DSA PRECONDITIONING FOR THE S_N EQUATIONS WITH STRICTLY POSITIVE SPATIAL DISCRETIZATION

A Thesis

by

DONALD EUGENE BRUSS

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

May 2012

Major Subject: Nuclear Engineering

DSA PRECONDITIONING FOR THE S_N EQUATIONS WITH STRICTLY POSITIVE SPATIAL DISCRETIZATION

A Thesis

by

DONALD EUGENE BRUSS

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Approved by:

Co-Chairs of Committee,	Jim E. Morel
	Jean C. Ragusa
Committee Members,	Jean-Luc Geurmond
	Marvin L. Adams
Head of Department,	Yassin A. Hassan

May 2012

Major Subject: Nuclear Engineering

ABSTRACT

DSA Preconditioning for the S_N Equations with Strictly Positive Spatial Discretization. (May 2012)

Donald Eugene Bruss, B.S. Nuclear Engineering, Oregon State University

Co-Chairs of Advisory Committee: Dr. Jim E. Morel
Dr. Jean C. Ragusa

Preconditioners based upon sweeps and diffusion-synthetic acceleration (DSA) have been constructed and applied to the zeroth and first spatial moments of the 1-D transport equation using S_N angular discretization and a strictly positive nonlinear spatial closure (the CSZ method). The sweep preconditioner was applied using the linear discontinuous Galerkin (LD) sweep operator and the nonlinear CSZ sweep operator. DSA preconditioning was applied using the linear LD S_2 equations and the nonlinear CSZ S_2 equations. These preconditioners were applied in conjunction with a Jacobian-free Newton Krylov (JFNK) method utilizing Flexible GMRES.

The action of the Jacobian on the Krylov vector was difficult to evaluate numerically with a finite difference approximation because the angular flux spanned many orders of magnitude. The evaluation of the perturbed residual required constructing the nonlinear CSZ operators based upon the angular flux plus some perturbation. For cases in which the magnitude of the perturbation was comparable to the local angular flux, these nonlinear operators were very sensitive to the perturbation and were significantly different than the unperturbed operators. To resolve this shortcoming in the finite difference approximation, in these cases the residual evaluation was performed using nonlinear operators "frozen" at the unperturbed local ψ. This was a Newton method with a perturbation fixup. Alternatively, an entirely frozen method always performed the Jacobian evaluation using the unperturbed nonlinear operators. This frozen JFNK method was actually a Picard iteration scheme. The

perturbed Newton's method proved to be slightly less expensive than the Picard iteration scheme.

The CSZ sweep preconditioner was significantly more effective than preconditioning with the LD sweep. Furthermore, the LD sweep is always more expensive to apply than the CSZ sweep. The CSZ sweep is superior to the LD sweep as a preconditioner. The DSA preconditioners were applied in conjunction with the CSZ sweep. The nonlinear CSZ DSA preconditioner did not form a more effective preconditioner than the linear DSA preconditioner in this 1-D analysis. As it is very difficult to construct a CSZ diffusion equation in more than one dimension, it will be very beneficial if the results regarding the effectiveness of the LD DSA preconditioner are applicable to multi-dimensional problems.

ACKNOWLEDGMENTS

Special thanks are due to Dr. Morel and Dr. Ragusa for their guidance and instruction over the two year scope of this project.

Funding for this project came from the Department of Homeland Security through the SHIELD (Smuggled HEU Interdiction through Enhanced anaLysis and Detectors) project.

NOMENCLATURE

LD	Linear Discontinuous Galerkin
CSZ	Consistent Set-to-Zero
DSA	Diffusion Synthetic Acceleration
S2SA	S_2 Synthetic Acceleration
S_N	Discrete Ordinates Method
GMRES	Generalized Minimum Residual Method
FGMRES	Flexible Generalized Minimum Residual Method

TABLE OF CONTENTS

Page

Page

LIST OF FIGURES

LIST OF TABLES

TABLE Page

TABLE Page

1. INTRODUCTION

A strictly positive spatial discretization for the S_N neutron transport equations was recently developed by Maginot, Morel, and Ragusa [1, 2]. Their Consistent Set-to-Zero (CSZ) method exactly solves the zeroth and first moments of the transport equations and reduces to the finite element linear discontinuous Galerkin (LD) method when the LD solution yields a strictly positive solution. Maginot et al. demonstrated that the CSZ method is a valuable alternative to existing strictly positive discretization schemes.

Previous research with the CSZ method was performed with a source iteration solution algorithm. Such algorithms suffer from arbitrarily slow convergence in diffusive problems, severely limiting their applicability. The purpose of this work is to define preconditioning techniques that enable the CSZ equations to be efficiently solved in diffusive problems.

1.1 Past Work

Non-physical solutions containing negative angular fluxes can be obtained in 1-D problems containing optically thick cells and in multidimensional problems under a wide variety of conditions. Maginot demonstrated that the CSZ method was comparable to other strictly positive techniques in certain problems and had several advantageous characteristics not shared by alternative methods in other problems [1, 2]. These results suggest that the CSZ method is worth further investigation. An acceleration method will be required to apply the CSZ method to diffuse problems.

Diffusion Synthetic Acceleration preconditioning was developed to overcome the arbitrarily slow convergence of source iteration schemes in diffusive problems with a

---

This thesis follows the style of the Journal of Computational Physics.

scattering ratio c close to one. By approximating the error after each source iteration with the diffusion equation, DSA can reduce the spectral radius ρ of any problem from approximately the scattering ratio to $\rho \leq 0.2247c$ [3, 4]. The application of DSA to source iteration algorithms allows these methods to be effectively applied to diffusive problems.

Krylov methods can converge significantly more rapidly than source iteration methods [3]. Source iteration schemes can be recast as a preconditioner and the resultant system solved with a Krylov method such as GMRES. A source iteration scheme utilizing DSA, recast as a preconditioner and solved with GMRES, should converge rapidly for even highly diffusive problems.

Applying the strictly positive CSZ method to a problem creates a nonlinear system that must be solved with a nonlinear solution technique. Maginot's research used a Newton iteration scheme inside of the source iteration scheme. Applying acceleration to this method would require a nonlinear DSA scheme; such schemes can be difficult to make robust. If instead the Newton iteration is outside the source iteration, standard linear DSA can be applied to the method. The Newton Krylov method preconditioned with DSA should converge significantly more rapidly than fixed-point methods in highly diffusive problems.

The CSZ method with this solution algorithm requires storing the full angular flux. Each Krylov vector will also be the length of the angular flux vector. The storage costs associated with this method will be large, but it should converge significantly more rapidly than a source iteration method.

1.2 New Research

The CSZ moment equations are reformulated as a preconditioned nonlinear system to be solved with GMRES. The scalar flux solution calculated with the CSZ method is compared to the scalar flux solution generated using the LD moment equations for comparison. The scalar flux solutions of the CSZ and LD moment

equations are expected to differ in optically thick regions. The cost associated with solving the nonlinear system of CSZ equations and the cost associated with solving the linear system of LD equations are compared to assess the relative cost of each method.

Two preconditioning schemes are applied to accelerate convergence of the CSZ method. Sweep preconditioning is first applied to the moment equations. DSA preconditioning is then applied after preconditioning with the sweep. The combination of these two methods should result in a very effective preconditioning scheme.

The CSZ moment equations are preconditioned with a sweep. The sweep preconditioning can be performed using either the linear LD or nonlinear CSZ operators. Because the CSZ operators are formed at each Newton step, preconditioning with a linear sweep is no less expensive than preconditioning with the CSZ sweep. In addition, the LD sweep is an approximation of the actual sweep and may be significantly less effective than preconditioning with the CSZ sweep. Nevertheless, it is of theoretical interest to determine the effectivness of this approach. The LD sweep preconditioner is therefore applied and compared to the CSZ sweep preconditioner.

DSA preconditioning is applied to the CSZ method to accelerate convergence in diffusive problems. In slab geometry, the S_2 transport equation is equivalent to the diffusion equation, so S_2 Synthetic Acceleration (S2SA) is equivalent to DSA. As the S_2 equations can be solved using the LD or CSZ operators, S2SA can be applied using the LD or CSZ equations. These two methods are referred to as LD S2SA and CSZ S2SA.

The equivalence between the S2SA and DSA equations in one dimension makes DSA easy to implement in slab geometry. However, no such equivalence exists in multi-dimensional problems. Although LD DSA preconditioning is possible in multi-dimensional problems, the complexity of deriving the equivalent diffusion equation using the CSZ operator makes CSZ DSA preconditioning problematic. CSZ DSA preconditioning is therefore compared to LD DSA preconditioning in diffusive one

dimensional problems to determine if the more complex CSZ DSA significantly outperforms the simpler LD DSA.

The remainder of this Thesis is divided into five sections. The moment equations and CSZ equations are derived in the first section. Source Iteration and Krylov solution techniques are outlined in the second section. A description of the sweep and DSA preconditioning strategies that are investigated is found in the third section. Test problems and computation results comprise the fourth section, followed by conclusions in the fifth section.

2. THE TRANSPORT EQUATION

2.1 The Boltzmann Equation

The neutron transport equation is derived from the Boltzmann equation for rarefied gasses. It appears:

$$\frac{1}{v}\frac{\partial\psi}{\partial t} = \int_0^\infty \int_{4\pi} \sigma_s \left(E' \to E, \vec{\Omega}' \to \vec{\Omega}\right) \psi(\vec{\Omega}', E') d\Omega' dE' - \vec{\Omega}\cdot\vec{\nabla}\psi - \sigma_t\psi + Q \quad (2.1)$$

where ψ is the time- and energy-dependent angular flux, v is the neutron velocity, $Q(\vec{r}, \vec{\Omega}, E, t)$ is the inhomogeneous source term, $\sigma_s\left(E' \to E, \vec{\Omega}' \to \vec{\Omega}\right)$ is the macroscopic differential scattering cross section at position $\vec{r}$ from energy E' to E and direction $\vec{\Omega}'$ to $\vec{\Omega}$, and $\sigma_t(\vec{r}, E)$ is the macroscopic cross section as a function of position and energy. Considering only steady-state problems with a single energy group in one-dimensional Cartesian coordinates with isotropic scattering cross sections and an isotropic external source simplifies the transport equation to

$$\mu\frac{\partial\psi(x,\mu)}{\partial x} + \sigma_t(x)\psi(x,\mu) = \frac{\sigma_s(x)}{4\pi}\phi(x) + \frac{Q(x)}{4\pi}, \quad (2.2)$$

with directional cosine μ, and the angle-integrated scalar flux ϕ defined as

$$\phi(x) = 2\pi\int_{-1}^{+1}\psi(x,\mu)d\mu. \quad (2.3)$$

2.2 Discretization in Angle and Space

The Discrete Ordinates (or S_N) approximation is used to discretize the transport equation in angle. Gauss-Legendre quadrature sets of angle and weight restrict the continuous angular flux to a finite number of directions. The angle-integrated scalar flux, ϕ, is approximated by summing the N weight and angle pairs:

$$\phi(x) \approx 2\pi\int_{-1}^{1}\psi(x,\mu)\partial\mu = \sum_{d=1}^{N} w_d\psi_d(x). \quad (2.4)$$

In slab geometry, the discrete ordinates form of the transport equation is a function of discrete direction d and position x. Assuming a source term $S(x, \mu_d)$ that

contains both the inhomogeneous and scattering sources, the one-dimensional S_N transport equation is

$$\mu_d \frac{\partial \psi(x, \mu_d)}{\partial x} + \sigma_t(x)\psi(x, \mu_d) = S(x, \mu_d). \tag{2.5}$$

Spatial discretization of the problem is straightforward. The slab is divided into cells with width h_i, with the index i denoting the i^{th} cell. The center of each cell is located at x_i, and the left and right edges of each cell are located at $x_{i-1/2}$ and $x_{i+1/2}$, respectively. Material properties are constant within each cell but can vary between cells. Thus the material properties carry index i, for the material property in the i^{th} cell.

2.3 The Spatial Moment Equations

The zeroth and first spatial moment equations are generated by multiplying the transport equation (2.5) by a basis function and integrating the resultant equation over the cell. The basis functions for the zeroth and first moment equations are

$$B_0 = 1 \tag{2.6a}$$

and

$$B_1 = \frac{1}{h_i}\left(x - x_i\right), \tag{2.6b}$$

respectively. This yields the following respective equations:

$$\frac{\mu_d}{h_i}\left(\psi_{d,i+1/2} - \psi_{d,i-1/2}\right) + \sigma_{t,i}\psi_{A,d,i} = S_{A,d,i}, \tag{2.7a}$$

$$\frac{3\mu_d}{h_i}\left(\psi_{d,i+1/2} - 2\psi_{A,d,i} + \psi_{d,i-1/2}\right) + \sigma_{t,i}\psi_{X,d,i} = S_{X,d,i}, \tag{2.7b}$$

where the average angular flux $\psi_{A,d}$ and the slope of the angular flux $\psi_{X,d}$ in cell i are respectively defined by:

$$\psi_{A,d,i} = \frac{1}{h_i}\int_{x_{i-1/2}}^{x_{i+1/2}} \psi(x, \mu_d)dx, \tag{2.8a}$$

and

$$\psi_{X,d,i} = \frac{6}{(h_i)^2} \int_{x_{i-1/2}}^{x_{i+1/2}} \psi(x, \mu_d)(x - x_i)dx. \tag{2.8b}$$

For each cell i in each direction d, the two moment equations in (2.7) include three unknowns: the outflow, $\psi_{d,i\pm1/2}$ ($\psi_{d,i+1/2}$ in positive directions and $\psi_{d,i-1/2}$ in negative directions); the average angular flux $\psi_{A,d,i}$, and the slope of the angular flux in each direction, $\psi_{X,d,i}$. A third equation is required to provide closure for this set of equations. Most closures relate the outflow to the average flux and slope of the flux within the cell. An infinite number of closures are possible, but the linear discontinuous Galerkin closure is perhaps the most widely used. In the next section we examine the LD scheme and Maginot's CSZ scheme in detail.

2.3.1 Linear Discontinuous Method

The Linear Discontinuous Galerkin (LD) closure for the moment equations assumes a linear flux distribution in every cell:

$$\psi(x)_{LD} = a_{LD} + \frac{2}{h_i}(x - x_i)\, b_{LD}. \tag{2.9a}$$

The value of the angular flux ψ at the interfaces between cells is defined in the following manner. In the positive directions ($\mu_d > 0$) Eq. (2.9a) applies for $x \in \left(x_{i-1/2}, x_{i+1/2}\right]$, and at the right face of the cell,

$$\psi(x_{i+1/2}) = a_{LD} + b_{LD}. \tag{2.9b}$$

The value of ψ at the left face of the cell is upwinded from the previous cell:

$$\psi(x_{i-1/2}) = \psi(x_{(i-1)+1/2}). \tag{2.9c}$$

In the negative directions, ($\mu_d < 0$), Eq. (2.9a) is valid for $x \in \left[x_{i-1/2}, x_{i+1/2}\right)$, with the value of ψ at the left face of the cell given by

$$\psi(x_{i-1/2}) = a_{LD} - b_{LD}, \tag{2.9d}$$

and the value at the right face of the cell is again upwinded from the previous cell:

$$\psi(x_{i-1/2}) = \psi(x_{(i+1)-1/2}). \tag{2.9e}$$

The LD closure thus defines the flux distribution across each cell $x \in \left[x_{i-1/2}, x_{i+1/2}\right]$

The definitions of ψ_A and ψ_X given in (2.8) yield

$$a_{LD} = \psi_A \tag{2.10a}$$

and

$$b_{LD} = \psi_X. \tag{2.10b}$$

The moment equations (2.7) can be written entirely in terms of a_{LD} and b_{LD} given the definitions for ψ_A, ψ_X found in Eq. (2.10), the value of the angular flux on the interfaces defined in Eq. (2.9) and the Gausian quadrature weight of each discreet direction w_d. The moment equations in the positive direction, then, are written:

$$\frac{\mu_d}{h_i}(a_{LD,d,i} + b_{LD,d,i}) + \sigma_t a_{LD,d,i} = \frac{\sigma_s}{2}\sum_{d=1}^{N}(a_{LD,d,i}w_d) + q_{A,d,i} + \psi_{d,(i-1)+1/2} \tag{2.11a}$$

$$\frac{3\mu_d}{h_i}(-a_{LD,d,i} + b_{LD,d,i}) + \sigma_t b_{LD,d,i} = \frac{\sigma_s}{2}\sum_{d=1}^{N}(b_{LD,d,i}w_d) + q_{X,d,i} + \psi_{d,(i-1)+1/2} \tag{2.11b}$$

The moment equations now include two unknown variables and two equations.

2.3.2 Consistent Set-to-Zero Method

The Consistent Set-to-Zero (or CSZ) method was recently developed by Maginot, Morel, and Ragusa [1, 2]. The CSZ closure for the moment equations modifies the linear discontinuous Galerkin closure in cases in which the LD closure yields negative outflows or a negative flux within the cell. If the LD method yields a positive solution at all points within a cell, the CSZ method is equivalent to the LD method ($\psi_{csz} = \psi_{LD}$) in that cell.

