

# LEAST-SQUARES FINITE-ELEMENT DISCRETIZATION OF THE NEUTRON TRANSPORT EQUATION IN SPHERICAL GEOMETRY

C. KETELSEN\*, T. MANTEUFFEL\* AND J. B. SCHRODER†

**Abstract.** The main focus of this paper is the numerical solution of the Boltzmann transport equation for neutral particles through mixed material media in a spherically symmetric geometry. Standard solution strategies, like the Discrete Ordinates Method (DOM), may lead to nonphysical approximate solutions. In particular, a point source at the center of the sphere yields undesirable *ray effects*. Posing the problem in spherical coordinates avoids ray effects and other non-physical numerical artifacts in the simulation process, at the cost of coupling all angles in the PDE setting. In addition, traditional finite element or finite difference techniques for spherical coordinates often yield incorrect scalar flux at the center of the sphere, known as flux dip, and oscillations near steep gradients.

In this paper, a least-squares finite element method with adaptive mesh refinement is used to approximate solutions to the non-scattering one-dimensional neutron transport equation in spherically symmetric geometry. It is shown that the resulting numerical approximations avoid flux dip and oscillations. The least-squares discretization yields a symmetric positive definite linear system which shares many characteristics with systems obtained from Galerkin finite element discretization of totally anisotropic elliptic PDEs. In general, standard Algebraic Multigrid (AMG) techniques fail to scale on non-grid-aligned anisotropies. In this paper, a new variation of Smoothed Aggregation (SA) is employed and shown to be essentially scalable. The effectiveness of the method is demonstrated on several mixed-media model problems.

**1. Introduction.** The main focus of this paper is the numerical solution of the Boltzmann transport equation for neutral particles in a spherically symmetric domain with mixed media. A First-Order System Least-Squares (FOSLS) finite element discretization is developed and shown to be accurate and to avoid flux dip near the center of the sphere and oscillations near steep fronts caused by material boundaries. The least-squares functional provides a locally sharp and globally accurate a posteriori error measure which is used in conjunction with an adaptive mesh strategy, ACE, that is based on Accuracy per Computational cost [7, 5]. FOSLS discretization of the non-scattering transport operator in spherically symmetric geometry yields a linear system that resembles an anisotropic diffusion operator and is, therefore, amenable to solution by a recently developed Smoothed Aggregation (SA) [25] algorithm specifically aimed at anisotropic diffusion problems [22]. In this context, the SA algorithm is shown to converge independently of the number of refinement levels. This paper will demonstrate that the FOSLS-ACE-SA approach provides an accurate and very efficient solution algorithm.

A broader motivation for this work is the application of FOSLS-ACE-SA to more general first-order hyperbolic conservation laws. In particular, numerical solution of Boltzmann transport in more general geometry using the Discrete Ordinates Method (DOM) and Diffusion Synthetic Acceleration (DSA) on irregular, isoparametric, or

---

\*Department of Applied Mathematics, 526 UCB, University of Colorado, Boulder, CO 80309, email: {ketelsen@colorado.edu, tmanteuf@colorado.edu}. This work was performed under the auspices of the U.S. Department of Energy under grant numbers DE-FC02-03ER25574 and DE-NA0002376, Lawrence Livermore National Laboratory under contract numbers B568677, and the National Science Foundation under grant number CBET-1249858.

†Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, P.O. Box 808, L-561, Livermore, CA 94551, email: schroder2@llnl.gov. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-JRNL-656198)

adaptively refined grids using the self-adjoint form of the transport operator involves solving a large number of problems of a character similar to those examined in this paper. Since SA algorithms are multilevel, they have potential for parallel scalability, thus overcoming a drawback to the conventional strategy of transport sweeps.

The next two subsections provide background on Boltzmann transport for neutral particles and provide broader motivation for this work.

**1.1. Boltzmann Transport.** The Boltzmann transport equation describes the movement of neutral particles through material media. The primary quantity of interest is the angular flux,  $\psi(\mathbf{x}, \Omega)$ , which represents the flux of particles through spatial location  $\mathbf{x}$ , moving in direction  $\Omega$ . In general Cartesian coordinates, the steady-state, single-energy, isotropic transport equation can be written as

$$\Omega \cdot \nabla \psi + \sigma_t \psi = \frac{\sigma_s}{4\pi} \int_{\mathcal{S}} \psi d\Omega + q, \quad (\mathbf{x}, \Omega) \in \mathcal{D} \times \mathcal{S}, \quad (1.1)$$

$$\psi = g, \quad \mathbf{x} \in \partial\mathcal{D}, \quad \mathbf{n} \cdot \Omega < 0, \quad (1.2)$$

where  $\mathcal{D}$  is the spatial domain and  $\mathcal{S}$  is the unit sphere. The first term in (1.1) represents the streaming of particles through the point  $\mathbf{x}$  in the direction  $\Omega$ , the second term represents particles that are lost due to a collision with the media, the third term represents particle rescattering, and  $q$  represents a particle source. The coefficients,  $\sigma_t$  and  $\sigma_s$ , are the total and scattering cross-sections, respectively, and, like the source term,  $q$ , may have spatial dependence. The function  $g$  defines inflow conditions on the boundary of the domain. In the general case, a solution process for the neutron transport equation involves discretization schemes for space and angle, and a solver for resulting linear system of equations.

