = A_l^m Y_{l+1}^m + B_l^m Y_{l-1}^m \quad (2.37)$$

$$\sqrt{1 - \mu^2} e^{i\varphi} = -C_l^m Y_{l+1}^{m+1} + D_l^m Y_{l-1}^{m+1} \quad (2.38)$$

$$\sqrt{1 - \mu^2} e^{-i\varphi} = E_l^m Y_{l+1}^{m-1} - F_l^m Y_{l-1}^{m-1}, \quad (2.39)$$

where

$$A_l^m = \sqrt{\frac{(l-m+1)(l+m+1)}{(2l+3)(2l+1)}} \quad B_l^m = \sqrt{\frac{(l-m)(l+m)}{(2l+1)(2l-1)}} \quad (2.40)$$

$$C_l^m = \sqrt{\frac{(l+m+1)(l+m+2)}{(2l+3)(2l+1)}} \quad D_l^m = \sqrt{\frac{(l-m)(l+m-1)}{(2l+1)(2l-1)}} \quad (2.41)$$

$$E_l^m = \sqrt{\frac{(l-m+1)(l-m+2)}{(2l+3)(2l+1)}} \quad F_l^m = \sqrt{\frac{(l+m)(l+m-1)}{(2l+1)(2l-1)}}. \quad (2.42)$$

Using these recurrence relations we arrive at the spherical harmonics equations [1]

$$\begin{aligned} \frac{1}{c} \frac{\partial I_l^m}{\partial t} + \frac{1}{2} \frac{\partial}{\partial x} (-C_{l-1}^{m-1} I_{l-1}^{m-1} + D_{l+1}^{m-1} I_{l+1}^{m-1} + E_{l-1}^{m+1} I_{l-1}^{m+1} - F_{l+1}^{m+1} I_{l+1}^{m+1}) \\ + i \frac{1}{2} \frac{\partial}{\partial y} (-C_{l-1}^{m-1} I_{l-1}^{m-1} + D_{l+1}^{m-1} I_{l+1}^{m-1} + E_{l-1}^{m+1} I_{l-1}^{m+1} - F_{l+1}^{m+1} I_{l+1}^{m+1}) \\ + \frac{\partial}{\partial z} (A_{l-1}^m I_{l-1}^m + B_{l+1}^m I_{l+1}^m) + \sigma_s I_l^m = \sigma_a \left(\frac{acT^4}{2\sqrt{\pi}} \delta_{l0} \delta_{m0} - I_l^m \right) + \sigma_s I_0^0 \delta_{l0} \delta_{m0}. \end{aligned} \quad (2.43)$$

The material temperature equation only involves the I_0^0 , due to the fact that the energy density is a constant times the I_0^0 moment. This equation is thus

$$C_v(T) \frac{\partial T}{\partial t} = \sigma_a (2\sqrt{\pi} I_0^0 - acT^4). \quad (2.44)$$

2.2.3 Truncating the Expansion: the closure issue

In Eqs. 2.43 we have left the expansion of I in spherical harmonics as an infinite series expansion. Obviously, that expansion must be made finite to arrive at a discretization of the angular variables. The simplest closure is the P_n closure³ which sets the I_{n+1}^m

³This terminology comes from transport theory in one dimension where the intensity can be expanded simply in Legendre polynomials (rather than the spherical harmonics) [10, 16]. In multiple dimensions this terminology is somewhat of a misnomer, however, in calling Eqs. 2.43 the P_n equations it is understood that we mean a spherical harmonics expansion.

moment (and all greater moments) equal to zero. This makes the sum in Eq. (2.30)

$$\sum_{l=0}^n \sum_{m=-l}^l = n^2 + 2n + 1, \quad (2.45)$$

and there are this many expansion coefficients, I_l^m . These expansion coefficients with $m \neq 0$ each have a real and imaginary part because the spherical harmonics functions do, and this nearly doubles the number of unknowns. But, there is a trick to reduce this amount by using the fact that the intensity is a real quantity [1]. Since this is the case, we can use the property of the spherical harmonics $\bar{Y}_l^m = (-1)^m Y_l^{-m}$. This means that we can eliminate the $m < 0$ coefficients and have only

$$2 \sum_{l=0}^n \sum_{m=1}^l + \sum_{l=0}^n = n^2 + 2n + 1 \quad (2.46)$$

unknowns.

Table 2.2.3 gives the number of unknowns for various expansion orders. Those listed all have n odd. This is because even-order expansions have a peculiar zero-eigenvalue that prevents particles from moving (this zero eigenvalue has an eigenvector that contains I_0^0 and therefore prevents part of the energy density to be advected through the problem) [1, 2]. Other authors while looking at steady state problems have attributed the poor performance of even-order expansions to the requirement of an odd number of boundary conditions which would require the introduction of an asymmetry [12] and the notion that an even-order expansion supposes that I is continuous at $\mu = 0$ [10]. It has also been noted that even-order expansions do worse than their one ordinate lower odd counterpart (i.e. P_2 is worse than P_1) [10, 16]. The existence of this zero eigenvalue, though, is probably the most salient argument why even-order expansions should be shunned for time dependent problems.

Other types of closures suppose some form for the I_{n+1}^m moment and use this postulated form to close the equations. Pomraning derived a closure for the P_n equations

Table 2.1: Number of Unknowns for various 3D P_n expansions

P_1	4
P_3	16
P_5	36
P_7	64
P_9	100
P_{19}	400

where the qualitative shape of the angular dependence was assumed [26]. For example, in problems where beam-like distributions are likely, the closure can be developed to model such a distribution. This work also shows how a closure can be developed that exactly captures the Case discrete modes [10]. Su and Pomraning revisited this work and found that in multi-material problems, thin materials can degrade the performance of a P_3 -like closure that captures the Case modes [27].

These types of closures are linear in the sense that the equations are closed in a way that does not depend on the solution. Various nonlinear closures exist that choose a closure based on the values of the P_n moments [1, 5, 28]. Most times these closures are used as two-moment closures (i.e. to close the P_1 equations). One problem with nonlinear closures is that they make the eigenvalues of the P_n equations solution dependent and therefore allow shocks in the radiation field – a clearly unphysical phenomenon. However, such closures can avoid some drawbacks in the spherical harmonics approach as will be discussed later.

In this thesis we will only deal with the classical P_n closure. The methods developed throughout will be applicable to any hyperbolic, linear closure of the spherical harmonics equations. Furthermore, with proper consideration, the methods could be extended to the nonlinear closures as was done by Brunner [1].

2.2.4 Two-Dimensional P_n Equations

In two-dimensional problems we can drop the dependence of the flux any one of the (x, y, z) variables. Inspecting Eq. (2.43) we notice that dropping the y -dependence decouples the imaginary and real parts of the spherical harmonics expansion and we only solve for the real unknowns. This nearly halves the number of unknowns to

$$\text{unknowns} = \frac{1}{2}(n^2 + 3n) + 1. \quad (2.47)$$

We can write the 2D P_n equations using matrices to contain the streaming terms; these matrices are the Jacobians of the 2D equations. These matrices will not be symmetric because of the trick we used to eliminate the negative m unknowns. The 2D equations in matrix form are

$$\frac{1}{c} \frac{\partial \vec{I}}{\partial t} + \mathbf{A}_x \frac{\partial \vec{I}}{\partial x} + \mathbf{A}_z \frac{\partial \vec{I}}{\partial z} + \sigma_s \vec{I} = \sigma_a \left(\frac{ac\Gamma^4}{2\sqrt{\pi}} \delta_{l0} \delta_{m0} - \vec{I} \right) + \sigma_s \vec{I} \delta_{l0} \delta_{m0}. \quad (2.48)$$

We have written \vec{I} as a vector of spherical harmonics moments of I grouped in blocks of constant m

$$\vec{I} = (I_0^0, I_1^0, \dots, I_n^0, I_1^1, \dots, I_n^1, \dots, I_n^n)^t, \quad (2.49)$$

and $\mathbf{A}_x, \mathbf{A}_z$ are the Jacobian matrices. In P_1 these matrices are

$$\mathbf{A}_x = \begin{pmatrix} 0 & 0 & -\sqrt{\frac{2}{3}} \\ 0 & 0 & 0 \\ -\sqrt{\frac{1}{6}} & 0 & 0 \end{pmatrix}, \quad (2.50)$$

$$\mathbf{A}_z = \begin{pmatrix} 0 & \sqrt{\frac{1}{3}} & 0 \\ \sqrt{\frac{1}{3}} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (2.51)$$

For the sake of expediency we can also write Eq. (2.48) using the form

$$\nabla \cdot \mathbf{F}(\vec{I}) = \mathbf{A}_x \frac{\partial \vec{I}}{\partial x} + \mathbf{A}_z \frac{\partial \vec{I}}{\partial z}. \quad (2.52)$$

This reinforces the notion that this term is the flux of information out of a given point.

2.2.5 One-Dimensional P_n Equations

For problems in slab geometry we can drop another spatial variable. Also, in one dimension, the intensity can be azimuthally symmetric. In this case we can drop all moments with $m \neq 0$. Hence, we choose to solve the P_n equations only in the z -direction and there are only $n + 1$ moments to solve for. This makes the P_1 Jacobian

$$\mathbf{A}_z = \begin{pmatrix} 0 & \sqrt{\frac{1}{3}} \\ \sqrt{\frac{1}{3}} & 0 \end{pmatrix}. \quad (2.53)$$

2.3 Other Methods

The spherical harmonics equations are not the only way to discretize the direction of flight variables. The other important approaches are the discrete ordinates method (S_n), flux-limited diffusion, and implicit Monte Carlo.

2.3.1 Discrete ordinates methods

The discrete ordinates methods discretize the angular variables by asserting that Eq. (2.1) holds for specific directions, $\{\hat{\Omega}_m\}$, and using a quadrature rule to approximate the integral of the intensity [12,13]. In the case of isotropic scattering, this method is

$$\frac{1}{c} \frac{\partial I_m}{\partial t} + \hat{\Omega}_m \cdot \nabla I_m + \sigma_s I_m = \frac{1}{4\pi} \int_{4\pi} d\hat{\Omega}' \sigma_s I(\cdot, \hat{\Omega}') + \sigma_a \left(\frac{acT^4}{4\pi} - I_m \right) \quad (2.54)$$

$$\int_{4\pi} d\hat{\Omega}' I(\cdot, \hat{\Omega}') \approx \sum_{m=1}^M w_m I_m. \quad (2.55)$$

This method has been very popular in steady-state nuclear reactor calculations and it has also been applied to thermal photon transport. One drawback to this method is the existence of ray effects in multi-dimensional problems. These ray effects are the

result of only using the intensity at discrete angles in the quadrature to find the energy density [12]. This can allow the energy density to have oscillations. These oscillations can have deleterious effects in coupled radiation hydrodynamics simulations where the radiation field influences the dynamics of material motion. These oscillations can cause material to develop “hot spots” where the peak of the oscillation is high, and hence cause the system to evolve incorrectly. S_n methods can also be troublesome in how the spatial variables are discretized. It is possible in many popular spatial discretizations to have negative energy densities or even have the method become unstable [29].

2.3.2 Flux-limited Diffusion

The diffusion method in 1D is derived by starting with the P_1 equations

$$\frac{1}{c} \frac{\partial I_0}{\partial t} + \frac{1}{\sqrt{3}} \frac{\partial I_1}{\partial z} = \sigma_a \left(\frac{acT^4}{2\sqrt{\pi}} - I_0 \right) \quad (2.56)$$

$$\frac{1}{c} \frac{\partial I_1}{\partial t} + \frac{1}{\sqrt{3}} \frac{\partial I_0}{\partial z} = -\sigma_t I_1. \quad (2.57)$$

Next, we make an assumption on the I_1 moment, namely that $\partial I_1 / \partial t$ is small

$$\frac{\partial I_1}{\partial t} \approx 0. \quad (2.58)$$

This assumption has been shown to be accurate to leading order in an asymptotic expansion of the transport equation [30,31]. The expansion is valid when collisions dominate the streaming of particles. Upon making this assumption we get Fick’s law

$$I_1 = -\frac{1}{\sigma_t \sqrt{3}} \frac{\partial I_0}{\partial z}, \quad (2.59)$$

and this leads to the diffusion equation

$$\frac{1}{c} \frac{\partial I_0}{\partial t} - \frac{1}{3} \frac{\partial}{\partial z} \frac{1}{\sigma_t} \frac{\partial I_0}{\partial z} = \sigma_a \left(\frac{acT^4}{2\sqrt{\pi}} - I_0 \right). \quad (2.60)$$

Oftentimes, the quantity $1/3\sigma_t$ is called the diffusion coefficient D .

This equation has a singular streaming term in vacuum. This is quite a shortcoming in photon transport where voids are often a part of the problem domain – in neutron transport in traditional nuclear reactors generally there are few vacuum regions. Even without voids, problems can arise. The diffusion equation allows the first moment of the intensity to be greater than the zeroth moment

$$I_0 < I_1 = -D \frac{\partial I_0}{\partial z}. \quad (2.61)$$

This is unphysical because it implies that the number of moving particles is greater than the number of particles that exist. This can be corrected by the use of a flux limiter. This method limits the size of D to keep the first moment smaller than the zeroth. In general most flux limiters cause D to be a function of I_0 and $\partial I_0 / \partial z$, and therefore also make the problem nonlinear. There are a number of variations on this theme and the performance of a particular flux limiter is problem dependent [28, 32].

2.3.3 Implicit Monte Carlo

The implicit Monte Carlo (IMC) method is a stochastic method of solving the coupled transport and material energy system. It is built on the foundation of the Monte Carlo method for linear transport [12]. Implicit Monte Carlo simulates sets of sample particles moving through the problem domain. By sampling more and more photons the correct intensity can be found. However, the material temperature equation is nonlinear, and some approximation must be made to handle the absorption and re-emission of photons during a time step. There are various approaches to stochastically simulating the material temperature equation [23]. One marked difference between IMC and linear Monte Carlo is the need for a spatial grid (to handle the variation of temperature) and a time step dependent error term. In linear transport Monte Carlo solutions are generally considered “truth” (modulo the sampling noise); in nonlinear problems the introduc-

tion of a grid and temporal truncation error makes such an assertion dubious, however, IMC results are often considered authoritative *faute de mieux*⁴.

An important issue with IMC methods is the noise in the solution due to the sampling of a finite number of particles. Such noise can seed instabilities in radiation hydrodynamics simulations. Specifically, in problems where magnetohydrodynamics is involved, the noise can initiate a Rayleigh-Taylor instability in the simulation [4].

⁴This phrase comes from the French “for want of something better”. IMC may not be perfect, but it is a common feeling that it is the best method available when computational cost is no object.

CHAPTER III

Upwind Finite Volume Discretization

A numerical method for hyperbolic partial differential equations can be said to be upwind if the discrete equations only allow information to propagate in the same directions as the underlying continuous differential equations. For example, in one dimensional scalar advection with positive speed of propagation an upwind method moves information from left to right and does not use information from the downstream direction¹. In a system of conservation laws, upwinding is done in regards to the directions of the characteristic variables.

The principle of upwinding is commonly applied in numerical methods for radiation transport - even if that term is not often used [33]. Transport sweeps in discrete ordinates (S_n) methods move information in an upwind direction and the characteristic directions are obvious in the S_n formulation. For the spherical harmonics equations, the process of upwinding is not so obvious.

We will now proceed to describe the numerical method for the two dimensional P_n equations. This discretization falls into the category of a finite volume scheme (although in 2D perhaps finite *area* would be more apropos). These types of schemes average the underlying equations over a computational cell and work in terms of these

¹High resolution schemes do use information from the downstream direction but only to reconstruct the slope inside a cell.

averaged quantities. On regular grids these methods become finite difference methods.

3.1 Cell-Averaged P_n Equations

To begin we take a generic two-dimensional spatial cell denoted by i and average the P_n equations (Eqs. (2.43-2.44)) over that cell. Performing this average on the P_n equations gives

$$\frac{1}{\Delta V} \int_{\Delta V} dV \left(\frac{1}{c} \frac{\partial \vec{I}}{\partial t} + \nabla \cdot \mathbf{F}(\vec{I}) + \sigma_t \vec{I} = \sigma_s \vec{I} \delta_{l0} \delta_{m0} + \sigma_a \delta_{l0} \delta_{m0} a c T^4 \right), \quad (3.1)$$

where in two dimensions the streaming term is

$$\nabla \cdot \mathbf{F}(\vec{I}) = \mathbf{A}_x \frac{\partial \vec{I}}{\partial x} + \mathbf{A}_z \frac{\partial \vec{I}}{\partial z}. \quad (3.2)$$

If we define the cell-averaged moments of the intensity,

$$\vec{I}_i \equiv \frac{1}{\Delta V} \int_{\Delta V} \vec{I} dV, \quad (3.3)$$

and assume T has a constant value in each cell, T_i , Eq. (3.1) becomes

$$\frac{1}{c} \frac{\partial \vec{I}_i}{\partial t} + \frac{1}{\Delta V} \oint_{\partial \Delta V} \mathbf{F}(\vec{I}) \cdot \hat{n} ds + \sigma_t \vec{I}_i = \sigma_s \vec{I}_i \delta_{l0} \delta_{m0} + \sigma_a \delta_{l0} \delta_{m0} a c T_i^4. \quad (3.4)$$

We have used the divergence theorem [34] to convert the volume integral over the streaming term in Eq. (3.1) to a surface integral. To treat the surface integral we break the surface into a series of faces and sum over each face given by

$$\oint_{\partial \Delta V} \mathbf{F}(\vec{I}) \cdot \hat{n} ds = \sum_{\gamma} F_{i,\gamma} \Delta s \cdot \hat{n}_{\gamma}, \quad (3.5)$$

where $F_{i,j}$ is the numerical approximation for the flux at cell interface γ , Δs is the length of the cell face, and \hat{n}_{γ} is the unit outward normal at interface γ . In the case of curved faces, this outward normal would be an average normal.

Eq. (3.4) is not complete because the flux terms are in terms of cell-edge quantities. We need a way to relate the cell-edge quantities to the cell-averaged quantities in the

other terms. A method of doing this in a physical way is the heart of the Riemann solver approach.

3.2 Upwinding The Equations Using A Riemann Solver

3.2.1 The flux between cells

Finding the flux at a given cell interface in terms of cell-averaged quantities can be done in various ways. One possible way would be to treat the cell interface as a Marshak boundary and write incoming/outgoing fluxes in terms of half-range integrals of the angular flux. This has been done in steady-state problems to with the P_1 equations to be used as an accelerator for S_n methods [35].

The method we will explore here is a Riemann solver based method. In this approach we exactly solve the the free-streaming P_n equations with a step-function initial condition. This problem is called a Riemann problem and we will now use it to finish our spatial discretization.

The flux at a particular cell interface in is given by

$$\hat{n}_\gamma \cdot F_{i,\gamma} = n_x \mathcal{A}_x \vec{I}_{i,\gamma} + n_z \mathcal{A}_z \vec{I}_{i,\gamma}. \quad (3.6)$$

Notice $\vec{I}_{i,\gamma}$ is a cell-edge value, and the rest of the variables in Eq. (3.4) are cell-centered quantities. Finding these cell-edge values in terms of cell-centered quantities in an upwinded fashion is the pith of the Riemann solver. If we define a coordinate χ , that is along the normal to a generic cell interface, the flux across that interface becomes

$$F_{i,\gamma} = \mathcal{A}_\chi \vec{I}_{i,\gamma}, \quad (3.7)$$

where we have defined $\mathcal{A}_\chi \equiv n_x \mathcal{A}_x + n_z \mathcal{A}_z$.

The problem we will now solve is the free-streaming problem in this direction

$$\frac{\partial \vec{I}}{\partial t} + \mathcal{A}_\chi \frac{\partial \vec{I}}{\partial \chi} = 0, \quad (3.8)$$




Figure 3.1: The initial conditions for the Riemann problem has two flat solutions discontinuously joined at $\chi = 0$. This effectively models the boundary between two cells.

with initial condition, as shown Fig. 3.1, given by

$$\vec{I}(\chi, 0) = \begin{cases} \vec{I}_l & \chi < 0, \\ \vec{I}_r & \chi > 0. \end{cases} \quad (3.9)$$

There is an exact solution to this problem which can be found in Refs. [1,2,5,36–38].

The solution gives the flux at $\chi = 0$ as

$$\mathcal{A}_\chi \vec{I} \Big|_{\chi=0} = \frac{1}{2} \mathcal{A}_\chi (\vec{I}_l + \vec{I}_r) - \frac{1}{2} |\mathcal{A}_\chi| (\vec{I}_r - \vec{I}_l). \quad (3.10)$$

In Eq. 3.10, the absolute value signs around a matrix signify that the matrix was evaluated using the absolute values of its eigenvalues – specifically

$$|\mathcal{A}_\chi| = \sum_k r_k^\vec{} |\lambda_k| l_k^\vec{}, \quad (3.11)$$

where $r_k^\vec{}$, $l_k^\vec{}$, and λ_k are the k^{th} right and left eigenvectors and eigenvalues of A_χ . This solution is an upwind solution in the sense that positive eigenvalues move information from the left to right. The idea of left and right loses its standard meaning on an unstructured mesh, but the idea is that moving in the positive direction – towards increasing χ – is moving from left to right. Eq. (3.10) can be thought of as a centered-difference scheme – the first term on the RHS in Eq. (3.10) – plus just the right amount of dissipation to make the centered difference scheme stable.




Figure 3.2: The two-dimensional Cartesian grid. A generic cell (i, j) is shaded in grey.

Now using Eq. (3.10) we have discretized the P_n equations in space. This gives a first order spatial method because the solution in each cell is assumed to be flat in the solution of the Riemann problem. We will now specialize this result to further illuminate the method.

3.2.2 Example: Two-dimensional Cartesian Grid

We take a Cartesian grid where we denote each cell and cell averaged quantities by an ordered pair (i, j) . Furthermore, we will use half ordinates to reference cell edges (e.g. the right boundary of cell (i, j) would be given by $(i + 1/2, j)$). For a graphical depiction of this grid see Fig. 3.2. Note that $\Delta x \neq \Delta z$, nor does the mesh need to be uniform.

For this geometry the fluxes across each face of cell (i, j) are given by

$$F_{i+1/2,j} = \frac{1}{2}\mathcal{A}_x \left(\vec{I}_{(i,j)} + \vec{I}_{(i+1,j)} \right) - \frac{1}{2}|\mathcal{A}_x| \left(\vec{I}_{(i+1,j)} - \vec{I}_{(i,j)} \right) \quad (3.12)$$

$$F_{i-1/2,j} = -\frac{1}{2}\mathcal{A}_x \left(\vec{I}_{(i-1,j)} + \vec{I}_{(i,j)} \right) + \frac{1}{2}|\mathcal{A}_x| \left(\vec{I}_{(i,j)} - \vec{I}_{(i-1,j)} \right) \quad (3.13)$$

$$F_{i,j+1/2} = \frac{1}{2}\mathcal{A}_z \left(\vec{I}_{(i,j)} + \vec{I}_{(i,j+1)} \right) - \frac{1}{2}|\mathcal{A}_z| \left(\vec{I}_{(i,j+1)} - \vec{I}_{(i,j)} \right) \quad (3.14)$$

$$F_{i,j-1/2} = -\frac{1}{2}\mathcal{A}_z \left(\vec{I}_{(i,j-1)} + \vec{I}_{(i,j)} \right) + \frac{1}{2}|\mathcal{A}_z| \left(\vec{I}_{(i,j)} - \vec{I}_{(i,j-1)} \right). \quad (3.15)$$

$$(3.16)$$

Given these fluxes, the net flux for cell (i, j) becomes

$$\sum_{\gamma} \hat{n}_{\gamma} \cdot F_{i,j} \Delta s = \frac{\Delta z}{2} \left(\mathcal{A}_x \left(\vec{I}_{(i+1,j)} - \vec{I}_{(i-1,j)} \right) - |\mathcal{A}_x| \left(\vec{I}_{(i+1,j)} - 2\vec{I}_{(i,j)} + \vec{I}_{(i-1,j)} \right) \right) \quad (3.17)$$

$$+ \frac{\Delta x}{2} \left(\mathcal{A}_z \left(\vec{I}_{(i,j+1)} - \vec{I}_{(i,j-1)} \right) - |\mathcal{A}_z| \left(\vec{I}_{(i,j+1)} - 2\vec{I}_{(i,j)} + \vec{I}_{(i,j-1)} \right) \right).$$

Using this net flux we can write the original cell-averaged thermal radiation P_n Eqs. (3.4)

as

$$\begin{aligned} \frac{1}{c} \frac{\partial \vec{I}_{(i,j)}}{\partial t} + \frac{1}{2\Delta x} \left(\mathcal{A}_x \left(\vec{I}_{(i+1,j)} - \vec{I}_{(i-1,j)} \right) - |\mathcal{A}_x| \left(\vec{I}_{(i+1,j)} - 2\vec{I}_{(i,j)} + \vec{I}_{(i-1,j)} \right) \right) \\ + \frac{1}{2\Delta z} \left(\mathcal{A}_z \left(\vec{I}_{(i,j+1)} - \vec{I}_{(i,j-1)} \right) - |\mathcal{A}_z| \left(\vec{I}_{(i,j+1)} - 2\vec{I}_{(i,j)} + \vec{I}_{(i,j-1)} \right) \right) \end{aligned} \quad (3.18)$$

$$+ \sigma_t \vec{I}_{(i,j)} = \sigma_s \vec{I}_{(i,j)} \delta_{l0} \delta_{m0} + \sigma_a \delta_{l0} \delta_{m0} a c T_{(i,j)}^4.$$

One aspect of these equations is that the x and z directions are decoupled; there is no flow in the $x - z$ direction. This could contribute some grid effects for flows that are not aligned to the grid.

3.2.3 Example: One-dimensional Equations

In one dimension we can simply use Eq. (3.18) and drop the flow in the x direction.

Then we have the equations

$$\frac{1}{c} \frac{\partial \vec{I}_j}{\partial t} + \frac{1}{2\Delta z} \left(\mathcal{A}_z \left(\vec{I}_{j+1} - \vec{I}_{j-1} \right) - |\mathcal{A}_z| \left(\vec{I}_{j+1} - 2\vec{I}_j + \vec{I}_{j-1} \right) \right) + \sigma_t \vec{I}_j \quad (3.19)$$

$$= \sigma_s \vec{I}_j \delta_{l0} \delta_{m0} + \sigma_a \delta_{l0} \delta_{m0} a c T_j^4. \quad (3.20)$$

3.2.4 Extra Dissipation

The flux between cells given by the Riemann solver in Eq. (3.10) has an issue for characteristics that have an eigenvalue of zero. Such modes are treated using a centered-difference approach, which can be unstable. To handle this phenomenon, Brunner [1] suggested that these zero eigenvalues be replaced by the smallest nonzero eigenvalue when $|\mathbf{A}_\chi|$ is constructed. This adds extra dissipation to these modes which adds extra-dissipation, and is equivalent to treating these modes using a Lax-Friedrichs solver.