In cases in which the LD definition of ψ is negative at any point within the cell, ψ_{csz} is set to zero across the negative regions. Let the distribution of the angular

flux within the cell is denoted $\widetilde{\psi}_{csz}$. The unknowns a_{LD} and b_{LD} in Eq. (2.9a) are replaced with a_{csz} and b_{csz}:

$$\widetilde{\psi}_{csz}(x) = a_{csz} + \frac{2}{h_i}(x - x_i)\, b_{csz}. \tag{2.12}$$

The CSZ angular flux, ψ_{csz}, is defined as equal to $\widetilde{\psi}_{csz}$ at all points where $\widetilde{\psi}_{csz}$ is positive and zero at all points where $\widetilde{\psi}_{csz}$ is negative.

$$\psi_{csz}(x) = \begin{cases} \widetilde{\psi}_{csz}(x) & \text{if} \quad \widetilde{\psi}_{csz}(x) \geq 0; \\ 0 & \text{if} \quad \widetilde{\psi}_{csz}(x) < 0. \end{cases} \tag{2.13}$$

With these definitions, ψ_{csz} is strictly positive across the cell.

The relationship between the linear distribution for the angular flux within the cell, ψ_{LD}, and the CSZ distribution for ψ within the cell, ψ_{CSZ}, is illustrated in Figure 2.1.

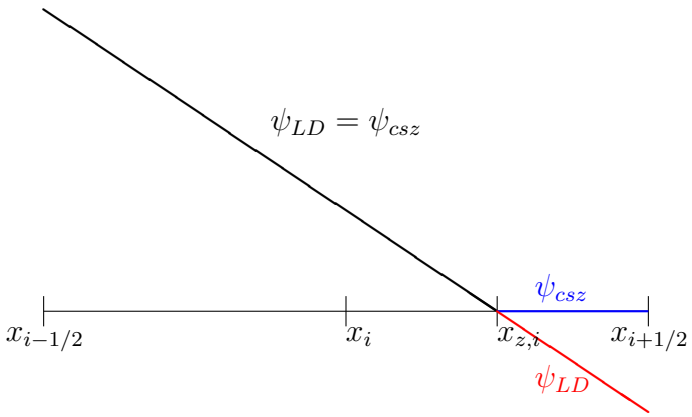

Fig. 2.1.: A graphical representation of ψ_{LD} and ψ_{csz} in a cell.

The average flux ψ_A and the slope of the flux ψ_X are determined by multiplying by the appropriate basis function and integrating across the portion of the cell where $\widetilde{\psi}_{csz}$ is positive. For example, in a cell in which ψ_{csz} is positive on the left edge of the

cell, equal to zero at some position $x_{z,i}$, and negative on the right edge of the cell as in Figure 2.1,

$$x_{z,d,i} = \frac{1}{2}\left(1 - \frac{a_{csz,d,i}}{b_{csz,d,i}}\right), \tag{2.14a}$$

$$\psi_{A,d,i} = \frac{1}{h_i}\int_{x_{d,i-1/2}}^{x_{z,d,i}} \psi_{csz,d,i}(x)dx, \tag{2.14b}$$

and

$$\psi_{X,d,i} = \frac{6}{(h_i)^2}\int_{x_{d,i-1/2}}^{x_{z,d,i}} \psi_{csz,d,i}(x)(x - x_i)dx. \tag{2.14c}$$

With these definitions for ψ_A and ψ_X, the CSZ definitions for the outflow, and the incoming flux upwinded from the previous cell, the moment equations (2.7) can be written entirely in terms of a_{csz} and b_{csz}. Note that ψ_A and ψ_X are no longer the primary unknowns in the CSZ method as they were in the LD method, where $a_{LD} = \psi_A$ and $b_{LD} = \psi_X$.

The CSZ formulation of the moment equations was shown by Maginot to eliminate negativities in one- and two- dimensional problems with significantly less cost than Exponential Discontinuous methods [1, 2]. Because the CSZ method defaults to the LD method in cells that do not have negativities in the linear method, the CSZ solution exactly matches the LD solution and LD solution cost in positive cells and only becomes more expensive as the number of cell-direction cases that include negativities increases.

3. SOLUTION TECHNIQUES FOR THE TRANSPORT EQUATION

A variety of techniques exist for solving the transport equation. Linear problems can be solved with a source iteration scheme or Krylov method. Nonlinear problems are solved with Newton's Method. The application of each of these schemes to the transport problem is outlined.

3.1 Source Iteration

The LD moment equations (2.7) derived earlier can be written in operator form:

$$\mathbf{L}\psi = \mathbf{SP}\psi + q. \tag{3.1}$$

In this form, ψ is the angular flux, the operator $\mathbf{L}$ includes the streaming and interaction terms, the operator $\mathbf{S}$ contains the scattering terms, $\mathbf{P}$ is the angle integration operator, and q is the inhomogeneous source term. The operator $\mathbf{L}$, in the positive directions ($\mu_d > 0$) in one dimensional slab geometry is defined as

$$\mathbf{L} = \begin{bmatrix} \mu_d + \sigma_t h & \mu_d \\ -3\mu_d & 3\mu_d + \sigma_t h \end{bmatrix}. \tag{3.2}$$

The operator $\mathbf{S}$ is a function of the macroscopic scattering cross section term, σ_s:

$$\mathbf{S} = \begin{bmatrix} \frac{\sigma_s}{2} \\ \frac{\sigma_s}{2} \end{bmatrix}. \tag{3.3}$$

The operator $\mathbf{P}$ generates the average and slope of the scalar flux, ϕ_A and ϕ_X, from the average and slope of the angular flux, $\psi_{A,d}$ and $\psi_{X,d}$, by summing across the d angular directions:

$$\mathbf{P}\psi = \begin{bmatrix} \sum_{d=1}^{N} w_d \psi_{A,d} \\ \sum_{d=1}^{N} w_d \psi_{X,d} \end{bmatrix}. \tag{3.4}$$

Finally, the source term q is also a vector:

$$q = \begin{bmatrix} q_A \\ q_X \end{bmatrix}. \tag{3.5}$$

A source iteration solution technique is natural to apply to the operator form of the moment equations because each of the terms on the right side of the expression is coupled in angle and the terms on the left side of the expression are uncoupled in angle. In a source iteration solution technique, the scattering and distributed source terms are lagged one iteration behind the transport term, as in Eq. (3.6):

$$\mathbf{L}\psi^{(\ell+1)} = \mathbf{SP}\psi^{(\ell)} + q. \tag{3.6}$$

If the initial guess for ψ is zero, then $\psi^{\ell+1}$, solved for with the sweep in iteration ℓ, contains all of the particles that have undergone, at most, ℓ scattering collisions.

In optically thin problems or problems with a strong absorption cross section, most particles undergo only a few collisions before being absorbed or leaking from the problem. Source iteration techniques converge rapidly for these types of problems. Particles in highly diffusive problems undergo a large number of collisions before leaking or being absorbed. As the scattering ratio c, defined as the ratio of the macroscopic scattering cross section to the macroscopic total cross section,

$$c = \frac{\sigma_s}{\sigma_t}, \tag{3.7}$$

approaches unity, the number of source iterations required to converge in optically thick systems can become arbitrarily large. It is therefore desirable to implement an acceleration scheme that is both efficient and unconditionally convergent. An *efficient* acceleration method converges with fewer total required computations than the unaccelerated method.

3.1.1 Sweep Preconditioning

The $\mathbf{L}$ operator is block lower triangular in each direction. It is therefore simple to precondition with $\mathbf{L^{-1}}$. This is called preconditioning with a sweep. After applying sweep preconditioning, equation (3.1) becomes:

$$\psi = \mathbf{L}^{-1}\left(\mathbf{SP}\psi + q\right). \tag{3.8}$$

When applied in a fixed-point iterative method, sweep preconditioning very effectively attenuates error modes in ϕ with a rapid spatial dependence. Error modes with a slow spatial dependence are not effectively managed with only sweep preconditioning [3, 4]. These types of slowly varying error modes have eigenvalues close to one in diffuse problems.

3.1.2 Diffusion Synthetic Acceleration Preconditioning

Diffusion Synthetic Acceleration (DSA) is a convergence acceleration scheme developed to compensate for the arbitrarily slow convergence of source iteration solution techniques in diffusive problems. Recall that the transport equation, in terms of the zeroth and first moment equations and in operator form, can be written as

$$\mathbf{L}\psi = \mathbf{SP}\psi + q \tag{3.9}$$

As the DSA method adds an additional step to the iterative solution process, the 'old' iterate is denoted with superscript (ℓ), the 'new' iterate from the source iteration step is denoted with $(\ell + 1/2)$, and the solution after the DSA step is denoted with superscript $(\ell + 1)$. With this notation, the source iteration form of the transport equation appears very similar to the way it appeared in (3.6):

$$\mathbf{L}\psi^{(\ell+1/2)} = \mathbf{SP}\psi^{(\ell)} + q \tag{3.10}$$

Adding and subtracting $\mathbf{S}P(\psi^{(\ell+1/2)})$ to equation (3.10) and algebraic rearrangement yields:

$$\mathbf{L}\psi^{(\ell+1/2)} = \mathbf{SP}(\psi^{(\ell)} + \psi^{(\ell+1/2)} - \psi^{(\ell+1/2)}) + q \tag{3.11}$$

Subtracting equation (3.11) from (3.9), an equation for the exact error in iteration $(\ell + 1/2)$, $\delta\psi^{(\ell+1/2)}$ is obtained:

$$\mathbf{L}\delta\psi^{(\ell+1/2)} - \mathbf{SP}\delta\psi^{(\ell+1/2)} = \mathbf{SP}(\psi^{(\ell+1/2)} - \psi^{(\ell)}) \tag{3.12}$$

Equation (3.12) exactly solves for the correction to $\psi^{(\ell+1/2)}$ to yield the exact solution, but is no simpler to solve than the original transport problem (3.9). Instead

of solving the transport equation to correct the solution ψ, a good approximation is to substitute the diffusion equation for the transport equation in Eq. 3.12 [3, 4]. This yields

$$-\frac{d}{dx}\frac{1}{3\sigma_t}\frac{d}{dx}\delta\phi^{(\ell+1/2)} + \sigma_a\delta\phi^{(\ell+1/2)} = \sigma_s\left[\phi^{(\ell+1/2)} - \phi^{(\ell)}\right]. \tag{3.13}$$

When applied to source iteration, DSA very effectively attenuates slowly varying error modes [3, 4]. Pairing a sweep and DSA effectively attenuates both rapidly-varying and slowly-varying error modes, and should therefore be a very effective pairing of preconditioners. The cost of solving the S_2 transport problem at each Krylov step is offset by greatly reducing the number of Krylov iterations required to solve the problem. Pairing DSA and sweep preconditioning should result in a very effective preconditioning scheme.

3.1.3 Consistent DSA and S_2 Synthetic Acceleration (S2SA)

DSA is generally not stable with optically thick cells with a scattering ratio $c \approx 1$ unless it is discretized in a manner consistent with the transport equation [5]. Although not guaranteed to fail if inconsistently discretized, the diffusion and transport equations generally must share the same discretization to obtain the fastest accelerated convergence. For this reason, in one dimension S_2 Synthetic Acceleration (S2SA) is often used instead of DSA. Because the S_2 and P_1 equations are analytically equivalent in one dimension, S2SA is equivalent to DSA [4]. This approach allows the existing discretization to be applied to the S_2 form of the S_N problem, resulting in identical spatial discretization. Although S2SA is somewhat more costly than DSA, the consistent discretization problem is avoided, resulting in an unconditionally convergent acceleration technique. In two- and three- dimensional problems, the S2SA and DSA methods are not equivalent [4].

The boundary conditions associated with the diffusion equation differ from the boundary conditions applied to S_2 transport. The diffusion equation is solved by

setting the ϕ to zero at an extrapolated distance from the problem boundary. The S_2 transport equations can be solved with vacuum boundary conditions, reflecting boundary conditions, or with an incident flux at the problem surface. Therefore even in one dimension S2SA is not exactly equivalent to solving the diffusion problem, but S2SA is approximately equal to DSA in one dimension. S2SA and DSA will be used synonymously throughout this paper.

3.1.4 Fourier Analysis of DSA

Fourier analysis is used to determine the spectral radius, ρ, of the source iteration solution technique for a model infinite-medium problem. The spectral radius characterizes the convergence rate of an iteration scheme, and is defined as the reduction in error between consecutive iterations:

$$\rho = \lim_{\ell \to \infty} \frac{\left\| \delta\psi^{(\ell+1)} \right\|}{\left\| \delta\psi^{(\ell)} \right\|}. \tag{3.14}$$

Optically thick problems solved with source iteration methods have a spectral radius approximately equal to the scattering ratio c [3]. As the optical thickness of a problem decreases, the leakage rate increases and the spectral radius decreases. The application of DSA preconditioning to source iteration problems reduces the spectral radius to $\rho \leq 0.2247c$ as the number of discrete angle directions N increases, an improvement of *at least* a factor of four for sufficiently large N [4, 5]. As DSA and S2SA are algebraically equivalent in one dimension, S2SA converges with the same spectral radius as DSA [4].

3.2 Krylov Methods

Krylov subspace methods are a class of iterative techniques for solving general linear systems. Effective source iteration methods can be used to define effective preconditioners that are easily applied to a Krylov solution technique, making Krylov meth-

ods a natural choice for accelerating source iteration problems. Numerous Krylov methods have been developed, but the basic Krylov scheme is common to all of them [6]. In general, Krylov methods can significantly accelerate the convergence of source iteration schemes [3].

3.2.1 The Krylov Subspace Method

Krylov methods are used solve problems of the form

$$\mathbf{A}\vec{x} = \vec{b}, \tag{3.15}$$

where $\mathbf{A}$ is an $N \times N$ matrix, $\vec{x}$ is a vector with N elements, and $\vec{b}$ is a source vector also of length N. The Krylov subspace of dimension m is denoted $K_m(\mathbf{A}, b)$ and has a basis made from the matrix $\mathbf{A}$ and the source vector $\vec{b}$ [3, 6]:

$$K_m(\mathbf{A}, \vec{b}) = span\{\vec{b}, \mathbf{A}\vec{b}, \mathbf{A}^2\vec{b}, ..., \mathbf{A}^{m-1}\vec{b}\}. \tag{3.16}$$

These basis vectors are labeled Krylov vectors. The solution to (3.15) is constructed as a linear combination of the Krylov vectors. It is simple to prove that the solution $\vec{x}$ to (3.15) can be constructed as a linear combination of the Krylov vectors [3]. First, the minimum polynomial of $\mathbf{A}$, $P_d(\mathbf{A})$, is defined as

$$P_d(\mathbf{A}) = a_0 + a_1\mathbf{A}^1 + a_2\mathbf{A}^2 + ... + a_d\mathbf{A}^d = 0. \tag{3.17}$$

Assuming that the matrix $\mathbf{A}$ is nonsingular, the polynomial can be scaled so that $a_0 = 1$. With further rearrangement,

$$I = -\left(a_1 I + a_2\mathbf{A} + ... + a_d\mathbf{A}^{d-1}\right)\mathbf{A}. \tag{3.18}$$

By multiplying from the right by $\mathbf{A}^{-1}$, equation (3.19) is formed,

$$\mathbf{A}^{-1} = -\sum_{i=0}^{d-1} a_{i+1}\mathbf{A}^i, \tag{3.19}$$

and equation (3.15) is solved with this definition of $\mathbf{A}^{-1}$:

$$\vec{x} = -\sum_{i=0}^{d-1} a_{i+1}\mathbf{A}^i\vec{b}. \tag{3.20}$$

It is clear that the solution to (3.15) is a linear combination of the Krylov vectors in $K_d(\mathbf{A}, \vec{b})$. The manner in which the polynomial $P_d(\mathbf{A})$ is determined depends upon the implementation of each particular Krylov scheme.

3.2.2 The Generalized Minimum Residual Method

The Generalized Minimum RESidual method, GMRES, was developed by Y. Saad and M.H. Schultz in 1986 [7]. GMRES differs from the standard Krylov method described previously in two main fashions. First, because the Krylov vectors are nearly linearly dependent, GMRES uses the Arnoldi process to generate an orthogonal basis of the Krylov subspace [3, 7]. As the solution to equation (3.15) lies within the Krylov subspace, it can obviously still be constructed by a linear combination of these orthogonal basis functions.

The second notable characteristic about GMRES is that the polynomial $P_d(\mathbf{A})$ is determined by solving a least-squares problem for the residual [7]. The solution for $\vec{x}$ is computed at each Krylov iteration m based upon the m basis functions that define the Krylov subspace and the solution to the least-squares problem, $P_d(\mathbf{A})$. The residual,

$$f(\vec{x}_m) = \vec{b} - \mathbf{A}\vec{x}_m, \tag{3.21}$$

is computed after each Krylov iteration until some (user-defined) convergence criteria is fulfilled.

3.2.3 Flexible GMRES

Flexible GMRES, or FGMRES, was developed by Youcef Saad in 1993 [8]. Flexible GMRES was designed for use in conjunction with right preconditioning methods

with the recognition that it is often beneficial to vary the preconditioner within the Krylov iterations. This results in an iteration matrix that varies between Krylov steps. This becomes necessary when the CSZ method is solved with a Jacobian-free method, as is explained in the section on Jacobian-free Newton Krylov methods.

FGMRES retains the positive characteristics of GMRES, including the minimization of the norm of the residual, and requires only a few extra computations per Krylov step if right preconditioning is not applied [8]. However, the memory requirement doubles because each Krylov vector is stored in its original and preconditioned state. A problem that is memory intensive to solve with GMRES requires twice as much memory to solve with FGMRES.