The use of the Discrete Ordinates Method (DOM) for angular discretization is ubiquitous in the transport community. Here, a finite set of directions,  $\{\Omega_j\}_{j=1}^N$ , are chosen, usually quadrature points of the integral in the scattering operator, and (1.1) is replaced by:

$$\Omega_j \cdot \nabla \psi_j(\mathbf{x}) + \sigma_t \psi_j(\mathbf{x}) = \sigma_s \phi(\mathbf{x}) + q_j(\mathbf{x}), \quad j = 1, \dots, N, \quad (1.3)$$

$$\phi(\mathbf{x}) = \sum_j \omega_j \psi_j(\mathbf{x}), \quad (1.4)$$

where  $\omega_j$  are quadrature weights,  $\phi(\mathbf{x})$  is the scalar flux, and  $\psi_j(\mathbf{x}) = \psi(\mathbf{x}, \Omega_j)$ .

The primary advantage of such methods, coupled with a standard spatial discretization method (e.g. Corner Balance, Weighted-Diamond Differences, Discontinuous Galerkin, etc [1]), is the decoupling of the angles in the steady-state streaming operator, the left-hand side in (1.3). A popular and effective solution approach, known as Diffusion Synthetic Acceleration (DSA), involves solving (1.3) for each angle independently, then updating the scalar flux with a diffusion equation [12, 16]. The DSA algorithm can be summarized as follows: given  $\phi^\ell$ , set

$$(\Omega_j \cdot \nabla + \sigma_t I) \psi_j^{\ell+1/2} = \sigma_s \phi^\ell + q, \quad j = 1, \dots, N, \quad (1.5)$$

$$\phi^{\ell+1/2} = \sum_j \omega_j \psi_j^{\ell+1/2}, \quad (1.6)$$

$$(-\nabla \cdot \frac{1}{3\sigma_t} \nabla + \sigma_a I) \delta^{\ell+1/2} = \sigma_s (\phi^{\ell+1/2} - \phi^\ell), \quad (1.7)$$

$$\phi^{\ell+1} = \phi^{\ell+1/2} + \delta^{\ell+1/2}, \quad (1.8)$$

where  $\sigma_a = \sigma_t - \sigma_s$  is the absorption cross-section.

The DSA algorithm requires the solution of (1.5), the so-called transport sweeps, for every angle,  $\Omega_j$ , at every iteration. The traditional choice of spatial discretization leads to a triangular or block-triangular linear systems that must be solved repeatedly in each angle [1, 16]. Solution of a triangular linear systems is inherently serial which has the potential to degrade performance on parallel architectures. A number of very effective parallel schemes have been developed to alleviate this difficulty [3, 19], yielding algorithms that scale with  $O(P^{\frac{1}{d}})$ , where  $P$  is the number of processors and  $d$  is the spatial dimension. Irregular grids, isoparametric grids, and adaptively refined grids degrade the performance of these algorithms.

A recent focus on the self-adjoint form of the transport operator, which leads to a least-squares finite element discretization of the sweep equation, (1.5), yields a linear system that shares many similarities with a standard finite element discretization of an anisotropic, second-order system [11]. The benefits of this approach is that it produces a linear system of equations that is symmetric positive definite (SPD) and has the potential to be solved efficiently in parallel by multilevel algorithms such as SA [22]. The demonstration of such a scheme is one focus of this paper.

The main drawback of the FOSLS method in the context of neutron transport is that the discretization is not locally conservative. However, the FOSLS solution converges rapidly to the true solution (especially when coupled with adaptive refinement) and, thus, satisfies approximate local conservation on a sufficiently refined mesh. The results in this paper are in a different context, but demonstrate the accuracy and efficiency of the FOSLS-ACE-SA approach

**1.2. Transport in Spherically Symmetric Geometry.** As stated above, this paper is focused on Boltzmann transport in spherically symmetric geometry (see figure (2.1)). In this setting, a drawback of the DOM method is that when a strong source is placed at the sphere center, unphysical oscillations of the discrete solution are observed far away from the source. These unphysical artifacts are called *ray effects* and typically necessitate the use of a large number of discrete directions [16, 14, 15].