3.2.5 Direction Splitting Issues

This method of upwinding the finite volume P_n equations does have a drawback in terms of how it treats the inherent multidimensionality of the underlying equations. The direction perpendicular to each face is split from the other directions in the system. This could cause effects of the grid to appear in the solution. The easiest way to see this is to think of a Cartesian grid with a wave traveling diagonally to the grid. In the set up we have described the wave has to travel across two cells to move in its diagonal direction, causing the wave to be smeared out transverse to the direction of flow. This phenomenon is diagrammed in Fig. 3.3: the red lines represent information propagating skew to the grid, the green lines demonstrate how the numerical method would move that information.




Figure 3.3: This cartoon shows how a dimensionally split numerical method would propagate information moving skew to the grid.

3.3 High-Resolution Correction

The discretization given by Eq. (3.10) is a first order method in Δx and Δz . This is due to the fact that it presupposes the values in each cell to constant. To make a higher order method we would like to reconstruct a slope inside each cell using the values from neighboring cells. Of course in doing so we would like to avoid creating artificial extrema in the solution. Unfortunately, the situation is not all that simple: Godunov's Theorem² [40] demonstrates that it is impossible to have a linear method high order method that does not introduce artificial oscillations. Hence, to achieve better than first order spatial accuracy we must reconstruct the slope in each cell using a nonlinear method.

Much work has been done on developing schemes that are better than first order in space and do not introduce oscillations. There is a body of literature [41–44] that has been devoted to methods of reconstructing the gradient on unstructured grids while

²Here I want to contribute my opinion of the correct pronunciation of this name. Seminars on the topic of hyperbolic equations are peppered with two distinct means of pronouncing Godunov. In its entry on the Russian historical figure and subject of an opera by Mussorgsky, *The American Heritage* dictionary [39] gives the preferred pronunciation as *göod'n - òf* over the common *gòd'n - òf*.

maintaining certain desirable properties such as exactly capturing the gradient for a linear function. Another tack are the essentially non-oscillatory (ENO) and weighted ENO (WENO) schemes [45, 46] which use high order functions to reconstruct the solution using an adaptive stencil.

In this thesis we will reconstruct the slope perpendicular to each face of a particular cell using the harmonic mean limiter and the minmod slope limiter. In both approaches the slope (in our rotated coordinate system that is perpendicular to the cell interface) of cell i is calculated

$$m_- = \frac{I_i - I_{i-1}}{\Delta\chi} \quad (3.21)$$

$$m_+ = \frac{I_{i+1} - I_i}{\Delta\chi}, \quad (3.22)$$

for each element of \vec{l} . In the harmonic mean approach the slope in cell i is set to [47]

$$m_i = \begin{cases} \frac{2m_+m_-}{m_+ + m_-} & m_-m_+ > 0 \\ 0 & \text{otherwise} \end{cases} \quad (3.23)$$

This approach assures that the interpolated value is between $I_{i\pm 1/2}$ is between I_{i+1} and I_{i-1} [1], hence the term “limiter”. One benefit of this limiter is that it smoothly transitions between using m_+ and m_- away from extreme points, where the slope is set to 0. Due to this smoothness property this limiter was used to in our nonlinear implicit scheme (see results in Sec. 9.1.

In the minmod method m_- and m_+ are compared and the slope in cell i is set to

$$m_i = \begin{cases} m_- & |m_-| < |m_+| \text{ and } m_-m_+ > 0 \\ m_+ & |m_+| < |m_-| \text{ and } m_-m_+ > 0 \\ 0 & \text{otherwise} \end{cases} \quad (3.24)$$

If the cell is not an extreme point, slope m_i is limited to be the minimum of the slope to either side of the cell. The minmod limiter has the benefit of being linear once the choice of stencil (i.e. whether to use m_- or m_+) is made. This feature will be exploited in development of our quasi-linear implicit method (see Sec. 5.2).

3.4 Justification of High Resolution Method

We have yet to show if our high resolution method is guaranteed to not introduce artificial oscillations. To show this we will first make some definitions. The total variation (TV) of a moment of the intensity, I , is defined as

$$TV(I) \equiv \sum_{j=0}^J \Delta x |I_j - I_{j-1}| \quad (3.25)$$

An important property for a numerical method for hyperbolic equations is, total variation diminishing (TVD). This says that the total variation of the numerical method does not increase between time steps

$$TV(I^{k+1}) \leq TV(I^k) \quad (3.26)$$

for each time step k . This property is important for two reasons, the true solutions to hyperbolic conservation laws are TVD [40] and if a method is TVD it will not introduce artificial oscillations – this would increase the TV.

We state without proof that the first order Riemann method is TVD. This can easily be seen through a modified equation analysis and show that the modified equation has a diffusive term but no dispersive term [40]. Since the Riemann method is TVD, we can define the following recipe for a high resolution TVD method. First, we take the previous time step's solution and reconstruct the solution using a slope reconstruction technique. Then use this reconstructed solution to obtain the next time step via the Riemann method. If the reconstruction technique is TVD then its combination

with the Riemann solver will be TVD. Both the minmod and harmonic mean limiters produce TVD reconstructions, making our high resolution method TVD and hence non-oscillatory.

In an implicit time integration scheme, the slope reconstruction would have to be done on the next time step's solution. This leads to a nonlinear system of equations that need to be solved at each time step. The next chapter introduces some implicit time integration schemes and later in Chapter V we will show how these nonlinear systems can be solved.

CHAPTER IV

Time Integration

There are two diseparate time scales in the P_n equations for thermal transport: the speed of light time scale and the material temperature time scale. The speed of light time scale is orders of magnitude shorter than the material temperature scale and often the details of radiation propagation at the speed of light scale are not important. Given this situation, it is natural for us to want to integrate the semi-discrete P_n equations over a time step that is large in regards to the radiation propagation time scale but small on the material temperature time scale. For this reason we have investigated implicit time integration methods without a time step size restriction.

In this chapter we will present a general approach to the time integration of the semi-discrete P_n equations (Eq. (3.4)). In the end we arrive at two different schemes: the backward Euler method, which is a first order unconditionally stable method, and a high-resolution second order method based on the Crank-Nicolson method, which users a limiter based approach to achieve a non-oscillatory scheme. Most of the theorems and proofs in this chapter deal with the scalar linear advection equation rather than a system of conservation laws, such as the P_n equations. This is due to the fact that the upwind discretization for this scalar equation is equivalent to our first order Riemann solver for a linear system of equations.

4.1 θ Notation for Time Integration

To discretize the time variable in Eq.(3.4) we average the time derivative over Δt and combine the other terms in a convex sum of their values at the current and next time step. The semi-discrete P_n equations become

$$\frac{\vec{I}_i^{n+1} - \vec{I}_i^n}{c\Delta t} = \theta L_i^{n+1} + (1 - \theta)L_i^n, \quad (4.1)$$

where

$$L_i^n = -\frac{1}{\Delta V} \sum_j F_{i,j}^n \Delta s \cdot \hat{n}_j - \sigma_t \vec{I}_i^n + \sigma_s \vec{I}_i \delta_{l0} \delta_{m0} + \sigma_a \delta_{l0} \delta_{m0} ac T_i^4, \quad (4.2)$$

the superscripts refer to values at the n^{th} time step and $\theta \in [0, 1]$. The material energy equation in this notation is given as

$$C_v \frac{T_i^{n+1} - T_i^n}{\Delta t \sigma_a} = \theta (2\sqrt{\pi} I_{0,i}^{0,n+1} - ac T_i^{n+1,4}) + (1 - \theta) (2\sqrt{\pi} I_{0,i}^{0,n} - ac T_i^{n,4}). \quad (4.3)$$

By choosing specific values of θ we recover some common methods of time integration. In the following sections we will discuss some of these methods and present those which we have implemented.

4.2 The Backward Euler Method

The backward Euler method [48] (also called the implicit Euler method) is an implicit method where $\theta = 1$. This method only uses information from the future time step to evaluate the non-time derivative terms. The method is unconditionally stable for any size time step on a properly discretized hyperbolic problem, as can be shown by a simple example.

Take a linear advection equation with positive speed

$$\frac{\partial u}{\partial t} + c \frac{\partial u}{\partial x} = 0, \quad (4.4)$$

that is discretized with the first order upwind (aka the Courant-Isaacson-Rees [49]) method¹ and the backward Euler method

$$u_j^{n+1} - u_j^n = -\frac{c\Delta t}{\Delta x}(u_j^{n+1} - u_{j-1}^{n+1}). \quad (4.5)$$

If we write the solution as a Fourier series in x , i.e. $u_i^n = \sum_m F_m(n\Delta t)A_m e^{imj\Delta x}$, Eq. (4.5) becomes

$$F_m((n+1)\Delta t) - F_m(n\Delta t) = -\frac{c\Delta t}{\Delta x}(F_m((n+1)\Delta t) - F_m((n+1)\Delta t)e^{-im\Delta x}), \quad (4.6)$$

for each Fourier mode. We now look at the growth rate of the solution between two time steps. If the growth rate between time steps is no bigger the one in magnitude, the method is stable. This growth rate, g , is given by

$$g_m = \frac{F_m((n+1)\Delta t)}{F_m(n\Delta t)}. \quad (4.7)$$

With this definition, Eq. (4.6) gives g as

$$g_m = \frac{1}{1 + \frac{c\Delta t}{\Delta x}(1 - \cos(m\Delta x) + i \sin(m\Delta x))}. \quad (4.8)$$

Hence, $|g| \leq 1$ and the method unconditionally stable, as is well known. One interesting point to note is that as Δt gets bigger the growth rate gets smaller (i.e. the time integration introduces more damping). This suggests that just because it is possible to take an arbitrarily large time step does not mean the solution will be properly resolved.

This analysis of the stability of the backward Euler method is specific to the upwind discretization. The growth rate would change for different spatial discretizations (i.e. centered-difference or Lax-Friedrichs).

Finally, we will show that the following theorem:

Theorem 1. *The backward Euler method is unconditionally TVD for the upwind discretization of the linear advection equation with positive velocity.*

¹This method is the scalar version of our first order Riemann solver.

Proof. The backward Euler method can be written as

$$u_j^{n+1} = u_j^n - \frac{c\Delta t}{\Delta x}(u_j^{n+1} - u_{j-1}^{n+1}). \quad (4.9)$$

We will use the total variation to define a norm [40]

$$\|u^n\|_{TV} = TV(u^n). \quad (4.10)$$

Rearranging Eq. (4.9) gives

$$\left(1 + \frac{c\Delta t}{\Delta x}\right) u_j^{n+1} = u_j^n + \frac{c\Delta t}{\Delta x} u_{j-1}^{n+1}. \quad (4.11)$$

Taking the TV norm of Eq. (4.11) and using the triangle inequality gives

$$\left(1 + \frac{c\Delta t}{\Delta x}\right) \|u^{n+1}\|_{TV} \leq \|u^n\|_{TV} + \frac{c\Delta t}{\Delta x} \|u^{n+1}\|_{TV}, \quad (4.12)$$

which implies $\|u^{n+1}\|_{TV} \leq \|u^n\|_{TV}$ and the method is TVD. \square

We have now shown that the backward Euler method is unconditionally stable and TVD. However, the method is only first order accurate; next, we will look at a method that is second order accurate.

4.3 Crank-Nicolson Method

The Crank-Nicolson method (also known as the trapezoid method) [50] is a second order time integration scheme obtained by setting $\theta = 0.5$. However, the Crank-Nicolson scheme is only TVD when the time step size satisfies [8]

$$\Delta t \leq 2 \frac{c\Delta t}{\Delta x}. \quad (4.13)$$

We can demonstrate this with the following theorem,

Theorem 2. *For the upwind method, the Crank-Nicolson method will be TVD if the time step size satisfies Eq. (4.13).*

Proof. First we write the Crank-Nicolson method

$$u_j^{n+1} = u_j^n - \frac{c\Delta t}{2\Delta x}(u_j^{n+1} - u_{j-1}^{n+1} + u_j^n - u_{j-1}^n). \quad (4.14)$$

We rewrite this equation as

$$\left(2 + \frac{c\Delta t}{2\Delta x}\right) u_j^{n+1} = 2u_j^n - \frac{c\Delta t}{\Delta x}(u_j^{n+1} - u_{j-1}^{n+1}) + \frac{c\Delta t}{\Delta x} u_{j-1}^n. \quad (4.15)$$

Taking the TV norm and using the triangle inequality we get

$$\|u_j^{n+1}\|_{TV} \leq \|u_j^n - \frac{c\Delta t}{2\Delta x}(u_j^{n+1} - u_{j-1}^{n+1})\|_{TV}. \quad (4.16)$$

Notice that

$$u_j^{n+1} = u_j^n - \frac{c\Delta t}{2\Delta x}(u_j^n - u_{j-1}^n) \quad (4.17)$$

is the forward Euler method with a time step of $0.5\Delta t$. We state without proof that the forward Euler method is TVD when the time step size, $\Delta\tau$, satisfies

$$1 \geq \frac{c\Delta\tau}{\Delta x}. \quad (4.18)$$

If Δt satisfies Eq. (4.13) then the forward Euler method is TVD and Eq. (4.16) becomes

$$\|u_j^{n+1}\|_{TV} \leq \|u_j^n\|. \quad (4.19)$$

However, if Eq. (4.13) is not satisfied then the forward Euler step will not be TVD and the Crank-Nicolson method will not be TVD.

□

This result shows that while Crank-Nicolson is second order accurate, it has a time step restriction to be TVD. The above theorem is a generalization of a similar TVD theorem for a general linear multi-step methods [8]. That theorem also shows that any method better than first order will have a time step restriction. Now we can turn to a high-resolution time integration scheme that combines the unconditionally TVD property of the backward Euler method and the second order convergence of the Crank-Nicolson method.

4.4 High-Resolution Time Integration

Recently, Duraisamy, et al. presented a limiter based method to combine the use of the backward Euler method and the Crank-Nicolson method to create a higher than first order TVD method [8]. This method is based on defining θ locally in space and choosing it to go between Crank-Nicolson in smooth regions and backward Euler in non-smooth regions.

To locally define θ we invoke a lemma developed by Huynh [51] that states: *Given $f(t_1)$ and $f(t_2)$ and $f'(t_1)$ or $f'(t_2)$, the resulting quadratic interpolant between t_1 and t_2 is monotone in $[t_1, t_2]$ if $f'(t_1), f'(t_2) \in [0, 2s]$, where $s = \frac{f(t_2)-f(t_1)}{t_2-t_1}$.* This lemma can be used for time integration of the first order upwind scheme. We define

$$f(t_1) = u_j^n \quad (4.20)$$

$$f(t_2) = u_j^{n+1} \quad (4.21)$$

$$f'(t_1) = L_j^n = \frac{\partial u^n}{\partial t} = \frac{1}{\Delta x}(u_j^n - u_{j-1}^n) \quad (4.22)$$

$$f'(t_2) = L_j^{n+1} = \frac{\partial u^{n+1}}{\partial t} = \frac{1}{\Delta x}(u_j^{n+1} - u_{j-1}^{n+1}) \quad (4.23)$$

$$s_j^{n+1/2} = \frac{u_j^{n+1} - u_j^n}{\Delta t} \quad (4.24)$$

We would like to adjust the value of θ locally to guarantee monotonicity as in the lemma. This will cause our method to smoothly transition between a Crank-Nicholson and backward Euler scheme at different points in the problem domain. To apply the lemma we first check to see if

$$L_j^{n+1} (L_j^{n+1} - 2s_j) \leq \epsilon \quad \text{and} \quad L_j^n (L_j^n - 2s_j) \leq \epsilon, \quad (4.25)$$

for ϵ small and positive. If Eq. (4.25) true then the assumptions of the lemma are satisfied and θ_j is set to 0.5. Otherwise we check to see if L_j^{n+1} or L_j^n has a different sign as s_j ,

i.e. if

$$L_j^n s_j < 0 \quad \text{or} \quad L_j^{n+1} s_j < 0. \quad (4.26)$$

In the case Eq. (4.26) is true θ_j is set to be 1 because the assumptions of the lemma are false, and the backward Euler method is used. If neither Eq. (4.25) nor Eq. (4.26) are satisfied then we are in a situation where at least one of the L_j 's is too large. In this case we set

$$\theta_j = 1.0 - \min \left[\frac{s_j}{L_j^{n+1} + \epsilon}, \frac{s_j}{L_j^n + \epsilon}, 0.5 \right]. \quad (4.27)$$

The above process of deciding θ_j is equivalent to what we did in developing a high resolution spatial scheme: we adjusted our interpolation step to make our solution monotone. We can show that this method will be linearly stable by applying a similar analysis as we did for the backward Euler method.

Theorem 3. *The local θ method is stable for any size time step on the first order upwind method [8].*

Proof. The local θ method for the upwinded linear advection equation is written as

$$u_j^{n+1} = u_j^n - \frac{c\Delta t}{\Delta x} (\theta_j u_j^{n+1} - \theta_{j-1} u_{j-1}^{n+1} + (1 - \theta_j) u_j^n - (\theta_{j-1}) u_{j-1}^n). \quad (4.28)$$

Now we write the spatial part of Eq.(4.28) as a Fourier series, and for a particular Fourier mode

$$\begin{aligned} F_m((n+1)\Delta t) &= F_m(n\Delta t) - \frac{c\Delta t}{\Delta x} (\theta_j F_m((n+1)\Delta t) - \theta_{j-1} F_m((n+1)\Delta t) e^{-im\Delta x} \\ &\quad + (1 - \theta_j) F_m(n\Delta t) - (\theta_{j-1}) F_m(n\Delta t) e^{-im\Delta x}). \end{aligned} \quad (4.29)$$

The growth rate is then given by

$$g_m = \frac{1 - \frac{c\Delta t}{\Delta x} ((1 - \theta_j) - (1 - \theta_{j-1}) e^{-im\Delta x})}{(1 + \frac{c\Delta t}{\Delta x} (\theta_j - \theta_{j-1} e^{-im\Delta x}))}. \quad (4.30)$$

The minimum value of the denominator occurs when $\theta_j - \theta_{j-1}e^{-im\Delta x} = -0.5$. This implies that

$$g \leq \frac{1 - \frac{c\Delta t}{2\Delta x}}{1 - \frac{c\Delta t}{2\Delta x}} \quad (4.31)$$
$$\leq 1$$

The maximum growth rate is 1 and the method is stable. □

Now that we have presented different time integration schemes we have fully discretized the P_n equations. In the next chapter we press on by describing how we implement and solve these discrete equations.

CHAPTER V

Implementation

The impetus behind developing space and time discretizations is to solve the underlying equations in large scale simulations. By large scale we mean hundreds of thousands or millions of unknowns. This fact we have kept in mind in designing our implementation. The most important part of this chapter is the discussion of the quasi-linear method of solving the nonlinear systems that arise from the high resolution discretization of the P_n equations. The basis of the idea is grounded in the fact that the streaming operator is linear and the nonlinearities were added only to make a high order scheme that is non-oscillatory. Our quasi-linear approach allows us to make nonlinearities exist in the temperature source only. Below we will discuss how we use modern solver technology to implement our discretization.

5.1 Implicit Equations

The time integration schemes we detailed in the previous chapter were all implicit schemes. To use an implicit approach we must solve a system of equations at each time step. We will write our equations in the form

$$\Psi^{n+1} + \theta(\mathcal{L}\Psi^{n+1} + \mathcal{F}(\Psi^{n+1})) = \Psi^n - (1 - \theta)(\mathcal{L}\Psi^n + \mathcal{F}(\Psi^n)). \quad (5.1)$$

The vector Ψ represents all the radiation and temperature unknowns in our system. For a Cartesian mesh this vector has a length of $N_x \times N_z \times (\frac{1}{2}(N_{mom}^2 + 3N_{mom}) + 2)$. These unknowns are the number of cells in the x and z directions (N_x and N_z), and the number of unknowns from the spherical harmonic moments, (see Eq. (2.47)) where N_{mom} is the order of the P_n expansion. There is one more unknown per spatial cell than in Eq. (2.47) because here we have included the temperature variable. The linear operator, \mathcal{L} , represents the first order upwind discretization. The nonlinear terms from the high-resolution correction along with the temperature nonlinearities are contained in \mathcal{F} . This system of equations can be solved using a Newton-Krylov solver (see §5.3). An alternative means of solution for our particular system of nonlinear equations can be solved using a quasi-linear approach which we will describe next.

5.2 Quasi-Linear High-Resolution Implementation

The quasi-linear approach will be developed for the P_n equation independent of the temperature equation. Later we will show how to bring in the temperature coupling using a full nonlinear treatment for those terms.

For the high-resolution spatial scheme using the backward Euler method without temperature coupling, the system of equations that must be solved at each time step is written as

$$\Psi^{n+1} + \mathcal{L}\Psi^{n+1} + \mathcal{F}(\Psi^{n+1}) = \Psi^n. \quad (5.2)$$

Recall, that $\mathcal{F}(\Psi)$ contains all the nonlinear terms from the spatial discretization. Given the nature of the nonlinearities of $\mathcal{F}(\Psi)$ it is hard to compute an approximate Jacobian due to the discontinuities in the derivative of $\mathcal{F}(\Psi)$. The harmonic mean approach of van Leer does help to alleviate this issue, as that limiter smoothly combines the slopes to the left and right of each cell and only has a discontinuous derivative at an extreme

Table 5.1: Quasi-Linear Algorithm

1.	Solve for Ψ^* in $\Psi^* + \mathcal{L}\Psi^* = \Psi^*$ using GMRES
2.	Build the matrix $\mathcal{M}(\Psi^*)$
3.	Solve $\Psi^{n+1} + \mathcal{L}\Psi^{n+1} + \mathcal{M}(\Psi^*)\Psi^{n+1} = \Psi^n$ using GMRES

point. The situation is worse in the minmod case where the method switches between the slope to the left and right of the cell and also to zero in a non-smooth way.

But, the minmod limiter has an advantage that we exploit. With the minmod method once it is known whether to interpolate in a cell and in which direction to interpolate, the limiter is effectively linear. Using this fact can write Eq. (5.2) for the minmod limiter in quasi-linear form

$$\Psi^{n+1} + \mathcal{L}\Psi^{n+1} + \mathcal{M}(\Psi^{n+1})\Psi^{n+1} = \Psi^n. \quad (5.3)$$

Here the matrix \mathcal{M} contains the “choices” from the minmod limiter; by analyzing Ψ^{n+1} the stencil for reconstruction is determined and \mathcal{M} can be constructed. By using this structure of the minmod method we can avoid using a traditional nonlinear solver. This is not the case for the harmonic mean approach as that method nonlinearly combines the data from surrounding cells to compute a slope.

By solving the linear part of the system Eq. (5.1) we can retrieve the information needed to determine how to reconstruct the slope in each cell. We then can make the function $\mathcal{F}(\Psi)$ linear and perform another linear solve to get the final answer for the time step. This procedure is given in Table 5.1.

To build the operator $\mathcal{M}(\Psi^*)$, the solution vector is analyzed in each cell to determine the m_+ , m_- and, therefore, how to interpolate in that cell: to the left, to the right, or not at all. Then using this information we have determined the stencil for the high resolution method.

The first solve using GMRES (see §5.4) uses the previous time-step’s result as the

initial guess while the final solve takes Ψ^* as the guess. While the algorithm involves solving two linear systems, the solution of the second system is made much easier by the use of this initial guess. Our implementation uses the fact that a Krylov method such as GMRES only needs to perform the action of a matrix times a vector. To perform this action we multiply \mathcal{L} by the vector and add to it $\mathcal{M}(\Psi^*)$ times the vector. Also, we use a matrix-free representation of $\mathcal{M}(\Psi^*)$, which saves the cost of building a new matrix at each time step.

We also remark that this algorithm is only the first step in a series of iterations. After solving for Ψ^{n+1} one could repeat steps two and three of the algorithm until a converged solution is found. Our results, from both analysis and numerics below, suggest that this is not necessary and that the three steps outlined in the algorithm are sufficient. In other words, one iteration is enough.

5.2.1 Analysis of the Quasi-Linear Approach

One natural question about the quasi-linear method is whether it is justified to use the solution to the first-order method to approximate the shape of the solution of the high resolution scheme. We will answer this question by looking at the phase velocity of monotonic initial conditions in each method.

To simplify the analysis, we will look at the method on a scalar advection equation given by

$$u_t + au_x = 0, \tag{5.4}$$

where we have denoted derivatives by subscripts. Without loss of generality we assume that $a > 0$ for the rest of this discussion. The first-order Riemann method (also known as the upwind method) discretizes this equation as

$$u_j^{n+1} = u_j^n - \frac{a\Delta t}{\Delta x} (u_j^{n+1} - u_{j-1}^{n+1}). \tag{5.5}$$

This discretization is a first-order approximation to the Eq. (5.4) but it is a second-order approximation [40] to

$$u_t + au_x = \frac{a\Delta x}{2} \left(1 + \frac{a\Delta t}{\Delta x} \right) u_{xx}. \quad (5.6)$$

This is known as the model equation for the first-order Riemann numerical method. The theory of using model equations to analyze numerical methods can be found in Refs. 52 and 53.

We now derive the model equation for the minmod based method to compare the properties of the high resolution method with the first-order method. To begin we assert that the slope in each cell is computed using the value in the cell to the left. This is general and we will point out where this choice influences our result. Given this fact, the method is written as

$$u_j^{n+1} = u_j^n - \frac{a\Delta t}{\Delta x} \left(\frac{3}{2}u_j^{n+1} - 2u_{j-1}^{n+1} + \frac{1}{2}u_{j-2}^{n+1} \right). \quad (5.7)$$

Expanding the numerical solution, u , in a Taylor series about $u(x_j, t^{n+1})$ then gives

$$u_t + au_x = \frac{\Delta t}{2}u_{tt} + \frac{a\Delta x^2}{3}u_{xxx} + O(\Delta x^3) + O(\Delta t^2). \quad (5.8)$$

To continue the derivation we will eliminate the u_{tt} term by taking the derivative of Eq. (5.8) w.r.t. t to get

$$\begin{aligned} u_{tt} &= -au_{tx} + \frac{\Delta t}{2}u_{ttt} + \frac{a\Delta x^2}{3}u_{txxx} + O(\Delta x^3) + O(\Delta t^2) \\ &= -a \left(-au_x + \frac{\Delta t}{2}u_{tt} + \frac{a\Delta x^2}{3}u_{xxx} \right)_x + O(\Delta t) + O(\Delta x^3) \\ &= a^2u_{xx} + \frac{a\Delta x^2}{3}u_{txxx} - \frac{a^2\Delta x^2}{3}u_{xxxx} + O(\Delta t) + O(\Delta x^3). \end{aligned} \quad (5.9)$$

Next, we operate on Eq. (5.8) with $\partial^3/\partial x^3$ and find that

$$u_{txxx} = -au_{xxxx} + O(\Delta t) + O(\Delta x^2). \quad (5.10)$$

Using Eqs. (5.9) and (5.10) in Eq. (5.8) the model equation becomes

$$u_t + au_x = \frac{a^2 \Delta t}{2} u_{xx} + \frac{a \Delta x^2}{3} u_{xxx} - \frac{a^2 \Delta x^2 \Delta t}{3} u_{xxxx} + O(\Delta t^2) + O(\Delta x^3). \quad (5.11)$$

This equation would have the same structure had we used a different stencil to reconstruct the slope (i.e. if the upwind method were used and the cell to the right were used to reconstruct the slope).