3.2.4 Krylov Form of the Transport Problem

Problems that have been structured in a form suitable for solution with a source iteration scheme can be easily reformatted for solution with a Krylov method. Recall that the source iteration form of the moment equations appears

$$\mathbf{L}\psi^{(\ell+1)} = \mathbf{S}P\psi^{(\ell)} + q \tag{3.22}$$

The Krylov method drops the iteration indexing as used in equation (3.22), above. The transport operator $\mathbf{L}$ is inverted and multiplied across the problem with a sweep. After minor algebraic rearrangement, the problem is in the form of equation (3.15),

$$\left(\mathbf{I} - \mathbf{L}^{-1}\mathbf{SP}\right)\psi = \mathbf{L}^{-1}q, \tag{3.23}$$

and can be solved with the application of a Krylov method.

3.2.5 DSA Applied to the Krylov Method

Acceleration techniques applied to source iteration methods are also effective as preconditioners for a Krylov solution technique. Diffusion Synthetic Acceleration can

be applied to the Krylov form of the problem in the same manner as it was applied to the source iteration problem. As when applying DSA to the source iteration problem, the the S_2 transport problem is used to approximate the difference between the exact solution and the solution at the current iteration.

3.2.6 DSA Degradation over Material Discontinuities

Warsa, Wareing, and Morel showed that DSA applied to source iteration suffers significant degradation when applied to multi-dimensional problems with strong material discontinuities [9]. The degradation of the spectral radius occurs even when fully consistent DSA is applied to the source iteration scheme and appears to be a shortcoming of the DSA method itself. Therefore, acceleration of source iteration problems with DSA is not guaranteed to be an efficient solution technique for all problems.

The application of DSA to a Krylov solution scheme for problems with material discontinuities avoids this degradation to a large extent. The source iteration scheme with DSA applied was found to be of comparable efficiency to unaccelerated GMRES. DSA preconditioning significantly improved the efficiency of GMRES, even when only partially-consistent DSA was applied [9]. The authors also noted that because source iteration methods can easily be replaced with a Krylov method, preconditioning a Krylov method with partially-consistent DSA is a very attractive solution technique for many transport problems that are not easily solved with source iteration.

3.3 Newton's Method

Newton's method is used to solve the nonlinear CSZ equations. At each Newton step, a linear technique (like the preconditioned Krylov method previously discussed) is employed to solve for the Newton step. The linear LD equations will be solved exactly with this linear solve, so applying Newton's method to a linear problem

results in the correct solution after a single iteration. The nonlinear CSZ equations will require multiple Newton iterations to converge.

Newton's method solves for the roots of a problem. In order to transform the moment equations into a problem in which the root is the angular flux that exactly satisfies the equations, the moment equations are written residual form. This is a trivial rearrangement of the equations that were to be solved with the Krylov method (3.23), yielding

$$F(\psi) = \left(\mathbf{I} - \mathbf{L}^{-1}\mathbf{SP}\right)\psi - \mathbf{L}^{-1}q. \tag{3.24}$$

When the moment equations are written in a residual format, as they are in (3.24), Newton's method searches for a solution ψ that has zero residual; this solution exactly solves the transport equation across all directions in all cells.

Newton's method is easily derived from a Taylor expansion of $F(\psi)$ about the current estimate of ψ, ψ^m (where m denotes the current Newton iteration for ψ):

$$F(\psi^{m+1}) = F(\psi^m) + F'(\psi^m)(\psi^{m+1} - \psi^m) + \text{higher order terms.} \tag{3.25}$$

If enough terms are included, $F(\psi^{m+1})$ will equal zero, meaning that ψ^{m+1} is the exact solution to the nonlinear problem. Setting the value of the residual at the $m+1$ step equal to zero and neglecting the higher order terms, equation (3.24) reduces to

$$0 = F(\psi^m) + F'(\psi^m)(\psi^{m+1} - \psi^m). \tag{3.26}$$

The difference between the current estimate of the angular flux, ψ^m, and the next Newton iterate, ψ^{m+1}, can also be written as the correction to the current iteration $\delta\psi^m$. Finally, the derivative of the residual equation in vector form is the Jacobian with respect to the current guess for psi, $J(\psi^m)$. With rearrangement, equation (3.26) results in a Newton's method formulation of our problem:

$$J(\psi^m)\delta\psi^m = -F(\psi^m), \tag{3.27a}$$

$$\psi^{m+1} = \psi^m + \delta\psi. \tag{3.27b}$$

Newton's method requires an initial guess for the angular flux, ψ^0. This initial guess is typically a zero vector. From this initial guess, the problem is iterated until a convergence criteria is met. The problem's convergence after each Newton Step is calculated by taking an L^2 norm of the residual and dividing it by an L^2 norm of the solution vector ψ. This convergence is compared to the convergence criteria ϵ; if the current Newton guess results in a relative residual smaller than the convergence criteria,

$$\frac{\|F(\psi^m)\|_2}{\|\psi^m\|_2} \leq \epsilon, \tag{3.28}$$

the solution is considered converged.

3.3.1 Jacobian-free Newton Krylov

Analytically determining the Jacobian matrix is costly in terms of both the number of computations carried out and amount of storage required. From the discussion of Krylov methods and Equation (3.27), it is clear that GMRES does not actually require the determination of the analytic Jacobian matrix. Instead, GMRES requires the action of the Jacobian matrix on a vector to build a Krylov subspace. The Jacobian matrix of a function $F(u)$ is defined as

$$J = \begin{bmatrix} \frac{\delta F_1}{\delta u_1} \frac{\delta F_1}{\delta u_2} \\ \frac{\delta F_2}{\delta u_1} \frac{\delta F_2}{\delta u_2} \end{bmatrix}. \tag{3.29}$$

During the Krylov iterations the action of the Jacobian matrix upon the Krylov vector v is determined:

$$Jv = \begin{bmatrix} v_1 \frac{\delta F_1}{\delta u_1} v_2 \frac{\delta F_1}{\delta u_2} \\ v_1 \frac{\delta F_2}{\delta u_1} v_2 \frac{\delta F_2}{\delta u_2} \end{bmatrix}. \tag{3.30}$$

The action of J on the vector v can be approximated by the definition of the derivative:

$$\frac{F(\psi + \epsilon v) - F(\psi)}{\epsilon} = \begin{bmatrix} \frac{F_1(u_1+\epsilon v_1, u_2+\epsilon v_2) - F_1(u_1,u_2)}{\epsilon} \\ \frac{F_2(u_1+\epsilon v_1, u_2+\epsilon v_2) - F_2(u_1,u_2)}{\epsilon} \end{bmatrix}. \tag{3.31}$$

A Taylor expansion of $F(u+\epsilon v)$ in equation (3.31), keeping only the first order terms, appears:

$$\frac{F(\psi+\epsilon v)-F(\psi)}{\epsilon} \approx \begin{bmatrix} \frac{F_1(u_1,u_2)+\epsilon v_1\frac{\delta F_1}{\delta u_1}+\epsilon v_2\frac{\delta F_1}{\delta u_2}-F_1(u_1,u_2)}{\epsilon} \\ \frac{F_2(u_1,u_2)+\epsilon v_1\frac{\delta F_2}{\delta u_1}+\epsilon v_2\frac{\delta F_2}{\delta u_2}-F_2(u_1,u_2)}{\epsilon} \end{bmatrix}, \tag{3.32}$$

which simplifies to

$$\begin{bmatrix} v_1\frac{\delta F_1}{\delta u_1}+v_2\frac{\delta F_1}{\delta u_2} \\ v_1\frac{\delta F_2}{\delta u_1}+v_2\frac{\delta F_2}{\delta u_2} \end{bmatrix} = Jv. \tag{3.33}$$

Equations (3.29) through (3.32) make it clear that Eq. (3.31) computes the action of the Jacobian matrix J upon the Krylov vector v without constructing the Jacobian matrix. The residual need only be evaluated at the unperturbed Newton guess ψ once per Newton step, so Eq. (3.31) requires one residual evaluation at each Krylov iteration. This is much preferable to calculating and storing the Jacobian matrix at each Krylov iteration.

3.3.2 Magnitude of the Perturbation ϵ

The Jacobian-free evaluation of Jv can be sensitive to the magnitude of the scalar perturbation ϵ. For very small perturbations compared to u and v, the derivative calculation is lost amid round-off errors. For very large perturbations, however, the local derivative is poorly approximated. D.A. Knoll and D.E. Keyes recommend determining the perturbation ϵ at each Krylov Iteration with equation (3.34):

$$\epsilon = \frac{\sqrt{(1+\|u\|)\epsilon_{machine}}}{\|v\|}, \tag{3.34}$$

where $\|u\|$ is the L^2 norm of the latest Newton guess for ψ, and $\|v\|$ is the L^2 norm of the Krylov vector v [6].

Although this method for determining ϵ can be successful for some problems, it is poorly suited for problems in which ψ spans many orders of magnitude. Equation (3.31) for approximating of the action of the Jacobian upon the Krylov vector Jv

is accurate for small perturbations. The value of ϵ determined with equation (3.34) may be small compared to ψ in most cells, but comparable to or significantly larger than ψ in strongly absorbing regions where ψ van be very small or zero. This creates a problem with approximating the CSZ operators $\mathbf{L}(\psi)$ and $\mathbf{P}(\psi)$ in the strongly absorbing regions, which are the regions that most require the use of the CSZ method.

Recall that the operators $\mathbf{L}(\psi)$ and $\mathbf{P}(\psi)$ are constructed based upon the ratio $\frac{b_{CSZ}}{a_{CSZ}}$. If the perturbation ϵv is of similar magnitude (or larger) than the vector ψ, the construction of $\mathbf{L}(\psi + \epsilon v)$ and $\mathbf{P}(\psi + \epsilon v)$ is determined by the Krylov vector instead of ψ. The perturbed residual $f(\psi + \epsilon v)$ is therefore significantly different than the unperturbed residual $f(\psi)$ and the Jacobian is poorly approximated in this region.

Although FGMRES allows for the use of a different preconditioner at each Krylov iteration (which results in a changing iteration matrix, as is encountered in these problems), FGMRES fails to converge for a variety of problems containing strongly absorbing regions. Two solutions were investigated for this problem.

3.3.3 Newton's Method with an ϵ Fixup

The value of ϵ is only similar in size to ψ in strongly absorbing regions. In these regions, ψ can be orders of magnitude smaller than ϵ. This causes a problem in the construction of the $\mathbf{L}(\psi + \epsilon v)$ and $\mathbf{P}(\psi + \epsilon v)$ operators. To address this issue, the ratio R of the perturbation to ψ is computed for each entry in ψ:

$$R_i = \frac{\epsilon v_i}{\psi_i}. \tag{3.35}$$

A fixup tolerance ζ_{fixup} is chosen to be several orders of magnitude smaller than one. In each case that the ratio R_i is larger than the fixup tolerance, the perturbation used in the construction of the $\mathbf{L}(\psi + \epsilon v)$ and $\mathbf{P}(\psi + \epsilon v)$ operators is set to zero. The construction and application of these CSZ operators is summarized in equations (3.36):

$$[\mathbf{L}(\psi_i + \epsilon v_i)]\,(\psi_i + \epsilon v_i) \approx [\mathbf{L}(\psi_i + c_i \epsilon v_i)]\,(\psi_i + \epsilon v_i), \tag{3.36a}$$

where

$$c_i = \begin{cases} 1 & \text{if} \quad \frac{\epsilon v_i}{\psi_i} < \zeta_{fixup}; \\ 0 & \text{if} \quad \frac{\epsilon v_i}{\psi_i} > \zeta_{fixup}. \end{cases} \tag{3.36b}$$

This method requires the use of Flexible GMRES because the iteration matrix changes at each Krylov iteration. Computational results generated with this method are presented in the section titled "Newton's Method compared to the Frozen Newton Method." The fixup applied in this method suggests a second method for constructing the CSZ operators and approximating the numerical Jacobian.

3.3.4 Picard Iterations and a "Frozen" Newton Krylov Method

Instead of the conditional fixup applied in equations (3.36), one could also always construct the CSZ operators using the unperturbed value of ψ. In this case, the Jacobian is "Frozen" at the latest Newton iteration for ψ and does not change with the Krylov vectors. This Frozen JFNK method appears

$$J(\psi^{(\ell)})\delta\psi = -f(\psi^{(\ell)}), \tag{3.37a}$$

$$\psi^{(\ell+1)} = \psi^{(\ell)} + \delta\psi. \tag{3.37b}$$

Although this resembles a Jacobian-free Newton Krylov method, by freezing the Jacobian at the most recent Newton guess and solving the frozen system, this method is more accurately described as a Picard Iteration scheme.

The Frozen JFNK method does not require FGMRES to converge the linear solves. However, because FGMRES reduces to the standard (albeit slightly more memory-intensive) GMRES algorithm if the iteration matrix is constant, FGMRES is always applied.

Computational results comparing the standard Newton method and the Frozen Newton method are presented in the section titled "Newton's Method compared to the Frozen Newton Method." Based upon these results, the Frozen Newton method is used throughout this paper.

4. PRECONDITIONING STRATEGIES DEFINED BY SOURCE ITERATION SCHEMES

4.1 Different Strategies for Solving the Transport Problem

Multiple source iteration schemes can be constructed by varying the details of how the sweep and DSA preconditioning are performed. Each of these source iteration schemes is then reexpressed as a preconditioned system and solved with a Krylov method. Sweep preconditioning is always applied. DSA preconditioning is optional, but can only be applied to the problem after the sweep. Because both the sweep and DSA preconditioning can be performed with either the LD or CSZ operator, several different preconditioning strategies are investigated.

The frozen JFNK method uses the most recent Newton guess for ψ to construct the CSZ equations. To clearly indicate this dependence, the CSZ operators are made functions of ψ^*, the most recent Newton iteration for ψ. In contrast, the LD operators are constant and are marked with "hat" notation. For example, the CSZ transport operator is denoted $\mathbf{L}(\psi^*)$ and the LD transport operator is denoted $\widehat{\mathbf{L}}$.

4.1.1 LD Solution

Sweeping the LD moment equations with the LD sweep yields the fully linear LD solution of the moment equations. This solution can include negative outflows and negative average angular fluxes. The LD solution is investigated to compare the scalar flux solution to the scalar flux solution calculated with the CSZ methods. The cost of the LD solution will reflect the cost of one Newton step; the CSZ methods will require multiple Newton steps to converge, and will therefore be several times more expensive than this method.

In operator form, the linear sweep appears

$$f(\psi) = \psi - \widehat{\mathbf{L}}^{-1}\left(\sigma_s \widehat{\mathbf{P}}\psi + Q\right). \tag{4.1}$$

Recall that $\mathbf{L}$ includes the streaming and interaction terms, σ_s is the macroscopic scattering cross section, $\mathbf{P}$ generates the scalar flux ϕ from the angular flux ψ, and Q is the inhomogeneous source term. The hat notation is used to denote the use of the LD equations in the matrix $\widehat{\mathbf{L}}$ and in the generation of the scalar flux from the angular flux with the operator $\widehat{\mathbf{P}}$.

4.1.2 Sweep Preconditioning the CSZ Equations

Nonlinear Sweep Applied to the Nonlinear Equations

The nonlinear CSZ sweep preconditioner is applied to the CSZ equations to generate a strictly positive angular flux solution. The sweeps are performed with the nonlinear CSZ operator $\mathbf{L}(\psi^*)$, which is formed for each cell in each direction. Recall that in cases in which the angular flux solutions are strictly positive, the nonlinear transport operator $\mathbf{L}(\psi^*)$ is identically equal to the linear operator $\widehat{\mathbf{L}}$. Applying the nonlinear sweep to the nonlinear equation always yields the identity matrix. This system is written

$$f(\psi) = \psi - [\mathbf{L}(\psi^*)]^{-1} \left(\sigma_s P(\psi^*)\psi + Q\right). \tag{4.2}$$

Linear Sweep Applied to the Nonlinear Equations

It is possible to apply the linear LD sweep to the nonlinear CSZ equations. Solving the CSZ equations always results in a strictly positive solution; the choice of preconditioner does not change this property. The nonlinear CSZ transport operator $\mathbf{L}(\psi^*)$ is constructed in each each direction in each cell and applied to ψ. The sweep is then performed with the linear transport operator $\widehat{\mathbf{L}}$. This strategy is called the linear sweep applied to the nonlinear equations and appears:

$$f(\psi) = \widehat{\mathbf{L}}^{-1} \left(\mathbf{L}(\psi^*)\psi - \sigma_S P(\psi^*)\psi + Q\right). \tag{4.3}$$

Preconditioning with the LD sweep requires both the evaluation of $\mathbf{L}(\psi^*)\psi$ and performing a sweep, making it computationally expensive. This method cannot be less expensive than sweeping with the CSZ operator, and it may be significantly less effective than the CSZ sweep, but it is of theoretical interest to assess the effectiveness of the LD sweep.

LD sweep preconditioning and CSZ sweep preconditioning should be similarly effective when applied to problems in which the transport operator is infrequently or very slightly modified by the CSZ definitions. However, as $\mathbf{L}(\psi^*)$ and $\widehat{\mathbf{L}}$ become more dissimilar, $\widehat{\mathbf{L}}^{-1}\mathbf{L}(\psi^*)$ becomes less like the identity matrix $\mathbf{I}$ and the LD sweep preconditioned method may require extra Krylov iterations to converge. Sweeping with the CSZ operator avoids this shortfall, making the CSZ sweep a more robust (and usually less expensive) preconditioner.