One solution to this difficulty is to formulate the transport problem in spherical coordinates. This creates a first-order coupling of the unknowns in angle and eliminates the ray effects (see equation (2.1)). However, since the angles do not decouple, transport sweeps do not separate, and one must solve for all angles of the solution simultaneously. The focus of this paper will be on the non-scattering operator, the left side of (2.1), denoted by  $T$ . This is equivalent to solving all the sweeps in (1.5). The diffusion solve, equivalent to (1.7), becomes an easily solved, one-dimensional equation.

The characteristic curves for this first-order hyperbolic operator,  $T$ , curve through the domain (see figure (3.1)). Applying a least-squares finite element discretization to this operator again yields linear systems that are similar to those that would arise from a Galerkin discretization of anisotropic diffusion. Now, however, the anisotropy curves through the domain along the characteristic curves.

The aim of this paper is to demonstrate the potential effectiveness of the FOSLS-ACE-SA approach on the non-scattering part of the neutron transport equation in spherical geometry. It will be demonstrated that the discretization, coupled with adaptive refinement, provides a faithful representation of the true solution while avoiding problematic numerical artifacts. It will also be demonstrated that the resulting linear systems can be solved with optimal algorithmic complexity by a recently developed algebraic algebraic multilevel method, a modification of SA. While a parallel

implementation of the solver is beyond the scope of this paper, the reader is directed towards literature discussing the parallelization of such methods [24].

We remark that this problem is an example of a broader class of problems, for example, first-order hyperbolic conservation laws. FOSLS discretization of such systems yields similar discrete linear systems and FOSLS-ACE-SA provides a viable solution technique [8, 9]. In particular, consider transport sweeps (1.5) in cartesian coordinates, but implemented on an irregular or isoparametric grid. Straight characteristics on such a grid yield discrete systems that resemble curved characteristics on a regular grid. These grids, along with adaptively refined or irregular, triangulated grids, are examples for which parallel sweeps become problematic. The approach here can be applied on such grids and yield the potential for both optimal algorithmic complexity and scalable parallel implementation.

The remainder of the paper is organized as follows. Section 2 of the paper describes the problem formulation in spherical geometry and some numerical pitfalls that are sometimes encountered in the solution process. Section 3 is concerned with the FOSLS methodology, finite element error bounds, and a discussion of the character of the linear system that results. In Section 4, exact solutions of a class of mixed media problems in spherical geometry will be presented. The numerical tests in Section 6 will address problems from this class. Section 5 contains a brief description of the recently developed algebraic multigrid method algorithm, a modification of SA, that will be used as the linear solver. In Section 6, FOSLS-ACE-SA approach will be applied to three test problems and the results will be discussed. Section 7 contains conclusions and future directions.

**2. Problem Formulation.** Consider the case of a spherical domain where the material properties of the domain and the source are radially symmetric. In this case, the angular flux is the same along any radial line in the spatial domain that passes through the origin (c.f. figure 2.1). In this case, the spatial dependence of the problem is reduced to a single dimension and the streaming angle,  $\Omega$ , is reduced to the unit circle. Denote the distance from the sphere center by  $\rho$  and let  $\mu$  be the cosine of the streaming angle,  $\theta$ .




Fig. 2.1: Spherical Domain

On the domain,  $\mathcal{R} = [0, 1] \times [-1, 1]$ , the 1D, spherically symmetric transport

equation in conservative form can be written as

$$\frac{\mu}{\rho^2} \frac{\partial}{\partial \rho} (\rho^2 \psi) + \frac{\partial}{\partial \mu} \left[ \frac{(1 - \mu^2) \psi}{\rho} \right] + \sigma_t \psi = \frac{\sigma_s}{2} \int_{-1}^1 \psi d\mu + q, \quad (2.1)$$

where the source term and cross-section are taken to be nonnegative but possibly discontinuous. In the remainder of the paper, we consider the non-scattering operator, omitting the scattering integral, and focus on the equation

$$T\psi = \frac{\mu}{\rho^2} \frac{\partial}{\partial \rho} (\rho^2 \psi) + \frac{\partial}{\partial \mu} \left[ \frac{(1 - \mu^2) \psi}{\rho} \right] + \sigma \psi = q. \quad (2.2)$$

Note that the total cross-section,  $\sigma_t$ , is now denoted simply as  $\sigma$  throughout the remainder of the paper. Vacuum boundary conditions are imposed on the inflow and reflective boundary conditions at the sphere's center:

$$\psi(1, \mu) = 0, \quad \mu \leq 0, \quad (2.3)$$

$$\psi(0, \mu) = \psi(0, -\mu), \quad \mu \geq 0. \quad (2.4)$$

Equation (2.2) can be considered a first-order equation in two dimensions and can be discretized using a variety of techniques [16]. Naive discretization techniques often produce a number of numerical artifacts. The first is the so-called *flux dip* that presents as a severe underestimation in the numerical approximation of the scalar flux near the sphere center, where the scalar flux is defined by

$$\phi(\rho) = \frac{1}{2} \int_{-1}^1 \psi(\rho, \mu) d\mu. \quad (2.5)$$

The second common numerical artifact is the appearance of non-physical oscillations of the scalar flux near material interfaces. Cartoons of these phenomena are depicted in figures 2.2 a) and b).