To compute the phase velocity a wave in the solutions to Eqs. (5.6) and (5.11) we will expand u in a Fourier series and look for solutions in the form $e^{i(kx - c(k)t)}$. Inserting this into the upwind model equation yields

$$c(k) = ak - \frac{ia^2 k \Delta x}{2} \left(1 + \frac{a \Delta t}{\Delta x} \right). \quad (5.12)$$

The phase velocity of a wave is given by $c(k)/k$, and for the upwind method is

$$v_p^{\text{up}} = a - \frac{ia^2 k \Delta t}{2} - \frac{iak \Delta x}{2}. \quad (5.13)$$

This says that every wave travels at the same speed but is dissipated to different degrees depending on the wave number k . Similarly, by inserting a Fourier mode into Eq. (5.11) we can obtain the phase velocity for the minmod method,

$$v_p^{\text{mm}} = a - \frac{ia^2 k \Delta t}{2} - \frac{ak^2 \Delta x^2}{3} + \frac{ia^2 \Delta x^2 \Delta t k^3}{3}. \quad (5.14)$$

The wave speed for the high resolution scheme is different for different wave numbers. This can be seen in the third term on the LHS of Eq. (5.14). This term is real and changes the phase velocity based on the value of k^2 . The different wave speeds would be problematic for solutions with discontinuities because these solutions would have large range values of k and hence demonstrate considerable dispersion. The Gibbs oscillations near extreme points that are seen in linear second order methods are a result of this dispersion. It is at these extreme points that we wish to “turn off” the slope reconstruction by

setting $m_i = 0$ and use the first order method. Beyond this dispersion effect, Eq. (5.11) an amount of dissipation that decreases at a rate of Δx^2 rather than the Δx convergence of the upwind method in Eq. (5.6).

The heart of our quasi-linear method is found in using solutions of the first-order upwind method to determine how to reconstruct the solution in each cell. The real phase velocities for the upwind and minmod methods are the same to order $k^2 \Delta x^2$. The evaluation of $\mathcal{M}(\Psi)$ is fairly insensitive to the values of Ψ used. This is due to the fact that in building this matrix, only relative comparisons between values in neighboring cells matter, not the actual values in each cell. Therefore, the model equations having the same real phase velocities to order Δx^2 is sufficient to approximate $\mathcal{M}(\Psi)$. The first order solution is “good enough” to get the shape and fix the minmod choice. In our computational results from actual calculations will demonstrate this point.

A natural question at this point is whether it would be practical to approximate the shape of the solution using the previous time step’s solution; this would require only one linear solve per time step. We claim that this, in fact, would be a poor idea for large time steps where information can travel across multiple cells. The location of extreme points in the solution would almost certainly be at a different place than those in the initial data. Moreover, in the case of a small time step an implicit method would not be necessary.

5.2.2 Quasi-linear time integration

For our high-resolution time integration scheme detailed in the previous chapter, we can apply the same approach as for the spatial scheme. We simply modify the algorithm to determine the value of θ for each cell, the set of which we denote by the vector, $\vec{\theta}$. In this case we are using the first order solution in both time and space to determine

how to correctly use the high order method. This method works in nearly the same the way as the approach for just the spatial scheme because the equations are linear once $\vec{\theta}$ is known. The results we present below for high resolution time integration use this approach to calculate each time step.

5.2.3 Temperature coupling

Up to this point we have neglected how the material temperature equation is integrated. This equation is nonlinear and our quasi-linear approach would work only if the material equation and temperature source terms were linearized. However, the temperature equation is local and only feeds directly into the transport equations for its own cell. For this reason, we treat the temperature fully nonlinearly. To do this we use a Newton-Krylov solver in steps one and three of the algorithm in Table 5.1. This nonlinear treatment could be modified to use a linearization of the temperature equation, in which case the procedure in Table 5.1 would appear to be identical. We have not explored linearizing the temperature terms.

5.3 Newton-Krylov Methods

Newton's method (or the Newton-Raphson method) is a means of solving nonlinear equations through successive linearizations [54, 55]. The method iterates on an initial guess. For the solution to the nonlinear equation $F(x) = 0$, Newton's method gives the i^{th} iterate as

$$x_i = x_{i-1} - F'(x_{i-1})^{-1}F(x_{i-1}). \quad (5.15)$$

For a system of equations, Newton's method appears to be similar

$$\vec{x}_i = \vec{x}_{i-1} - \mathbf{F}'(\vec{x}_{i-1})^{-1}F(\vec{x}_{i-1}), \quad (5.16)$$

but the derivative of $F(\vec{x})$ is the Jacobian matrix of the function defined element by

element as

$$\mathbf{F}'(\vec{x})_{i,j} = \frac{\partial F_i}{\partial x_j}(x). \quad (5.17)$$

To use the Newton method for systems we must evaluate the product of the inverse of the Jacobian and the result of the function on the previous iterate. To evaluate this we must solve the linear system

$$\mathbf{F}'(\vec{x}_{i-1})\vec{s} = -F(\vec{x}_{i-1}). \quad (5.18)$$

Methods that use an approximate means to find \vec{s} , hence the product we needed, are called inexact Newton methods [55].

One notable example of an inexact Newton method are Newton-Krylov methods [55,56]. These use a Krylov based iterative solver such as GMRES (see §5.4). Krylov methods approximate the solution to a linear system within a specified tolerance. Inexact Newton methods can also approximate the Jacobian of F using finite differences. This is done in such a way that the Jacobian matrix need not be built (see §5.4.2).

5.4 GMRES

The generalized minimum residual method (GMRES) [57] is an iterative method for solving non-symmetric linear systems of the form $\mathbf{A}\vec{x} = \vec{b}$. GMRES is a Krylov subspace method, meaning that it uses the subspace

$$K_i(\mathbf{A}, \vec{r}_0) = \text{span}\{\vec{r}_0, \mathbf{A}\vec{r}_0, \mathbf{A}^2\vec{r}_0, \dots\}, \quad (5.19)$$

where $\vec{r}_0 = \vec{b} - \mathbf{A}\vec{x}_0$ is the residual from the initial guess \vec{x}_0 in the process of finding the solution. GMRES gives for the i^{th} iteration the value of \vec{x}_i as the vector in the space $\vec{x}_0 + K_i$ that minimizes $\|\vec{b} - \mathbf{A}\vec{x}_i\|_2$. This minimization is performed by solving a least squares problem of size $(i + 1) \times i$ which is inexpensive for small i . This procedure is

especially useful for solving large systems because the matrix \mathbf{A} need not be stored, only how \mathbf{A} acts on a vector needs to be known.

GMRES will converge to the solution $\vec{x} = \mathbf{A}^{-1}\vec{b}$ in at most n iterations where n is number of rows and columns of \mathbf{A} . However, one drawback of GMRES is that the Krylov subspaces, K_i must be stored and this could lead to a memory issue. However, there is a fix to this dilemma in *restarted* GMRES. In this method after some predetermined number of iterations, the old Krylov subspaces are purged and the last solution vector is used as \vec{x}_0 to restart the method. This method will not be guaranteed to converge and when to restart is a decision that must be made on a problem by problem basis. The use of GMRES has been previously explored for solving problems in neutron transport [58, 59]. It was found that GMRES is a robust solution method and that it shows strong promise when using a preconditioner.

5.4.1 Preconditioning

In preconditioning a linear system another matrix, \mathbf{M} is applied to the system:

$$\mathbf{M}^{-1}\mathbf{A}\vec{x} = \mathbf{M}^{-1}y. \quad (5.20)$$

If $\mathbf{M} = \mathbf{A}$, then we would have the solution immediately without doing an GMRES iterations. If this were the case then we would not need to be worrying about an iterative solver. When $\mathbf{M} = \mathbf{I}$, then we have done nothing. The power in preconditioning lies in choosing $\mathbf{M} \approx \mathbf{A}$ so that hopefully fewer iterations will be needed to solve the system, but in such a way that \mathbf{M}^{-1} is easy to compute. Algebraically, a good choice of \mathbf{M} will make the spectral radius \mathbf{A} smaller, but the notion that \mathbf{M}^{-1} approximates the inverse of \mathbf{A} will suffice for picking a preconditioner for many applications.

There are two types of preconditioners one can use: algebraic preconditioners and physics-based preconditioners. Algebraic preconditioners approximate the operator \mathbf{A}

in such a way that an approximate inverse can be cheaply found. One important type of this preconditioner is the incomplete LU (ILU) factorization. This factors the matrix \mathbf{A} into the form

$$\mathbf{A} = \mathbf{LU} - \mathbf{R} \quad (5.21)$$

where \mathbf{L} and \mathbf{U} are lower and upper triangular matrices respectively and \mathbf{R} is the error in the factorization. The error is neglected in defining the preconditioner, $\mathbf{M} = \mathbf{LU}$. This factorization is generally done in such a way that \mathbf{LU} has nearly the same nonzero structure as \mathbf{A} , i.e. where \mathbf{A} has a zero, so does \mathbf{LU} .

Variations on the theme of ILU factorizations include the ILU(0) preconditioner in which the sparsity of \mathbf{LU} is exactly the same as \mathbf{A} . Another approach is the ILU(τ) method [60] which sets an element of \mathbf{LU} to zero if that element's magnitude is less than τ times the norm of the row. An ILU($\tau = 0$) method would compute the actual LU decomposition of the matrix and require storage of a possibly dense matrix. A large value of τ would make the LU decomposition inaccurate and a less effective preconditioner.

The other type of preconditioner is the physics-based preconditioner. This type of preconditioner approximates the inverse of \mathbf{A} using a physical approximation to \mathbf{A} . For example, if \mathbf{A} represents a discretized P_3 transport system, a physics based preconditioner could be an operator that is just the streaming part of the P_3 transport system. The benefits from a physics-based preconditioner are seen when the preconditioner is easily inverted and represents a good approximation to the physics of the full operator. For instance, this was the observation when P_1 or diffusion operators were used to precondition the transport operator in discrete ordinates codes [35,61]. Another important type of physics-based preconditioner is the multigrid approach, which uses a series of different size meshes as preconditioners [62]. In fact the use of P_1 as a preconditioner

for and S_n calculation can be thought of as a multigrid-in-angle preconditioner.

5.4.2 Matrix Free

Another important property of the GMRES algorithm is that it does not require the matrix \mathbf{A} to be stored. All that is required is the action of \mathbf{A} on a vector. This is due to the way the Krylov subspace is created – by repeatedly applying the operator to a vector. This can reduce the amount of memory needed and can allow the method designer more freedom. This feature will be used when we develop our quasi-linear solver and is used by inexact Newton methods to avoid specifying an analytic Jacobian, but instead allowing its action to be computed by a cheap finite difference approximation.

5.5 Trilinos

The numerical solver library Trilinos from Sandia National Laboratories [63] was used to invoke the solver methods described in this chapter. While offering much more, we primarily used the Epetra package (a collection of matrix construct/storage classes), the AztecOO package (a linear solver package which has GMRES) and the NOX package which implements a Newton-Krylov method.

For the Newton-Krylov method the user can stipulate a convergence tolerance or maximum number Newton iterations. The convergence tolerance can be a norm of the residual or a norm of the relative change in the solution. The norm can be length scaled also if the user desires. Unfortunately, NOX does not have the functionality to have tolerances for different parts of the solution vector. This is a problem for our solution vectors which contain radiation intensities and temperatures. For a temperature of 300eV the value of I_0^0 is on the order of 10^{19} . In such a case we would like to converge the temperatures to much less than an eV and the moments of the intensity to about 10^7 . With NOX we must set the convergence criteria to be 10^7 for the entire vector.

In practice we have not seen evidence that this drawback of NOX has degraded our solutions.

Also, with NOX one can specify the line search method. That is, rather than taking a full Newton step, a shorter step can be taken if the full Newton step increases the residual [55]. One type of line search is the polynomial line search where the step length is reduced by fitting a polynomial to the residual gained by taking smaller steps and choosing an appropriate step length. The step length is only reduced if the full Newton step increases the residual. In most problems we have observed that the full Newton step does not increase the residual, however, we turn on the polynomial line search method as a cheap hedge.

In NOX we use a matrix-free approximate Jacobian, this allows us to avoid computing and storing the Jacobian for our system. Using this approximate Jacobian a linear system of equations is solved using GMRES. NOX also allows the user to specify an operator to precondition this system with. As a physics based preconditioner, we use the first order Riemann discretization with the temperatures assumed to be constant to precondition. This is a reasonable choice because for time steps that are small compared to the temperature time scale, this approximation would be valid. We do not compute the action of the first order matrix to precondition. Rather than computing the action of the inverse of the first order matrix – which would require a linear solve, an $ILU(\tau)$ factorization of the first order matrix is computed and this is used to approximate the inverse of the first order matrix. If \mathbf{A} is the first order matrix, instead of using \mathbf{A}^{-1} as the preconditioner we use the incomplete LU factorization, $(\mathbf{LU})^{-1}$, as the preconditioner. The incomplete LU inverse is very cheap to compute and much less expensive than the linear solver required to compute the action of the inverse of the first order matrix.

CHAPTER VI

Boundary Conditions

In transport theory boundary conditions are such that the *incoming* radiation must be specified at a boundary. Even in the fully continuous case for the P_n equations the boundary conditions represent an approximation to the transport boundary condition. The transport boundary conditions are given by an incoming intensity

$$I(\vec{x}, \hat{\Omega}) = B(\hat{\Omega}) \quad \text{for } \hat{n} \cdot \hat{\Omega} < 0, \quad (6.1)$$

where \hat{n} is the outward normal of the boundary. One might be tempted to just take spherical harmonic moments of the incoming intensity, B , and use these as the P_n boundary condition. However, B is only specified for incoming directions and the exiting intensity is not known. One approach to take is to take the spherical harmonic moments of B only over the incoming directions. This tack is called the Marshak boundary condition [10, 13, 16]. Another approach is to take the value of B along the characteristics of the P_n equations and use those values to enforce the boundary conditions. This is called the Mark boundary condition [10, 13, 16].

The Mark boundary condition allows the specification of an incoming intensity. However, often we may want a boundary that reflects particles instead of allowing them to pass. Such reflecting boundaries are well known for one-dimensional P_n methods, but

multiple dimensional reflecting boundary conditions are, at best, not widely disseminated. These multi-dimensional reflecting boundary conditions will be derived below.

6.1 Mark Boundary Condition

While it is possible to implement a Marshak Boundary condition in method that uses an upwind discretization, the partial fluxes that need to be specified at a boundary require a completely different treatment of the boundary cells [1]. The Mark boundary condition can be implemented within the framework of the Riemann solver without completely recharacterizing the flux between a boundary cell and the actual problem boundaries.

The Mark boundary condition specifies the incoming intensity along particular directions. The movement of information along characteristics is automatic with the Riemann solver. In fact, the Riemann solver was designed to move information upwind along the characteristics so implementing the Mark boundary condition is natural. To implement the Mark boundary conditions we use a technique known as ghost cells. This involves placing a distribution that has the desired incoming angular dependence “just outside” the physical system. It is not necessary to compute the appropriate incoming flux from the boundary cell (as is necessary in the Marshak boundary condition). This is because the Riemann solver already computes the flux across a cell face in the proper way. In the ghost cell boundary condition, only the information that is supposed to flow across a system boundary does. Since the impetus behind developing a Riemann solver for radiative transfer was the desire to have a method that moves information in an upwind manner, we need not make an exception at the boundary.

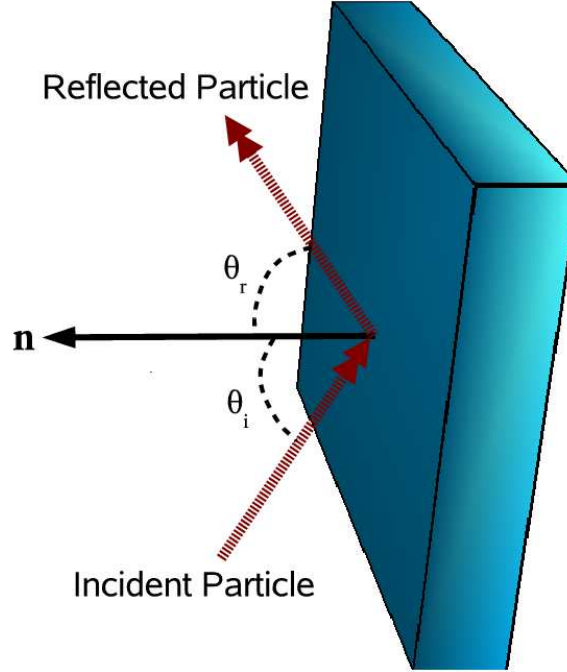


Figure 6.1: The reflected angle equals the incident particle angle for a reflecting boundary.

6.1.1 Vacuum boundary condition

The vacuum boundary condition can be easily implemented with ghost cells. For example, in a cell at the left edge of the system we set the value of $\vec{I}_{\text{ghost}} = \vec{0}$ (\vec{I}_{ghost} is the value in the ghost cell). Then the flux across the boundary is (c.f. Eq. ??)

$$F_{\text{boundary}} = \frac{1}{2} \left(n_x \mathcal{A}_x \left(\vec{I}_{\text{interior}} \right) + n_y \mathcal{A}_y \left(\vec{I}_{\text{interior}} \right) + n_z \mathcal{A}_z \left(\vec{I}_{\text{interior}} \right) \right) - \frac{1}{2} \left(n_x |\mathcal{A}_x| \left(\vec{I}_{\text{interior}} \right) + n_y |\mathcal{A}_y| \left(\vec{I}_{\text{interior}} \right) + n_z |\mathcal{A}_z| \left(\vec{I}_{\text{interior}} \right) \right). \quad (6.2)$$

Vacuum boundaries for other locations are similarly treated.

6.2 Reflecting Boundary Condition

The reflecting boundary condition causes the system boundary to reflect particles impinging on the boundary interface. In Fig. 6.1 this process is illustrated. The particles hitting the boundary are to be reflected with the same angle relative to the outward

normal to the interface. The relation between θ_i and θ_r , taking into account direction, is $\theta_r = \pi - \theta_i$. For an incident particle with direction $\hat{\Omega}_i$, the angle between the outward normal and $\hat{\Omega}_i$ is given by $\hat{\Omega}_i \cdot \hat{n} = \cos \theta_r$ and the reflected angle is then given by $\hat{\Omega}_r \cdot \hat{n} = \cos(\pi - \theta_i) = -\cos \theta_i$. Therefore the intensity at the boundary is symmetric in the cosine of the polar angle with respect to \hat{n}

$$I(\cos \theta_i)_{\cos \theta_i > 0} = I(-\cos \theta_i)_{\cos \theta_i < 0}. \quad (6.3)$$

The incident direction of a particle in terms of Cartesian coordinates (x, y, z) is

$$\hat{\Omega}'_i = \left(\sqrt{1 - \mu^2} \cos \phi, \sqrt{1 - \mu^2} \sin \phi, \mu \hat{k} \right) \quad (6.4)$$

The reflecting boundary condition for general inward normal is a linear combination of the boundary conditions for each of the x , y , and z directions. We will now derive the boundary conditions for each of these directions.

Z direction

If the outward normal, \hat{n} , equals \hat{k} the boundary condition becomes

$$I(\mu, \phi)_{\mu > 0} = I(-\mu, \phi), \quad (6.5)$$

because $\hat{\Omega}_i \cdot \hat{n} = \mu$. The intensity at the boundary is an even function of μ . Therefore, the moments of this intensity are given by

$$\vec{I}_{l \text{ boundary}}^m = \int_0^{2\pi} d\phi \int_{-1}^1 d\mu Y_l^m \bar{I}(\mu, \phi) = \begin{cases} \vec{I}_l^m & l + m \text{ even} \\ 0 & l + m \text{ odd} \end{cases}. \quad (6.6)$$

This is because the associated Legendre functions are odd in μ for $l + m$ odd and they are even if $l + m$ is even.

To impose this boundary condition using ghost cells we will set the boundary flux to be

$$\vec{F}_{l \text{ Boundary}}^m = \frac{1}{2} \left(\vec{I}_{l \text{ interior}}^m + \vec{I}_{l \text{ ghost}}^m \right), \quad (6.7)$$

with

$$\vec{I}_{l\text{ghost}}^m = \begin{cases} \vec{I}_{l\text{interior}}^m & l + m \text{ even} \\ -\vec{I}_{l\text{interior}}^m & l + m \text{ odd} \end{cases}. \quad (6.8)$$

Note that this is equivalent to making our scheme a centered-difference (and not up-winded) scheme at the boundary.

X direction

For $\hat{n} = \hat{i}$, the incident angle is given by

$$\cos \theta_i = \sqrt{1 - \mu^2} \cos \phi. \quad (6.9)$$

This makes the outgoing (exiting the system) direction $\cos \phi > 0$ or $\phi \in [-\pi/2, \pi/2]$, and the reflected direction $\cos \phi < 0$ or $\phi \in [\pi/2, 3\pi/2]$. The boundary condition is

$$I(\mu, \phi)_{\phi \in [-\pi/2, \pi/2]} = I(\mu, \pi - \phi). \quad (6.10)$$

To determine what the value of I_l^m should be at the boundary let us examine the integrals that define the moment

$$\begin{aligned} I_l^m &= \int_{-1}^1 d\mu \int_0^{2\pi} d\phi \bar{Y}_l^m I(\mu, \phi) \\ &= \int_{-1}^1 d\mu \int_{-\pi/2}^{\pi/2} d\phi \bar{Y}_l^m I(\mu, \phi) + \int_{-1}^1 d\mu \int_{\pi/2}^{3\pi/2} d\phi \bar{Y}_l^m I(\mu, \pi - \phi). \end{aligned} \quad (6.11)$$

The real part of the second integral can be simplified using

$$\int_{\pi/2}^{3\pi/2} d\phi K_l^m P_l^m(\mu) \cos(m\phi) I(\mu, \pi - \phi) = \int_{-\pi/2}^{\pi/2} d\phi K_l^m P_l^m(\mu) \cos(m(\pi - \phi)) I(\mu, \phi), \quad (6.12)$$

where K_l^m is the normalization constant, and the identity

$$\cos(m(\pi - \phi)) = \begin{cases} \cos(m\phi) & m \text{ even} \\ -\cos(m\phi) & m \text{ odd} \end{cases}, \quad (6.13)$$

to get

$$\Re(I_l^m) = 0 \quad \text{for } m \text{ odd} \quad (6.14)$$

The imaginary part is handled by examining

$$-\int_{\pi/2}^{3\pi/2} d\phi K_l^m P_l^m(\mu) \sin(m\phi) I(\mu, \pi - \phi) = -\int_{-\pi/2}^{\pi/2} d\phi K_l^m P_l^m(\mu) \sin(m(\pi - \phi)) I(\mu, \phi), \quad (6.15)$$

and the identity

$$\sin(m(\pi - \phi)) = \begin{cases} -\sin(m\phi) & m \text{ even} \\ \sin(m\phi) & m \text{ odd} \end{cases}, \quad (6.16)$$

giving

$$\Im(I_l^m) = 0 \quad \text{for } m \text{ even.} \quad (6.17)$$

The conditions of Eqs. 6.14 and 6.17 are satisfied using ghost cells by setting the flux at the boundary to

$$\vec{F}_{l \text{ Boundary}}^m = \frac{1}{2} \left(\vec{I}_{l \text{ interior}}^m + \vec{I}_{l \text{ ghost}}^m \right), \quad (6.18)$$

with

$$\vec{I}_{l \text{ ghost}}^m = \begin{cases} \vec{I}_{l \text{ interior}}^m & m \text{ even} \\ -\vec{I}_{l \text{ interior}}^m & m \text{ odd} \end{cases}. \quad (6.19)$$

Y direction

For $\hat{n} = \hat{j}$, the incident angle is given by

$$\cos \theta_i = \sqrt{1 - \mu^2} \sin \phi. \quad (6.20)$$

This makes the outgoing direction $\sin \phi > 0$ or $\phi \in [0, \pi]$, and the incoming direction $\sin \phi < 0$ or $\phi \in [\pi, 2\pi]$. The boundary condition is

$$I(\mu, \phi)_{\phi \in [0, \pi]} = I(\mu, \pi - \phi). \quad (6.21)$$

To determine what the value of I_l^m should be at the boundary let us examine the integrals that define the moment

$$\begin{aligned} I_l^m &= \int_{-1}^1 d\mu \int_0^{2\pi} d\phi \bar{Y}_l^m I(\mu, \phi) \\ &= \int_{-1}^1 d\mu \int_0^\pi d\phi \bar{Y}_l^m I(\mu, \phi) + \int_{-1}^1 d\mu \int_\pi^{2\pi} d\phi \bar{Y}_l^m I(\mu, \pi - \phi). \end{aligned} \quad (6.22)$$

The real part of the second integral can be simplified using

$$\int_\pi^{2\pi} d\phi K_l^m P_l^m(\mu) \cos(m\phi) I(\mu, \pi - \phi) = \int_0^{2\pi} d\phi K_l^m P_l^m(\mu) \cos(m(2\pi - \phi)) I(\mu, \phi), \quad (6.23)$$

and the identity

$$\cos(m(2\pi - \phi)) = \cos(m\phi), \quad (6.24)$$

which says that the real part of \vec{I}_l^m should be the same in the ghost cell. The imaginary part is handled by examining

$$- \int_\pi^{2\pi} d\phi K_l^m P_l^m(\mu) \sin(m\phi) I(\mu, \pi - \phi) = - \int_0^\pi d\phi K_l^m P_l^m(\mu) \sin(m(\pi - \phi)) I(\mu, \phi), \quad (6.25)$$

and the identity

$$\sin(m(2\pi - \phi)) = -\sin(m\phi), \quad (6.26)$$

giving

$$\Im(I_l^m) = 0 \quad \text{for all } m. \quad (6.27)$$

Condition 6.27 is enforced using ghost cells through setting

$$\vec{F}_{l \text{ Boundary}}^m = \frac{1}{2} \left(\vec{I}_{l \text{ interior}}^m + \vec{I}_{l \text{ ghost}}^m \right), \quad (6.28)$$

with

$$\vec{I}_{l \text{ ghost}}^m = \Re(\vec{I}_{l \text{ interior}}^m). \quad (6.29)$$

6.2.1 Albedo Boundary Condition

An albedo boundary condition is used to model the partial return of radiation energy that has escaped the system. Specifically, the returned radiation will have lost all the angular information that the exiting radiation had. For example, if one were modeling the radiation field inside a hohlraum for inertial confinement fusion, the gold walls of the hohlraum might be treated as an albedo boundary, i.e. the walls isotropically return some fraction of the radiation entering the walls.