Notes on Sweep Preconditioning

The initial guess for ψ at the start of the first Newton step is always zero. Therefore the CSZ equations will always reduce to the LD equations during the first Newton step. After the first Newton step, the CSZ equations will continue to reduce to the LD equations in cells and directions that do not contain negativities. It is possible for a problem to have a strictly positive final solution but contain significant negativities during the iterative process.

The choice of sweep preconditioner does not change the solution of the CSZ equations. Both the LD and CSZ sweep preconditioners will yield the same solution, but the LD sweep preconditioner may require additional Krylov iterations to converge. In addition to requiring at least as many Krylov iterations to converge as the CSZ sweep preconditioner, the LD sweep preconditioner requires more multiplications per iteration than the CSZ sweep preconditioner.

4.2 DSA Preconditioning

DSA preconditioning can be applied to the Krylov iteration scheme after preconditioning with a sweep. The LD or CSZ definitions for the angular flux can be applied while performing the DSA preconditioning. DSA always solves the S_2 problem for the correction to the angular flux $\delta\psi$. In operator form, the generation of the right hand side of the S_2 problem requires angle integration with the $\mathbf{P}$ operator followed by the generation of the source term from $\delta\phi$ with the operator $\mathbf{H}$. A sparse matrix $\mathbf{S_2}$ is constructed containing the streaming and interaction terms for all directions in all cells. The sparse $\mathbf{S_2}$ matrix is of rank $(2 \times m)$, where m is the number of cells in the problem. The S_2 problem is of the same form as the S_N problem,

$$\mathbf{S_2}\delta\psi = \mathbf{HP}\delta\psi, \tag{4.4}$$

but the S_2 problem is of much smaller rank than the S_N problem.

This problem is reexpressed as a preconditioned system in the same manner as equation (3.22). MATLAB's direct solver MLDIVIDE is used to invert the sparse matrix $\mathbf{S_2}$ and solve for the S2SA correction to the S_N problem. The correction to the scalar flux calculated by solving the S_2 problem is projected to N directions and added to the vector ψ.

The DSA preconditioning step can be performed using the LD or CSZ equations, leading to two different strategies for applying DSA preconditioning. The effectiveness of each DSA preconditioner will be assessed to determine if the CSZ DSA preconditioner is a significantly more effective preconditioner than LD DSA.

4.2.1 DSA Source Term

After performing the sweep in each Krylov iteration, a residual associated with the angular flux $\delta\psi$ is passed to the DSA step. The source term for DSA is dependent upon a residual associated with the scalar flux $\delta\phi$. This angle-integrated residual is

calculated from $\delta\psi$ with the operator $\mathbf{P}$, but can be performed using the linear LD definitions or the nonlinear CSZ definitions for ψ. The use of the LD definitions for to calculate $\delta\phi$ from $\delta\psi$ is denoted $\widehat{\mathbf{P}}$. The use of the CSZ definitions to calculate $\delta\phi$ is denoted with the operator $\mathbf{P}(\psi^*)$.

The generation of $\delta\phi$ is an intermediate step between the sweep preconditioning and DSA preconditioning in each Krylov iteration. The DSA source term is written in terms of ψ_A and ψ_X, not a_{csz} and b_{csz}. The ψ vector generated with the LD sweeps is already in terms of ψ_A and ψ_X, but the CSZ sweep preconditioner generates a ψ vector in terms of a_{csz} and b_{csz}. The CSZ integrating operator $\mathbf{P}(\psi^*)$ will be called to generate ϕ_A and ϕ_X when the CSZ sweep preconditioner is used. Equations (4.5, 4.6, 4.7, and 4.11) reflect this notation.

4.2.2 Linear (LD) DSA

The most straightforward method that can be used to apply DSA preconditioning to the Krylov iterations is to apply fully linear (LD) S2SA. In this solution technique, the LD definitions of ψ are used throughout the construction of the $\mathbf{S_2}$ matrix. It is observed that the restriction from N to two angular directions strongly dampens negativities in most cells, resulting in a positive S_2 LD solution in those cell. Most negativity-containing cells are only negative in the steepest angular directions and yield strictly positive LD distributions when restricted to only two directions. Very optically thick cells may contain LD negativities in a considerable number of angular directions and can still result in negative solutions to the LD S_2 problem.

The LD S2SA matrix is a sparse matrix labeled $\widehat{\mathbf{S_2}}$. This matrix is constructed once and reused at each Krylov iteration. The S_2 source term is calculated from the S_N vector ψ with the operator H. The S_2 solution generated with DSA is projected to N directions using the scalar flux ϕ and the current J with the operator $\mathbf{T}$.

Linear DSA can be applied after the linear or nonlinear sweep. When applied to the linear equations preconditioned with the linear sweep to produce the LD solution (4.1), the Krylov formulation of the transport problem with DSA appears:

$$f(\psi) = \left(I + \mathbf{T}\widehat{\mathbf{S_2}}^{-1} H\sigma_s\widehat{\mathbf{P}}\right)\left[\psi - \widehat{\mathbf{L}}^{-1}\left(\sigma_s\widehat{\mathbf{P}}\psi + Q\right)\right]. \tag{4.5}$$

LD DSA can be applied to the CSZ equations after preconditioning with the LD sweep or the CSZ sweep. The effectiveness of the LD sweep preconditioner when paired with LD DSA preconditioning will be compared to the CSZ sweep preconditioner with LD DSA preconditioning. The application of LD DSA preconditioning in conjunction with the LD sweep preconditioner is written:

$$f(\psi) = \left(I + \mathbf{T}\widehat{\mathbf{S_2}}^{-1} H\sigma_s\widehat{\mathbf{P}})\right)\left[\widehat{\mathbf{L}}^{-1}\left(\mathbf{L}(\psi^*)\psi - \sigma_S\mathbf{P}(\psi^*)\psi + Q\right)\right], \tag{4.6}$$

and the application of LD DSA and the CSZ sweep is written:

$$f(\psi) = \left(I + \mathbf{T}\widehat{\mathbf{S_2}}^{-1} H\sigma_s\mathbf{P}(\psi^*)\right)\left[\psi - \left[\mathbf{L}(\psi^*)\right]^{-1}\left(\sigma_s\mathbf{P}(\psi^*)\psi + Q\right)\right]. \tag{4.7}$$

The effectiveness of the LD sweep preconditioner will be assessed by comparing the LD and CSZ sweep preconditioners without further preconditioning, and by comparing both methods when coupled with LD DSA.

4.2.3 Nonlinear (CSZ) DSA

The S2SA equations can be formed using the nonlinear CSZ definitions of the angular flux to solve for the correction to the angular flux in terms of δa_{csz} and δb_{csz} instead of $\delta\psi_A$ and $\delta\psi_X$. The matrix $\widehat{\mathbf{S_2}}$, constructed once as the problem is initialized, was constructed to apply to a vector containing $\delta\psi_A$ and $\delta\psi_X$; therefore, after inverting $\widehat{\mathbf{S_2}}$ with MATLAB's direct method MLDIVIDE, the vector $\delta\psi$ is in terms of $\delta\psi_A$ and $\delta\psi_X$. A coefficient matrix $\mathbf{C}(\psi^*)$ is introduced to perform the action of the operator $P(\psi^*)$ to the vector ψ:

$$\psi_{A,X} = \mathbf{C}(\psi^*)\psi_{a_{csz},b_{csz}}. \tag{4.8}$$

The nonlinear matrix $\mathbf{C}(\psi^*)$ is constructed at each Newton iteration based upon the value of ψ at the current Newton iteration (ψ^*). The CSZ $\mathbf{S_2}(\psi^*)$ matrix is then generated from the LD $\mathbf{S_2}$ matrix:

$$\mathbf{S_2}(\psi^*) = \widehat{\mathbf{S_2}}\mathbf{C}(\psi^*). \tag{4.9}$$

This operation must be completed at each Newton iteration, making this method more computationally expensive than LD DSA. However, both the matrix $\widehat{\mathbf{S_2}}$ and $\mathbf{C}(\psi^*)$ are sparse matrices, and the number of inner Krylov iterations usually greatly exceeds the number of outer Newton iterations.

The S_2 problem is solved using the nonlinear $\mathbf{S_2}(\psi^*)$ matrix,

$$\mathbf{S_2}(\psi^*)\delta\psi = \mathbf{H}\sigma_s\mathbf{P}(\psi^*)\delta\psi, \tag{4.10}$$

and the solution $\delta\psi$ is projected from two to N directions and added to the original vector ψ with the operator T.

Nonlinear CSZ DSA is applied to the moment equations preconditioned with the CSZ sweeps. This preconditioned system appears as follows:

$$f(\psi) = \left(I + \mathbf{TS_2}(\psi^*)^{-1}H\sigma_s\mathbf{P}(\psi^*)\right)\left[\psi - [\mathbf{L}(\psi^*)]^{-1}\left(\sigma_s\mathbf{P}(\psi^*)\psi + Q\right)\right]. \tag{4.11}$$

The effectiveness of CSZ DSA and LD DSA preconditioning, each applied in conjunction with CSZ sweep preconditioning, will be assessed for a variety of problems. While the CSZ DSA preconditioner should be more effective than the LD DSA preconditioner, it is desirable that both will form very effective preconditioners.

Although CSZ DSA should be more effective than LD DSA when applied to problems containing strong negativities, the difference between the two preconditioning strategies may be relatively minor. The steepest angular directions are most likely to contain negativities. Cells that contain negativities in the steeper angular directions may not contain negativities when reduced to two angular directions. The restriction from N to two angular directions may significantly reduce the cells containing negativities. Because the CSZ equations reduce to the LD equations when the LD

equations have strictly positive solutions, the LD and CSZ DSA methods may be very similar.

5. TEST PROBLEMS AND RESULTS

Test problems are presented to verify that the CSZ method and the various preconditioners are functioning properly and to assess the relative effectiveness of the preconditioning strategies. The results are presented in three parts. First, the solution of the LD moment equations and the CSZ moment equations are compared to illustrate the characteristics of the CSZ method. Second, the two sweep preconditioning strategies are assessed to compare the effectiveness of the LD sweep and the CSZ sweep. Finally, the two DSA preconditioning strategies are compared to determine if CSZ DSA significantly outperforms LD DSA, or if the two strategies are comparable.

5.1 The CSZ Method

Two simple problems are presented to illustrate the beneficial attributes of the CSZ method. These problems compare the solution of the LD moment equations (4.1) and the CSZ moment equations (4.2).

5.1.1 CSZ as a Strictly Positive Method

A simple problem is presented to highlight the strictly positive nature of the CSZ method. The LD and CSZ methods are most clearly contrasted in optically thick cells. A homogeneous slab of strongly absorbing material is divided into three cells. The total slab width is 1cm, with $\sigma_t = \sigma_a = 30.0cm^{-1}$, so that each cell has an optical thickness of 10 mean free paths (MFPs). An isotropic flux normalized to produce a unit current is incident upon the left face of the slab, and a vacuum boundary condition is imposed on the right face of the slab. The scalar flux at the midpoint of each cell as calculated with the LD and CSZ methods and S_2 quadrature is presented in Figure 5.1.

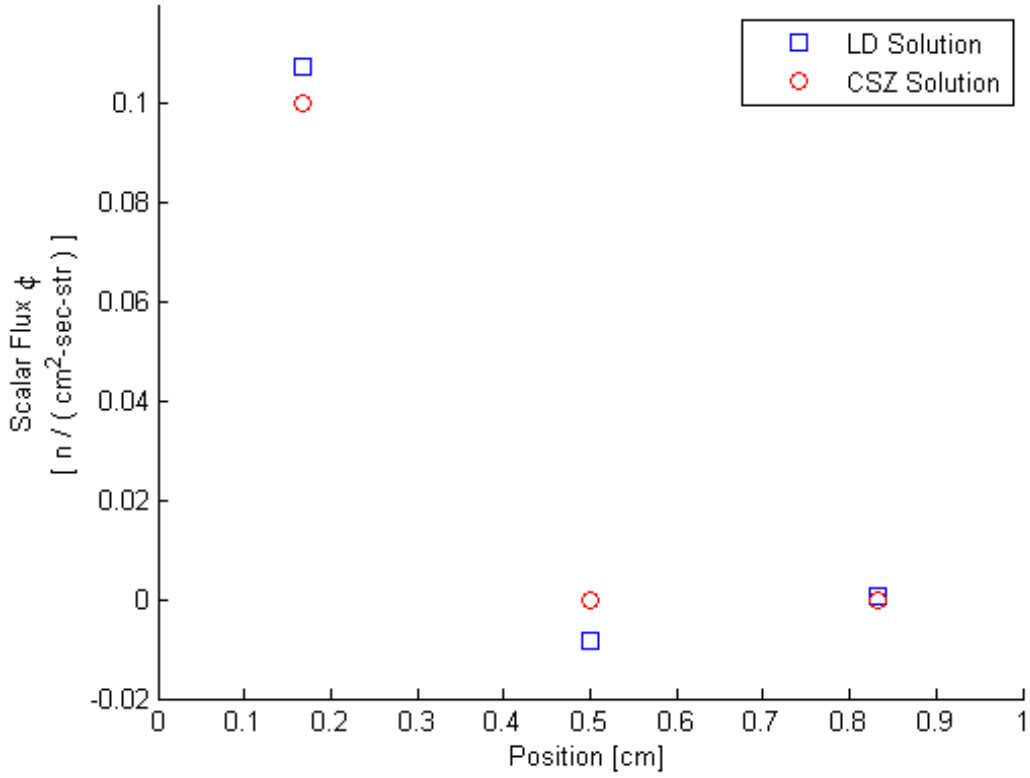

Fig. 5.1.: Comparison of average scalar flux calculated with LD and CSZ across a homogeneous absorber with optically thick cells.

Figure 5.1 clearly illustrates the strictly positive nature of the CSZ method. Notice that the LD solution method generates a negative scalar flux solution in the second cell and a positive solution in the third cell, while the CSZ method produces a zero solution in the second and third cells. The LD solution includes a negativity in the first cell resulting in a negative outflow. A negative outflow is equivalent to a positive inflow, so in the first cell the LD value for ϕ is larger than the CSZ scalar flux in the first cell. The negative inflow into the second cell results in a negative ϕ in the second cell. The sign of the angular flux changes from negative to positive in the second cell, yielding a positive outflow and a positive scalar flux in the third cell.

The analytic solution for one dimensional pure absorbers is an exponential flux shape within the slab. In these optically thick cells, the LD method yields a positive average scalar flux but a negative outflow for positive incident radiation. Alternatively, negative incoming radiation yields a negative average scalar flux but a positive outflow. These oscillations are damped but nonphysical. The CSZ method generates a positive average scalar flux in the first cell and a zero solution in the remaining cells. This underestimates the exact solution by some amount in the second and third cells. Neither the LD nor the CSZ method accurately reproduce the analytic solution, but the CSZ method does eliminate non-physical attributes that can characterize solutions generated with the LD method.

This problem highlights the difference between the LD and CSZ methods in optically thick cells. Although mesh sizes should be chosen to avoid extremely thick cells, even fine meshes can be very optically thick in the glancing or grazing directions. The CSZ equations' strictly positive nature is therefore most pronounced in the glancing directions in absorbing media.

5.1.2 The CSZ Method in Strictly Positive Problems

The CSZ method reduces to the LD method at any Newton step in which the LD solution is strictly positive. These problems are termed "strictly positive;" that is, the LD solution contains no negativities and the CSZ equations reduce to the LD equations. The reduction of the CSZ equations to the LD equations preserves many of the positive attributes of the LD method that make it one of the most commonly used spatial discretizations.

A strictly positive problem is presented to compare the LD and CSZ solution methods. The problem consists of a homogeneous slab five centimeters in length, divided into fifty cells, with a total absorption cross section $\sigma_a = 1.00 cm^{-1}$ and scattering cross section $\sigma_s = 0.99 cm^{-1}$. Vacuum boundary conditions are imposed

on the left and right face of the slab, and a distributed unit source is present ($q = 0.10\frac{neutron}{cm^3-s-biradian}$). S_{16} angular quadrature was used in the solution of the problem.

The scalar flux solution at the center of each cell is presented in Figure 5.2.

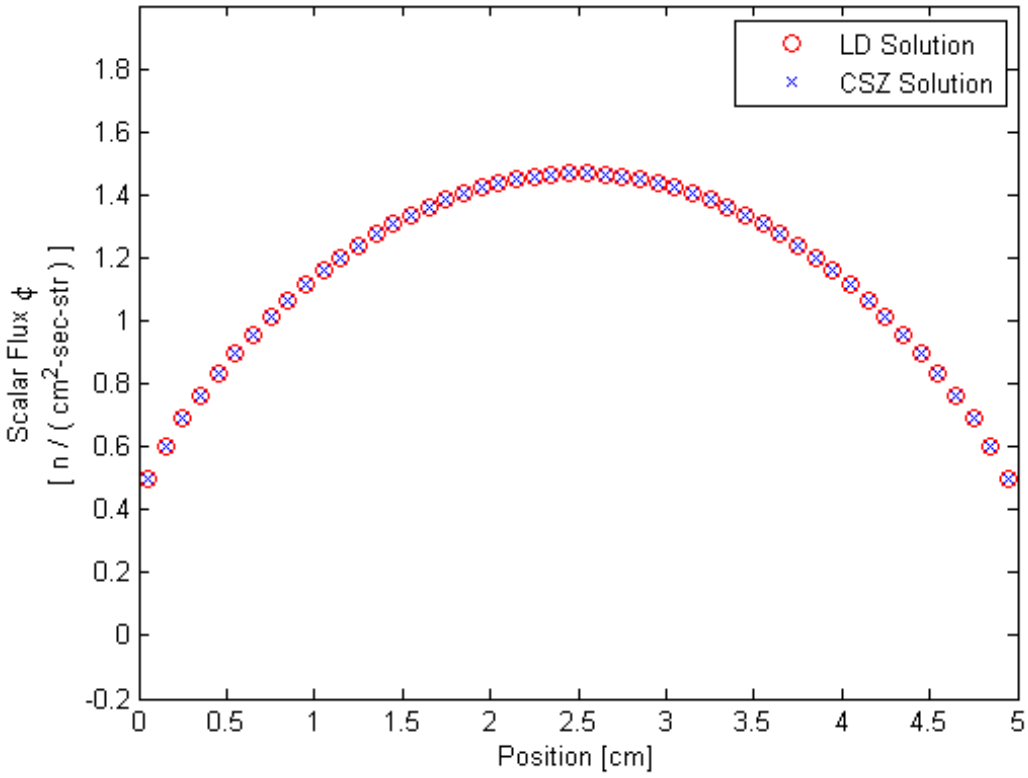

Fig. 5.2.: Average scalar flux in a homogeneous scatterer with a unit distributed source, as calculated with the LD and CSZ equations.