Fig. 2.2: Common Numerical Artifacts

**3. Least Squares Discretization.** In nonconservative form and scaled by  $\rho$ , equation (2.2) can be written as

$$\rho T\psi = \rho\mu \frac{\partial\psi}{\partial\rho} + (1 - \mu^2) \frac{\partial\psi}{\partial\mu} + \rho\sigma\psi = \rho q \quad (3.1)$$

We recast the (3.1) as a minimization of the residual in the  $L^2$  norm. That is, find  $\psi^*$  such that

$$\psi^* = \arg \min_{\psi \in \mathcal{V}} \mathcal{G}(\psi; q) := \arg \min_{\psi \in \mathcal{V}} \|\rho(T\psi - q)\|_2^2, \quad (3.2)$$

where

$$\mathcal{V} = \{\psi : \rho\psi \in L^2(\mathcal{R}) : \rho T\psi \in L^2(\mathcal{R}), \psi \text{ satisfies (2.3), (2.4)}\}. \quad (3.3)$$

For convenience we define the vector  $b = [\rho\mu, (1 - \mu^2)]^T$  and rewrite the model problem in simplified form as

$$\rho T\psi = b \cdot \nabla\psi + \rho\sigma\psi = \rho q, \quad (3.4)$$

where the gradient vector is defined with respect to  $\rho$  and  $\mu$  in the usual way as  $\nabla = [\partial/\partial\rho, \partial/\partial\mu]^T$ .

Continuity and coercivity of the functional  $\mathcal{G}(\psi; 0)$  can be established by relating back to the original geometry,  $\mathcal{D} \times \mathcal{S}$ , as described in (1.1) and (1.2), where  $\mathcal{D}$  is now a sphere of radius 1.0. Given  $\psi(\rho, \mu) \in \mathcal{V}$ , consider the corresponding spherically symmetric function,  $\Psi(\mathbf{x}, \Omega) \in L^2(\mathcal{D} \times \mathcal{S}^2)$ . It is easy to show that

$$\|\Psi\|_{L^2(\mathcal{D} \times \mathcal{S})}^2 = 8\pi^2 \|\rho\psi\|_{L^2(\mathcal{R})}^2. \quad (3.5)$$

Now, let  $\mathcal{T}$  be the streaming operator on the left side of (1.1). A similar calculation yields

$$\|\mathcal{T}\Psi\|_{L^2(\mathcal{D} \times \mathcal{S})}^2 = 8\pi^2 \|\rho T\psi\|_{L^2(\mathcal{R})}^2. \quad (3.6)$$

Thus, equation (3.1) posed on  $\mathcal{R}$  inherits the properties of the more general equation (1.1) posed in  $\mathcal{D} \times \mathcal{S}$ . Using (3.5), (3.6), together with results from [8] and [18], one can establish the Poincare inequality,

$$\|\rho\psi\| \leq C \|\rho T\psi\|, \quad \forall \psi \in \mathcal{V}. \quad (3.7)$$

(We have reverted to omitting the subscript on the norm, which is assumed to be the  $L^2$ -norm on  $\mathcal{R}$ .) Since  $\mathcal{D}$  is a sphere of radius 1.0, we have  $C = 1.0$ ,

Continuity and coercivity of the functional,  $\mathcal{G}(\psi, 0)$ , also follows from the same arguments; there exist positive constants,  $c$  and  $C$ , such that, for every  $\psi \in \mathcal{V}$ ,

$$c (\|\rho\psi\|_2^2 + \|b \cdot \nabla\psi\|_2^2) \leq \mathcal{G}(\psi; 0) \leq C (\|\rho\psi\|_2^2 + \|b \cdot \nabla\psi\|_2^2). \quad (3.8)$$

Define the norm on  $\mathcal{V}$  to be

$$\|\psi\|_{\mathcal{V}} := (\|\rho\psi\|_2^2 + \|b \cdot \nabla\psi\|_2^2)^{1/2}. \quad (3.9)$$

Equations (3.7) can be used to show that  $\mathcal{V}$  is a Hilbert space under the  $\mathcal{V}$ -norm.

Minimization (3.2) is equivalent to the following weak form:

$$\text{find } \psi^* \in \mathcal{V} \text{ s.t. } \langle \rho T \psi, \rho T v \rangle = \langle \rho q, \rho T v \rangle, \quad \forall v \in \mathcal{V}. \quad (3.10)$$

Equation (3.8) guarantees that (3.10) has a unique solution for every  $q$  such that  $\rho q \in L^2(\mathcal{R})$ .