For a P_n boundary condition, the albedo can be treated as a reflecting boundary in the I_0^0 moment multiplied by the albedo value, α . The other moments have a reflected value of zero.

$$\vec{I}_{l\text{ghost}}^m = \begin{cases} \alpha \vec{I}_{0\text{interior}}^0 & l = 0, m = 0 \\ 0 & \text{otherwise} \end{cases} . \quad (6.30)$$

CHAPTER VII

Limits of the P_n Equations

The behavior of the transport equation has some important limits that we would like our numerical method to respect. One of these limits is the diffusion limit, where scattering (or absorption/ reëmission¹) dominate. In this limit we would like to have our numerical method be robust when a diffusion length is resolved, rather than a mean free path. The other extreme is the free-streaming limit, which presents a problem for P_n methods in multiple dimensions. Below we will present a novel analysis that the standard Riemann solver does not limit to a discrete diffusion equation. Beyond this we will demonstrate that the P_n equations give *analytic* solutions with negative energy densities and we point to some possibilities to address these issues.

7.1 The Diffusion Limit

In this section we examine how Riemann solvers for the P_n equations behave in problems of a “diffusive” nature, that is, in problems where scattering interactions dominate. At steady-state these problems are described by an elliptic diffusion equation, which is an approximation to the full transport equation in the limit of asymptotically

¹The use of the diæresis over the second e in reëmission is technically correct. This symbol denotes that the vowels are part of two different syllables and should not be pronounced as a diphthong [64]. The *New Yorker* is particularly fond of using a diæresis in words where it regularly omitted, such as coöperate.

small absorption and small sources. In linear radiation transport the steady-state solution to diffusive problems varies on a length scale characterized by the diffusion length $L = 1/\sqrt{3\sigma_t\sigma_a}$ which is much larger than a mean-free-path $1/\sigma_t$ when $\sigma_a \ll \sigma_t$. As is well known, time dependent diffusion is governed by a parabolic equation, and unlike hyperbolic problems, parabolic equations are characterized by infinite propagation speed of information. They are also characterized by dissipation: they produce smooth solutions even from non-smooth sources and initial data.

This presents a challenge: Riemann solvers are specifically designed to solve hyperbolic equations and to upwind information flow at finite speed in a minimally dissipative way. Riemann solvers will give good results for cells that are on the order of $1/\sigma_t$ or smaller, but to resolve the solution in a diffusive problem a cell on the order of L should be sufficient. For robustness we would hope that a Riemann solver would capture this diffusive solution on a spatial grid that only resolves the diffusion length L . A method that has this robustness property is said to have a diffusion limit.

We will show that the standard Riemann solver for the P_n equations does *not* have a diffusion limit, and because of this a mean-free path must be resolved to obtain a good steady solution even in diffusive problems. An asymptotic analysis of the problem reveals that the main impediment to a diffusion limit in a Riemann solver is the numerical dissipation present in the scheme. This grid-size-dependent dissipation is present in a Riemann solver to make it upwinded and thus maintain the proper physics of finite propagation speed, and hence stability. However, the collisional interactions that in transport physics dominate the diffusion limit represent another type of dissipation. The number of particles travelling in a given direction is reduced by collisions. We propose a method that recovers the diffusion limit by systematically scaling out the numerical dissipation in regions where scattering collisions dominate. The result is a

hybrid method that is a Riemann solver when the grid resolves a mean-free-path, but which becomes a centered difference method when the grid is large.

7.1.1 Diffusion properties of the P_n equations

Since we study the spherical harmonic equations, we wish to understand their properties in the diffusion limit.

Let us begin by exploring a diffusion-like solution of the transport equation.² In one dimension with isotropic scattering, the grey transport equation is

$$\frac{1}{c} \frac{\partial I}{\partial t} + \mu \frac{\partial I}{\partial x} + \sigma_t I = \frac{\sigma_s}{2} \int_{-1}^1 I d\mu' + \frac{Q}{4\pi}. \quad (7.1)$$

Suppose that we have a linear-in-space source $Q = qx$ for some constant q , and we seek a solution of the transport equation that is linear in space x and linear in the direction cosine μ . Substituting $I(x, \mu) = ax + b\mu$ into Eq. (7.1) we find

$$\mu a + \sigma_t(ax + b\mu) = x\sigma_s a + x \frac{q}{4\pi}. \quad (7.2)$$

From this we find $a = q/(4\pi\sigma_a)$ and $b = -a/\sigma_t$, and thus

$$I(x, \mu) = \frac{q}{4\pi\sigma_a} \left[x - \frac{\mu}{\sigma_t} \right] \quad (7.3)$$

is a special solution of the one-speed transport equation with constant cross-sections and a linear-in-space isotropic source. This solution is not strictly positive, however this is a result of a having negative source for $x < 0$.

This solution is interesting because it satisfies Fick's law: in particular, it has

$$\phi = 2\pi \int_{-1}^1 I d\mu = \frac{q}{\sigma_a} x \quad \text{a linear in space scalar flux, and} \quad (7.4)$$

$$J = 2\pi \int_{-1}^1 \mu I d\mu = -\frac{q}{3\sigma_a\sigma_t} = -\frac{1}{3\sigma_t} \frac{\partial \phi}{\partial x} \quad \text{a constant current satisfying Fick's law.} \quad (7.5)$$

²This is a simple exercise in the method of manufactured solutions, a technique first suggested by Lingus [65]

Because of this property it has been suggested that a numerical scheme for the transport equation should capture this solution [66, 67] at least in the diffusion limit. A method that can produce these linear solutions is said to be linear solution preserving (LSP). We will see later that the Riemann solver for the P_n equations is LSP, and so is the modified Riemann solver that we propose in this paper.

7.1.2 Linear solution of the P_n equations

We should like to know if the P_n equations capture the linear-in-space, linear-in-direction cosine solution just described in Eq. (7.3). The 1D P_n equations are

$$\begin{aligned} \frac{1}{c} \frac{\partial I_0}{\partial t} + \frac{\partial}{\partial x} (B_1 I_1) + \Sigma_a I_0 &= \frac{Q}{2\sqrt{\pi}} \\ \frac{1}{c} \frac{\partial I_l}{\partial t} + \frac{\partial}{\partial x} (A_{l-1} I_{l-1} + B_{l+1} I_{l+1}) + \Sigma_t I_l &= 0 \quad l = 1 \dots n \quad (7.6) \\ I_{n+1} &= 0. \end{aligned}$$

For that linear source problem, we can immediately compute that

$$I_0 = \frac{qx}{2\sigma_a\sqrt{\pi}} \quad (7.7)$$

$$I_1 = -\frac{q}{2\sigma_a\sigma_t\sqrt{3\pi}} \quad (7.8)$$

$$I_l = 0 \quad l > 1. \quad (7.9)$$

Noting that $A_0 = B_1 = 1/\sqrt{3}$, and substituting these into Eqs. (7.6) along with $Q = qx$ we discover that they do indeed satisfy the P_n equations. Therefore the P_n equations have the special solution that represents Fick's law.

In order for a numerical scheme to capture this solution it is necessary only that the derivative in the zeroth moment equation map a constant to zero (so it must be at least zeroth order accurate) and the derivative in the first moment equation must map a linear function to a constant of the correct value (so it must be at least first order accurate). These not very demanding conditions will guarantee that a method for the

P_n equations will capture the linear-in-space and linear-in-direction cosine solution of the P_n equations. We will note below that the Riemann solver for the P_n equations does have these properties.

We also note in passing that because of the rotational invariance of the spherical harmonic equations, they will in fact have an entire family of linear-in-space and linear-in-direction cosine solutions resulting from rotations of the special x -dependent solution just displayed.

7.1.3 Asymptotic analysis of the P_n equations

To investigate the diffusive limit of the P_n equations, and numerical methods for them, we want to examine their solution when scattering dominates over absorption namely, $\sigma_t \gg \sigma_a$, and when time variation is negligible, $\partial I_l / \partial t \approx 0$. To do this we divide σ_t by a small, positive parameter ϵ and multiply σ_a , Q and $\partial / \partial t$ by ϵ as well, resulting in

$$\begin{aligned} \frac{\epsilon}{c} \frac{\partial I_0}{\partial t} + \frac{\partial}{\partial x} (B_1 I_1) + \epsilon \sigma_a I_0 &= \epsilon \frac{Q}{2\sqrt{\pi}} \\ \frac{\epsilon}{c} \frac{\partial I_l}{\partial t} + \frac{\partial}{\partial x} (A_{l-1} I_{l-1} + B_{l+1} I_{l+1}) + \frac{\sigma_t}{\epsilon} I_l &= 0 \quad l = 1 \dots n \\ I_{n+1} &= 0. \end{aligned} \quad (7.10)$$

We also then postulate a asymptotic expansion for I_l given by

$$I_l \sim \sum_{j=0}^{\infty} \epsilon^j I_l^{(j)}(x, t), \quad \epsilon \rightarrow 0. \quad (7.11)$$

Next, we present a theorem on the asymptotic behavior of the P_n equations that we will want to recapture with our modified Riemann solver.

Theorem 4. *Let $\sigma_t > 0$. Then for the asymptotic expansion Eq. (7.11) to satisfy the scaled P_n equations, Eq. (7.10), we must have $I_l^{(j)} = 0$ for $l > j$. In other words, $I_l = O(\epsilon^l)$.*

Furthermore, the solution satisfies Fick's law at leading order,

$$I_1^{(1)} = -A_0 \frac{\partial I_0^{(0)}}{\partial x} \quad (7.12)$$

and

$$\frac{1}{c} \frac{\partial I_0^{(0)}}{\partial t} - \frac{\partial}{\partial x} \frac{A_0 B_1}{\sigma_t} \frac{\partial I_0^{(0)}}{\partial x} + \sigma_a I_0^{(0)} = \frac{Q}{2\sqrt{\pi}} \quad (7.13)$$

Proof by induction. Substituting Eq. (7.11) into Eq. (7.10) yields

$$\begin{aligned} \sum_{j=1}^{\infty} \epsilon^j \frac{1}{c} \frac{\partial I_0^{(j)}}{\partial t} + \sum_{j=0}^{\infty} \epsilon^j \frac{\partial}{\partial x} \left(B_1 I_1^{(j)} \right) + \sum_{j=1}^{\infty} \epsilon^j \sigma_a I_0^{(j-1)} &= \epsilon \frac{Q}{2\sqrt{\pi}} \\ \sum_{j=1}^{\infty} \epsilon^j \frac{1}{c} \frac{\partial I_l^{(j)}}{\partial t} + \sum_{j=0}^{\infty} \epsilon^j \frac{\partial}{\partial x} \left(A_{l-1} I_{l-1}^{(j)} + B_{l+1} I_{l+1}^{(j)} \right) + \sum_{j=-1}^{\infty} \epsilon^j \sigma_t I_l^{(j+1)} &= 0 \quad l = 1 \dots n \\ \sum_{j=0}^{\infty} \epsilon^j I_{n+1}^{(j)} &= 0. \end{aligned}$$

Gathering terms of order ϵ^{-1} yields $I_l^{(0)} = 0$ for $l > 0$. Gathering terms of order ϵ^0 produces

$$\frac{\partial}{\partial x} \left(B_1 I_1^{(0)} \right) = 0 \quad (7.14)$$

$$\frac{\partial}{\partial x} \left(A_{l-1} I_{l-1}^{(0)} + B_{l+1} I_{l+1}^{(0)} \right) + \sigma_t I_l^{(1)} = 0 \quad l = 1 \dots n \quad (7.15)$$

$$I_{n+1}^{(0)} = 0. \quad (7.16)$$

Since $I_1^{(0)} = 0$, Eq. (7.14) is automatically satisfied. Eq. (7.15) for $l = 1$ implies $A_0 \partial I_0^{(0)} / \partial x + \sigma_t I_1^{(1)} = 0$, which is Fick's law. Also since $I_l^{(0)} = 0$ for $l > 0$, Eq. (7.15) for $l > 1$ implies $I_l^{(1)} = 0$ for $l > 1$.

Repeating the exercise for terms of order ϵ^1 we will discover that

$$\frac{1}{c} \frac{\partial I_0^{(0)}}{\partial t} + \frac{\partial}{\partial x} \left(B_1 I_1^{(1)} \right) + \sigma_a I_0^{(0)} = \frac{Q}{2\sqrt{\pi}} \quad (7.17)$$

$$\frac{1}{c} \frac{\partial I_l^{(0)}}{\partial t} + \frac{\partial}{\partial x} \left(A_{l-1} I_{l-1}^{(1)} + B_{l+1} I_{l+1}^{(1)} \right) + \sigma_t I_l^{(2)} = 0, \quad l = 1 \dots n \quad (7.18)$$

$$I_{n+1}^{(1)} = 0. \quad (7.19)$$

Equation (7.17) combines with Fick's law to give us the time-dependent diffusion equation for $I_0^{(0)}$. From Eq. (7.18) with $l > 2$, and using $I_l^{(1)} = 0$ for $l > 1$, we see that $I_l^{(2)} = 0$ for $l > 2$. Thus we begin to build up an induction on the order j to show that $I_l^{(j)} = 0$ for $l > j$.

Suppose that there is a value j such that for all values $j' \leq j$ we know $I_l^{(j')} = 0$ for $l > j'$. We need only consider $j > 2$ since we have already established this for $j \leq 2$.

Gathering terms of order ϵ^j we have

$$\frac{1}{c} \frac{\partial I_0^{(j-1)}}{\partial t} + \frac{\partial}{\partial x} (B_1 I_1^{(j)}) + \sigma_a I_0^{(j-1)} = 0 \quad (7.20)$$

$$\frac{1}{c} \frac{\partial I_l^{(j-1)}}{\partial t} + \frac{\partial}{\partial x} (A_{l-1} I_{l-1}^{(j)} + B_{l+1} I_{l+1}^{(j)}) + \sigma_t I_l^{(j+1)} = 0, \quad l = 1 \dots n \quad (7.21)$$

$$I_{n+1}^{(j)} = 0. \quad (7.22)$$

Using $I_l^{(j-1)} = 0$ for $l > j - 1$ and $I_{l-1}^{(j)} = 0$ for $l - 1 > j$, and hence $I_l^{(j)} = 0$ and $I_{l+1}^{(j)} = 0$ for $l - 1 > j$, with Eq. (7.21) for $l > j + 1$ therefore implies that $\sigma_t I_l^{(j+1)} = 0$, and hence $I_l^{(j+1)} = 0$ for $l > j + 1$. This completes the induction. □

Note that $A_0 B_1 = 1/3$, and also recalling that $\phi = 2\sqrt{\pi} I_0$ we see from Eq. (7.13) that the scalar flux ϕ satisfies the time dependent diffusion equation to leading order in ϵ . This proof shows that the diffusion limit of the P_n equations is connected to the angular moments of order l being of order ϵ^l , and in the correct Fick's law arising at first order in the expansion.

7.1.4 Diffusion properties of the Riemann discretization

We now want to explore the Riemann solver discretization in the thick diffusion limit. To begin, we write the first-order in space Riemann discretization of the P_n equations

in one-dimension with a uniform grid

$$\frac{1}{c} \frac{\partial}{\partial t} \vec{I}_i + \mathbf{A} \frac{(\vec{I}_{i+1} - \vec{I}_{i-1}))}{2\Delta x} - |\mathbf{\Lambda}| (\vec{I}_{i+1} - 2\vec{I}_i + \vec{I}_{i-1}) = -\mathbf{S}\vec{I}_i + \vec{Q}_i(t). \quad (7.23)$$

Here \vec{I}_i is a vector of angular moments ascending through l averaged over spatial cell i , and

$$|\mathbf{\Lambda}| = \sum_{k=0}^n \vec{r}_k |\lambda_k| \vec{l}_k. \quad (7.24)$$

We first consider the linear-in-space and linear-in-direction cosine solution, which in the discrete form on a uniform mesh of spacing Δx should be

$$I_{0,i} = \frac{q i \Delta x}{2\sigma_a \sqrt{\pi}} \quad (7.25)$$

$$I_{1,i} = -\frac{q}{2\sigma_a \sigma_t \sqrt{3\pi}} \quad (7.26)$$

$$I_{l,i} = 0 \quad l > 1. \quad (7.27)$$

We wish to see if this solution satisfies the discrete P_n equations (c.f. Eq. (7.23)). For this solution we immediately have $(I_{0,i+1} - I_{0,i-1})/2\Delta x = q/(2\sigma_a \sqrt{\pi})$, $(I_{l,i+1} - I_{l,i-1})/2\Delta x = 0$ for $l > 0$, and $\vec{I}_{i+1} - 2\vec{I}_i + \vec{I}_{i-1} = 0$. With these observations it is easy to conclude that this linear-in-space and linear-in-direction cosine solution is an exact solution of the Riemann discretized P_n equations. This really was inevitable from the first order accuracy of the discretization.

However, even though the Riemann discretized P_n equations have this exact diffusion-like solution (which exactly satisfies Fick's law, you will recall), it does not have a good diffusion limit. Introducing the same scaling in ϵ as for Theorem 1 we write the Rie-

mann discretized equations as

$$\begin{aligned} & \frac{\epsilon d\vec{I}_i}{c dt} + \mathbf{A} \frac{(\vec{I}_{i+1} - \vec{I}_{i-1})}{2\Delta x} - |\mathbf{\Lambda}| \frac{(\vec{I}_{i+1} - 2\vec{I}_i + \vec{I}_{i-1})}{2\Delta x} \\ &= - \begin{pmatrix} \epsilon\sigma_a & 0 & 0 & \dots \\ 0 & \sigma_t/\epsilon & 0 & \dots \\ 0 & 0 & \sigma_t/\epsilon & \\ \vdots & \vdots & & \ddots \end{pmatrix} \vec{I} + \begin{pmatrix} \epsilon Q/(2\sqrt{\pi}) \\ 0 \\ 0 \\ \vdots \end{pmatrix}, \end{aligned} \quad (7.28)$$

We once again use the asymptotic expansion $I_{l,i} \sim \sum_{j=0}^{\infty} \epsilon^j I_{l,i}^{(j)}(t)$ and have the following unfortunate theorem which says the Riemann solver has a poor diffusion limit.

Theorem 5. *Using $I_{l,i} \sim \sum_{j=0}^{\infty} \epsilon^j I_{l,i}^{(j)}(t)$ in Eq. (7.28) with $\sigma_t > 0$ we must have $I_1^{(0)} = 0$ and $I_{0,i+1}^{(0)} - 2I_{0,i}^{(0)} + I_{0,i-1}^{(0)} = 0$, so $I_{0,i}^{(0)}$ does not satisfy a discrete diffusion equation.*

Proof. Considering first terms of order ϵ^{-1} , we get contributions only from the right hand side and when $\sigma_t \neq 0$; these terms imply $I_{l,i}^{(0)} = 0$, and in particular $I_1^{(0)} = 0$, yielding the first claim of the theorem.

Moving on to terms of order ϵ^0 we have

$$\mathbf{A} \frac{\vec{I}_{i+1}^{(0)} - \vec{I}_{i-1}^{(0)}}{2\Delta x} - |\mathbf{\Lambda}| \frac{(\vec{I}_{i+1}^{(0)} - 2\vec{I}_i^{(0)} + \vec{I}_{i-1}^{(0)})}{2\Delta x} = \begin{pmatrix} 0 & 0 & 0 & \dots \\ 0 & \sigma_t & 0 & \dots \\ 0 & 0 & \sigma_t & \\ \vdots & \vdots & & \ddots \end{pmatrix} \vec{I}_i^{(1)}. \quad (7.29)$$

This appears to have many terms, but in fact we already know that only $I_{0,i}^{(0)}$ is non-zero, and all other $I_{l,i}^{(0)} = 0$ for $l > 0$, and so most of the terms on the left are zero. For the first equation, the I_0 equation, the first term on the LHS of Eq. (7.29) only contains I_1 terms hence, under our scaling this first equation becomes

$$|\mathbf{\Lambda}|_{0,0} \frac{(I_{0,i+1}^{(0)} - 2I_{0,i}^{(0)} + I_{0,i-1}^{(0)})}{2\Delta x} = 0 \quad (7.30)$$

where $|\Lambda|_{0,0}$ denotes the first row, first column of $|\Lambda|$, which corresponds to $l = 0$, explaining the zero indices. So, in order to establish the theorem we need only show that $|\Lambda|_{0,0} \neq 0$.

Going back to Eq. (7.24), $|\Lambda| = \sum_{k=0}^n \vec{r}_k |\lambda_k| \vec{l}_k$, Brunner & Holloway [2, 36] have perviously derived the eigenvectors and eigenvalues of \mathbf{A} , and from these results one can construct $|\Lambda|_{0,0}$ and see that it is non-zero. Alternately, noting that $B_{l+1} = A_l$ we see that \mathbf{A} is symmetric, hence $\vec{r}_k = \vec{l}_k$. Every term in the sum for $|\Lambda|$ is therefore non-negative, and if the first element of \vec{r}_k corresponding to a non-zero eigenvalue λ_k is non-zero, then $|\Lambda|_{0,0} > 0$. This is easy to discover from the structure of \mathbf{A} in 1D, which has a zero diagonal and non-zeros on the first super and sub-diagonals. If the first element of an eigenvector is zero then, for a non-zero eigenvalue, the second element is zero. And if the first and second elements are zero, the third must be, and so on down the line. Hence the first element of the eigenvector cannot be zero, so $|\Lambda|_{0,0} > 0$ and the theorem is proved. \square

So, at first order in ϵ we discover the equation for the leading order scalar flux $I_{0,i}$ is

$$I_{0,i+1} - 2I_{0,i} + I_{0,i-1} = 0, \quad (7.31)$$

hence the *Götterdämmerung*³ of the standard Riemann solver in the diffusion limit. This equality tells us that the leading order terms will not satisfy the correct diffusion equation. We see this be noting that Eq. (7.31) is (within a constant factor) a finite-difference Laplacian. This Laplacian being zero tells us the leading order terms are linear in space and satisfy an erroneous diffusion equation $\nabla^2 I_0^{(0)} = 0$.

³Properly translated into English, *Götterdämmerung*, means “twilight of the gods” and denotes the turbulent and complete downfall of a regime or institution. The word is the mistranslation into German of the Old Norse *ragnarok* (which means “fate of the gods”) and its most famous usage is by Richard Wagner as the title for the finale of *The Ring of the Nibelung*.

7.1.5 Intermediate Diffusion Limit

The intermediate diffusion limit has the cell thickness Δx equal to about one mean free path. In the notation from the above analysis this means $h \rightarrow \epsilon h$, which gives the scaling

$$\frac{\epsilon}{c} \frac{d\vec{I}_i}{dt} + \mathbf{A} \frac{(\vec{I}_{i+1} - \vec{I}_{i-1})}{2\epsilon h} - \frac{1}{2\epsilon h} |\mathbf{\Lambda}| (\vec{I}_{i+1} - 2\vec{I}_i + \vec{I}_{i-1}) = - \begin{pmatrix} \epsilon\sigma_a & 0 & 0 & \dots \\ 0 & \sigma_t/\epsilon & 0 & \dots \\ 0 & 0 & \sigma_t/\epsilon & \\ \vdots & \vdots & & \ddots \end{pmatrix} \vec{I} + \begin{pmatrix} \epsilon Q/(2\sqrt{\pi}) \\ 0 \\ 0 \\ \vdots \end{pmatrix}$$

We then proceed by postulating a smooth solution to the P_n equation such that

$$I_{l,i}(t, \epsilon) = g_l(x_i, t, \epsilon, h). \quad (7.32)$$

We then expand $g_l(x, t, \epsilon)$ as a Taylor series in $h\epsilon$ (hiding the dependence on time) as

$$I_{l,i+1} = g(x_i, \epsilon) + h\epsilon \frac{dg_l}{dx}(x_i, \epsilon) + \frac{h^2\epsilon^2}{2} \frac{d^2g_l}{dx^2}(x_i, \epsilon). \quad (7.33)$$

Inserting Eq. (7.33) into Eq. (7.1.5) we get

$$\frac{\epsilon}{c} \frac{d\vec{g}_i}{dt} + \mathbf{A} \frac{d\vec{g}_i}{dx} - \frac{1}{2\epsilon h} |\mathbf{\Lambda}| \left(\epsilon^2 h^2 \frac{d^2\vec{g}_i}{dx^2} \right) = - \begin{pmatrix} \epsilon\sigma_a & 0 & 0 & \dots \\ 0 & \sigma_t/\epsilon & 0 & \dots \\ 0 & 0 & \sigma_t/\epsilon & \\ \vdots & \vdots & & \ddots \end{pmatrix} \vec{I} + \begin{pmatrix} \epsilon Q/(2\sqrt{\pi}) \\ 0 \\ 0 \\ \vdots \end{pmatrix}, \quad (7.34)$$

We now inspect the equations in the system given by Eq. (7.34) in their long form which is

$$\frac{\epsilon}{c} \frac{dg_{0,i}}{dt} + B_1 \frac{dg_{1,i}}{dx} - \frac{\epsilon h}{2} |\mathbf{\Lambda}_0| \left(\frac{d^2\vec{g}_i}{dx^2} \right) = -\epsilon\sigma_a I_{0,i} + \epsilon \frac{Q}{2\sqrt{\pi}}, \quad (7.35)$$

$$\frac{\epsilon}{c} \frac{dg_{1,i}}{dt} + A_{l-1} \frac{dg_{l-1,i}}{dx} + B_2 \frac{dg_{0,i}}{dx} - \frac{\epsilon h}{2} |\mathbf{\Lambda}_l| \left(\frac{d^2\vec{g}_i}{dx^2} \right) = -\frac{1}{\epsilon} \sigma_t I_{l,i}, \quad l > 0 \quad (7.36)$$

$$I_{n+1} = 0, \quad (7.37)$$

where we have denoted the k^{th} row of $|\Lambda|$ by $|\Lambda_l|$. Expanding $I_{l,i}$ into a power series as before we get for the $O(1/\epsilon)$ terms

$$I_{l,i}^{(0)} = 0, \quad l > 0 \quad (7.38)$$

and from the $O(1)$ equations

$$\frac{\partial}{\partial x} I_{1,i}^{(0)} = 0 \quad (7.39)$$

$$A_{l-1} \frac{dg_{l-1,i}^{(0)}}{dx} + B_{l+1} \frac{dg_{l+1,i}^{(0)}}{dx} = -\sigma_t I_{l,i}^{(1)}, \quad l > 0 \quad (7.40)$$

For Eq. (7.40) we can simplify further by using the fact that since $I_{l,i}^{(0)}$ is zero for $l > 0$, its derivatives will also be zero. This gives us the result

$$\frac{-A_0}{\sigma_t} \frac{dg_{0,i}^{(0)}}{dx} = I_{l,i}^{(1)}, \quad (7.41)$$

$$I_{l,i}^{(1)} = 0, \quad l > 1. \quad (7.42)$$

Now we look at the $O(\epsilon)$ terms in the first equation in Eq. (7.34). These are

$$\frac{1}{c} \frac{dg_{o,i}^{(0)}}{dt} - \left(\frac{A_0 B_1}{\sigma_t} + \frac{h}{2} |\Lambda_{0,0}| \right) \frac{d^2 g_{0,i}^{(0)}}{dx^2} + \sigma_a I_{0,i}^{(0)} = \frac{Q}{2\sqrt{\pi}} + \frac{h}{2} \sum_{l=1}^N |\Lambda_l| \left(\frac{d^2 g_{l,i}}{dx^2} \right). \quad (7.43)$$

This equation is “diffusion-like” insofar as it looks like the diffusion equation with an incorrect diffusion coefficient and a different source term. This result is better than that of the thick diffusion limit. Eq. (7.43) does include material interaction terms and the time derivative term. However, it still is not the correct diffusion equation.