The algorithm reports that the CSZ equations reduce to the LD equations in each cell in each direction. The fully LD solution algorithm and both methods for sweeping the CSZ equations converge to the same solution, within a Newton tolerance of 1E-6, in a single Newton step requiring eleven Krylov iterations to reach a tolerance of 1E-8.

The CSZ method requires several more computations in each cell in each directions than the LD method (the cost of constructing the $L(\psi^*)$ matrix), but the CSZ equations reduce to the LD equations in this problem. This means that the the CSZ method is always at least slightly more expensive than the LD method, but that it also preserves the positive features of the LD method in cells and directions that do not contain negativities.

5.2 Sweep Preconditioning Strategies for the CSZ Equations

Two methods for applying sweep preconditioning to the CSZ equations are examined. Sweeping with the LD operator is always more expensive than sweeping with the CSZ operator, and the effectiveness of the LD sweep may degrade in problems containing significant negativities. Nevertheless, it is interesting to assess the effectiveness of the LD sweep as a preconditioner. Several problems are studied to assess the relative effectiveness of the LD and CSZ sweep preconditioners.

5.2.1 Grazing Radiation Problem

A problem containing significant negativities is solved to compare the LD and CSZ sweep preconditioners. The problem is solved with S_{16} quadrature, with incident radiation in the steepest positive direction on the left face of a slab. Vacuum boundary conditions are imposed on the right face of the slab. The slab has a scattering cross section $\sigma_s = 9.999cm^{-1}$, an absorption cross section $\sigma_t = 10.000cm^{-1}$, a length of $5cm$, and is divided into ten cells. Each cell has an optical thickness of $5.00\frac{mfp}{cell}$. The problem is repeated with a length of $10cm$, which increases the optical thickness of each cell to $10.0\frac{mfp}{cell}$.

The solution of the CSZ equations is independent of the preconditioners. The CSZ scalar flux solution and the fully LD scalar flux solution are plotted together

for the $5cm$ problem to compare the difference between the CSZ and LD solutions in Figure 5.3.

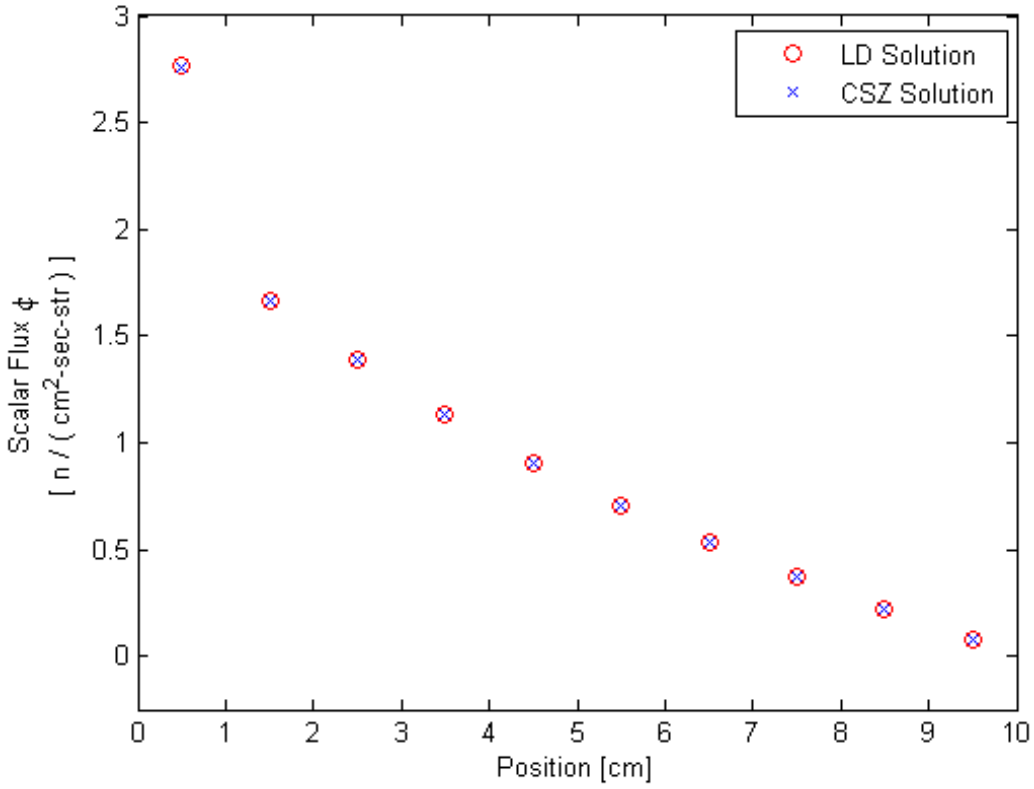

Fig. 5.3.: Average scalar flux in a homogeneous scatterer with grazing incident radiation.

The two sweep preconditioning strategies applied to the CSZ equations yield solutions that match to the problem tolerance, as expected. The LD solution differs from the CSZ solution by only approximately one percent in the first two cells and a very small amount in the remaining eight cells. Starting with an initial guess of zero, two Newton iterations were required for the problem to converge to a tolerance of 1E-6. The total number of Krylov iterations required for each preconditioning

strategy to converge is presented in Table 5.1 for a problem thickness of $5cm$ (as plotted in Fig 5.3) and $10cm$.

Table 5.1: A comparison of the number of Krylov iterations required for the grazing radiation problem to converge with the two sweep preconditioning strategies.

Preconditioner	**Iterations**	
	10 MFP/cell	20 MFP/cell
Linear Solution	24	24
LD Sweep	114	85
CSZ Sweep	120	74

The LD sweep and the CSZ sweep preconditioned methods required a similar number of iterations to converge for both slab thicknesses. The thicker slab required slightly fewer Krylov iterations when preconditioned with the LD sweep preconditioner, but the thinner problem required fewer iterations when preconditioned with the CSZ sweep. For this problem, the difference between the LD and CSZ sweep preconditioners was an order of magnitude smaller than the total number of Krylov iterations required by either sweep-preconditioned method to converge.

The LD scalar flux solution is very similar to the CSZ scalar flux solution, but the CSZ method requires significantly more iterations to converge than the LD method. Recall that each of the preconditioning strategies is equivalent to a source iteration scheme. Although the final solution is similar whether the LD moment equations or CSZ moment equations are solved, the intermediate source iteration solutions are quite different. The initial guess for the Newton step is a zero vector, so the CSZ equations reduce to the LD equations. The vector ψ after one Newton iteration includes a significant negativity in the direction of the incident radiation in the first cell and minor negativities in several directions in the furthest right cell. The value

for ψ after the second Newton iteration satisfies the convergence criteria and again contains a significant negativity in the direction of the incident radiation in the first cell and minor negativities in the final cell.

The numerical Jacobian was constructed for the method with LD sweep preconditioning and for the method with CSZ sweep preconditioning at each Newton step. The Eigenvalues associated with the Jacobian matrix were calculated with MATLAB's "eigs" function. The condition number, defined as the ratio of the largest and smallest eigenvalues, are presented in Table 5.2 for each sweep preconditioned method at each Newton step. Recall that at the first Newton step, the CSZ equations reduce to the LD equations, so the sweep preconditioners are identical.

Table 5.2: The condition number at each Newton step for the two sweep preconditioned methods.

Preconditioner	**Condition Number**		
	Step 1	**Step 2**	**Step 3**
LD Sweep	2.3728E3	2.3733E3	2.3734E3
CSZ Sweep	2.3728E3	2.3729E3	2.3729E3

There is not a significant difference in the condition numbers at each Newton step between the two sweep preconditioners. The grazing radiation problem does not suggest a significant difference between the two sweep preconditioning strategies.

5.2.2 A Reed-like Problem

A second problem is examined to compare the LD sweep preconditioner and the CSZ sweep preconditioner. Reed's problem is a five-region heterogeneous problem containing strong absorbers, highly diffusive regions, distributed sources, and a vac-

uum region [10]. The left face of the problem has an isotropic unit incident flux. A vacuum boundary condition is applied to the problem's right face. The cross sections used originally by Reed were adjusted to increase the diffusivity of the problem, and to increase the negativities in the strongly absorbing region. For this region, the problem is referred to as a "Reed-like" problem.

Description of the Reed-like Problem

The Reed-like problem has a length of $8cm$. The first region is a pure absorber two centimeters in length, with $\sigma_t = 50cm^{-1}$ and a distributed source of strength $Q = 50\frac{neutrons}{cm^3-s-biradian}$. The second region is one centimeter in length, contains no distributed source, and is a pure absorber with $\sigma_t = 5cm^{-1}$. The third region is two centimeters of vacuum. A fourth region one centimeter in length is highly diffusive with $\sigma_s = 19.99cm^{-1}$, $\sigma_t = 20.00cm^{-1}$, $Q = 0.10\frac{neutrons}{cm^3-s-biradian}$. Finally, the fifth region is two centimeters in length, highly diffusive ($\sigma_s = 19.99cm^{-1}$, $\sigma_t = 20.00cm^{-1}$), and contains no distributed source. The problem geometry is presented graphically in Table 5.3, with the material properties in the top portion of the table and the region lengths at the bottom of the table.

Table 5.3: Geometry of the Reed-like problem, where cross sections have units of cm^{-1}, distributed sources have units of *particles / (cm*^{3}*-sec-biradian)*, and lengths have units of *cm*.

$\sigma_s = 0$ $\sigma_t = 50$ $Q = 50$	$\sigma_s = 0$ $\sigma_t = 5$ $Q = 0$	$\sigma_s = 0$ $\sigma_t = 0$ $Q = 0$	$\sigma_s = 19.99$ $\sigma_t = 20.00$ $Q = 0.1$	$\sigma_s = 19.99$ $\sigma_t = 20.00$ $Q = 0$
$L = 2$	$L = 1$	$L = 2$	$L = 1$	$L = 2$

The strongly absorbing region contains strong negativities requiring the CSZ equations. The problem is split into 80 cells and solved with S_{16} quadrature. The Krylov iterations are solved to a tolerance of 1E-8, and the Newton Iterations are solved to a tolerance of 1E-6.

Solution of the Reed-like Problem

The scalar flux solution the Reed-like problem, solved with the CSZ equations, is presented in Figure 5.4.

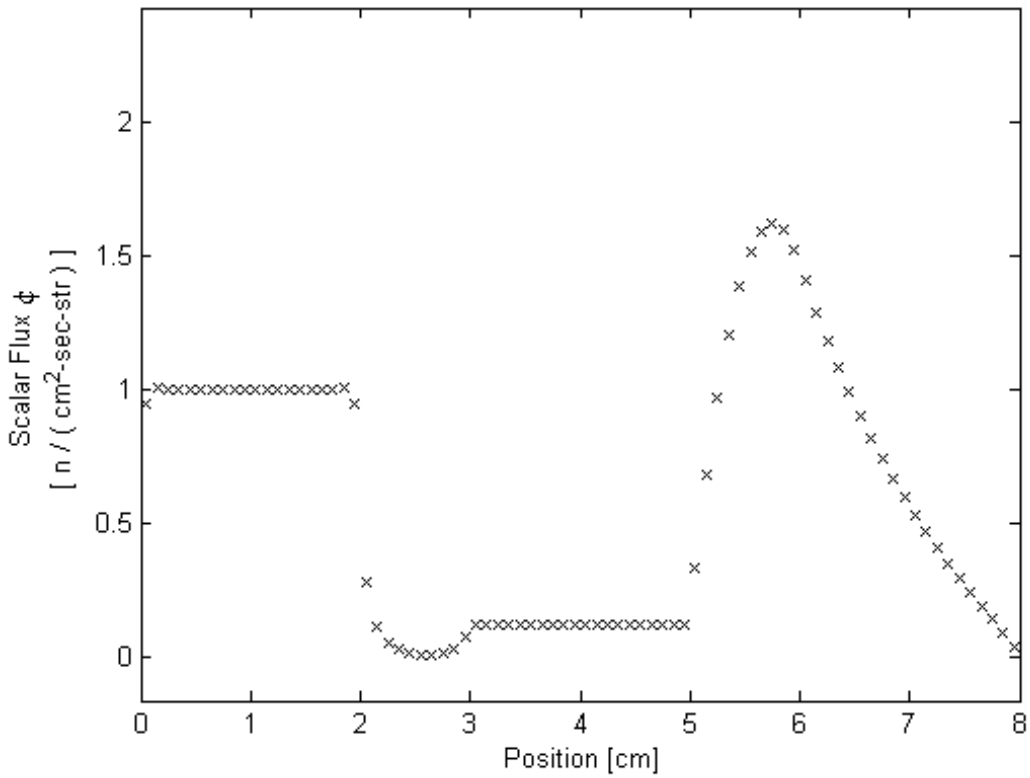

Fig. 5.4.: Scalar flux solution of a Reed-like problem.

The Strongly Absorbing Region of the Reed-Like Problem

Unlike the grazing radiation problem, the LD and CSZ solutions of the Reed-like problem differ significantly in the strongly absorbing region. The analytic solution in the purely absorbing region is eaily determined. Radiation is incident on this region on both the left and right face. In each direction, the analytic solution decreases exponentialy with space,

$$\psi_d(x) = \psi_d(0) e^{\frac{-x}{\mu}\sigma_t}. \tag{5.1}$$

The average angular flux in each cell i (with left and right boundary $x_{i-1/2}$ and $x_{i+1/2}$, respectively) $\psi_{A,d,i}$, is determined by integrating the analytic angular flux across cell i,

$$\psi_{A,d,i} = \frac{1}{h} \int_{x_{i-1/2}}^{x_{i+1/2}} e^{\frac{-x}{\mu}\sigma_t}, \tag{5.2}$$

for each of the N directions. The average scalar flux, ϕ_A, is determined by summing the average angular fluxes and their corresponding weights.

The error analysis was conducted for meshes with 3, 5, 10, 15, 20, and 30 cells in the strongly absorbing region. In each case, the boundary conditions were set by the incoming radiation from the solution to the 80 cell Reed-like problem. The analytic solution for the average scalar flux across the entire problem is plotted with the average scalar flux in each cell as computed with the LD and CSZ equations. These solutions are presented in Figure 5.5.

The error in the LD and CSZ scalar flux solutions is readily calculated by comparing the analytic and numerical solutions. The size of the cells in the strongly absorbing region is adjusted while holding the incident radiation on the left and right faces constant to compare the error in the LD and CSZ solutions for a variety of cell thicknesses. The L^2 norm of the relative error associated with the LD and CSZ solutions is plotted for a variety of mesh sizes in Figure 5.6.