Next, (3.10) is restricted to a finite dimensional subspace,  $\mathcal{V}^h \subset \mathcal{V}$ . Let  $\xi_k^h$ , for  $k = 1 \dots N$ , be a basis for  $\mathcal{V}^h$ . Then,  $\psi^h$  can be found by solving the familiar linear system,

$$A \mathbf{u} = \mathbf{f}, \quad (3.11)$$

where

$$(A)_{ij} = \langle \rho T \xi_j^h, \rho T \xi_i^h \rangle \quad \text{and} \quad \mathbf{f}_i = \langle \rho q, \rho T \xi_i^h \rangle. \quad (3.12)$$

From this definition it is easy to see that the stiffness matrix,  $A$ , is symmetric positive definite.

To get an idea of the likelihood of the success of an AMG method on the resulting linear system it is helpful to look at the leading term in the formal normal of the formulation; that is, consider  $(\rho T)^*(\rho T)$ . In the present case, the leading term of  $(\rho T)^*(\rho T)$  looks like an totally anisotropic diffusion term of the form

$$-\nabla \cdot [ b b^T ] \nabla. \quad (3.13)$$

The lines of strong anisotropy are the characteristic curves of the operator  $b \cdot \nabla \psi$ . The characteristics are everywhere tangent to the vector  $b$  and are given by the equation

$$\rho = \sqrt{\frac{\mu_0^2 - 1}{\mu^2 - 1}}, \quad (3.14)$$

where  $\mu_0$  parametrizes the inflow boundary,  $\{(1, \mu) : -1 \leq \mu \leq 0\}$ . Several of these curves are plotted in figure 3.1. It is clear from the picture that the anisotropy varies significantly over the domain and, on an orthogonal grid, are generally non-grid-aligned. In the case when the anisotropy is grid-aligned, a successful AMG method aggregates (or semicoarsens) unknowns in the direction of strong anisotropy. It has been shown in a number of papers that standard AMG solvers perform poorly when the anisotropy is non-grid-aligned (c.f. [6, 22]). One goal of this paper is to demonstrate an AMG algorithm, in this case an SA algorithm, that achieves optimal complexity on such systems. (See section 5).

In the case of smooth solutions, finite element convergence bounds for the FOSLS discretization can be established in the standard fashion. Assume that the finite dimensional subspace,  $\mathcal{V}^h$ , involves piecewise polynomials of degree  $p$  and admits interpolations operator,  $\mathcal{I}^h$ , with the approximation property

$$\|\psi - \mathcal{I}^h \psi\|_k \leq C_{k,s} h^{1+s-k} \|\psi\|_{1+s}, \quad (3.15)$$

for  $0 \leq k \leq s \leq p$ . Then, the FOSLS solution satisfies the bound

$$\|\rho T(\psi - \psi^h)\| \leq \|\rho T(\psi - \mathcal{I}^h \psi)\| \leq C \|\psi - \mathcal{I}^h \psi\|_1 \leq C h^s \|\psi\|_{1+s}, \quad (3.16)$$




Fig. 3.1: Characteristic Curves

for  $0 < s \leq p$ , where  $C$  is a generic constant.

Solutions to hyperbolic equations of the form (3.1) can be discontinuous along characteristics in the presence of discontinuous  $\sigma$  and  $q$ . In test problems 2 and 3 below, the solutions are  $C^0$ , but are only in  $H^\alpha$  for  $\alpha < 1$ . For the grid-aligned case, it can be shown that

$$\|\rho T(\psi - \mathcal{I}^h \psi)\| \leq Ch^\alpha \|\psi\|_\alpha, \quad (3.17)$$

for  $\alpha \leq p$ , where  $C > 0$  is a grid-independent constant. The exact bound for the non-grid-aligned case remains an open question. However, the theoretical limit for the grid-aligned case and other results offer some insight. In [23], an interpolation operator is established that satisfies the bound

$$\|\rho(\psi - \mathcal{I}^h \psi)\| \leq \|\psi - \mathcal{I}^h \psi\| \leq Ch^s \|\psi\|_s, \quad (3.18)$$

for  $0 < s \leq p$ . The Poincare inequality (3.7) yields

$$\|\rho(\psi - \mathcal{I}^h \psi)\| \leq C \|\rho T(\psi - \mathcal{I}^h \psi)\|. \quad (3.19)$$

Thus, the functional cannot converge any faster than the  $L^2$  norm of the error.

These are upper bounds on the error and specific problems may converge more quickly. That is the case in the test problems in section 6. It will be shown below that the non-smooth behavior is confined to the material interface in test 2, and along a special characteristic in test 3. For this reason, adaptive refinement is used. In this context, a more useful measure of the accuracy obtained per computation cost is to consider the error as a function of the number of degrees of freedom used. Since the SA algorithm has optimal complexity, the computational cost is directly related to the number of degrees of freedom.