7.1.6 Modified Riemann solver in the diffusive limit

Riemann solvers were designed to add just the right amount of dissipation to make the advective terms of a problem upwinded and stable. They treat an idealized problem (one in which there are no sources or sinks) exactly and use the solution to this problem

to determine the amount of flow across a cell interface. In problems where the advection of information dominates this is the correct approach. Yet, in many transport problems there are source terms, namely collisional interactions and inhomogeneous sources. When the advection dominates these source terms (i.e. when a mean free path is resolved), the Riemann solver's added dissipation is the correct amount. However, when a mean free path is not resolved, the dissipation is incorrect.

To address this problem we suggest that the Riemann dissipation be scaled out as the cell size relative to a mean free path grows. In particular we suggest that the dissipation matrix $|\Lambda|$ be multiplied by $[1 + (\sigma_s \Delta x)^2]^{-1}$, where $\sigma_s = \sigma_t - \sigma_a$ is the scattering cross-section. This scaling allows the dissipation to be largely unchanged when the cell size is smaller than a scattering mean free path, but also reduces the dissipation acutely when the cell size is larger than the scattering mean free path. This has the effect of effectively making $|\Lambda|_{0,0} = O(\epsilon^2)$ as $\epsilon \rightarrow 0$ in the proof of Theorem 2, and thereby this ‘‘diffusion correction’’ removes the problem revealed there.

This scaling factor obeys

$$\frac{1}{1 + (\sigma_t/\epsilon - \epsilon\sigma_a)^2 \Delta x^2} \sim \begin{cases} 1 & \epsilon \rightarrow \infty \\ \frac{\epsilon^2}{\sigma_s^2 \Delta x^2} & \epsilon \rightarrow 0. \end{cases} \quad (7.44)$$

Using this scaling the order $1/\epsilon$ equations still yield

$$I_{l,i}^{(0)} = 0 \quad (7.45)$$

for $l > 0$. But the order 1 equations now state

$$I_{1,i}^{(1)} = \frac{-1}{2\Delta x \sigma_t \sqrt{3}} \left(I_{0,i+1}^{(0)} - I_{0,i-1}^{(0)} \right), \quad (7.46)$$

which recalls Fick's law. Finally, the order ϵ equations give

$$\frac{1}{c} \frac{dI_{0,i}^{(0)}}{dt} - \frac{1}{3\sigma_t} \left(\frac{I_{0,i+2}^{(0)} - 2I_{0,i}^{(0)} + I_{0,i-2}^{(0)}}{4\Delta x} \right) + \sigma_a I_{0,i}^{(0)} = \frac{Q_i}{2\sqrt{\pi}}. \quad (7.47)$$

This is a discrete diffusion equation with the correct diffusion coefficient $D = 1/3\sigma_t$. The effect of the scaled dissipation is to convert the solver from an upwinded Riemann solver when computational cells are on the order of a mean-free-path or smaller, into a cell-centered diffusion solver when cells are many scattering mean-free-paths thick.

It should be noted that this diffusion equation is discretized on a mesh that is of size $2\Delta x$, rather than on the mesh of size Δx . As written this limit is therefore yielding two diffusion equations, one on even numbered mesh cells, and one on odd numbered cells. This arises because in the first-order Riemann solver all quantities are effectively cell-centered. However, as we will show in the results of the next section, we have not seen a problem with this in practice because the nonlinear interpolation used in a high-resolution Riemann solver does couple neighboring cells in this limit.

We also note that the scaling of the Riemann dissipation term could have been of the form $[1 + (\frac{\sigma_t}{c}\Delta x)^2]^{-1}$ based on σ_t rather than σ_s . Our thinking, however, was that for a mesh that contains large cells in a strong absorber we should continue to upwind the solution, rather than allow it to become a centered difference scheme. However, in problems of thermal radiative transfer the absorption/reëmission process behaves as effective scattering. In this case the scaling factor using σ_t would be appropriate.

7.2 Free Streaming Limit and Negative Solutions

In one dimension both the transport and P_n equations have simple solutions in the free streaming limit. For an initial condition of isotropic radiation released at the origin, $I_0^0(0, \mu, 0) = Q$ the solutions are

$$I_0^0(z, \mu, t) = \frac{Q}{2(\frac{z}{\mu} - t)} h\left(\frac{z}{\mu} - t\right) \quad \text{transport} \quad (7.48)$$

$$I_0^0(z, \mu, t) = \frac{1}{n+1} \sum_{i=0}^n \delta(z - \lambda_i t) \quad P_n, \quad (7.49)$$

where the λ_i are the eigenvalues of the P_n equations. In the limit on $n \rightarrow \infty$ the P_n solution agrees with the transport solution, however, for any finite number of moments there will be a series of delta functions moving out from the origin. These simple solutions can show us how the P_n equations have trouble in multi-dimensional problems. To see these problems we will transform this plane solution to a point solution. We can make such a transformation because the P_n equations are rotationally invariant and linear.

The plane to point transform is found by inverting the transformation from a point solution to a planar solution. Given a point solution, $I_{0, \text{point}}^0$, it is possible to find the solution from summing up a plane of those solutions. Let z be the distance from the plane and r be the distance from the point, then the planar solution is given by

$$I_{0, \text{plane}}^0(z) = \int_r^\infty dr 2\pi r I_{0, \text{point}}(r). \quad (7.50)$$

By inverting this equation we can find the value of a point solution from a plane solution:

$$I_{0, \text{point}}^0(r) = \frac{-1}{2\pi r} \left. \frac{\partial I_{0, \text{plane}}(z)}{\partial r} \right|_{z=r}. \quad (7.51)$$

Applying Eq. (7.51) formula to Eq. (7.49) *will give negative values for I_0^0* because the derivative of a δ function is negative in the sense of distributions [68]. This quantity is a constant multiple of the energy density, so the energy density will also be negative.

There are three ingredients that make the P_n equations give negative energy densities. The plane to point transform was valid because the P_n equations are rotationally invariant and linear. The delta functions which led to negatives exist because the P_n equations are hyperbolic and have a finite number of degrees of freedom in the angular variables. This means that if we want to guarantee that the solution given by a P_n method will be positive, we must give up one of these four properties. Obviously,

we cannot give up finite dimensionality if we want to do computer simulations; so there are really only three possibilities. This is unfortunate because the transport equation is linear, rotationally invariant and hyperbolic.

Other methods satisfy this quid pro quo in different ways. Diffusion methods give up hyperbolicity allowing information to travel infinitely fast. This eliminates the delta functions and gives diffusion positivity. Discrete ordinates methods give the same delta function solutions as those in Eq. (7.49). However, the discrete ordinates equations are not rotationally invariant so the transform in Eq. (7.51) does not apply⁴. This lack of rotational invariance leads to the ray effects seen in discrete ordinates solutions [13]. This makes the choice of dropping the rotational invariance of the P_n approach undesirable.

It would be possible to drop linearity through a nonlinear closure. The standard P_n equations are closed by assuming $I_{n+1}^m = 0$. A nonlinear closure would attempt to construct this moment using the other moments that are known. One potentially promising closure is the Minerbo maximum entropy closure which determines the unknown moment using physical arguments regarding the equilibration of the intensity in angle [1, 5]. It should be noted that linear closures derived based on some assumed form for the intensity [26] will not resolve our problem unless they give up hyperbolicity.

Another possible remedy could be to make a diffusion-like approximation in the highest moments in the P_n equations. This would lose the hyperbolic character in the P_n equations, yet how this would actually effect the physics is unknown. It is possible that this ad hoc approach would only seriously affect the higher moments in the P_n equations without harming the lower moments (which are the moments that we

⁴Discrete ordinates solutions can sometimes have negative values for the energy density, but this is a numerical effect caused by some spatial discretizations [29], whereas the P_n equations have these negatives in the spatially continuous limit.

care the most about).

CHAPTER VIII

Verification and Test Problems

Solving the time dependent thermal transport equations is a difficult problem, the number of degrees of freedom is large and the system is nonlinear. These nonlinearities in the temperature variable make all but the most trivial problems tractable analytically. This is unfortunate because it is rather useful to have analytic solutions to test a code against. With an analytic test one can show that a code than claims to solve the P_n equations does indeed solve those equations. This type of test is called verification and is an excellent way to check a code for bugs or other unexpected shortcomings. Below we will derive an analytic solution to the thermal P_1 equations and later we will use this problem to test our code. We will also present some other radiation transport test problems.

8.1 Analytic Solutions to the Thermal P_1 Equations

Despite the difficulting of solving thermal transport problems, using a cleverly form of the heat capacity (C_V) as a function of temperature, it is possible to find linearized equations that can be solved. Using this form, grey diffusion solutions [69,70] have been found for various problems. Also, using a modified P_1 and P_2 approximation for grey transport Pomraning and Shokair found the value of the radiation and temperature fields

at the free-surface and the total radiation and material energies of a half-space problem [71]; Ganapol and Pomraning found these quantities for the full transport equations [72]. The transport solution for the radiation and temperature fields throughout a problem domain was found by Su and Olson [73].

One might wonder why so much ink has been spilled to solve radiative transfer problems that involve an assumed (and unrealistic¹) form of the heat capacity. This work has not been an exercise in solipsism but rather represents a the need to have analytic solutions to verify the codes written to solve the nonlinear radiation transport equations. Given that a code should have the ability to take an arbitrary heat capacity for a material, these solutions can be used to verify numerical codes (i.e. a diffusion code should converge to the analytic diffusion solution). Furthermore, given an analytic solution to the transport equation we can test the accuracy and applicability of various approximations to the transport solution.

Below we will develop the first time-dependent P_1 (telegrapher's equation) solutions for thermal radiative transfer. We do this by first developing the Green's function for the P_1 equations and then using this solution to "build up" solutions to other problems. This solution will allow us to verify that indeed our method and code solves the P_1 equations as claimed. Beyond this capacity, the analytic solutions will give us insight into the behavior of P_1 solutions as compared to transport solutions.

8.1.1 Normalized P_1 equations

We will normalize the P_1 equations to make our discussion as general as possible. Also, we will write the equations in energy density and radiation flux form instead of spherical harmonic form since these are a more common form of the equations. In this

¹The heat capacity of a material is nearly constant as a function of temperature, except in the region where the material is ionizing [20].

notation the P_1 equations are

$$\frac{1}{c} \frac{\partial \mathcal{E}}{\partial \tau} + \frac{1}{c} \frac{\partial \mathcal{F}}{\partial z} + \sigma_a \mathcal{E} = a \sigma_a T^4 + S, \quad (8.1)$$

$$\frac{1}{c} \frac{\partial \mathcal{F}}{\partial \tau} + \frac{c}{3} \frac{\partial \mathcal{E}}{\partial z} + \sigma_a \mathcal{F} = 0, \quad (8.2)$$

where the radiation flux is defined as

$$\mathcal{F} = 2\pi \int_{-1}^1 d\mu \mu I(\mu), \quad (8.3)$$

and the energy density is \mathcal{E} . These equations are different than those used in the analysis on the P_1 by Shokair and Pomraning [71]. They dropped the time derivative in 8.2 in their analysis.

The equation that describes the coupling of material energy to radiation energy,

$$\frac{C_v}{c} \frac{\partial T}{\partial t} = \sigma_a (\mathcal{E} - aT^4), \quad (8.4)$$

with C_v denoting the heat capacity of the material. Eqs. (8.1), (8.2), and (8.4) give a closed system of equations that govern the interplay between radiation and material energies.

To proceed we make a common assumption on the behavior of the heat capacity of the material to make Eq. (8.4) linear in T^4 , namely we stipulate that

$$C_v = \alpha T^3. \quad (8.5)$$

This assumption on the heat capacity was used by Pomraning [69], Su and Olson in their development of a diffusion benchmark [70] and in a transport benchmark [73]. This assumption allows us to bring to bear the classic tools of mathematical physics on the P_1 system of equations.

Before continuing we will normalize the P_1 system using

$$x \equiv \sigma_a z, \quad \epsilon \equiv \frac{4a}{\alpha}, \quad t \equiv \epsilon c \sigma_a \tau, \quad (8.6)$$

$$\mathcal{E} \equiv \frac{E}{aT_r^4}, \quad U \equiv \frac{T^4}{T_r^4}, \quad \mathcal{F} \equiv \frac{F}{aT_r^4}, \quad Q = \frac{S}{aT_r^4}. \quad (8.7)$$

Given these definitions the P_1 system becomes

$$\epsilon \frac{\partial \mathcal{E}}{\partial t} + \frac{1}{c} \frac{\partial \mathcal{F}}{\partial x} = -(\mathcal{E} - U) + Q, \quad (8.8)$$

$$\epsilon \frac{\partial \mathcal{F}}{\partial t} + \frac{c}{3} \frac{\partial \mathcal{E}}{\partial x} + \mathcal{F} = 0, \quad (8.9)$$

$$\frac{\partial U}{\partial t} = (\mathcal{E} - U). \quad (8.10)$$

The boundary and initial conditions that we will use throughout this paper are

$$E(\pm\infty, t) = 0, \quad E(x, 0) = U(x, 0) = 0. \quad (8.11)$$

8.1.2 Green's function

We will now seek the Green's function for the system given by Eqs. (8.8)-(8.10) by specifying the source $Q = \delta(x)\delta(t)$. Upon finding this Green's function we will be able to construct the solution from more complex sources.

To begin finding the Green's function we first Laplace transform Eqs. (8.8)-(8.10) using the convention

$$\hat{f}(s) = \int_0^\infty dt e^{-st} f(t), \quad (8.12)$$

and we obtain

$$\epsilon s \hat{\mathcal{E}} + \frac{1}{c} \frac{\partial \hat{\mathcal{F}}}{\partial x} = -c_a (\hat{\mathcal{E}} - \hat{U}) + \delta(x), \quad (8.13)$$

$$\epsilon s \hat{\mathcal{F}} + \frac{c}{3} \frac{\partial \hat{\mathcal{E}}}{\partial x} + \hat{\mathcal{F}} = 0, \quad (8.14)$$

$$s \hat{U} = c_a (\hat{\mathcal{E}} - \hat{U}). \quad (8.15)$$

Solving Eq. (8.15) for \hat{U} we can eliminate that equation from the system changing Eq. (8.13) into

$$\epsilon s \hat{E} + \frac{1}{c} \frac{\partial \hat{F}}{\partial x} = - \left(1 - \frac{1}{1+s} \right) \hat{E} + \delta(x). \quad (8.16)$$

Applying a Fourier transform, given by

$$\tilde{f}(k) = \int_{-\infty}^{\infty} dx e^{ikx} f(x), \quad (8.17)$$

to Eqs. (8.14) and (8.16) yields

$$\epsilon s \bar{E} - \frac{ik}{c} \bar{F} = - \left(1 - \frac{1}{1+s} \right) \bar{E} + 1 \quad (8.18)$$

$$\epsilon s \bar{F} - \frac{ikc}{3} \bar{E} + \bar{F} = 0, \quad (8.19)$$

where for simplicity in notation we have denoted functions that have been both Laplace and Fourier transformed using an overbar. Upon solving Eq. (8.19) for \bar{F} and substituting the solution into Eq. (8.18) we arrive at

$$\epsilon s \bar{E} + \frac{k^2}{3(\epsilon s + 1)} \bar{E} = - \left(1 - \frac{1}{1+s} \right) \bar{E} + 1. \quad (8.20)$$

This gives the value of \bar{E} as

$$\bar{E}(k, s) = \frac{3(\epsilon s + 1)}{\left(3\epsilon s(\epsilon s + 1) + 3(\epsilon s + 1) - \frac{3(\epsilon s + 1)}{1+s} + k^2 \right)}. \quad (8.21)$$

Before performing any inverse transforms, we factor the denominator of Eq. (8.21)

$$\bar{E}(k, s) = \frac{3(\epsilon s + 1)}{(k + i\Lambda)(k - i\Lambda)}, \quad (8.22)$$

where we have defined

$$\Lambda^2 \equiv 3\epsilon s(\epsilon s + 1) + 3(\epsilon s + 1) - \frac{3(\epsilon s + 1)}{1+s}. \quad (8.23)$$

The inverse Fourier transform of Eq. (8.22) is given by

$$\hat{E}(x, s) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dk \frac{3(\epsilon s + 1)e^{-ikx}}{(k + i\Lambda)(k - i\Lambda)}. \quad (8.24)$$

This integral can be evaluated using the residue theorem [74] and the fact that the roots of a quadratic equation are either both real or complex conjugates. When $x > 0$ the path of integration taken in the complex plane is a semicircle in the lower-half plane enclosing the pole at $k = -i\Lambda$ and using the residue theorem

$$\hat{E} = \frac{1}{2\pi} \lim_{R \rightarrow \infty} \left(\int_{-R}^R dk \frac{3(\epsilon s + 1)e^{-ikx}}{(k + i\Lambda)(k - i\Lambda)} + \int_0^\pi d\theta \frac{3(\epsilon s + 1)e^{-iRe^{i\theta}x}}{\sigma_t (Re^{i\theta} + i\Lambda)(Re^{i\theta} - i\Lambda)} \right). \quad (8.25)$$

The integral over θ in Eq. (8.25) limits to zero as $R \rightarrow \infty$, making

$$\hat{E}(x, s) = \frac{3(\epsilon s + 1)e^{-\Lambda x}}{2\Lambda}, \quad x > 0. \quad (8.26)$$

For $x < 0$ the path of integration is a semicircle below the real axis and \hat{E} becomes

$$\hat{E}(x, s) = \frac{3(\epsilon s + 1)e^{+\Lambda x}}{2\Lambda}, \quad x < 0. \quad (8.27)$$

We combine Eqs. (8.26) and (8.27) to get

$$\hat{E}(x, s) = \frac{3(\epsilon s + 1)e^{-\Lambda|x|}}{2\Lambda}. \quad (8.28)$$

Our task now is to perform an inverse Laplace transform on Eq. (8.28). We are now at a crossroads in our journey; this transform is not tractable analytically (to the best of our knowledge) and we have two choices to proceed: we could numerically perform the inversion to obtain the value of E at particular values of t , this tack had been used in the past [70,73], or we could specify specific values of ϵ to simplify Eq. (8.28). Given a value of unity for ϵ the value of Λ changes and Eq. (8.28) becomes

$$\hat{\mathcal{E}}(x, s) = \frac{3(s + 1)e^{-\sqrt{3s(s+2)}|x|}}{2\sqrt{3s(s+2)}}. \quad (8.29)$$

To aid in the computation of the inverse Laplace transform we notice that Eq. (8.29) can be written as $\sqrt{3}/2(f(s) + sf(s))$ where

$$f(s) \equiv \frac{e^{-\sqrt{3s(s+2)}|x|}}{\sqrt{s(s+2)}}. \quad (8.30)$$

This tells us that $E(x, t) = \sqrt{3}/2(f(t) + f'(t))$. The inverse Laplace transform of $f(s)$ can be found in an exhaustive table of Laplace transforms [75] to be

$$f(t) = e^{-t} I_0(\sqrt{t^2 - 3x^2}) h(t - \sqrt{3}|x|). \quad (8.31)$$

Here we have denoted the modified Bessel function of the first kind of order zero with I_0 and the Heaviside step function by h . The derivative of $f(t)$ is

$$\frac{\partial f}{\partial t} = e^{-t} \left(-I_0(\sqrt{t^2 - 3x^2}) h(t - \sqrt{3}|x|) + \frac{t I_1(\sqrt{t^2 - 3x^2})}{\sqrt{t^2 - 3x^2}} h(t - \sqrt{3}|x|) \right) \quad (8.32)$$

$$+ I_0(\sqrt{t^2 - 3x^2}) \delta(t - \sqrt{3}|x|) \Big). \quad (8.33)$$

Therefore, we have the Green's function for the P_1 equations with $\epsilon = 1$

$$G(x, y, t) = \frac{\sqrt{3}}{2} e^{-t} \left(\frac{t I_1(\sqrt{t^2 - 3(x-y)^2})}{\sqrt{t^2 - 3(x-y)^2}} h(t - \sqrt{3}|x-y|) \right) \quad (8.34)$$

$$+ I_0(\sqrt{t^2 - 3(x-y)^2}) \delta(t - \sqrt{3}|x-y|) \Big). \quad (8.35)$$

In the Green's function we note some important properties. The finite propagation speed allowed by the P_1 equations (due to their hyperbolic nature) is manifest in the delta and Heaviside functions; the solution has a front of uncollided source particles with a speed of $1/\sqrt{3}$ moving away from the origin, with a decaying modified Bessel function behind the front. These features can be seen in Fig. 8.1 which shows the Green's function solution from a pulsed source at the origin (i.e. $y = 0$) at $t = 10$. The "strength" of the delta functions in this figure is fairly weak because a large majority of the particles have been absorbed and reemitted and very few are still streaming straight from the source. The strength of the delta function becomes $\exp(-t)$, hence, most of the energy in the solution is in the reemitted particles.

8.1.3 Finite source in time and space

A well known benchmark problem presented in the past has a unit source where $|x| < 0.5$ and $t < 10$. A transport solution and diffusion solution were presented by




Figure 8.1: The Green's function solution to the P_1 equations at $t = 10$.

Su and Olson [70, 73]. To present a P_1 solution to this problem we will integrate the Green's function Eq. (8.34) over space and time yielding

$$\begin{aligned}
 E(x, t) &= \frac{\sqrt{3}}{2} \int_{-0.5}^{0.5} dy \int_0^t dt' e^{-t'} \left(\frac{t' I_1(\sqrt{t'^2 - 3(x-y)^2})}{\sqrt{t'^2 - 3(x-y)^2}} h(t' - \sqrt{3}|x-y|) \right. \\
 &\quad \left. + I_0(\sqrt{t'^2 - 3(x-y)^2}) \delta(t' - \sqrt{3}|x-y|) \right) \quad (8.36) \\
 &= \frac{\sqrt{3}}{2} \int_{-0.5}^{0.5} dy \left(e^{-\sqrt{3}|x-y|} h(t - \sqrt{3}|x-y|) \right. \\
 &\quad \left. + \int_0^t dt' \frac{t' I_1(\sqrt{t'^2 - 3(x-y)^2})}{\sqrt{t'^2 - 3(x-y)^2}} h(t' - \sqrt{3}|x-y|) \right).
 \end{aligned}$$

This function can be computed and plotted using Mathematica in a straightforward manner. In Figs. 8.2 - 8.3 we plot the P_1 analytic solution along with the transport solutions [73].




Figure 8.2: Analytic solutions to the P_1 and transport equations for the finite source problem.

8.2 Point Source

To transform the solution from a pulsed plane source to a point source in 3D we can employ the transform

$$G_{\text{point}}(r, t) = -\frac{1}{2\pi r} \frac{\partial G_{\text{plane}}}{\partial r}. \quad (8.37)$$




Figure 8.3: The solutions shown in Fig. 8.2 on a logarithmic scale.

Upon using this transform, the solution function for a point source at the origin becomes

$$\begin{aligned}
 G_{\text{point}}(r, r', t) = & \frac{3\sqrt{3}e^{-t}I_1(\sqrt{t^2 - 3r^2})\delta(t - \sqrt{3}r)}{2\pi\sqrt{t^2 - 3r^2}} + \frac{3e^{-t}I_1(\sqrt{t^2 - 3r^2})\delta(t - \sqrt{3}r)}{2\pi r\sqrt{t^2 - 3r^2}} \\
 & + \frac{3\sqrt{3}e^{-t}(\sqrt{t^2 - 3r^2}I_0(\sqrt{t^2 - 3r^2}) + I_2(\sqrt{t^2 - 3r^2}))h(t - \sqrt{3}r)}{4\pi(t^2 - 3r^2)^{3/2}} \\
 & - \sqrt{3}e^{-t}t \frac{2I_1(\sqrt{t^2 - 3r^2})h(t - \sqrt{3}r)}{4\pi(t^2 - 3r^2)^{3/2}} + \frac{3e^{-t}I_0(\sqrt{t^2 - 3r^2})\delta'(t - \sqrt{3}r)}{2\pi r} \quad (8.38)
 \end{aligned}$$

8.2.1 Line source

By summing an appropriate arrangement of point sources we can construct a line source benchmark. This line source can be used to test our two-dimensional P_1 code for

radiative transfer. To construct this line source we employ the transform

$$E_{\text{line}}(\rho, t) = 2 \int_0^\infty dz E_{\text{point}}(\sqrt{\rho^2 + z^2}, t), \quad (8.39)$$

where the variable ρ is the distance from the line source. The integral can be more easily evaluated if we change variables to

$$R \equiv \sqrt{3}\sqrt{z^2 + \rho^2} \quad dz = \frac{Rdr}{\sqrt{3}\sqrt{R^2 - 3\rho^2}}. \quad (8.40)$$

After making this transform the terms that contain delta functions can be evaluated after some manipulation and we arrive at

$$E_{\text{line}}(\rho, t) = \frac{3te^{-t}(t^2 - 3\rho^2 - 2)h(t - \sqrt{3}\rho)}{4\pi(t^2 - 3\rho^2)^{\frac{3}{2}}} - \frac{3e^{-t}I_0(\sqrt{t^2 - 3\rho^2})}{2\pi\sqrt{R^2 - 3\rho^2}}\delta(t - \sqrt{3}\rho) \Big|_{R=\sqrt{3}\rho} \\ + \int_{\sqrt{3}\rho}^\infty \frac{3te^{-t}R(I_0(\sqrt{t^2 - R^2}) + I_2(\sqrt{t^2 - R^2}))h(t - R)}{4\pi(t^2 - R^2)\sqrt{R^2 - \rho^2}} dR \quad (8.41)$$

$$- \int_{\sqrt{3}\rho}^\infty \frac{3Re^{-t}I_1(\sqrt{t^2 - 3R^2})h(t - \sqrt{3}R)}{2\pi(t^2 - R^2)^{\frac{3}{2}}\sqrt{R^2 - 3\rho^2}} dR. \quad (8.42)$$

The integrals of Bessel functions over a finite domain must be evaluated numerically. These integrals need to be evaluated carefully due to the fact that there is a square-root singularity at the lower endpoint.