In the cases in which the absorbing region is divided into twenty or thirty cells, the CSZ equations do not engage and the LD and CSZ solutions match. In the meshes

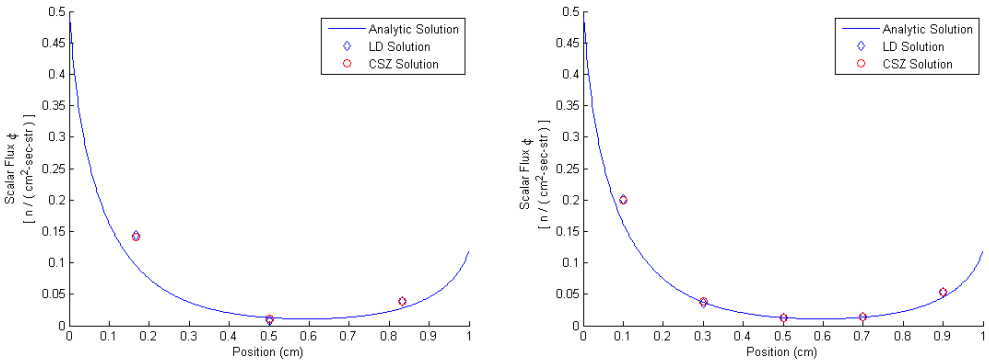

(a) 3 Cells

(b) 5 Cells

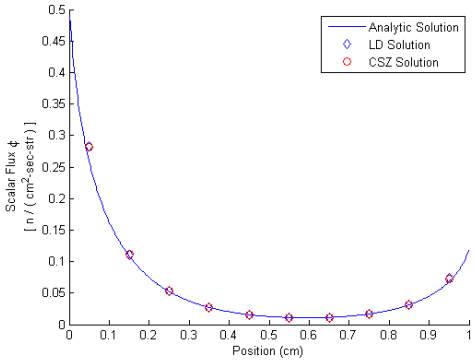

(c) 10 Cells

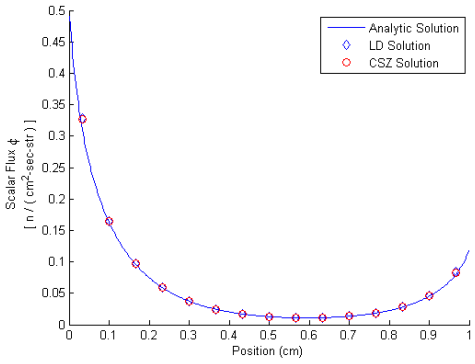

(d) 15 Cells

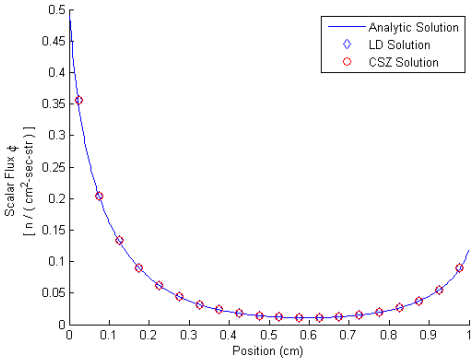

(e) 20 Cells

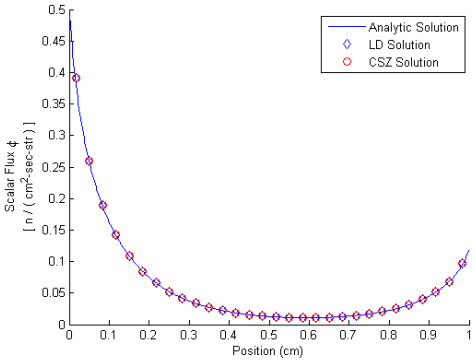

(f) 30 Cells

Fig. 5.5.: Relative error in the LD and CSZ numeric solutions of the purely absorbing region of the Reed-like problem for different mesh sizes.

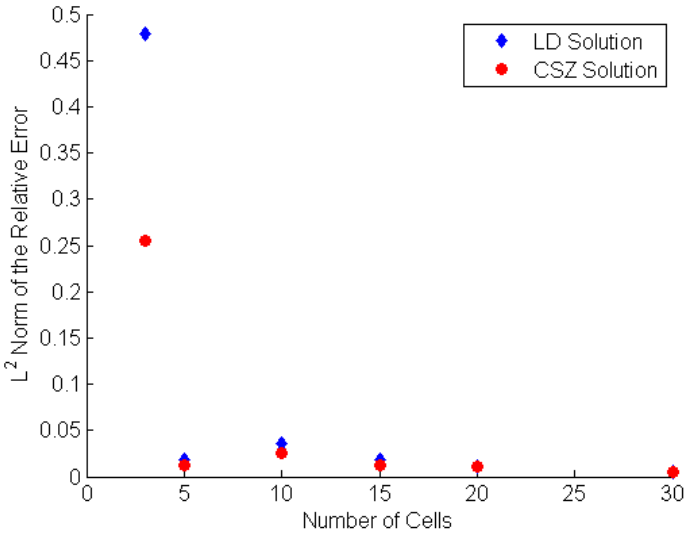

Fig. 5.6.: L^2 Norm of the relative error associated with the LD and CSZ solutions to the Reed-like problem.

with fewer than twenty cells, the CSZ solution technique has a smaller relative error than the LD solution. These results confirm that the CSZ solution is more accurate than the LD solution, and that this effect is especially apparent in optically thick cells in one dimension.

Comparing the LD and CSZ Sweep Preconditioners

The number of Krylov iterations required for the Reed-like problem to converge with LD sweep preconditioning and CSZ sweep preconditioning are presented in Table 5.4. The number of Krylov iterations required to generate the fully LD solution is presented as well to assess the relative cost of the CSZ solution.

It is worth noting that the CSZ sweep preconditioner required one more Newton step than the LD sweep preconditioned system, but required just more than half as many Krylov iterations to converge. The linear problem required a single Newton

Table 5.4: A comparison of the number of Krylov iterations required for a Reed-like problem to converge with the two sweep preconditioning strategies.

Preconditioning Strategy	**Newton Steps**	**Krylov Iterations**
Linear Solution	1	46
LD Sweep	5	298
CSZ Sweep	6	161

step, as expected. The condition number at each Newton step for the two sweep preconditioning methods is reported in Table 5.5.

Table 5.5: The condition number at each Newton step for the two sweep preconditioned methods.

Preconditioner	**Condition Number**					
	Step 1	**Step 2**	**Step 3**	**Step 4**	**Step 5**	**Step 6**
LD Sweep	7.30E2	7.30E2	2.03E6	2.84E11	1.42E5	n/a
CSZ Sweep	7.30E2	7.30E2	7.30E2	7.30E2	7.30E2	7.30E2

The condition number associated with the CSZ sweep preconditioner is constant across the Newton steps. The condition number associated with the LD sweep preconditioner, however, does not remain constant, and becomes extremely large in the fourth Newton step. The LD sweep preconditioner results in a very ill-conditioned Jacobian matrix; it is not surprising that the LD sweep preconditioner required many more iterations to converge than the method using the CSZ sweep preconditioner.

The CSZ sweep preconditioner is much more effective than the LD sweep preconditioner when applied to this Reed-like problem. This problem contains significantly more negativities than the grazing radiation problem, and these negativities persist into the final solution. The CSZ sweep clearly outperforms the LD sweep as a preconditioner with this problem, as expected, and the investigation into the DSA preconditioners will be performed using the CSZ sweep preconditioner.

5.3 The LD DSA Preconditioner

A test problem is presented to confirm that the LD DSA preconditioner is working properly. Two different tests are run using the same problem geometry to test the preconditioner in two different ways.

5.3.1 LD DSA in a Strictly Positive Problem

In one dimension, Diffusion Synthetic Acceleration is equivalent to solving the S_2 problem, as discussed in the section introducing DSA. Therefore, applying DSA is equivalent to solving the S_2 problem at each Krylov iteration. Applying DSA to a problem being solved with S_2 quadrature results in convergence in a single Newton iteration containing a single Krylov iteration, because the DSA step calculates the exact correction to solve the problem. An addition Krylov iteration is required to assess convergence, so LD DSA applied to the LD equations will converge after a total of two Krylov iterations.

By choosing a test problem in which the LD equations are strictly positive at each Newton iteration, the CSZ equations will always reduce to the LD equations. Therefore, LD DSA applied to the CSZ equations in this problem will result in convergence in two iterations. The CSZ DSA preconditioner also reduces to the LD DSA preconditioner in a strictly positive problem. Therefore, the application of

either LD or CSZ DSA to either the LD or CSZ equations will reduce to a fully LD problem and converge in a single iteration.

A highly diffusive homogeneous problem with a distributed source is used to confirm that the LD DSA preconditioner works correctly. The problem is ten centimeters in length, divided into one hundred cells, with a unit distributed source, a unit absorption cross section, and a scattering ratio of $c = 0.999$; this problem was examined previously to compare the LD and CSZ sweep preconditioners. The solution to this problem was presented in Figure 5.2. S_2 quadrature is used to test the DSA preconditioner.

The number of Krylov iterations required to converge to a Krylov tolerance of 1E-8 and a Newton tolerance of 1E-6 is presented in Table 5.6. Three different combinations of the LD equations, the LD sweep, and the CSZ sweep are presented to confirm that the CSZ equations reduce to the LD equations and the CSZ sweep reduces to the LD sweep in strictly positive problems. Recall that the CSZ sweep should not be applied to the LD equations, that combination is therefore excluded in Table 5.6.

Table 5.6: The number of Krylov iterations required for the highly diffuse problem to converge with S_2 quadrature with DSA preconditioning.

Equations	**LD Sweep**	**CSZ Sweep**
LD Equations	12	n/a
CSZ Equations	12	12

The CSZ DSA preconditioner similarly reduces to the LD DSA preconditioner in strictly positive problems. Table 5.7 presents the number of Krylov iterations required for the problem to converge to the same tolerances as Table 5.6 when solved with the LD equations, the LD sweep preconditioner, and LD DSA and for the CSZ

equations preconditioned with the CSZ sweep and both LD and CSZ DSA. As in Table 5.6, CSZ preconditioners are not applied to the LD equations, so only the LD DSA preconditioner is applied to the LD equations.

Table 5.7: The number of Krylov iterations required for the highly diffuse problem to converge with S_2 quadrature with DSA preconditioning.

Equations	**LD DSA**	**CSZ DSA**
LD Eqs & LD Sweep	2	n/a
CSZ Eqs & CSZ Sweep	2	2

The scalar flux solution calculated with the LD and CSZ equations are identical, further confirming that the CSZ equations are properly reducing to the LD equations in strictly positive problems. Because the methods require the same number of iterations to converge, Tables 5.6 and 5.7 confirm that the CSZ DSA preconditioner reduces to the LD DSA preconditioner in strictly positive problems. Finally, because each problem converged with only a single Krylov iteration with DSA applied, this result suggest that the LD DSA preconditioner is working correctly.

5.3.2 The Diffusion Limit

The diffusion limit of the highly diffuse problem is examined to verify that the LD DSA preconditioner is functioning properly. As a problem becomes more and more diffuse, the number of Krylov iterations required to converge should increase when only sweep preconditioning is applied. DSA preconditioning, however, is most effective in highly diffuse problems; therefore, the number of Krylov iterations required to converge should decrease with increasing problem diffusivity when DSA preconditioning is applied.

In slab geometry, is is simple to test the thick diffusion limit of a problem. Recall the form of the diffusion equation:

$$-\frac{d}{dx}\frac{1}{3\sigma_t}\frac{d}{dx}\phi(x)+\sigma_a\phi(x)=Q(x). \tag{5.3}$$

Multiplying the diffusion equation by a constant ϵ does not change the solution of the diffusion equation,

$$\phi-2L_D\frac{\partial\phi}{\partial x}=0, \tag{5.4}$$

with the diffusion length L_D defined

$$L_D=\left(\frac{1}{3\sigma_t\sigma_a}\right)^{1/2}. \tag{5.5}$$

The boundary conditions applied to the diffusion equation on domain $[0,L]$ are in terms of an extrapolated length $\widetilde{L}$:

$$\phi(-\widetilde{L})=\phi(L+\widetilde{L})=0. \tag{5.6}$$

In this analysis, vacuum boundary conditions are applied instead of the value of ϕ at an extrapolated boundary condition.

If the cross sections in the diffusion equation are scaled with ϵ as follows,

$$\sigma_t=\frac{\sigma_t}{\epsilon}, \tag{5.7a}$$

$$\sigma_a=\epsilon\sigma_a, \tag{5.7b}$$

$$\sigma_s=\frac{\sigma_t}{\epsilon}-\epsilon\sigma_a, \tag{5.7c}$$

and the distributed source q is scaled to make the relative magnitude of the solution independent of the scaling ϵ,

$$q=\epsilon q, \tag{5.7d}$$

it is clear that the diffusion equation is invariant with respect to ϵ. This is the scaling used in the asymptotic analysis of the diffusion limit.

The diffusion length is invariant with respect to ϵ. A problem thickness can be described in terms of the number of diffusion lengths L_D it contains. The optical

thickness in mean free paths is defined by the total cross section σ_t and the problem thickness in centimeters, presented here as L_{cm} for clarity. With the total cross section modified by ϵ, the problem thickness in mean free paths is given by:

$$L_{MFP} = \frac{\sigma_t}{\epsilon} L_{cm} \tag{5.8}$$

Thus it is possible to maintain a constant thickness in terms of the diffusion length but vary the optical thickness of a problem. As ϵ becomes very small, the absorption cross section becomes very small, the scattering cross section becomes very large, and the the cells become very optically thick. These conditions make the diffusion equation a valid approximation for the transport equation. In fact, an analysis of the transport equation with these scaled cross sections proves that for $\epsilon << 1$, the transport equation reduces to the diffusion equation [3].

The effectiveness of the LD DSA scheme can be analyzed as a function of the optical thickness of the problem. The highly diffusive problem presented earlier (Figure 5.2) will be reconsidered for this analysis. Because this problem does not retain its strictly positive nature for all optical thicknesses, only the LD equations preconditioned with the LD sweep are examined so that only the effect of the LD DSA is considered. The material is a homogeneous, highly diffusive slab five centimeters in length, divided into fifty cells, with a distributed source and vacuum boundary conditions. The invariant problem thickness is 10.0 diffusion lengths, and the thickness in terms of mean free paths given by Eq. (5.8). S_{16} quadrature was used in the solution of this problem. The average scalar flux in the center of each cell is presented in Figure 5.7.

The shape of the scalar flux solutions in Figure 5.7 indicate that as the optical thickness per cell increases, the solution converges to the solution of the diffusion problem. Table 5.8 presents the number of Krylov iterations required to converge for varying problem thicknesses in terms of mean free paths per cell.

As the optical thickness per cell increases, the sweep preconditioning strategy requires additional iterations to converge. The number of Krylov iterations required

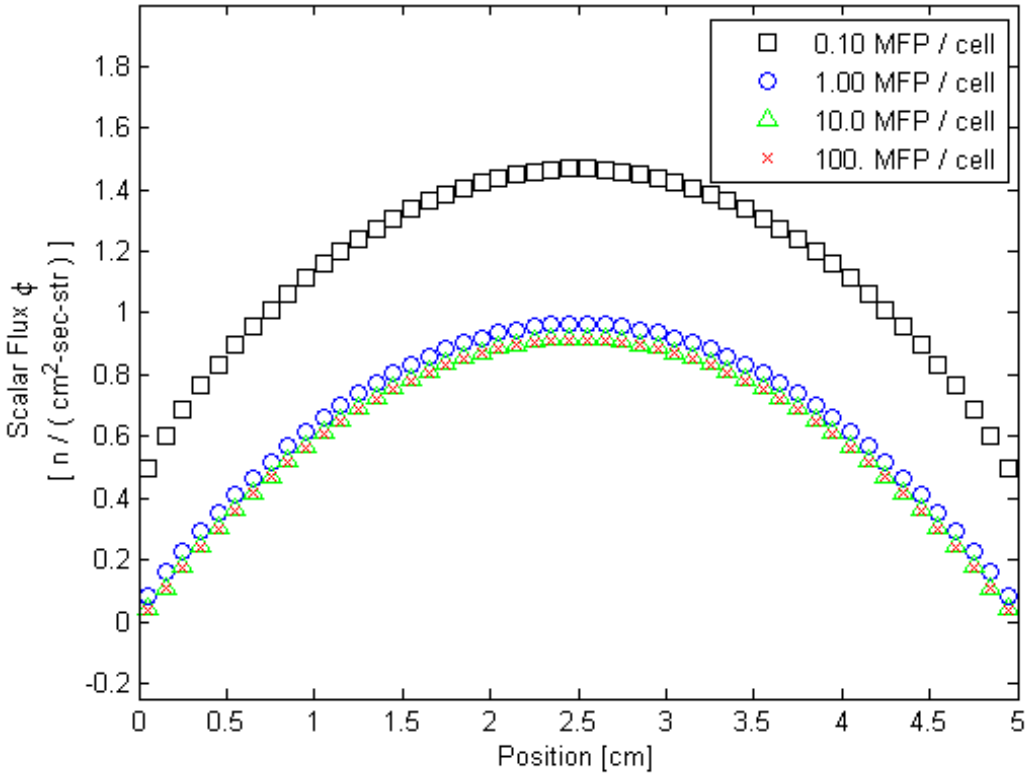

Fig. 5.7.: Average scalar flux in the diffusion limit.

Table 5.8: A comparison of the number of Krylov iterations required to converge with different combinations of sweep and DSA preconditioners.

	0.10 $\frac{MFP}{cell}$	1.00 $\frac{MFP}{cell}$	10.0 $\frac{MFP}{cell}$	100. $\frac{MFP}{cell}$
$1-c$	3.330E-3	3.330E-5	3.330E-7	3.330E-9
Sweep Preconditioning	11	30	71	98
DSA Preconditioning	8	7	5	4

with DSA preconditioning decreases as the optical thickness per cell increases, as

the diffusion approximation becomes increasingly accurate. This analysis further suggests that the DSA preconditioning appears to be working correctly.

5.3.3 Spectral Radius Investigation

As discussed in the section about Fourier Analysis of DSA, the spectral radii of source iteration schemes and of source iteration schemes preconditioned with DSA are well known. The preconditioned methods examined in this work are each based upon a source iteration scheme. Returning each preconditioned method to a source iteration form allows a comparison of the spectral radius of each preconditioned method and the expected value for that method.

The spectral radius ρ of the sweep preconditioning strategy is approximately equal to the scattering ratio c. The application of DSA preconditioning reduces the spectral radius to $\rho \leq 0.2247c$ in an infinite model problem, as discussed previously in the section titled Fourier Analysis of DSA. The spectral radius of each of the sweep and DSA preconditioning strategies are easily calculated with equation 3.14, repeated here for simplicity:

$$\rho = \lim \frac{\left\| \delta\psi^{(\ell+1)} \right\|}{\left\| \delta\psi^{(\ell-1)} \right\|} \tag{5.9}$$

Equation (5.9) requires the change between successive iterations in a source iteration method.

The relationship between the spectral radius of the sweep and DSA preconditioning methods and the optical thickness of a problem can be analyzed with the diffusion limit. For simplicity, the single region highly diffusive problem described in the previous section is used to determine the spectral radius of each method. The LD equations with LD DSA preconditioning and S_8 angular quadrature were used for this analysis.