**4. Mixed Media Solution.** The mixed-media test case in which the source,  $q$ , and cross-section,  $\sigma$ , are piecewise-constant with discontinuity at some point,  $\rho_{mid}$ , in the domain is examined in the numerical tests in section 6. In the spherical setting, this corresponds to an inner sphere and outer shell with distinct material and source properties. For example, one might have

$$\sigma(\rho) = \begin{cases} \sigma_1 & \text{for } 0 \leq \rho \leq \rho_{mid} \\ \sigma_2 & \text{for } \rho_{mid} \leq \rho \leq 1 \end{cases}, \quad q(\rho) = \begin{cases} q_1 & \text{for } 0 \leq \rho \leq \rho_{mid} \\ q_2 & \text{for } \rho_{mid} \leq \rho \leq 1 \end{cases}. \quad (4.1)$$



The exact solution to this problem can be written explicitly for three different portions of the domain as follows:

For  $0 \leq \rho \leq \rho_{mid}$  ,

$$\psi(\rho, \mu) = \frac{q_1}{\sigma_1} \left(1 - e^{-\sigma_1(\rho\mu+s)}\right) + \frac{q_2}{\sigma_2} e^{(\sigma_2-\sigma_1)s-\sigma_1\rho\mu} (e^{-\sigma_2 s} - e^{-\sigma_2 t}) .$$

For  $\rho_{mid} \leq \rho \leq 1$  and  $-1 \leq \mu \leq \sqrt{1 - (\rho_{mid}/\rho)^2}$  ,

$$\psi(\rho, \mu) = \frac{q_2}{\sigma_2} \left(1 - e^{-\sigma_2(\rho\mu+t)}\right) .$$

For  $\rho_{mid} \leq \rho \leq 1$  and  $\sqrt{1 - (\rho_{mid}/\rho)^2} \leq \mu \leq 1$  ,

$$\psi(\rho, \mu) = \frac{q_1}{\sigma_1} e^{\sigma_2 s - \sigma_2 \rho \mu} (1 - e^{-2\sigma_1 s}) + \frac{q_2}{\sigma_2} e^{2(\sigma_2 - \sigma_1)s - \sigma_2 \rho \mu} (e^{-\sigma_2 s} - e^{-\sigma_2 t}) + \frac{q_2}{\sigma_2} \left(1 - e^{-\sigma_2(\rho\mu-s)}\right) ,$$

where the distinct portions of the domain are separated by the curves

$$s = \sqrt{\rho_{mid}^2 - \rho^2(1 - \mu^2)} \quad \text{and} \quad t = \sqrt{1 - \rho^2(1 - \mu^2)} .$$

While the exact solution for this problem with discontinuous cross-section and source is itself continuous, there are two areas in the domain which exhibit sharp changes in the solution due to discontinuities in the first derivative. These sharp changes occur along the material discontinuity,  $\rho = \rho_{mid}$ , and also along the characteristic tangent to the material discontinuity (c.f. Figure 4.1).




Fig. 4.1: Osculating Characteristic

At the point of osculation,  $(\rho_{mid}, 0)$ , the behavior of  $\psi$  in the  $\rho$  direction is determined by the degree of osculation. In this case, as  $\rho \rightarrow \rho_{mid}$  from below, a square root behavior leads to an unbounded derivative,

$$\frac{d\psi}{d\rho}(\rho_{mid}, 0) \simeq (\rho - \rho_{mid})^{-1/2} . \quad (4.2)$$

In the absence of absorption ( $\sigma = 0$ ), this behavior is carried along the characteristic (to the left in figure 4.1) without modification. With absorption, the behavior is attenuated, depending on the strength of the absorption. In test 3,  $\sigma = 0.5$  in the

outer shell and the behavior persists until the outflow boundary (see figure 6.6a). In test 2,  $\sigma = 100$  in the outer shell and the behavior disappears quickly (see figure 6.3a). In either case, the behavior along the osculating characteristic implies  $\psi \in H^\alpha$  for  $\alpha < 1.0$ . The numerical results below demonstrate that finite element convergence is much faster for test 2 than for test 3.

In order to accurately capture these sharp changes in the solution, it is necessary to have a very fine grid resolution along the material interface in test 2 and along the osculating characteristic in test 3. However, away from these two curves the solution is relatively smooth and does not require fine resolution. A natural way to capture this behavior without wasting significant computational effort is to use adaptive mesh refinement. Fortunately, accurate AMR is relatively easy to obtain in the FOSLS framework. A local version of the upper bound (continuity) in (3.8) can be used to show that the local FOSLS functional,

$$\epsilon_\tau = \|\rho T\psi - \rho q\|_\tau,$$

provides a locally sharp estimation of the error in element  $\tau$ . The lower bound (3.8) shows that the functional is also a globally reliable error estimator. An efficient adaptive refinement strategy based on the FOSLS functional can be obtained by choosing to refine elements with large local FOSLS functional so as to maximize the accuracy obtained per increase in computational cost. This refinement strategy is called ACE and presented in detail in [5, 7].