Results from the line source are shown in Fig. 8.4. This figure shows the energy density from the absorbed and reemitted (i.e. the delta functions are not shown). Just behind the wave front the energy density is negative. This negativity decays in time and at $ct = 10$ is nearly gone. Early in time P_1 is doing a poor job of capturing the angular dependence of the intensity. As more photons are absorbed and emitted isotropically, the P_1 approximation is more viable and the negatives go away.

8.3 Linear Exact Solutions

Problems without temperature feedback can also be used to test the transport features of our P_n method. Brunner developed P_1 solutions to two such problems, a plane




Figure 8.4: The collided energy density from the line source problem at time $ct = 5, 10$.

source [5] in a scattering medium and a line source in vacuum [2], both pulsed at time zero. For the plane source there is also a transport solution [76]. These P_1 solutions again contain by delta functions traveling away from the origin at speed $c/\sqrt{3}$. The plane solution is plotted in Fig. 8.5. Both problems present a numerical method with a mare's nest of difficulties. Early in time the P_1 solution is a pair of delta function spikes moving away from the origin, but as time progresses a smooth solution representing the collided particles grows in the middle of the domain. Hence, a numerical scheme must capture both steep gradients (the spikes) and smooth regions. Also the method must move the uncollided particles at the appropriate speed.

8.4 Other Test Problems

There are some problems without analytic solutions that will demonstrate how the P_n method we have developed will perform in more realistic simulations. One such

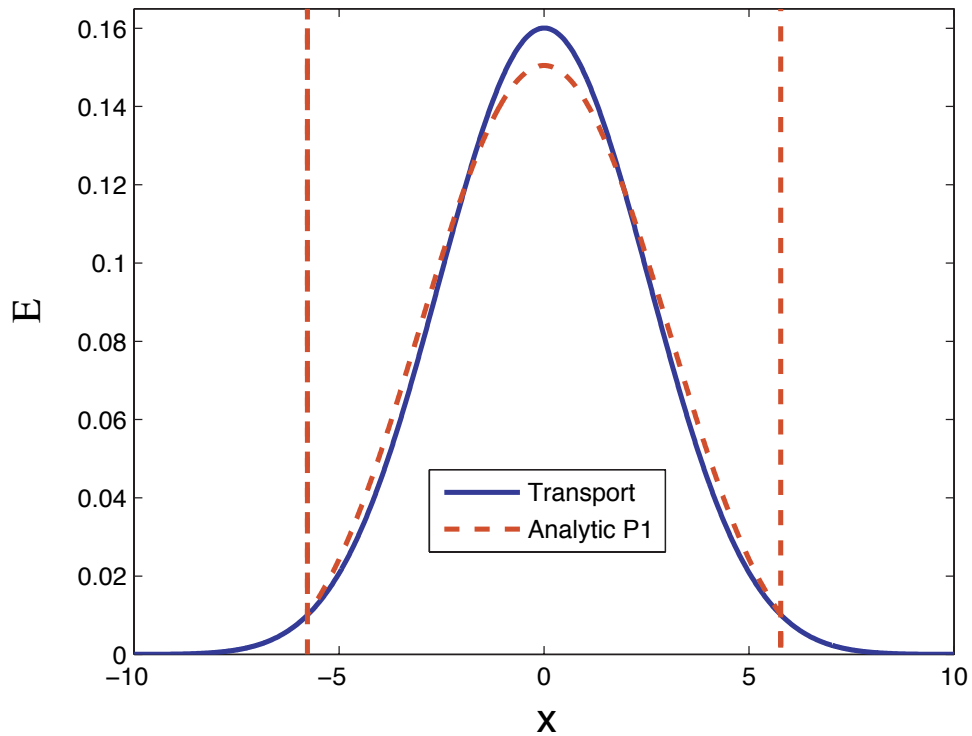


Figure 8.5: The P_1 and transport solutions to the plane, pulsed source problem.

problem is called the “M” problem. This problem is a variation on the so-called crooked pipe problem (aka the tophat problem) [77]. In our version an M shaped vacuum is cut out of a dense material. A schematic of this problem is shown in Fig. 8.6. We will run this problem in two modes: the first is in normalized units with $c = 1$ where the blue region is a purely scattering medium of $\sigma_s = 20$ making this a problem of linear transport. The other mode has thermal feedback with $c = 2.998 \times 10^8$ m/s, $C_v = 5 \times 10^{10}$ J/kg-keV and $\sigma_a = 5000$ m $^{-1}$. In both instances radiation is input isotropically through the middle leg of the problem.

A one-dimensional linear problem that we will examine is Reed’s problem [78]. This is a multi-material problem and its set up is detailed in Table 9.1.1. Also, there is a vacuum boundary on the left and a reflecting boundary on the right. This problem




Figure 8.6: The “M” problem schematic: the maize colored regions are vacuum and the blue regions are dense material.

features very distinct material types and provides a way to assess a method’s ability to handle the interfaces between dissimilar media.

The final test problem we will show is problem meant to roughly represent a hot wire in a Z pinch array heating another wire across a vacuum. This problem (see Fig. 9.28) has a central block of material heated to $T = 300$ eV and another block starting at $T = 1$ eV, at the beginning of the problem there is no radiation. The blocks are substantial with $C_v = 5 \times 10^{10}$ J/kg-keV and $\sigma_a = 5000$ m⁻¹. The reflecting boundaries quadruple the actual size of the problem – there are four of the initially cooler blocks.

Table 8.1: Material Layout in Reed's Problem

Scattering Region	Vacuum	Absorber	Strong Source
$\sigma_a = 0.1$		$\sigma_a = 5$	$\sigma_a = 50$
$\sigma_s = 0.9$		$\sigma_s = 0$	$\sigma_s = 0$
$S = 0$		$S = 0$	$S = 50$
$x \in (0, 2)$ $x \in (2, 4)$	$x \in (4, 5)$	$x \in (5, 6)$	$x \in (6, 8)$




Figure 8.7: The problem of one wire heating another across a vacuum.

CHAPTER IX

Numerical Results

Now we will study some results from calculations using the numerical methods introduced in the previous chapters. The presentation of the results will mimic the development of the implicit Riemann solvers. First, we will begin with some results from linear transport (i.e. problems where there is no material temperature). These results begin with a demonstration of the effectiveness of an implicit method on the nonlinear slope reconstruction terms of the high resolution Riemann solver; this is followed by a discussion of results in the diffusion limit. Next, the results from the quasi-linear method are presented. This is followed by results from thermal radiative transfer (aka nonlinear transport). Finally, we will show results from high resolution time integration. Throughout this chapter the high resolution spatial scheme will be used (unless otherwise noted) along with backward Euler time integration (except in the last section).

9.1 Linear Transport Results

Linear transport is an important part of a thermal transport calculation because in many codes the material temperature and the transport operators are split. This means that the two equations are solved independently, possibly by using a guess at the

temperature, solving the transport equation and iterating on the temperature guess. In fact, most of the literature on transport methods presents results for linear transport, in part due to the fact that in this split case a linear transport equation is solved.

9.1.1 Implicit Riemann solver using a Newton method

The high resolution Riemann solver method adds a nonlinear term to the linear P_n equations; because of this the effectiveness of an implicit scheme on such equations was not a foregone conclusion. A nonlinear equation solver would have to be used to solve a nonlinear algebraic system at each time step. Newton based methods are the most straightforward approach to solve such equations, so the following results use the harmonic mean limiter and a Newton method to solve the implicit equations. In all our linear calculations the speed of light is set to one. This makes the scalar flux, ϕ , equal to the energy density.

One-dimensional problems

The first problem we will examine is the plane pulse problem in a purely scattering medium (see Sec. 8.3); Fig. 9.1 shows some results. In this figure $CFL = \frac{c\Delta x}{\Delta t}$ and $CFL > 1$ would be unstable for an explicit time integration method. Even time steps that allow particle to travel across fifty spatial cells in one time unit show good agreement with the analytic P_1 solution. The numerical solutions 5 seconds after the pulse demonstrate a smoothing of the spikes in the analytic P_1 solution. It is clear that the solutions are converging in Δt to the P_1 solution, as would be expected for a method solving that P_1 equations in discretized form. However, the front of the wave of particles with speed $\pm\sqrt{(1/3)}$ – which is captured in the analytic solution by a delta function – is retarded from the proper speed with large time steps. This is not an artifact from the Riemann solver, rather it arises from the time integration scheme, as can be seen in Ref. 1. In




Figure 9.1: The high resolution P_1 solutions between $x = -3.5$ and 3.5 at $t = 5$.

this reference results for the same problem using explicit time integration have the wave fronts moving at the appropriate speeds. Increasing the size of the time step beyond the CFL limit moves particles at the wrong speeds. This effect is seen in Fig. 9.1, in which we see that as Δt is reduced (CFL decreases for Δx fixed) the wave-front tracks the correct location more accurately. At later times there are fewer uncollided particles (the delta functions are weaker). This makes the analytic P_1 solution easier to compute and this effect is seen in Fig. 9.2. In this figure, even a large time step, $\Delta t = 1$, performs very well in capturing the analytic solution.

The high resolution and first order spatial schemes were compared on a series of tests using the P_7 equations to solve the plane pulse problem. These results at $t = 1$ are shown in Fig. 9.3. The P_7 solutions have a series of “spikes” representing the particles travelling at speeds characteristic of the P_7 equations, much like the delta functions of the P_1 equations. In the series of plots in Fig. 9.3 the high resolution method captures

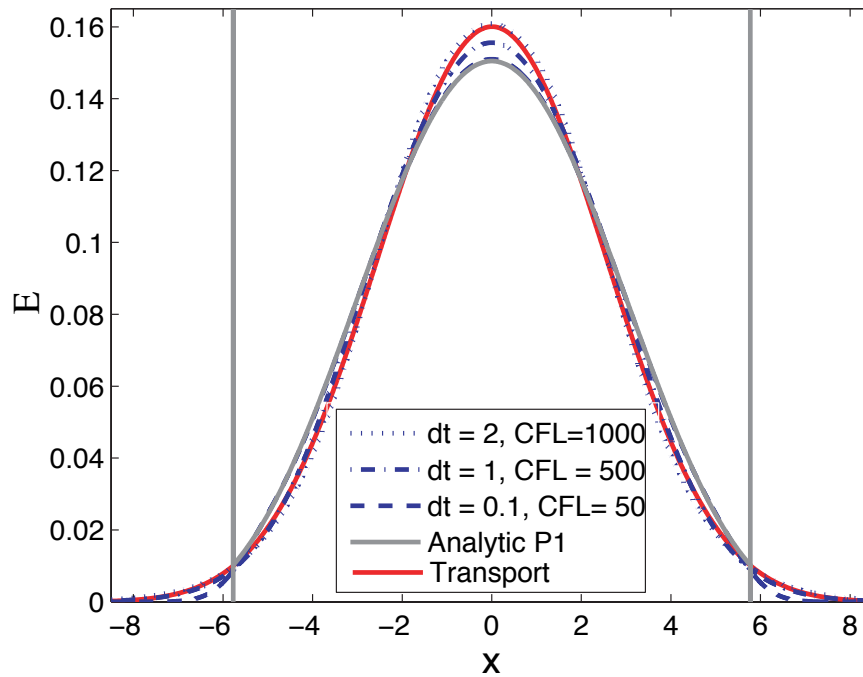


Figure 9.2: The high resolution P_1 solutions at $t = 10$. The analytic P_1 solution and the $\Delta t = 0.1$ solution are coincident in the middle region.

these features more accurately for both coarse and fine grids. Furthermore, the high resolution solution with $N_x = 150$ is better than the first order solution with $N_x = 300$. This improvement with fewer cells will be an important effect in multidimensional problems.

Using our time dependent method, we solve steady state problems using what is colloquially referred to as “time-stepping to steady state.” An initial condition is given and this condition is propagated forward in time until the solution converges (i.e. the difference in solution vectors between successive time steps is sufficiently small). The results for steady state calculations using the high resolution spatial scheme on Reed’s problem from Sec. 8.4 are given in Fig. 9.4. This figure depicts P_7 solutions obtained using the high resolution spatial scheme and $\Delta t = 10$ for various Δx values. Here we see that our method is able to handle the sharp changes in solution behavior at the material

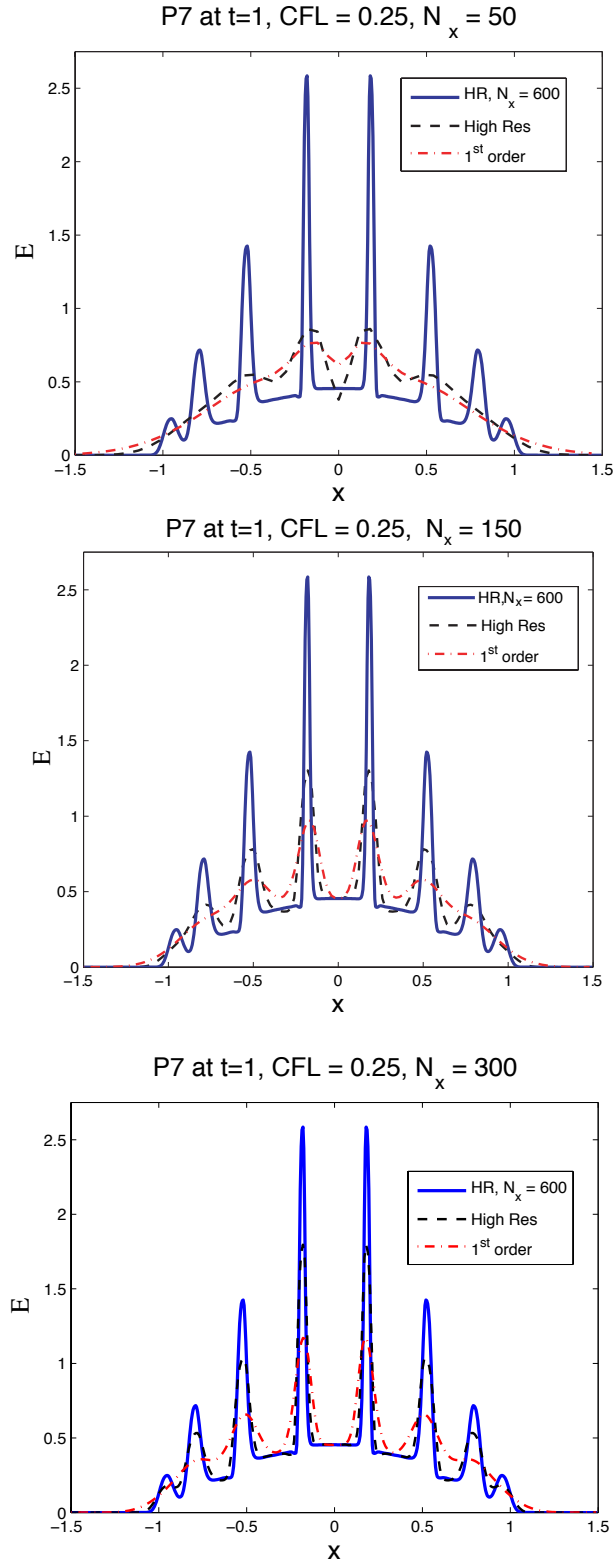


Figure 9.3: Comparison of high resolution and first-order solutions to the P_7 equations with different numbers of spatial grid points.




Figure 9.4: Steady state results from Reed's problem

interfaces and the solution is fairly insensitive to increasing the spatial grid size. Also, in Fig. 9.4 the S_8 solution using the multidimensional multibalance method [66] is shown. The agreement of the solutions suggests that the implicit Riemann solver is converging to the correct solution.

Figure 9.5 shows the solution to Reed's problem at $t = 1$. The initial condition had all moments set to zero and steady state solution “grows” out of this. The only region where the solution is sensitive to the size of the time step is the strong scattering region. This suggests that the scattering source becomes important with large Δt and affects the numerical results in this area.

Two-dimensional problems

The analog of the plane pulse test problem for two dimensions is a pulsed line source. We will take a P_1 analytic solution for a line source pulsed in a vacuum and use it to compare to our numerical solutions. The numerical solution for this problem is shown in




Figure 9.5: The scalar flux for different values of Δt ; cell width is $\Delta x = 0.05$.

Fig. 9.6. This figure shows the effects of the grid on the solution. At points skew to the Cartesian grid, the solution is slightly different than at points where information could arrive by only passing orthogonally through the grid. We can also use this problem to compare the high resolution method to the first order scheme. Fig. 9.7 shows results from both spatial discretizations. Similar to the results from the one-dimensional problem, the high resolution method with fewer points gives a better result than the first order method with more points. In the figure, the result using a grid of $N_x = N_z = 50$ has a quarter of the number of cells as the grid with $N_x = N_z = 100$. This makes the high resolution method even more important in multiple dimensions. The effect of the time step size on the solution is shown in Fig. 9.8. As the time step size is increased, the propagation of the wave front is smoothed out. This smoothing effect is common with the backward Euler method and is a result of the diffusion term in the model equation, Eq. (5.5), for the method which is proportional to $c^2 \Delta t^2$. In this problem where a delta function propagates in a vacuum, even rather small time steps introduce noticeable time integration error.




Figure 9.6: P_1 results from the pulsed line source problem with $N_x = N_z = 50$ at $t=1$. The linear color scale represents the energy density.

Diffusion correction

We now turn to results detailing the effectiveness of the diffusion correction suggested in Sec. 7.1.6. Both steady state and time-dependent problems will be used to explore the properties of the diffusion-corrected Riemann solver with the high resolution spatial scheme. Figure 9.9 presents results from a P_5 steady state calculation both with and without the diffusion correction. The problem has an incident beam at $x = 0$, a strong absorbing region from $x = 0$ to $x = 2$ and a strong scattering region from $x = 2$ to $x = 7$. In the scattering region, where the diffusion approximation is valid, the uncorrected method (Fig. 9.9(b)) gives different solutions for different mesh spacings. With the correction added to the dissipation term, the solution does not vary significantly with changes in the cell size. Even for cells that are very large compared to a




Figure 9.7: Comparison of P_1 results from high resolution and first order methods on the pulsed line source problem.

mean-free-path the method with a proper diffusion limit still yields correct results as seen in Fig. 9.9(a).

Figure 9.10 shows a problem of a uniform source embedded in a diffusive material with vacuum boundary conditions. Only resolving the diffusion length, the modified Riemann solver produces a nearly identical solution to the result calculated with a mesh that resolves a mean free path. The standard Riemann solver causes the height of the solution “plateau” and the boundary layer to be incorrect when the mean free path is not resolved.

Another steady state problem used to test the diffusion correction is a modified version of Reed’s problem [78]. The problem was modified to make the diffusive region in the problem optically thick. This was in an attempt to gauge the ability of the diffu-






Figure 9.8: The effect of the time scale on P_1 results for the pulsed line source problem. All of these were computed using the high resolution spatial scheme.

sion correction in problems with a variety of materials. The results in Fig. 9.11 show that in the diffusion correction does indeed improve the calculated flux in the diffusive region. Moreover, the solution in the strong source and absorber regions is nearly identical. There does appear to be an issue in the void region: the scalar flux is too high. When we presented the diffusion correction we did not discuss how it would work in a multi-material situation. In its current formulation the correction is not conservative at the interface between a diffusive region and a streaming region. This would explain the higher value in the vacuum region.

The diffusion correction was also tested in time-dependent problems. One problem used to test the correction places a plane pulse of particles at the center of a medium dominated by scattering ($\sigma_t = 10, \sigma_s = 9.9$). The differences between the corrected



(a) The solution using the diffusion-corrected Riemann solver



(b) The solution using the standard Riemann solver without the diffusion correction

Figure 9.9: The P_5 , steady state solution with incident beam on the left, and two regions: a strong absorber and a strong scatterer




Figure 9.10: The P_1 , steady-state solution to a uniform source problem with $Q = \frac{1}{2\pi}$ and $\sigma_a = 0.01, \sigma_t = 10$

solution and the standard Riemann solver are noteworthy. Figures 9.12 and 9.13 show the P_7 and P_1 solutions to this problem at $t = 35$ after the pulse with $\Delta t = 0.5$. The lack of a diffusion limit in the unaltered Riemann solver leads to very different solutions with different spatial grids. Figures 9.12(b) and 9.13(b) show that the width of the pulse in the solution artificially spreads as Δx increases in the solution without the diffusion correction. The fact that the diffusion correction behaves the same both in the P_1 approximation and the P_7 approximation is demonstrated by these figures. This is manifest in the fact that the standard Riemann solver solutions are similar for both P_1 and P_7 and the diffusion corrected solutions are similar for both angular approximations (i.e. Fig. 9.12(b) is similar to Fig. 9.13(b) and Fig. 9.12(a) is similar to Fig. 9.13(a)).

Table 9.1: Material Layout in Reed's Problem

Scattering Region	Vacuum		Absorber	Strong Source
$\sigma_a = 0.1$ $\sigma_s = 9.9$			$\sigma_a = 5$ $\sigma_s = 0$ $S = 0$	$\sigma_a = 50$ $\sigma_s = 0$ $S = 50$
	$S = 0$ $\Delta x = 2$	$S = 1$ $\Delta x = 2$	$\Delta x = 1$	$\Delta x = 2$

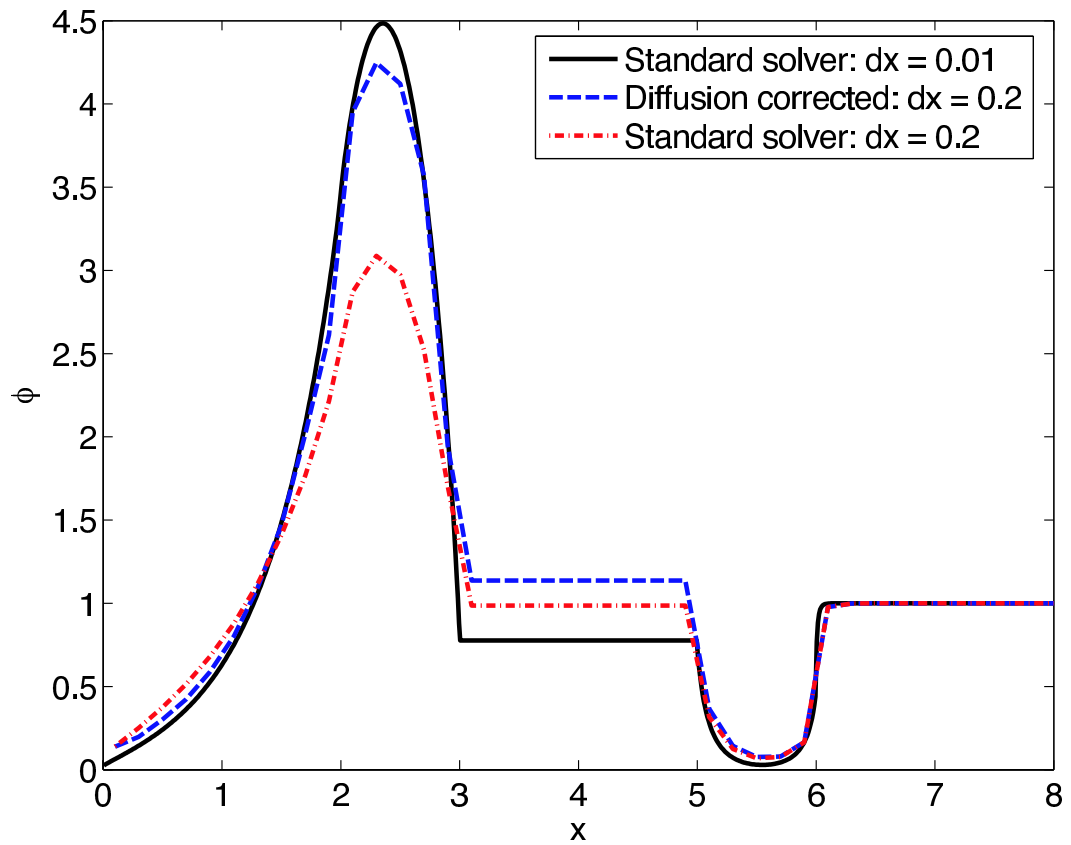




Figure 9.11: The solution to the modified Reed's problem.