Because DSA is a very effective preconditioner for highly diffusive problems, the spectral radius analysis is performed with an initial (non-zero) guess for ψ and zero

source terms. The solution is zero. A random initial guess for ψ ensures that all of the error modes are present in the calculation. Iterations can continue until ψ is almost arbitrarily small to ensure that enough iterations are performed to measure the spectral radius.

Table 5.9 compares the scattering ratio c and the spectral radius of the different preconditioning strategies as a function of the optical thickness of the problem. The problem is held at a constant thickness of ten diffusion lengths. These results were generated by reformatting the Krylov preconditioners as source iteration methods as discussed above.

Table 5.9: The spectral radius ρ of the LD Sweep and LD DSA preconditioners as a function of problem thickness.

	0.10 $\frac{MFP}{cell}$	1.00 $\frac{MFP}{cell}$	10.0 $\frac{MFP}{cell}$	100. $\frac{MFP}{cell}$
$1-c$	3.330E-3	3.330E-5	3.330E-7	3.330E-9
ρ_{Sweep}	0.92476	0.99873	0.99993	0.99998
ρ_{DSA}	0.21603	0.21234	0.14456	0.12792

As the optical thickness of the problem increases, the spectral radius of the LD sweep preconditioned source iteration scheme approaches the scattering ratio. The spectral radius of the DSA-preconditioned scheme approaches $\rho \approx 0.2247c$ for high S_N order. This analysis suggests that the LD DSA operator is working correctly.

5.4 LD DSA Applied to the LD and CSZ Sweep

The LD DSA preconditioner can be paired with either the LD sweep preconditioner or the CSZ sweep preconditioner. While the CSZ sweep preconditioner greatly outperformed the LD sweep preconditioner for some problems, it may be interest-

ing to see if the application of LD DSA to the two different sweep preconditioners produces very different results.

Because the two sweep preconditioners were comparable when applied to the grazing radiation problem, the Reed-like problem is examined. The total number of Krylov iterations required to converge with just the LD sweep preconditioner, the LD sweep preconditioner with LD DSA, just the CSZ sweep preconditioner, and the CSZ sweep preconditioner with LD DSA are reported in Table 5.10. By reporting all four of these values, the effectiveness of LD DSA can be assessed with both sweep preconditioners as well as the relative effectiveness of both sweep preconditioning strategies.

Table 5.10: The number of Krylov iterations required for the Reed-like problem to converge with LD and CSZ sweep preconditioning and LD DSA preconditioning.

Preconditioner	**Sweep Only**	**Sweep and LD DSA**
LD Sweep	298	141
CSZ Sweep	161	22

The application of LD DSA to the CSZ sweep is highly effective. Applying LD DSA preconditioning in conjunction with the LD sweep reduced the number of Krylov iterations required by a factor of two, but the CSZ sweep required roughly this many iterations to converge without DSA preconditioning. Applying DSA preconditioning to the CSZ sweep reduced the required number of iterations by a factor of more than seven. The CSZ sweep preconditioner clearly outperforms the LD sweep preconditioner for this problem.

The eigenvalues of the Jacobian matrix corresponding to each method are calcualted by forming the numerical Jacobian at each Newton step. Recall that an effective source iteration scheme has eigenvalues close to zero. Because of the form

of the preconditioned solution schemes, these eigenvalues are the eigenvalues of the identity matrix minus the Jacobian matrix; ie, one minus the eigenvalues. Therefore, an effective Krylov scheme will have eigenvalues close to one, and eigenvalues close to zero are detrimental to the method.

The eigenvalues of the LD sweep preconditioned system and the LD sweep with LD DSA preconditioning are plotted on the real and imaginary axis for each Newton step in Figure 5.8.

The eigenvalues are scattered along the real axis after preconditioning with the LD sweep. Many of the eigenvalues appear to be close to zero, and none of them are close to one. In the first Newton step, the LD DSA preconditioner effectively moves all of the eigenvalues away from zero and towards one. However, the LD DSA preconditioner is much less effective in each of the remaining Newton steps; the smallest eigenvalue remains very close to zero. It is not surprising that the LD DSA preconditioner is very effective in the first Newton step, because the CSZ equations reduce to the LD equations in the first Newton step.

The condition numbers of the LD sweep and the LD sweep paired with LD DSA at each Newton step are presented in Table 5.11.

Table 5.11: The condition number at each Newton step for the two LD sweep and LD sweep plus LD DSA preconditioning strategies.

Preconditioner	**Condition Number**				
	Step 1	**Step 2**	**Step 3**	**Step 4**	**Step 5**
LD Sweep	7.30E2	7.30E2	2.03E6	2.84E11	1.42E5
LD Sweep and LD DSA	1.24E0	2.02E2	9.99E3	7.65E5	3.65E9

Although LD DSA preconditioning did reduce the total number of Krylov iterations required for the system to converge when the LD sweep preconditioner is

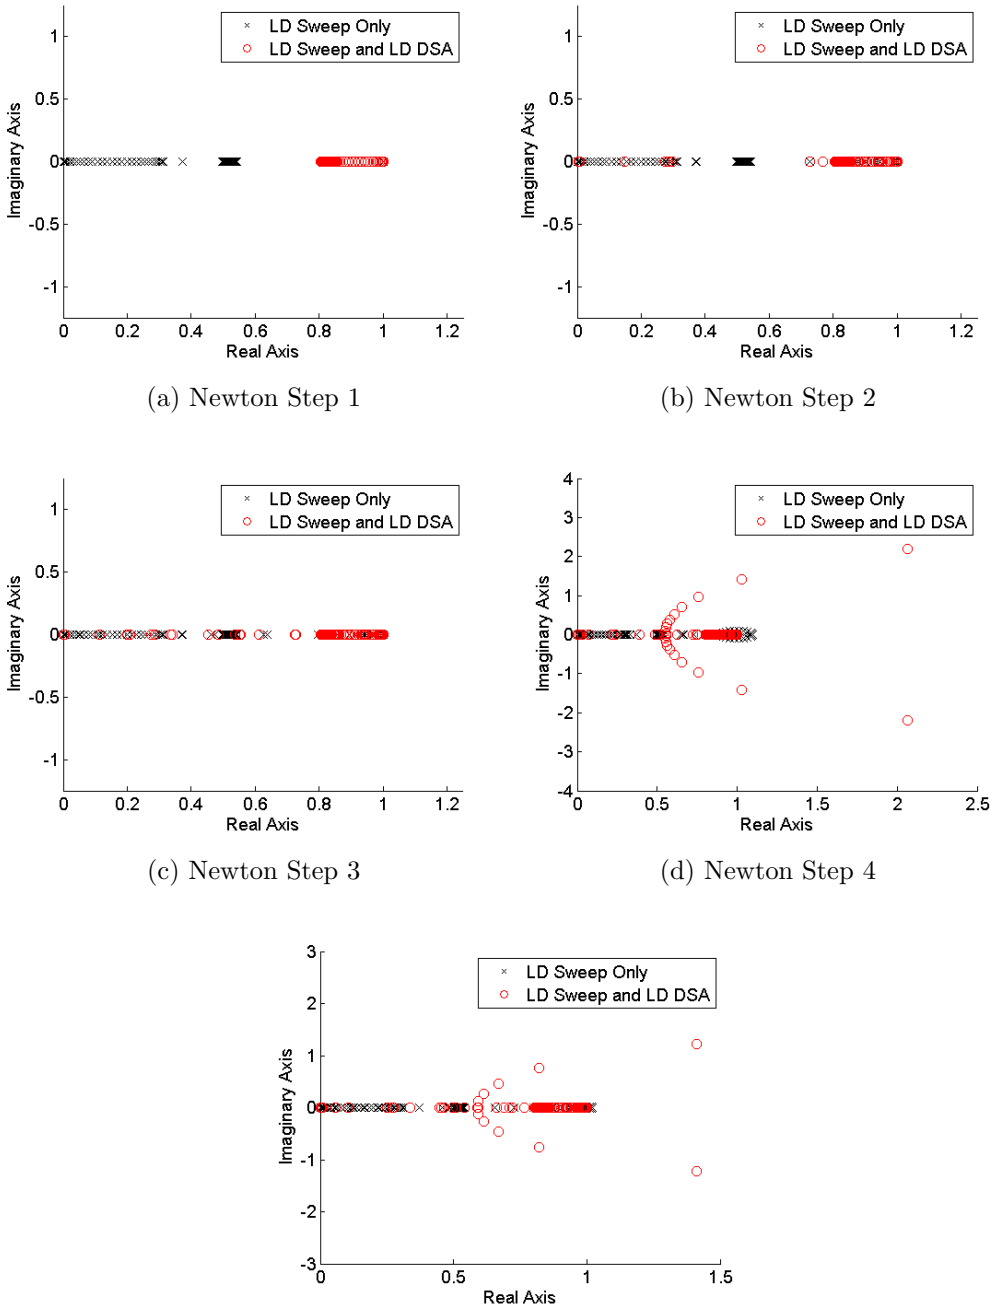

(a) Newton Step 1

(b) Newton Step 2

(c) Newton Step 3

(d) Newton Step 4

(e) Newton Step 5

Fig. 5.8.: Eigenspectrum of the LD sweep preconditioned Reed-like problem at each Newton step. Note the change in the scale of the axis for the fourth and fifth Newton steps.

applied, it is clear from the eigenspectra and the condition numbers that LD DSA preconditioning is not a very effective preconditioner when paired with the LD sweep preconditioner. LD DSA preconditioning should move the eigenvalues away from zero towards one, but one eigenvalue is remaining very close to zero. LD DSA preconditioning does not compensate for the LD sweep preconditioner's poor performance, and the CSZ sweep is a much more effective preconditioner.

As determined before, the CSZ sweep is strictly superior to the LD sweep. The CSZ sweep is less expensive than the LD sweep, requires fewer iterations to converge when only sweep preconditioning is applied, and DSA preconditioning is much more effective with the CSZ sweep than the LD sweep. Only the CSZ sweep will be considered for the remainder of the test problems.

5.5 Comparing the LD and CSZ DSA Preconditioners

The sweep preconditioners and the LD DSA preconditioner have been demonstrated to work properly for a variety of simple test problems. The CSZ sweep preconditioning significantly outperformed the LD sweep preconditioner, so the CSZ sweep preconditioner will be used to compare the LD and DSA preconditioners. Two test problems were designed to contain significant negativities and highly diffuse regions. The LD and CSZ DSA preconditioners can be compared to determine if the more difficult to implement CSZ DSA preconditioner significantly outperforms the simpler LD DSA preconditioner.

5.5.1 Metrics for Comparing the DSA Preconditioning Strategies

The primary metric by which the different preconditioning methods can be compared is the number of Krylov iterations required for the problem to converge to a Krylov tolerance of 1E-8 and a Newton tolerance of 1E-6. The computations associated with each Newton step are almost negligible compared to the computations

associated with each Krylov iteration. Therefore, the number of Krylov iterations required by each method forms the primary basis for comparing the methods.

The number of computations required per Krylov iteration varies between each preconditioning scheme. Because of this, the CPU time required for the problem to converge is considered in addition to the number of Krylov iterations. This will provide an estimate of the total cost of each method, although it has several inherent disadvantages. First, because this code is not optimized, these results will be only approximate. Second, because this project was written in MATLAB instead of C++ or FORTRAN, it is impossible to neglect background operations inherent to MATLAB. Finally, the amount of time required to converge varies between successive iterations.

Each preconditioned system is solved a large number of times to reduce the variance in the timing results. MATLAB loads the problem during the first iteration, so the first few iterations are discarded. The average time required to solve the problem and a standard deviation of that time are calculated over two hundred active iterations. These calculations were performed on the Texas A&M Nuclear Engineering cluster "Grove." An idle compute node was tasked with the timing analysis and remained otherwise unoccupied for the duration of the timing analysis to reduce or eliminate interference between tasks. Although the timing metric is not exact, it will be useful to determine if any of the preconditioning strategies are significantly more expensive than the competing strategies.

Finally, the eigenspectra and condition numbers of the two preconditioned systems may suggest which DSA preconditioner is more effective. If one DSA preconditioning strategy more effectively clusters the eigenvalues about one, it may be reflected in the number of Krylov iterations required for the problem to converge. The eigenspectra and condition number, the number of Krylov iterations required to converge, and the CPU time required to solve the problem are each related measures of the effectiveness of the preconditioning strategy. These three measures will

determine if the CSZ DSA preconditioner is significantly more effective than the LD DSA preconditioner.

5.5.2 Grazing Radiation

The grazing radiation problem examined previously is well suited for comparing the two DSA preconditioning strategies. The steep angle of the incident radiation on the left face causes a significant negativity in the first cell. The problem is highly diffusive, with a scattering ratio $c = 0.999$, making DSA preconditioning necessary. The problem is five centimeters in length, $\sigma_s = 9.999cm^{-1}$, $\sigma_t = 10.000cm^{-1}$, has a vacuum boundary condition on the right face, and is divided into ten cells, each with a thickness of ten mean free paths. Sixteen angular directions are used in this solution. The solution to the grazing radiation problem was presented as Figure 5.3.

Table 5.12 compares the number of Krylov iterations required for the LD DSA and CSZ DSA preconditioned strategies to converge to a Newton tolerance of 1E-6 and a Krylov tolerance of 1E-8. Table 5.13 compares the CPU time required for each combination of preconditioning strategies to converge to the same tolerance.

Table 5.12: Number of Krylov iterations required to solve the grazing radiation problem.

Preconditioner	**Krylov Iterations**
Sweep Only	74
LD DSA	21
CSZ DSA	21

Table 5.13: CPU time required for each preconditioning strategy to solve the grazing radiation problem. Times are presented in seconds plus or minus one standard deviation.

Preconditioner	**CPU time (s)**
Sweep Only	0.5209 ± 0.0064
LD DSA	0.2065 ± 0.0054
CSZ DSA	0.2074 ± 0.0072

The standard deviation (in seconds) associated with each method is reported alongside the time required to solve the preconditioned system. The standard deviations are on the order of five percent of the average time.

The eigenspectra of the preconditioned systems are plotted at each Newton step in Figure 5.9.

Finally, the condition number of the iteration matrix is reported in Table 5.14.

Table 5.14: The condition number at each Newton step for the two LD sweep and LD sweep plus LD DSA preconditioning strategies.

Preconditioner	**Condition Number**		
	Step 1	**Step 2**	**Step 3**
Sweep Only	2.37E3	2.37E3	2.37E3
LD DSA	1.18	1.21	1.21
CSZ DSA	1.18	1.21	1.21

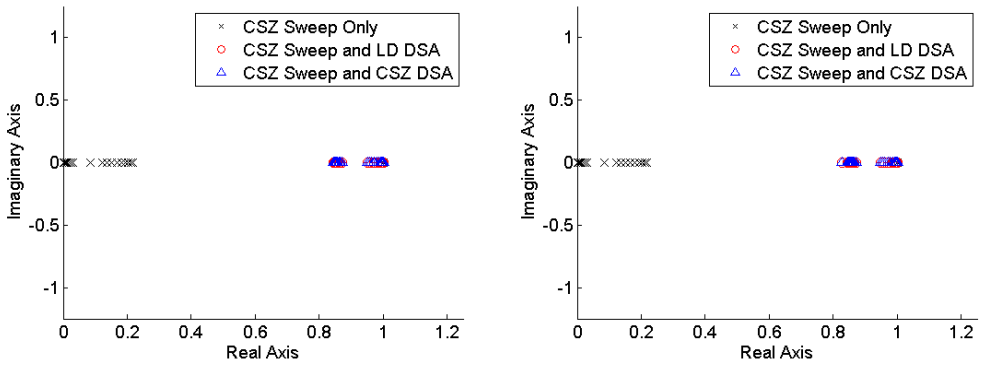

(a) Newton Step 1

(b) Newton Step 2

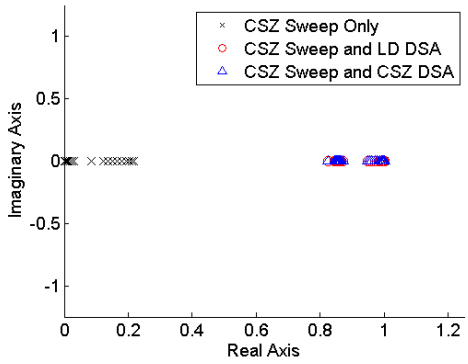

(c) Newton Step 3

Fig. 5.9.: Eigenspectrum of the preconditioned grazing radiation problem at each Newton step.

The application of LD DSA preconditioning significantly reduces the number of Krylov iterations and CPU time required for the problem to converge. The number of iterations required to converge was reduced by a factor of more than three, and the CPU time required was reduced by a factor greater than two. As demonstrated previously, LD DSA preconditioning significantly reduces the cost of solving the CSZ equations in diffuse problems.

CSZ DSA preconditioning achieved performance almost exactly equivalent to preconditioning with LD DSA. The difference in CPU time required between the two methods is smaller than the standard deviation in the data, and both methods required the same number of Krylov iterations to converge. The eigenspectra produced with both DSA preconditioning strategies is nearly identical, and the condition numbers are the same. CSZ DSA preconditioning is not more beneficial than LD DSA preconditioning in this problem.

5.5.3 Reed-like Problem

The Reed-like problem introduced previously was designed to include severe negativities and highly diffusive regions. The problem geometry was presented in Table 5.3, and the scalar flux solution was presented in Figure 5.4. Strong negativities are encountered in the highly absorbing regions on the left of the problem, and DSA acceleration is beneficial because of the highly diffusive regions on the right of the problem.