**5. Smoothed Aggregation Algebraic Multigrid.** AMG methods are a popular and effective approach for solving the sparse linear systems that arise from discretizing partial differential equations because they are scalable or optimal and solve a linear system with  $n$  degrees of freedom with  $O(n)$  work. However while AMG was designed for elliptic problems, we noted in Section 3 that the case of non-grid-aligned anisotropic diffusion is not adequately addressed by existing AMG methods and the matrices here resemble such highly non-grid-aligned operators. Therefore, we apply the recently developed SA solver described in [22], which was shown to be scalable for a range of model non-grid-aligned anisotropic diffusion problems.

We first give an overview of how multigrid solves the linear system  $Ax = b$ , where  $x, b \in \mathbb{R}^n$ ,  $A \in \mathbb{R}^{n \times n}$  and  $A$  is symmetric and positive definite (SPD). Multigrid consists of two main parts, relaxation and the coarse grid correction. The coarse grid correction uses interpolation operators (often called prolongation and restriction) to transfer vectors and matrices between the fine and coarse “grids”, denoted by  $\mathbb{R}^n$  and the lower dimensional (coarse) vector space  $\mathbb{R}^{n_c}$ . The prolongator  $P : \mathbb{R}^{n_c} \rightarrow \mathbb{R}^n$  and we make the common choice of restriction equal  $P^T$ .

Multigrid is optimal when relaxation and coarse grid correction are complementary. For a classic model diffusion problem, relaxation is simple, like Gauss-Seidel, and effectively reduces high frequency (high energy) error. The remaining low frequency error is then accurately captured by  $P$  and moved to the coarse grid where it is eliminated. In the general setting, this low frequency error is referred to as “algebraically smooth” and must be captured accurately by interpolation. The result is that all error components are uniformly damped with an  $O(n)$  cost. A single two-grid

iteration is defined in (5.1), where  $x$  is some initial guess.

$$\text{Relax on } Ax = b \tag{5.1a}$$

$$\text{Correct } x \leftarrow x + P(P^T AP)^{-1} P^T (b - Ax) \tag{5.1b}$$

$$\text{Relax on } Ax = b \tag{5.1c}$$

$$\text{Return } x \tag{5.1d}$$

For optimality, (5.1) is applied recursively to the correction step to solve  $P^T APx_c = P^T(b - Ax)$ , until the coarse grid problem is of trivial size. Multigrid V-cycles correspond to only one recursive call at the correction step, while F-cycles implement a V-cycle iteration at each coarse grid, i.e., coarse grids are visited more frequently, but still in a scalable way. An F( $k, m$ )-cycle refers to  $k$  relaxation sweeps in (5.1a) and  $m$  relaxation sweeps in (5.1c).

AMG methods automatically construct the hierarchy and associated coarse level matrices  $P^T AP$  during the so-called setup phase. The solve phase refers to (5.1). In order to assess optimality of our approach, we measure computational effort of both phases using work units, which are equivalent to one matrix-vector product with the finest level matrix. When the work units required to solve a problem does not change regardless of problem size, we can say that the method is  $O(n)$  or scalable.

The strategy described in [22] examines the three major parts to a SA method for the case of model rotated anisotropic diffusion problems and discerns that existing coarsening strategies<sup>1</sup> and relaxation methods are sufficient for scalable performance, but that existing interpolation methods are not. To augment  $P$ , classical SA is modified to accommodate much wider interpolation stencils where the weights are determined by the energy-minimization method from [21]. Classical SA essentially does one energy-minimizing step (with weighted-Jacobi and a narrow interpolation stencil) to find a low-energy, complementary interpolation operator. Here, many energy-minimization steps and a wider stencil are used to find a complementary interpolation operator. The result is a more robust method that has a higher setup cost, but a lower overall time to solution.

As described in [22], the coarsening strategy relies on the advanced strength measure developed in [20]. Coarsening uses strength to determine the coupling between degrees of freedom, and decide whether one degree of freedom can serve as a good representation on the coarse level for a group of degrees of freedom on the fine level.