(a) The solution from the modified Riemann solver




(b) The standard Riemann solver solution

Figure 9.12: The P_7 solution at $t = 35$ after the initial pulse of particles



(a) The solution from the modified Riemann solver



(b) The standard Riemann solver solution

Figure 9.13: The P_1 solution at $t = 35$ after the initial pulse of particles

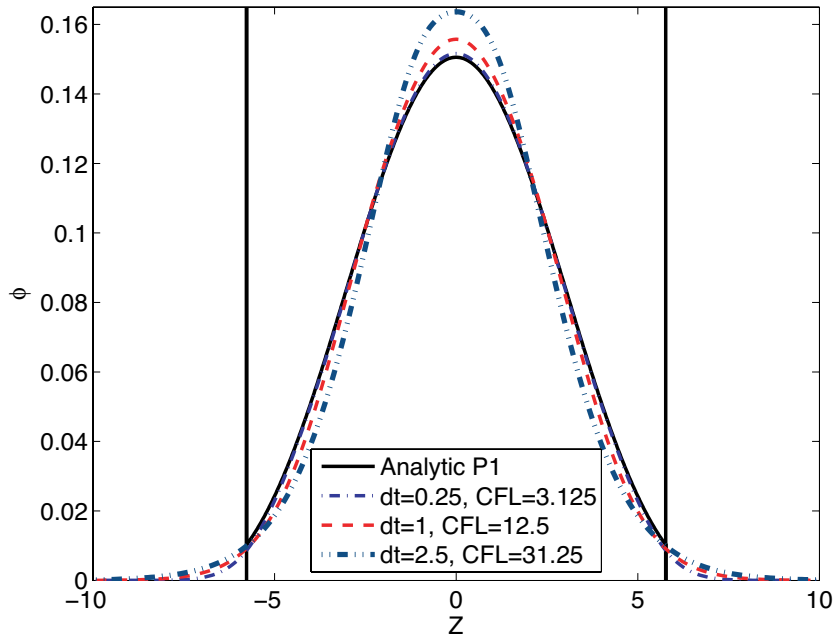


Figure 9.14: The P_1 scalar flux for the first test problem at $T = 10$ after the initial pulse for different Δt , solved using the quasi-linear formulation.

9.2 Quasi-linear time integration results

The quasi-linear approach to solving the nonlinear system of equations arising from the high resolution method was justified via analysis in Sec. 5.2.1. Now we will see if this approach bears fruit in actual numerical simulations. In Fig. 9.14 results from the quasi-linear method are shown. The solutions are converging to the analytic solution as they did when using the Newton based high resolution method.

We will now gauge the performance of the quasi-linear method against a full nonlinear calculation using a fully nonlinear Newton-Krylov high resolution method and using the first-order method. To accomplish this we have plotted the error for the plane pulse test problem as a function of spatial resolution (see Fig. 9.15) and the CPU time as a function of error (see Fig. 9.16) In both of these figures, the L-2 norm error was computed using the region just inside the delta functions in the analytic solution. This was




Figure 9.15: The error as a function of spatial resolution for P_1 scalar flux with $\Delta t = 0.05$ at $T = 10$ as compared with the analytic solution. Note the errors for the quasi-linear and Newton-based nonlinear method are essentially identical.




Figure 9.16: The CPU time as a function of error for P_1 scalar flux with $\Delta t = 0.05$ at $T = 10$ as compared with the analytic solution.

done because the error at a delta function would be infinite for any numerical method. In Fig. 9.15 we see that both the fully nonlinear and the quasi-linear method have second order error convergence whereas the first-order method achieves slightly better than first order accuracy (order 1.33). Moreover, we see that the errors in the nonlinear and quasi-linear methods are essentially identical. This suggests that for this problem there is no loss of accuracy in using the quasi-linear method versus the full nonlinear solve. Finally, the error for the first-order method with the highest spatial resolution ($\Delta x = 0.025$) is still not as small as for the high resolution methods with $\Delta x = 0.08$.

The superior error convergence shown in the quasi-linear method is a heartening result, however, it would be for naught if the time inherent in performing the extra linear solve was prohibitively long. One might imagine a circumstance where the first-order method with a much finer mesh would have a faster calculation time and the same

error as a high resolution method with a coarser mesh. Fig. 9.16 addresses this concern by comparing CPU time with error.

The quasi-linear and nonlinear methods have the CPU time and error related by an inverse power law while the first-order method has no such simple relation. This is most likely due to the memory requirements for the problem size required to make the error small using the first-order method. To produce results with small error, the number of spatial cells must be large. In smaller problems the entire mesh can fit in the computer's cache and the calculation is completed quickly. The sharp increase in CPU time near an error of 0.001 could be explained by the problem becoming too large to fit in the computer's cache. The high resolution method does not need such large meshes to get small errors and does not demonstrate this effect. In comparison, for relatively large errors the first-order method is faster than the high resolution methods. The first-order method ceases to perform better at an error of 0.001 where the quasi-linear method finishes faster. In all cases tested the fully nonlinear method is the slowest for a given error.

To demonstrate the effectiveness of the quasi-linear implicit high resolution Riemann solver, we will compare some P_n results with implicit Monte Carlo (IMC) results on the "M" problem from Sec. 8.4. By comparing results from an IMC computation [79, 80] with the implicit Riemann solver we can appraise the effectiveness of the method with large time steps. In this problem of without temperature feedback, the IMC method becomes exact up to statistical error.

In Fig. 9.17 results for the scalar flux using the quasi-linear method and the P_7 approximation are compared with IMC results. In these figures the color scale is proportional to $\log_{10} \phi$ and limited to 3.5 orders of magnitude. With a CFL number of 5, the P_7 solution compares well with the IMC solution. As the time step size is increased

the features of the solution are blurred, as we see in the results from $\Delta t = 0.25$ (CFL = 25). This spreading of the solution can be explained using the model equation for this method (c.f. Eq. (5.11)). The larger time step introduces larger dissipation (proportional to Δt^2) and this effect can be seen by comparing the P_7 solutions with different time steps.




A telling comparison can be obtained by looking at the scalar flux along the lineout at $x = 0$ in Fig. 9.18. We see the effect of the larger time step is to make the numerical solution lower in the middle of the problem. This result of the damping in the model equation is quite apparent with the solution having $\Delta t = 1$. Here we see the maximum time integration error for this situation because that solution was arrived at in only one time step. Finally, in this region and time in the calculation, the P_7 solution is good approximation to the true solution, even when computed using a CFL number of 25.

9.3 Thermal Transport Results

9.3.1 One-dimensional problems

Now we look at problems where there is a material temperature providing nonlinear feedback to the radiation intensity. The first problem we will examine is the Su-Olson finite source problem from Sec. 8.1.3, a problem thermal transport with $C_v \propto T^3$. The P_1 analytic solution was given in that section and we could use this solution to verify that our method is converging to the P_1 solution.

Fig. 9.19 shows the numerical solution and analytic P_1 solutions to this problem. In this figure the radiation flux (see Eq. (8.3)) is denoted by F_z . The analytic solution has a kink (i.e. discontinuity in the derivative of E) at $z \approx 0.077$. This kink is due to the edge of the source at $z = -0.5$ propagating in the positive z direction. This kink is an artifact of the sharp transient created when the source was turned on at $t = 0$. To the right of

(a) P_7 $\Delta x = 0.01$, $\Delta t = 0.05$ (CFL = 5)(b) P_7 $\Delta x = 0.01$, $\Delta t = 0.25$ (CFL = 25)

(c) IMC

Figure 9.17: Results for the “M” problem at $t = 1$ after the source is turned on.




Figure 9.18: The scalar flux for the “M” problem along $x = 0$ at $t = 1$ after the source is turned on.




Figure 9.19: The numerical energy density and radiation flux, F_z , at $ct = 1$ compared with the P_1 analytic solution for the Su-Olson benchmark problem.

this kink the material does not know that there is a source going back to $z = -0.5$; in hydrodynamics we might say that these points are not sonically connected to the source edge. There is a similar, though less noticeable kink appears at $z \approx 1.077$. Such kinks are difficult to capture using a numerical method and this appears in our numerical solution. Beyond this expected behavior near a kink, the numerical solution shows a jump at the edge of the source ($z = 0.5$). This jump is not present in the analytic solution; this numerical artifact is a result of the dissipation added to the numerical method becomes zero in the I_0 equation. Consider the discrete P_1 equations in one dimension in the free streaming limit with first order spatial differencing

$$\frac{1}{c} \frac{I_{0,i}^{n+1} - I_{0,i}^n}{\Delta t} + \frac{1}{2\Delta z\sqrt{3}}(I_{1,i+1}^{n+1} - I_{1,i-1}^{n+1}) - \frac{1}{2\Delta z\sqrt{3}}(I_{0,i+1}^{n+1} - 2I_{0,i}^{n+1} + I_{1,i-1}^{n+1}) = 0 \quad (9.1)$$

$$\frac{1}{c} \frac{I_{1,i}^{n+1} - I_{1,i}^n}{\Delta t} + \frac{1}{2\Delta z\sqrt{3}}(I_{0,i+1}^{n+1} - I_{0,i-1}^{n+1}) - \frac{1}{2\Delta z\sqrt{3}}(I_{0,i+1}^{n+1} - 2I_{0,i}^{n+1} + I_{1,i-1}^{n+1}) = 0. \quad (9.2)$$

In Eq. (9.1) the dissipation is proportional to the second derivative of I_0 . At the edge of the source there is an inflection point and hence no dissipation. If this happened at a place where the radiation flux (which is proportional to I_1) were smooth there would be no consequence from the loss of dissipation. However, at the edge of the source there is a cusp in the radiation flux and a jump in the flow between cells in the I_0 equation. This numerical artifact does go away, but slowly. Fig. 9.20 shows the converge to the analytic solution as the grid is refined. Up to a point, the convergence as N_z increases is greater than one. We would not expect full second order convergence in the problem due to the kinks in the solution. As the mesh gets further refined the convergence rate slackens to less than one. In this region we are adding more and more spatial cells to remove the jump at the source edge when only points near the jump affect the solution.

Finally, this problem shows the benefit of having analytic solutions to compare to. If we did not have this analytic solution, we would likely think that this behavior at the




Figure 9.20: Logarithm of the Error between high resolution solution and analytic solution for the Su-Olson benchmark as a function of the logarithm of the number of grid points.

edge of the source were actually in the analytic solution.

Using this same problem we can examine how the P_n equations limit to transport. This is done by comparing different P_n approximations with fine spatial grids and time steps ($\Delta x = 0.015$ and $\Delta t = 0.01$) to the analytic transport solution found by Su and Olson [73]. Figs. 9.21 and 9.22 show such comparisons; therein we can see that P_7 is sufficient to capture the transport solution in this problem. We also see that P_1 does a poor job, yet P_3 performs noticeably better than P_1 . In Fig. 9.22 we can see the effect of the slow P_1 wave speed. P_1 only allows information to propagate at $c/\sqrt{3}$ and on the logarithmic scale we can see that P_1 does not propagate information far enough into the problem.

The next one-dimensional problem we will look at is problem where the opacities depend nonlinearly on temperature. This problem was suggested by Olson et al. and




Figure 9.21: Comparison of several P_n solutions with the transport solution for the Su-Olson benchmark.




Figure 9.22: Comparison of several P_n solutions with the transport solution for the Su-Olson benchmark on a logarithmic scale.

has a slab of purely absorbing material irradiated on the left boundary by an isotropic light source [32]. The heat capacity of the material is a constant function of temperature and the opacities are governed by

$$\sigma_a = \sigma_0 T^{-p}. \quad (9.3)$$

Olson et al. only examined the case of $p = 3$. This is actually a special case as it makes the material temperature time scale independent of temperature. This time scale, τ , is the e folding time for T^4 ,

$$\tau = \frac{C_v}{4ac\sigma_a T^3}. \quad (9.4)$$

When $p = 3$, this time scale is constant because the T^3 in σ_a cancels the T cubed term. The problem is still nonlinear, but the nonlinearity only appears in the transport equation opacity and the temperature equation is linear [32]. We will look at the $p = 3$ case and also the harder $p = 5$ case to be assured that our method can handle such problems that are nonlinear in opacity, temperature, and radiation.

P_7 results for this problem with $p = 3$ and 5 are shown in Figs. 9.23 and 9.24. In these figures the dashed lines are the material temperature and the solid lines are the radiation temperature ($T_r = \sqrt[4]{E/a}$). The initial material temperature was set to $T = 10^{-1.25}$ as in Ref. 32. P_7 was chosen because it performed well in the problem with $C_v \propto T^{-3}$ and in this new problem the additional wrinkle is the nonlinearity, which should not contribute to additional angular complexity. Also, in previous solutions of this problem it was noted that even P_3 provided “a quite adequate solution” [32]. In Fig. 9.23 the effect of different time step and cell sizes on the solution is portrayed. Early in time a small time step and cell size are required to accurately capture the front of the wave. The largest time step used placed the wave front slightly too far from the boundary source for all times except $ct = 7$. In this case the time step is about the order of the time scale




Figure 9.23: P_7 solutions to the variable opacity test problem with $p = 3$; the dashed lines are the material temperature and the solid lines are the radiation temperature.




Figure 9.24: P_7 solutions to the variable opacity test problem with $p = 5$; the dashed lines are the material temperature and the solid lines are the radiation temperature.

of the material heating, approximately $\tau = 1$ s. Despite the time discretization errors of this time step, the method remained stable and produced reasonable results even for this unreasonable time step. The “medium” sized time step ($\Delta t = 10^{-2}$) solution in Fig. 9.23 is nearly perfectly coincident with the fine grid/small time step solution.

For the case of $p = 5$, as seen in Fig. 9.24, a finer spatial grid was used because the mean free path of a photon is much smaller in this problem. In this problem the time scale ranges from about 0.5 in the hot regions all the way down to 5×10^{-7} where the material is cold. Our solutions are converged in terms of time integration errors despite not resolving the very fastest time scale in the cold material. This is probably a benefit of treating the temperature terms (and thereby the opacities) nonlinearly. Also, the shoulder of the wavefront is where the character of the problem changes from streaming dominated to absorption dominated and the temperature below the shoulder is higher; in this region of the problem the material time scale is resolved.

9.3.2 Two-dimensional results

The first test in two dimensions we will look at the P_1 line source problem in medium with $C_v \propto T^3$. This problem has an analytic solution as derived in Sec. 8.2.1. Although this problem is one-dimensional in the distance from the line, ρ , we can test our two-dimensional code by using a Cartesian grid to model the line. The two-dimensional solution is shown in Fig. 9.25 and the solutions are compared to the analytic solution in Figs. 9.26 and 9.27. At $t = 5$ the numerical method is having a difficult time capturing the analytic solution. This is due to the fact that the numerical method is trying to resolve the delta functions containing the uncollided particles. At this point in time there are still a moderate number of uncollided photons and the numerical solution is smoothing out this delta function, thereby corrupting the solution containing




Figure 9.25: The energy density from the P_1 solution to the thermal line source problem at time $t = 5$.




Figure 9.26: Comparison of energy density from P_1 numerical and analytic solutions at $t = 5$.




Figure 9.27: Comparison of energy density from P_1 numerical and analytic solutions at $t = 10$.

collided photons. Also, at this point in time there is an artifact from the approximation of having temperature be constant in a spatial cell. This appears as the bump near $\rho = 0$. The temperature is very sharp there as the initial line source of radiation quickly turns into a localized temperature spike. This temperature spike is difficult to resolve. This spike cannot conduct away, but must instead radiate. Despite these difficulties early in time, the numerical solution captures the analytic solution very well later at $t = 10$ in Fig. 9.27.

The wire problem detailed in Sec. 8.4 is a good indicator of how both our method and the P_n approach in general will work on realistic problems with materials heating each other across a vacuum. This problem uses a constant heat capacity, making the temperature feedback nonlinear. This particular instance of the problem had the heated block's material fed 300 eV over one nanosecond. In Fig. 9.28 energy density from different orders of P_n are shown. For this figure, the radiation temperature was calculated





(a) P_1 $N_x = 60$ $\Delta t = 0.1$ ns(b) P_3 $N_x = 100$ $\Delta t = 0.1$ ns(c) P_7 $N_x = 60$ $\Delta t = 0.1$ ns(d) P_{11} $N_x = 40$ $\Delta t = 0.05$ ns

Figure 9.28: Radiation temperature (eV) at 2 nanoseconds for the mock wires problem using different P_n orders. All problems were run on a Cartesian mesh with $N_x = N_z$.

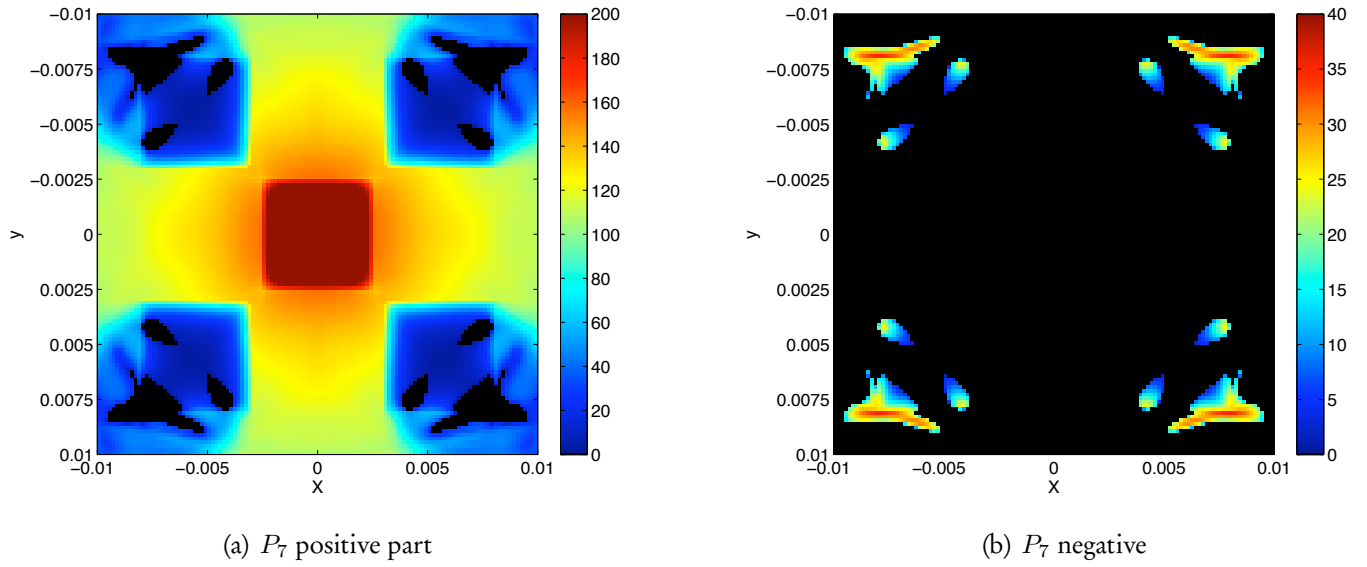


Figure 9.29: Comparison of the positive and negative parts of the radiation temperature for the P_7 solution at 2 nanoseconds.




Figure 9.30: Material temperatures (eV) at 10 nanoseconds for P_5 and P_7 . The colorscale is such that any temperature -1 eV or below is black.

to be

$$T_r = \text{sign}(E) \sqrt[4]{\frac{|E|}{a}}, \quad (9.5)$$

because the numerical solutions had negative energy densities. This figure has any radiation temperature zero or less set to be black. This will give us an estimate of how well a given order can capture the shadow behind the cold blocks. The P_1 solution has no shadow behind the cold block, showing that in P_1 photons unphysically wrap around the cold block. P_3 shows this phenomenon as well but to a lesser extent; a shadow does exist in this solution. Both the P_7 and the P_{11} solutions have a shadow behind the cold block. With the P_7 solution there is some light at temperature of 60 – 89 eV hitting the back of the block. In Fig. 9.28 all of the solutions with $n > 1$ also have a negative part to the radiation temperature in the shadow region. P_1 has no negatives because it has no shadow. The magnitude of these unphysical negatives goes down as n is increased. The P_3 solution has a maximum negative of 63 eV while the P_{11} result has a maximum of 31 eV. The separated positive and negative parts of the P_7 solution are shown in Fig. 9.29.

The real importance of the negative radiation energy density is how it affects the material temperatures. If the radiation does not cause the material temperature to go negative, then the unphysical nature of the negatives is more of a quirk than a serious drawback. The material temperatures are shown in Fig. 9.30. In this figure both P_7 and P_5 have regions of negative material temperature. The P_7 minimum temperature is about 2 eV and the P_5 negative is almost 24 eV.

Finally, it should be pointed out that in all of these figures (Figs. 9.28 - 9.30) implicit time stepping was heavily utilized; every solution had the photons travelling across the problem domain multiple times in a time step. To justify taking such large time steps we can look at the effect of large time steps on the solution. A study was done on an instance of the wires problem where the hot block was started at 300 eV rather than




Figure 9.31: Effect of time step size on the P_7 radiation energy density (J/m^3) at 8.5 ns.




Figure 9.32: Effect of time step size on the P_7 material temperature (eV) at 8.5 ns.

ramped up. This is a more difficult problem to simulate because of the large difference in radiation and material temperatures in the beginning of the problem. Figs. 9.31 and 9.32 show the effect of time step size on the P_7 solution with $N_x = N_z = 60$. These figures show the lineout along the diagonal from the center of the heated block through the cold block and into the shadow region. In this problem the time scale of the material temperature was as short as about 0.1 ns. These figures show that both the radiation and material temperature fields are converged in time when time steps capture this material temperature time scale are taken. This is the case for the radiation field even when the time scale of radiation propagation is greater than 100 times shorter than the time step.

The duct problem with temperature feedback is shown in Figs. 9.33 - 9.34. This problem, as detailed in Sec. 8.4, has a 300 eV isotropic radiation source entering the middle leg of an “M” shaped duct. This nonlinear version of the problem produces negative energy densities whereas the purely scattering problem did not. Fig. 9.33 shows the effect of different P_n approximations on the solution at an early time. The P_3 solution has too much energy turning the corner of the duct; the P_7 and P_{11} solutions show less of this effect and display a cone of radiation near the top of the duct. All solutions demonstrate wave effects. In the P_3 solution there is a local maximum at the top of the middle leg of the duct. The higher order in angle solutions also demonstrate wave effects, yet these appear in the outer legs of the duct. These wave effects are less evident in the solution at later times as seen in Fig. 9.34. At later times we also notice how P_3 creates a “mushroom” of photons near the top of the duct while the P_7 solution maintains a less rounded shape. In terms of material heating, the P_3 solution has significant artificial heating of the outer legs of the duct. The P_7 solution shows much less of this effect. This difference could be important in a problem where the heating of the wall affected some other physical process (e.g. in a radiation-hydrodynamics simulation). Finally,

this problem shows the difficulty of multi-dimensional transport simulations. The P_{11} solution was limited in the spatial resolution possible by the amount of memory in the computer because the method has only been implemented for serial machines. The P_{11} solution required approximately 150,000 unknowns and overwhelmed the capacity of a single machine.

9.4 High Resolution Time Integration

The final results that we shall examine deal with the use of the high resolution time integration scheme. These results will show why such an approach is necessary, and illustrate the behavior of the method. First, we will look at some P_1 results from a one-dimensional problem. In this problem there is a vacuum with $c = 1$ where initially the scalar intensity, $\phi = E/c$, and the radiation flux, F_Z (see Eq. (8.3)) are given by a square pulse

$$\phi(z, 0) = F_z(z, 0) = 1 \quad \text{for } z \in [0.25, 0.75]. \quad (9.6)$$

The solution for the scalar flux in this problem is a traveling square pulse moving with velocity $1/\sqrt{3}$. In Fig. 9.35 we see some numerical results for this problem at time $t = 1$. In all the numerical results there are 200 cells in the z direction and the time step size is $t = 0.125$. The Crank-Nicolson (trapezoidal method) has obvious issues with artificial oscillations around the sharp edges of the peak, whereas the backward Euler method markedly smooths out the numerical solution. The high resolution local θ method does a better job of capturing the exact solution without numerical oscillations.

This method was also explored using a two-dimensional square pulse in vacuum. The pulse was isotropic (i.e. $F_z = F_x = 0$) and of the form

$$E_r(x, z, 0) = 1 \quad \text{for } z \in [-0.5, 0.5] \text{ and } x \in [-0.5, 0.5]. \quad (9.7)$$




Figure 9.33: Radiation and material temperature (eV) at 0.9 ns for the nonlinear duct problem using different P_n orders. All problems were run on a Cartesian mesh with a time step of $\Delta t = 0.05$ ns.




Figure 9.34: Radiation and material temperature (eV) at 6 ns for the nonlinear duct problem from P_3 and P_7 .







Figure 9.35: Results from the 1D square pulse problem at $t = 1$ seconds and $\Delta t = 0.125$ and $\text{CFL} = 12.5$ for different methods compared with the exact solution.

To compare results with different time step sizes a reference solution using the Crank-Nicolson method and a fine time step was computed. Fig. 9.36 shows results from the different methods with a time step yielding $\text{CFL} \approx 7.5$. The Crank-Nicolson method shows oscillations in the energy density and gives the solution the likeness of a Hopi woven cloth. The backward Euler method did the opposite and smeared out the effect of the corners of the square. The high resolution method was much closer to the reference solution than the backward Euler method and avoided the oscillations of the Crank-Nicolson method. The next two figures, Figs. 9.37 and 9.38, show the solution along the line $x = 0$. The oscillations from the Crank-Nicolson method mar the solution with time step of $\Delta t = 0.5$ and the solution in the middle of the problem goes far below the reference solution. These oscillations go away in the solution with the smaller time step in Fig. 9.38. This smaller time step has $\text{CFL} = 3$ which is beyond the TVD limit for the Crank-Nicolson method, although no oscillations appear. The backward Euler


solution is too smooth with the larger time step and does not capture the peaks and valleys of the reference solution with a smaller time step. Finally, the high resolution method does better than the Crank-Nicolson method in the large time step case in terms of not having any oscillations. However, the wave front is further from the peak of the reference solution in the high resolution method. In the small time step solution, there is still a difference between the Crank-Nicolson method and the high resolution method. This is due to the fact that the Crank-Nicolson method is not guaranteed to be nonoscillatory with this time step.

(a) Reference solution: Crank-Nicolson Method $\Delta t = 0.01$ 

(b) Crank-Nicolson Method



(c) High Resolution Method



(d) Backward Euler

Figure 9.36: The radiation energy density for different methods with $\Delta t = 0.25$ compared with the reference solution at $t = 1.5$ seconds. All solutions have $\Delta x = 0.033$.




Figure 9.37: Results from the 2D square pulse problem at $t = 1.5$ seconds and $\Delta t = 0.5$ for different methods compared with the reference solution.




Figure 9.38: Results from the 2D square pulse problem at $t = 1.5$ seconds and $\Delta t = 0.1$ for different methods compared with the reference solution.

CHAPTER X

Conclusions and Future Work

10.1 Implicit time integration

One of the main goals of this work was to explore the use of implicit schemes to solve an upwinded discretization of the P_n equations. Implicit time integration is important because the radiation propagation times scale is much shorter than the material energy time scale. Also, at long times the negativities of the P_n expansion are less prevalent – as our P_1 analytic solution shows. Our implicit approach solved the fully coupled, nonlinear radiation - material energy equations. With the implicit method we were able to solve thermal radiation transport problems on the time scale of the material energy equation rather than on a time scale limited by the CFL limit. These results were the first use of an implicit Riemann method for thermal radiation transport.