The algorithm reports that the CSZ equations do not reduce to the LD equations in roughly one to two percent of the total cell and direction cases in both the sweep preconditioning step and the CSZ DSA preconditioning step in all but the first Newton step; recall that the CSZ equations reduce to the LD equations in the first Newton step because the initial guess for ψ is zero. In each of the remaining Newton steps, the CSZ equations do not reduce to the LD equations in cells near the boundaries of the strongly absorbing region.

The number of iterations required for the two DSA preconditioning strategies to converge is presented in Table 5.15. The results calculated using the only sweep preconditioning and not DSA preconditioning are presented as well to assess the effectivness of the DSA preconditioners.

Table 5.15: Number of Krylov iterations required for the DSA preconditioning strategies to solve the Reed-like problem.

Preconditioner	**Krylov Iterations**
Sweep Only	161
LD DSA	22
CSZ DSA	22

The CPU time required for the two DSA preconditioning strategies to converge is presented in Table 5.16. Notice that as in the grazing problem timing results (Table 5.13), the standard deviation is on the order of one to two percent of the average time required for each method to converge.

Table 5.16: CPU time required for each combination of preconditioning strategies to solve the Reed-like problem. Times are presented in seconds plus or minus one standard deviation.

Preconditioner	**CPU Time (s)**
Sweep Only	9.2618 ± 0.0475
LD DSA	3.1570 ± 0.0173
CSZ DSA	3.1769 ± 0.0174

The eigenspectra at each Newton step are plotted in Figure 5.10.

The condition number at each Newton step is reported in Table 5.17.

Table 5.17: The condition number of the Reed-like problem at each Newton step for the CSZ sweep and the CSZ sweep paired with the LD or CSZ DSA preconditioner.

Preconditioner	**Condition Number**					
	Step 1	**Step 2**	**Step 3**	**Step 4**	**Step 5**	**Step 6**
Sweep only	730	730	730	730	730	730
LD DSA	1.24	1.24	1.24	1.24	1.24	1.24
CSZ DSA	1.24	1.24	1.24	1.24	1.24	1.24

The application of the CSZ sweep and the LD or CSZ DSA preconditioners to the CSZ equations produced similar trends to those observed in the grazing radiation problem. As before, both DSA preconditioners required the same number of Krylov iterations to converge. The difference in the CPU time required for the two DSA preconditioning methods to converge was similar in magnitude to the standard deviation in the timing data. The eigenspectra and condition numbers generated with both DSA preconditioning strategies are extremely close.

Pairing DSA preconditioning with the sweep preconditioner reduced the number of iterations required by more than a factor of seven, and the CPU time required to converge was reduced by almost a factor of three. DSA preconditioning is very beneficial for problems with highly diffusive regions, but CSZ DSA does not outperform LD DSA even in problems containing strong negativities.

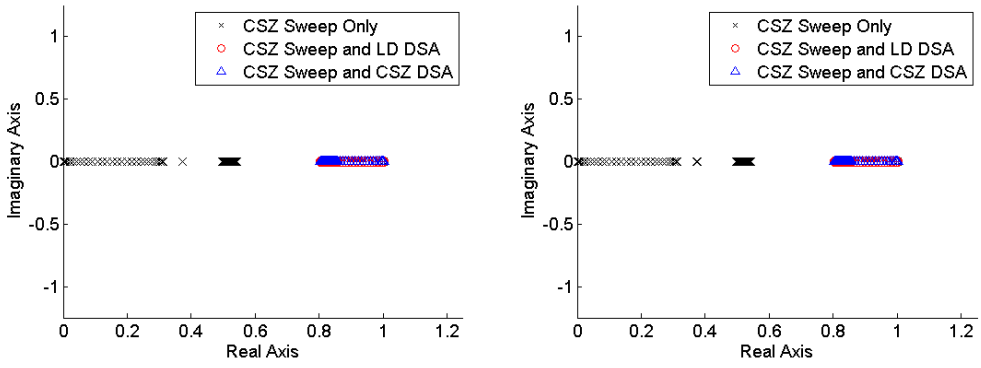

(a) Newton Step 1

(b) Newton Step 2

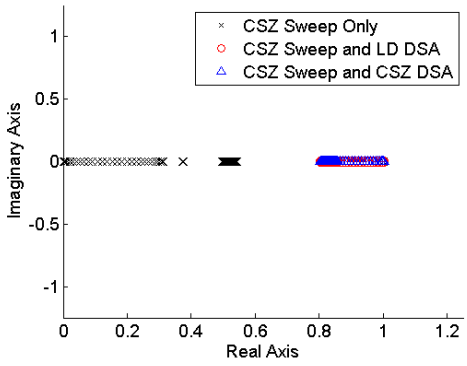

(c) Newton Step 3

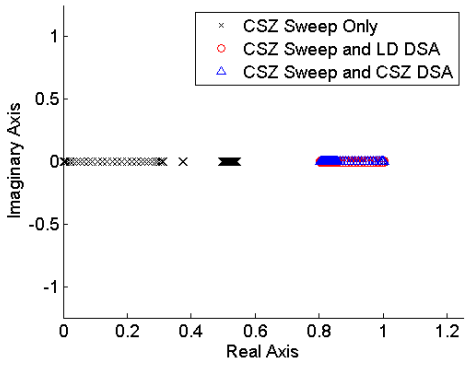

(d) Newton Step 4

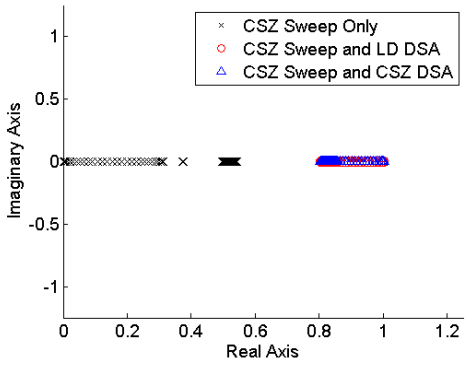

(e) Newton Step 5

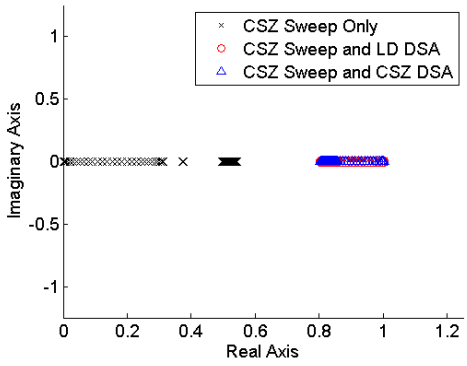

(f) Newton Step 6

Fig. 5.10.: Eigenspectrum of the preconditioned Reed-like problem at each Newton step. The spectrum after only the CSZ sweep is bound between 1.371E-3 and 1.0000; after applying either form of DSA, the spectrum is bound between 0.8045 and 1.0000.

5.6 Comparing the LD and CSZ Methods

Finally, it is interesting to compare the total cost assosiated with solving the Reed-like problem with both the LD and CSZ equations to assess the relative cost of the CSZ method. LD DSA is applied to both methods. Table 5.18 compares the number of Newton and Krylov iterations required for to solve the Reed-like problem to a Newton tolerance of 1E-6 and a Krylov tolerance of 1E-8.

Table 5.18: Number of Newton and Krylov iterations requred to solve the Reed-like problem with the LD and CSZ methods, to Newton tolerance of 1E-6 and a Krylov tolerance of 1E-8.

Method	**Newton Iterations**	**Krylov Iterations**
LD (with LD DSA)	1	7
CSZ (with LD DSA)	6	22

Table 5.19 presents the same data for a Newton tolerance of 1E-12 and a Krylov tolerance of 1E-14 to determine if tightening the convergence tolerance yields a similar result.

Table 5.19: Number of Newton and Krylov iterations requred to solve the Reed-like problem with the LD and CSZ methods, to Newton tolerance of 1E-12 and a Krylov tolerance of 1E-14.

Method	**Newton Iterations**	**Krylov Iterations**
LD (with LD DSA)	2	12
CSZ (with LD DSA)	7	71

The CSZ method is always more expensive than the LD method for problems in which the LD solution yields negativities. The LD method is a linear problem and converges (to a tolerance of at least 1E-6) in a single Newton step. This greatly reduces the number of Krylov iterations needed to solve the LD problem, and this is reflected in Tables 5.18 and 5.19. The CSZ method is always several times more expensive than the LD solution. Maginot examined the cost of the CSZ method and determined it to be comparable to other nonlinear methods for generating strictly positive solutions to the transport equation [1, 2].

5.7 Comparing Newton's Method and the Frozen Newton's Method

The Reed-like problem is used to compare the Frozen JFNK method used thus far in the test problems to the true JFNK method outlined in the section titled "Newton's Method with an ϵ Fixup." The CSZ equations preconditioned with LD DSA were chosen for this analysis based upon the results comparing the LD and DSA preconditioners. Table 5.20 reports the number of Newton and Krylov iterations required for the Frozen JFNK and the true Newton methods to converge to a Newton tolerance of 1E-6 and a Krylov tolerance of 1E-8. The Newton method is examined for various fixup tolerances, ζ_{fixup}. Fixup tolerances larger than 1E-2 caused FGMRES to stagnate and are not included.

These computations were repeated for a Newton tolerance of 1E-12 and a Krylov tolerance of 1E-14. These results are reported in Table 5.21.

The number of entries in the Krylov vector that required fixup when constructing the CSZ operators varied from about half the entries for the fixup tolerance of 1E-5 to about 11% of the entries for a fixup tolerance of 1E-3. In all cases, the final Krylov step required no fixup at all, as is expected.

These results suggest that applying the Jacobian-free Newton's method with the epsilon fixup yields slightly better results than the Frozen Jacobian-free Newton method. However, the reduction in cost is not very significant. For this reason, the

Table 5.20: Number of Newton and Krylov iterations requred for the Frozen JFNK and Newton JFNK method to converge to a Newton tolerance of 1E-6 and a Krylov tolerance of 1E-8.

Method	**Newton Iterations**	**Krylov Iterations**
Frozen JFNK	6	22
JFNK, ζ_{fixup} = 1E-5	6	22
JFNK, ζ_{fixup} = 1E-4	6	22
JFNK, ζ_{fixup} = 1E-3	5	24

Table 5.21: Number of Newton and Krylov iterations requred for the Frozen JFNK and Newton JFNK method to converge to a Newton tolerance of 1E-12 and a Krylov tolerance of 1E-14.

Method	**Newton Iterations**	**Krylov Iterations**
Frozen JFNK	7	71
JFNK, ζ_{fixup} = 1E-5	7	58
JFNK, ζ_{fixup} = 1E-4	7	58
JFNK, ζ_{fixup} = 1E-3	6	65

frozen JFNK method was used to generate results for this paper, and is suggested for use with the CSZ operators.

6. CONCLUSIONS

The application of two preconditioning techniques to the Consistent Set-to-Zero S_N spatial discretization developed by Maginot, Morel, and Ragusa was investigated. The CSZ equations were first reformulated to be solved with a sweep-preconditioned Krylov method. Diffusion Synthetic Acceleration was then applied as a preconditioner for the Krylov system. Both the Krylov solver FGMRES and DSA preconditioning were demonstrated to be compatible with the CSZ equations.

The CSZ equations can be effectively reexpressed as a preconditioned system and solved with a Krylov method. Krylov methods generally converge significantly faster than source iteration methods. Source iteration can be reexpressed as a sweep preconditioner relatively easily. Effective iteration techniques recast as preconditioners are generally effective in a Krylov solve. These attributes make Krylov schemes straightforward to implement in applications built around a source iteration method.

Two different sweep preconditioning strategies were investigated. While preconditioning the CSZ equations with an LD sweep is possible, it appears that a CSZ sweep is a much more effective preconditioner. LD sweep preconditioning was demonstrated to be significantly less effective than the CSZ sweep preconditioner for some problems, as expected when the methods were formulated. The most complicated problems demonstrated the largest difference between the methods, suggesting that the CSZ sweep preconditioner should be chosen for any further research.

Diffusion Synthetic Acceleration preconditioning using the LD S2SA equations (LD DSA) significantly accelerated the solution of the CSZ equations. Optically thick model problems required between two and seven times fewer iterations to converge with the application of LD DSA preconditioning in addition to sweep preconditioning. LD DSA preconditioning paired with CSZ sweep preconditioning was significantly more effective than LD DSA preconditioning paired with LD sweep preconditioning.

This further suggests that the CSZ sweep preconditioner is superior to the LD sweep preconditioner for the CSZ equations.

Preconditioning with the CSZ S2SA equations (CSZ DSA) was not more effective than preconditioning with LD DSA. Both DSA preconditioners significantly reduced the cost associated with solving problems containing strongly absorbing regions and highly diffusive regions. The CSZ DSA preconditioner was not significantly more effective than the LD DSA preconditioner for those problems.

In one dimension, the restriction from N angular directions to S_2 quadrature suppresses LD negativities because these negativities generally occur only in the steepest angular directions. For this reason the CSZ DSA preconditioner reduces to the LD DSA preconditioner in all but a few cases even in problems containing strongly absorbing regions. Several test problems were designed to have significant LD negativities in the DSA step. CSZ DSA was not significantly more effective than LD DSA even in these problems.

It would be very beneficial if this result regarding the comparable performance of the LD and CSZ DSA preconditioners in slab geometry applies to diffusion preconditioning in multiple spatial dimensions. In slab geometry, the S_2 equations are equivalent to the diffusion equation, making CSZ diffusion straightforward to derive. However, no simple analogous relationship exists in multiple spatial dimensions. This makes DSA more difficult to apply, making CSZ DSA preconditioning very difficult to apply.

The numerical approximation of the Jacobian of this system can be difficult because ψ spans many orders of magnitude. The frozen Jacobian-free Newton Krylov method was found to be much simpler and only slightly more expensive than a true Jacobian-free Newton Krylov method. Future researchers should consider the difficulties that can arise in computing a numerical approximation of the Jacobian with the CSZ operators and be aware of the two approaches for resolving this issue outlined in this paper.

The Consistent Set-to-Zero S_N spatial discretization appears to be well suited to problems that contain both strongly absorbing and highly diffusive regions. These types of problems can contain significant negativities and benefit greatly from DSA preconditioning. The CSZ sweep preconditioner paired with the LD DSA preconditioner demonstrated significant acceleration in these types of problems and should be considered for further research in multiple dimensions.

REFERENCES

[1] P. G. Maginot, J. E. Morel, J. C. Ragusa, A new nonlinear spatial S_N discretization method for 1-D and 2-D cartesian geometries, Journal of Computational Physics.

[2] P. G. Maginot, A nonlinear positive extension of the linear discontinuous spatial discretization of the transport equation, Master's thesis, Texas A&M University, College Station, TX (December 2010).

[3] Y. Y. Azmy, E. Sartori, Nuclear Computational Science: A Century in Review, Springer, New York, NY, 2010.

[4] M. L. Adams, E. W. Larsen, Fast iterative methods for discrete-ordinates particle transport calculations, Progress in Nuclear Energy 40 (2002) 3–159.

[5] E. W. Larsen, Unconditionally stable diffusion-synthetic acceleration methods for the slab geometry discrete-ordinates equations part i: Theory, Nuclear Science and Engineering 82 (1982) 47–64.

[6] D. A. Knoll, D. E. Keyes, Jacobian-free Newton-Krylov methods: a survey of approaches and applications, Journal of Computational Physics 193 (2) (2004) 357–397.

[7] Y. Saad, M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM Journal of Statistical Computing 3 (1986) 856–869.

[8] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, SIAM Journal of Statistical Computing 14 (1993) 461–469.

[9] J. S. Warsa, T. A. Wareing, J. E. Morel, Krylov iterative methods and the degraded effectiveness of diffusion synthetic acceleration for multidimensional S_N calculations in problems with material discontinuities, Nuclear Science and Engineering 147 (2004) 218–248.

[10] W. H. Reed, New difference schemes for the neutron transport equation, Nuclear Science and Engineering 45 (1971) 309–314.

[11] E. W. Larsen, J. E. Morel, Asymptotic solutions of numerical transport problems in optically thick, diffusive regimes ii, Journal of Computational Physics 83 (1989) 212–236.

[12] C. Börgers, E. W. Larsen, M. L. Adams, The asymptotic diffusion limit of a linear discontinuous discretization of a two-dimensional linear transport equation, Journal of Computational Physics 98 (1992) 285–300.

[13] P. Barbucci, F. Di Pasquantonio, Exponential supplementary equations for sn methods: The one-dimensional case, Nuclear Science and Engineering 63 (1977) 179–187.

[14] Y. Saad, Iterative Methods for Sparse Linear Systems, 2nd Edition, Society for Industrial and Applied Mathmatics, Philadelphia, PA, 2003.

VITA

Donald Eugene Bruss received a Bachelor of Science degree in nuclear engineering from Oregon State University in June of 2009. He began graduate school at Texas A&M University in August of 2009 in pursuit of a Masters of Science degree in nuclear engineering, to be awarded in May of 2012. He will begin working on a PhD in nuclear engineering in Spring 2012. His areas of interest include predictive simulation of multiphysics problems, finite element and Monte Carlo methods, and parallelizable algorithms for these problems.

Mr. Bruss can be contacted by mail through the Texas A&M University Nuclear Engineering Department, 3133 TAMU, College Station, TX, 77843. His university email address is BrussD@neo.tamu.edu.