An implicit method for a high resolution spatial scheme requires the solution of a system of nonlinear equations at each time step. We were able to reduce the computational cost of solving these equations by using a two-step, quasi-linear approach. The quasi-linear approach was developed out of the fact that the underlying P_n equations are linear in radiation propagation and that the nonlinearities were added to get a high resolution spatial scheme. By exploiting similarities between first order and high resolution solutions we were able to reduce the computational cost of adding a high resolution

method without sacrificing accuracy. We showed through analysis and numerical results that using the first order solution to determine the stencil for the high resolution scheme is sufficient to suppress oscillations and get high order error convergence. This might lead to the use of such a quasi-linear approach to add non-oscillatory spatial schemes to time dependent S_n methods. Also, other fields that solve linear hyperbolic problems could benefit from this discovery.

10.1.1 High resolution time integration

Using the quasi-linear approach we were able to demonstrate the first radiation transport results from the nascent area of high resolution implicit time integration. The preliminary results shown in this thesis suggest this could be an area of future research. A firm conclusion we can draw from our analysis is that the Crank-Nicolson method should be used when the time step is within the TVD limit. However, beyond this limit a high resolution method needs to be used to get higher than first order accuracy without oscillations.

10.2 Accuracy of P_n expansion

For the Su-Olson benchmark problem we were able to study the convergence of P_n to transport in the one-dimensional case. Results showed that there is a leap in accuracy in going from P_1 and P_3 . The P_7 solution captured the transport solution well. Using P_7 we then satisfactorily solved two radiation wave propagation problems where the opacity was a nonlinear function of the material temperature with time steps large compared with the mean free time of a photon.

Despite our robust solutions to these nonlinear one-dimensional problems, our results in two dimensions did show some problems with the P_n method. In two dimensions we also produced robust results with time steps large compared with the radiation time

scale. However, even at long times the negative energy densities of the P_n solutions persisted and caused the material temperature to become negative. We showed why these negatives arise in analytic solutions to the P_n equations and how this shortcoming could be explained by the P_n method having rotational invariance, linearity, and hyperbolicity. Based on this argument we can point to future work based on the P_n methods, namely to find a way to give up one of these properties to avoid negative energy densities. Partially due to the work of this thesis, another student, Tiberius Moran, is taking up this mantle by investigating the use of a nonlinear closure to eliminate the negatives in P_n .

The negativities in energy density are not necessarily a mortal blow to the possibility of using the P_n equations for production purposes. There are simple, ad hoc ways of dealing with quantities becoming unphysically negative in large code projects [4]. These can take the form of a temperature floor that does not allow the material temperature to become negative. It has yet to be determined if such a floor would cause problems in a coupled simulation using the P_n method. This option, while destroying energy conservation, would not affect the negative radiation energy densities in vacuum regions.

While there cannot be a guarantee of positivity for the P_n equations for a finite expansion, for many problems there will be an order of expansion that reduces the negatives to a negligible amount. Some important future work would be to understand how these negatives go away as the order of the expansion is increased. For example, in the thermal duct problem will P_{15} reduce the negatives to an acceptable amount? At this point little is known theoretically about the magnitude of negativities in the P_n equations as n is increased.

10.3 Future implementation issues

This leads to another area of future work: implementing our P_n method in parallel. To fully explore the behavior of P_n methods in problems of industrial importance, the size of the problem will necessitate the utilization of massively parallel architecture. To perform a simulation using a 3D method with P_{15} and 100 frequency groups would require 40,000 unknowns per spatial element. This is before introducing the extra unknowns from a finite element description. In each spatial cell there could be more unknowns than in the entire problem for an average calculation done in this thesis. Luckily, our method was implemented with this in mind and since we used the Trilinos library the task to move to parallel might not be too onerous.

Beyond the parallelization a useful future project would be to implement our method on non-orthogonal grids. The method was developed with this in mind and is extensible to such a case. Results on such grids would be the first such results from upwind P_n methods for thermal radiation transport. This work would also answer important questions regarding the method's robustness in difficult scenarios such as the Kershaw "z" mesh, high aspect ratio cells, and "bowtie" cells.

10.3.1 Diffusion limit

Satisfying a diffusion limit is the ineluctable modality¹ of a practical radiation transport scheme – methods that do not satisfy the diffusion limit are often not robust enough to handle the mesh constraints of real problems. We showed that the stan-

¹This phrase is borrowed from James Joyce's opening of the Proteus chapter in *Ulysses* [81] – "The ineluctable modality of the visible" – this uses the first definition of modality given by the Oxford English Dictionary (OED) [82]: "Those aspects of a thing which relate to its mode, or manner or state of being, as distinct from its substance or identity; the non-essential aspect or attributes of a concept or entity. Also: a particular quality or attribute denoting the mode or manner of being of something." In the context of diffusion limits we evoke the sixth definition, "A symptom, procedure, or other factor which aids diagnosis."

standard Riemann solver approach does not limit to a diffusion equation in the case of thick cells and a collision dominated medium. The failure in the diffusion limit is a result of the dissipation added by the Riemann solver. By scaling away this dissipation we were able to show how to correct this situation. However, our scaling resulted in a stencil that decoupled the solution in neighboring cells and contributed to difficulties at material interfaces. These effects would be more pronounced in 2D and for this reason, the diffusion correction was not pursued in multiple dimensions. Furthermore, our correction has issues with conservation at the interface of a diffusive region. Future research into numerical methods for P_n will have to address these issues for the P_n method to be a viable method for large scale thermal radiation transport simulations. It is likely that going to a discontinuous Galerkin finite element approach will correct the method in the diffusion limit. This was necessary to make S_n methods robust in the diffusion limit [33, 83].

10.4 Benchmark solutions

Finally, we have introduced the first benchmark solutions for P_1 thermal radiation transport. These solutions will be especially important as P_n methods become more used for thermal transport. Having these exact solutions will help in the process of verification of computer codes that solve the P_n equations. These benchmarks may be especially useful because we have developed them in one and two dimensions. However, there is more work to be done in this regard. Our two-dimensional solution is for a pulsed delta function source. It would be much more useful to have a solution to a finite source problem where the source remains on for a finite time. This would be an analog of the Su-Olson slab source problem in two dimensions and would be more useful for testing 2D P_n codes as well as exploring the P_1 approximation.

10.5 Coda

While our results show that there is still important work to be done in the realms of numerics and analysis of the P_n equations, these methods still show promise. P_n methods are at sizeable disadvantage to S_n methods because the latter has such a body of research regarding analysis and implementation. We have made significant progress in the implementation of the P_n methods with the quasi-linear approach and implicit scheme. Furthermore, our analysis showed where the real problems with P_n are and suggested ways to fix them. The work of this thesis coupled with research to address the outstanding issues will make P_n methods a viable option for large scale simulation of thermal radiation transport.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] Thomas A. Brunner. *Riemann Solvers for Time-Dependent Transport Based on the Maximum Entropy and Spherical Harmonics Closures*. PhD thesis, University of Michigan, Ann Arbor, Michigan, 2000.
- [2] Thomas A. Brunner and James Paul Holloway. Two dimensional time dependent Riemann solvers for neutron transport. *J. Comput. Phys.*, 210(1):386–399, 2005.
- [3] S. Chandrasekhar. *Radiative Transfer*. Dover Publications, Inc., 1960.
- [4] Thomas A. Brunner. personal communication.
- [5] Thomas A. Brunner and James Paul Holloway. One-dimensional Riemann solvers and the maximum entropy closure. *J. Quantitative Spectroscopy and Radiative Transfer*, 69:543–566, 2001.
- [6] M. Eaton, C. C. Pain, C.R.E. de Oliveira, and A. Goddard. A high-order Riemann method for the Boltzmann transport equation. In *Nuclear Mathematical and Computational Sciences: A Century in Review; A Century Anew*, Gatlinburg, Tennessee, April 2003. American Nuclear Society.
- [7] C.C. Pain, M.D. Eaton, J. Bowsher, R.P. Smedley-Stevenson, A.P. Umpleby, C.R.E. de Oliveira, and A.J.H. Goddard. Finite element based Riemann solvers for time-dependent and steady-state radiation transport. *Transport Theory and Statistical Physics*, 32(5-7):699–712, 2003.
- [8] Karthikeyan Duraisamy, James D. Baeder, and Jian Guo Liu. Concepts and application of *time-limiters* to high resolution schemes. *Journal of Scientific Computing*, 19(1-3), 2003.
- [9] Stewart Harris. *Introduction to the Theory of the Boltzmann Equation*. Dover Publications Inc., 1971.
- [10] Kenneth M. Case and Paul F. Zweifel. *Linear Transport Theory*. Addison-Wesley, Reading, Massachusetts, 1967.
- [11] Dimitri Mihalas and Barbara Weibel-Mihalas. *Foundations of Radiation Hydrodynamics*. Dover Publications Inc., 1999.

- [12] E.E. Lewis and W.F. Miller. *Computational Methods of Neutron Transport*. John Wiley and Sons, 1984.
- [13] George I. Bell and Samuel Glasstone. *Nuclear Reactor Theory*. Robert E. Kreiger Publishing, Malabar, Florida, 1970.
- [14] J.H. Tait. *Neutron Transport Theory*. Mathematical Physics Series. American Elsevier Publishing Company, Inc., 1964.
- [15] B. Davison. *Neutron Transport Theory*. International Series of Monographs on Physics. Oxford University Press, 1957.
- [16] Gerald C. Pomraning. *The Equations of Radiation Hydrodynamics*. Pergamon Press, Oxford, 1973.
- [17] D. Salzmänn. *Atomic Physics in Hot Plasmas*. Oxford University Press, 1998.
- [18] Ya B. Zel'dovich and Yuri P. Raizer. *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena*. Dover, 2002.
- [19] Max Planck. Zur theorie des gesetzes der energieverteilung im normalspektrum. *Verhandl. Deutsch. phys. Ges.*, 2(237), 1900.
- [20] R. Paul Drake. *High-Energy-Density Physics: Fundamentals, Inertial Fusion, and Experimental Astrophysics*. Springer, 2006.
- [21] Encyclopædia Britannica. Kaspar Hauser – Encyclopædia Britannica, 2006. [Online; accessed 11-October-2006].
- [22] Ezequiel Berdichevsky. personal communication.
- [23] Jeffrey D. Densmore and Edward W. Larsen. Asymptotic equilibrium diffusion analysis of time-dependent Monte Carlo methods for grey radiative transfer. *J. Comput. Phys.*, 199:175–204, 2004.
- [24] Thomas Murray MacRobert. *Spherical harmonics; an elementary treatise on harmonic functions, with applications*. Methuen and Co., 1947.
- [25] Philip R. Wallace. *Mathematical Analysis of Physical Problems*. Dover Publications Inc., 1972.
- [26] G.C. Pomraning. A generalized P-N approximation for neutron transport problems. *Nukleonik*, 6:348–356, 1964.
- [27] Bingjing Su and G.C. Pomraning. A modified p_3 -like approximation to the transport equation. *Nucl. Sci. and Eng.*, 124:309–319, 1996.
- [28] Bingjing Su. Variable Eddington factors and flux limiters in radiative transfer. *Nucl. Sci. and Eng.*, page 281, March 2001.

- [29] Kirk A. Mathews. On the propagation of rays in discrete ordinates. *Nucl. Sci. and Eng.*, 132:155, 1999.
- [30] Edward W. Larsen and Joseph B. Keller. Asymptotic solution of neutron transport problems for small mean free paths. *Journal of Mathematical Physics*, 15(1), January 1974.
- [31] G. J. Habetler and B.J. Matkowsky. Uniform asymptotic expansions in transport theory with small mean free paths, and the diffusion approximation. *Journal of Mathematical Physics*, 16(4), April 1975.
- [32] Gordon L. Olson, Lawrence H. Auer, and Michael L. Hall. Diffusion, P_1 , and other approximate forms of radiation transport. *J. Quantitative Spectroscopy and Radiative Transfer*, 64:619–634, 2000.
- [33] R.B. Lowrie and J.E. Morel. Issues with high-resolution Godunov methods for radiation hydrodynamics. *J. Quantitative Spectroscopy and Radiative Transfer*, 69:475–489, 2001.
- [34] George Green. *Essay on the Mathematical Theory of Electricity and Magnetism*. Self published, Nottingham, 1828.
- [35] James S. Warsa, Todd A. Wareing, and Jim E. Morel. Fully consistent diffusion synthetic acceleration of linear discontinuous S_n transport discretizations on unstructured tetrahedral meshes. *Nucl. Sci. and Eng.*, 141:236–251, 2002.
- [36] Thomas A. Brunner and James Paul Holloway. Two dimensional time dependent Riemann solvers for neutron transport. In *Proceedings of the 2001 ANS International Meeting on Mathematical Methods for Nuclear Applications*, Salt Lake City, Utah, September 2001. American Nuclear Society.
- [37] Ryan McClarren, James Paul Holloway, Thomas A. Brunner, and Thomas Mehlhorn. Implicit Riemann solvers for the P_n equations. In Frank Graziani, editor, *Computational Methods in Transport*, Lake Tahoe, California, September 2005. Institute for Pure and Applied Mathematics, Springer-Verlag.
- [38] Ryan G. McClarren, James Paul Holloway, Thomas A. Brunner, and Thomas A. Mehlhorn. An implicit Riemann solver for the time-dependent P_n equations. In *International Topical Meeting on Mathematics and Computation, Supercomputing, Reactor Physics and Nuclear and Biological Applications*, Avignon, France, September 2005. American Nuclear Society.
- [39] Joseph P. Pickett, et al., editor. *The American Heritage Dictionary of the English Language*. Houghton Mifflin Company, fourth edition, 2004.
- [40] Randall J. Leveque. *Numerical Methods for Conservation Laws*. Birkhäuser Verlag, Boston, USA, 1992.

- [41] Timothy J. Barth and Dennis C. Jespersen. The design and application of upwind schemes on unstructured meshes. In *27th Aerospace Sciences Meeting*, number AIAA-89-0366, Reno, Nevada, January 1989. American Institute of Aeronautics and Astronautics.
- [42] David C. Slack, D.L. Whitaker, and Robert W. Walters. Time integration algorithms for the two-dimensional Euler equations on unstructured meshes. *American Institute of Aeronautics and Astronautics Journal*, 32(6):1158–1166, June 1994.
- [43] Russ D. Rausch, Henry T. Y. Yang, and John T. Batina. Spatial adaption procedures on unstructured meshes for accurate unsteady aerodynamic flow computation. In *AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, number AIAA-91-1106, Baltimore, Maryland, April 1991. American Institute of Aeronautics and Astronautics.
- [44] D.G. Holmes and S.D. Connell. Solution of the 2d navier-stokes equations on unstructured adaptive grids. In *AIAA 9th Computational Fluid Dynamics Conference*, number AIAA-89-1932. American Institute of Aeronautics and Astronautics, June 1989.
- [45] Ami Harten and Stanley Osher. Uniformly high-order accurate nonoscillatory schemes, I. *SIAM Journal of Numerical Analysis*, 24(2), 1987.
- [46] Xu-Dong Liu, Stanley Osher, and Tony Chan. Weighted essentially nonoscillatory schemes. *Journal of Computational Physics*, 115:200–212, 1994.
- [47] Bram van Leer. Towards the ultimate conservative difference scheme. II. Monotonicity and conservation combined in a second-order scheme. *J. Comput. Phys.*, 14:361–370, 1974.
- [48] John Bonnycastle. *An Introduction to Algebra*. W.E. Dean, New York, revised and enlarged edition, 1851.
- [49] Edwige Godlewski and Pierre-Arnaud Raviart. *Numerical Approximation of Hyperbolic Systems of Conservation Laws*. Springer, 1996.
- [50] Erwin Kreysig. *Advanced Engineering Mathematics*. John Wiley and Sons, eighth edition, 1999.
- [51] H. Huynh. Accurate cubic interpolation. *SIAM J. Numer. Anal.*, 30:57–100, 1993.
- [52] G. Hedstrom. Models of difference schemes for $u_t + u_x = 0$ by partial differential equations. *J. Comput. Phys.*, 29:969–977, 1975.
- [53] R.F. Warming and B.J. Hyett. The modified equation approach to the stability and accuracy analysis of finite-difference methods. *J. Comput. Phys.*, 14:159–179, 1974.
- [54] Kendall Atkinson. *Elementary Numerical Analysis*. John Wiley and Sons, second edition, 1993.

- [55] C.T. Kelley. *Iterative Methods for Linear and Nonlinear Equations*. Frontiers in Applied Mathematics. Society for Industrial and Applied Mathematics, Philadelphia, USA, first edition, 1995.
- [56] William Gropp, David Keyes, Lois Curfman McInnes, and M. D. Tidriri. Globalized Newton-Krylov-Schwarz algorithms and software for parallel implicit CFD. *The International Journal of High Performance Computing Applications*, 14(2):102–136, 2000.
- [57] Yousef Saad. *Iterative Methods for Sparse Linear System*. PWS Publishing Company, 1995.
- [58] Brian Guthrie, James Paul Holloway, and Bruce W. Patton. GMRES as a multi-step transport sweep accelerator. *Transport Theory and Statistical Physics*, 28(1), 1999.
- [59] Bruce W. Patton and James Paul Holloway. Application of preconditioned GMRES to the numerical solution of the neutron transport equation. *Annals of Nuclear Energy*, 29:109–136, 2002.
- [60] Yousef Saad. Ilut: a dual threshold incomplete lu factorization. *Numerical Linear Algebra with Applications*, 1(4):387–402, 1994.
- [61] James S. Warsa, Todd A. Wareing, and Jim E. Morel. Krylov iterative methods and the degraded effectiveness of diffusion synthetic acceleration for multidimensional S_n calculations in problems with material discontinuities. *Nucl. Sci. and Eng.*, 147:218–248, 2004.
- [62] William L. Briggs, Van Emden Henson, and Steve F. McCormick. *A Multigrid Tutorial*. Society for Industrial and Applied Mathematics, 2000.
- [63] Michael Heroux, Roscoe Bartlett, Vicki Howle Robert Hoekstra, Jonathan Hu, Tamara Kolda, Richard Lehoucq, Kevin Long, Roger Pawlowski, Eric Phipps, Andrew Salinger, Heidi Thornquist, Ray Tuminaro, James Willenbring, and Alan Williams. An Overview of Trilinos. Technical Report SAND2003-2927, Sandia National Laboratories, 2003.
- [64] Bryan A. Garner. *The Oxford Dictionary of American Usage and Style*. Oxford University Press, 2000.
- [65] C. Lingus. Analytical test case’s [sic] for neutron and radiation transport codes. In *Second Conference on Transport Theory*, pages 655–659, Los Alamos, New Mexico, January 1971. United States Atomic Energy Commission - Division of Technical Information.
- [66] Heath L. Hanshaw. *The Multidimensional Multiple Balance Method for S_n Radiation Transport*. PhD thesis, University of Michigan, Ann Arbor, 2005.

- [67] Edward W. Larsen. Infinite medium solutions of the transport equation, Sn discretization schemes, and the diffusion approximation. *Transport Theory and Statistical Physics*, 32(5-7):623–643, 2003.
- [68] Gerald B. Folland. *Fourier Analysis and Its Applications*. Brooks/Cole Publishing Company, 1992.
- [69] G.C. Pomraning. The non-equilibrium Marshak wave problem. *J. Quantitative Spectroscopy and Radiative Transfer*, 21:249–261, 1979.
- [70] Bingjing Su and Gordon L. Olson. Benchmark results for the non-equilibrium Marshak diffusion problem. *J. Quantitative Spectroscopy and Radiative Transfer*, 56(3):337–351, 1996. Diffusion Benchmark.
- [71] I. R. Shokair and G. C. Pomraning. Boundary conditions for differential approximations. *J. Quantitative Spectroscopy and Radiative Transfer*, 25:325, 1981.
- [72] B.D. Ganapol and G. C. Pomraning. The non-equilibrium marshak wave problem: A transport theory solution. *J. Quantitative Spectroscopy and Radiative Transfer*, 29:311, 1983.
- [73] Bingjing Su and Gordon L. Olson. An analytic benchmark for non-equilibrium radiative transfer in an isotropically scattering medium. *J. Quantitative Spectroscopy and Radiative Transfer*, 24(13):1035–1055, 1997.
- [74] Ruel V. Churchill and James Ward Brown. *Complex Variables and Applications*. McGraw Hill, New York, 5 edition, 1990.
- [75] Milton Abramowitz and Irene A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York, ninth dover printing, tenth gpo printing edition, 1964.
- [76] Barry D. Ganapol. Solution of the one-group time-dependent neutron transport equation in an infinite medium by polynomial reconstruction. *Nucl. Sci. and Eng.*, 92:272–279, 1986.
- [77] N. A. Gentile. Implicit monte carlo diffusion, an acceleration method for monte carlo time-dependent radiative transfer simulations. *J. Comput. Phys.*, 172:543, 2001.
- [78] W. Reed. New difference schemes for the neutron transport equation. *Nucl. Sci. and Eng.*, 46:31–39, 1971.
- [79] N. A. Gentile, Malvin Kalos, and Thomas A. Brunner. Obtaining identical results on varying numbers of processors in domain decomposed particle Monte Carlo simulations. In Frank Graziani, editor, *Computational Methods in Transport*, Lake Tahoe, California, September 2005. Institute for Pure and Applied Mathematics, Springer-Verlag.

- [80] Thomas A. Brunner, Todd J. Urbatsch, T. M. Evans, and N. A. Gentile. Comparison of four parallel algorithms for domain decomposed implicit Monte Carlo. *J. Comput. Phys.*, 212(2):527–539, March 2005.
- [81] James Joyce. *Ulysses*. Vintage, New York, Gabler (1986) edition, 1922.
- [82] John Simpson and Edmund Weiner, editors. *Oxford English Dictionary*. Oxford University Press, second edition, 1989.
- [83] Marvin L. Adams. Discontinuous finite element transport solutions in thick diffusive problems. *Nucl. Sci. and Eng.*, 137:298–333, 2001.

Index

- A Rebours*, 1
- albedo boundary condition
 using ghost cells, 66
- angular flux, 15
 related to intensity, 15
- backward Euler method, 39
- blackbody radiation, 12
- Boltzmann equation
 Linear, 9
- Case modes, 21
- Courant-Isaacson-Rees method, 40
- cross-section, 9, 10
 differential, 9, 10
- deterministic methods, 2
- diæresis, 67
- diffusion length, 68
- diphthong, 67
- discrete ordinates methods, 3, 23, 24
- Eliot, T.S., ii
- energy density
 in terms of scalar flux, 15
 in terms of spherical harmonics, 18
 material energy density, 13
 radiation energy density, 13
- extra dissipation, 33
- Fick's law, 69, 70
- flux-limited diffusion, 3, 24, 25
- "Four Quartets", ii
- ghost cells, 60
- GMRES, 48, 49, 54, 55, 57, 58
 restarted, 55
- Godunov's Theorem, 34
- Götterdämmerung*, 76
- grid effects, 33
- Hauser, Kaspar, 15
- heat capacity, 13, 85
 T^3 form, 86
- Hopi weaving, 137
- Huysmans, J.K., 1

- Implicit Monte Carlo, 4
- intensity, 9
 - in terms of angular flux, 15
- Joyce, James, 144
- Kirchoffs law, 12
- Krylov Methods, *see* GMRES
- Legendre polynomials, 17
 - associated Legendre functions, 17
- linear solution preserving, 70
- local thermodynamic equilibrium, 11
- “M” problem, 96, 117, 133
- Mark boundary condition, 59
- Marshak boundary condition, 29, 59, 60
- material energy equation, 14
- model equation analysis, 50
- Mussorgsky, Modest, 34
- The New Yorker*, 67
- Newton’s method, 53
 - inexact, 54
 - Newton-Krylov method, 54
- NOX, 57
- opacity, *see* cross-section
- P_1 Green’s function
 - line source, 94, 126
 - plane source, 90
 - point source, 93
 - plane to point transform, 81, 93
 - point to line transform, 94
- preconditioner
 - algebraic, 56
 - incomplete LU, 56, 58
 - physics based, 56, 58
- Ragnarok, 76
- ray effects, 3, 23
- Reed’s problem, 96, 102
 - modified, 107
- reflecting boundary condition, 61
- Riemann problem, 29
- Riemann solver, 5, 29
 - intermediate diffusion limit, 78
 - LSP property, 74
 - modified dissipation, 79
 - thick diffusion limit, 76
- Ring of the Nibelung, The*, 76
- scalar flux, 15
 - related to energy density, 15
- slope limiter, 35

- harmonic mean, 35
 - minmod, 35
- specific heat, 13
- spherical harmonics, 11, 17
 - definition, 16
 - intensity expanded in, 18
 - moments of transport equation, 18
 - recursion relations, 18
- spherical harmonics methods, 4
 - avoidance of even-order expansion, 20
 - classical closure, 19
 - LSP property, 70
 - number of unknowns, 19, 20, 22
 - other closures, 21
- stochastic methods, 2
- Su-Olson benchmark, 91, 118
- telegrapher's equation, 85
- total variation, 36, 41, 42
- total variation diminishing, 36, 40, 42
- transport equation, 9
 - grey, 11
 - spherical harmonics moments of, 18
- Ulysses*, 144
- vacuum boundary condition
 - using ghost cells, 61
- Wagner, Richard, 76
- wave effects, 4
- wire heating problem, 97, 128

ABSTRACT

Spherical Harmonics Methods for Thermal Radiation Transport

by

Ryan G. McClarren

Chair: James Paul Holloway

An implicit, spherical harmonic (P_n) method for solving thermal transport problems is developed. The method uses a high resolution Riemann solver to produce an up-winded discretization. The high resolution scheme introduces nonlinearities to the radiation transport operator to avoid the creation of artificial oscillations in the solution. A quasi-linear approach to solving this nonlinear system of equations is developed. Through analysis and numerical results it is shown that the quasi-linear approach does suppress artificial oscillations and gives better than first order accuracy. The time integration methods considered are the backward Euler method and a high resolution time integration method. Also, reflecting boundary conditions for the P_n equations in three-dimensions are presented. It is shown that the standard Riemann solver is not robust in the diffusion limit. A fix is suggested that scales out the dissipation added by the Riemann solver as spatial cells become optically thick. The free-streaming limit of the P_n equations is explored and it is shown why in multiple-dimensions the P_n solutions can have negative energy densities. The Green's function for the one-dimensional P_1

thermal transport equations with $C_v \propto T^3$ is derived. The Green's function is used to create the P_1 solution to a common benchmark and to a problem of an infinite, pulsed line source. The implicit method was able to produce robust results to thermal transport problems in one and two dimensions. The implicit approach allowed the numerical method to take time steps on the longer material energy time scale rather than the speed of light time scale. In two dimensional problems the P_n solutions contained negative radiation energy densities. These negatives caused the material temperature to become negative as well.