The method from [20] extracts strength information from locally smoothed  $\delta$ -functions. Here, 4 local Jacobi relaxation sweeps are used. The strength values reflect how accurately algebraically smooth error can be approximated between degrees of freedom. As such, the strength drop-tolerances are all greater than 1.0 (as opposed to classical strength drop-tolerances which are between 0.0 and 1.0). Here, a drop-tolerance of 8.0 is used, which indicates that all entries within a factor of 8 of the strongest (i.e., smallest) connection are considered strong. See [20] for more detail. The energy minimization parameters use a wide interpolation stencil based on distance-six strength-of-connection information (corresponding to  $k = 6$  in the algorithm from [21]) and a total of 15 energy-minimization steps.

The notable modifications to the solver from [22] used here are as follows. Diagonally dominant rows are not coarsened because relaxation alone is sufficient to reduce the error. This is important for efficiency when  $\sigma$  is large. Additionally, the

---

<sup>1</sup>The coarsening strategy decides which degrees of freedom should exist on the next coarser level and form  $\mathbb{R}^{n_c}$

near nullspace modes used to help construct interpolation include the linears and the constant, as opposed to only the constant, which was used in [22]. The addition of this extra algebraically smooth information was found to be important to the solver's performance. These near nullspace modes are commonly used by multigrid as input and serve as constraints on interpolation as modes that must be interpolated well. These modes are usually determined by the dominant differential part of the weak form of the governing PDE, which is also the case here.

While the current implementation of the method from [22] is in serial, parallelization is not a serious issue. The algorithm relies on existing linear algebra operations already present in most parallel multigrid packages, i.e., sparse matrix-matrix multiplication and LAPACK routines. The main algorithmic substitution required would be using a parallel relaxation method, such as hybrid Gauss-Seidel.

Last, we note that while the method from [22] addressed AMG scalability for difficult model problems, the matrices here present additional challenges known to create issues for AMG, most importantly AMR and coefficient jumps. So, we do not expect performance to be as perfectly scalable as for the model problems tested in [22].

**6. Numerical Results.** Numerical experiments were carried out on three test problems of increasing difficulty. In all cases, we take  $(\rho, \mu) \in [0, 1] \times [-1, 1]$ . They are

$$\text{Test 1: } \sigma = 10^{-6}, \quad q = 1, \quad \text{for } 0 \leq \rho \leq 1,$$

$$\text{Test 2: } \sigma(\rho) = \begin{cases} 10^{-8} & \text{for } 0 \leq \rho \leq \rho_{mid} \\ 100 & \text{for } \rho_{mid} \leq \rho \leq 1 \end{cases}, \quad q(\rho) = \begin{cases} 1 & \text{for } 0 \leq \rho \leq \rho_{mid} \\ 0 & \text{for } \rho_{mid} \leq \rho \leq 1 \end{cases},$$

$$\text{Test 3: } \sigma(\rho) = \begin{cases} 1.0 & \text{for } 0 \leq \rho \leq \rho_{mid} \\ 0.5 & \text{for } \rho_{mid} \leq \rho \leq 1 \end{cases}, \quad q(\rho) = \begin{cases} 4.0 & \text{for } 0 \leq \rho \leq \rho_{mid} \\ 1.0 & \text{for } \rho_{mid} \leq \rho \leq 1 \end{cases},$$

where the material jump in Tests 2 and 3 occurs at  $\rho_{mid} = 0.4$ .

Test 1 is a simple, single-media problem with constant cross-section and source throughout the domain. Tests 2 and 3 are mixed media problems corresponding to concentric spheres with different material properties within. Test 2 represents a void surrounded by an optically thick shell with a unit source only within the void. Test 3 corresponds to a moderately thick shell surrounding a slightly thicker center.

For the three test problems we will measure the suitability of the FOSLS discretization is measured by considering the convergence of the discretization, the performance of the SA, preconditioned Conjugate Gradient solver on the associated linear system, and the character of the scalar flux  $\phi$  compared to the known exact solution.

In the mixed-media case we are particularly interested in seeing that the scalar flux avoids oscillations around the material jump and flux dip at the center of the sphere. The convergence of the discretization is measured both in the least-squares functional,

$$\epsilon_F = \|T\psi^h - \rho q\|_2,$$

and in the weighted  $L^2$  norm of the error with the known solution,

$$\epsilon_2 = \|\rho(\psi^h - I^h\psi)\|_2,$$

where  $\psi^I$  is the exact solution evaluated at the mesh grid points. Quadratic finite elements are used in each test along with uniform mesh refinement and adaptive mesh refinement where indicated.

**Test 1:** Figures 6.1a and 6.1b show the approximate angular and scalar flux, respectively. Figures 6.2a and 6.2b show the convergence of the approximate solution in the functional norm and  $L^2$ -norm, respectively. Note that in the single-media case the method converged very quickly with just uniform refinement with orders  $\mathcal{O}(h^{0.75})$  and  $\mathcal{O}(h^{1.25})$  for the functional and the  $L^2$ -norm, respectively.




Fig. 6.1: Test 1 - Solutions

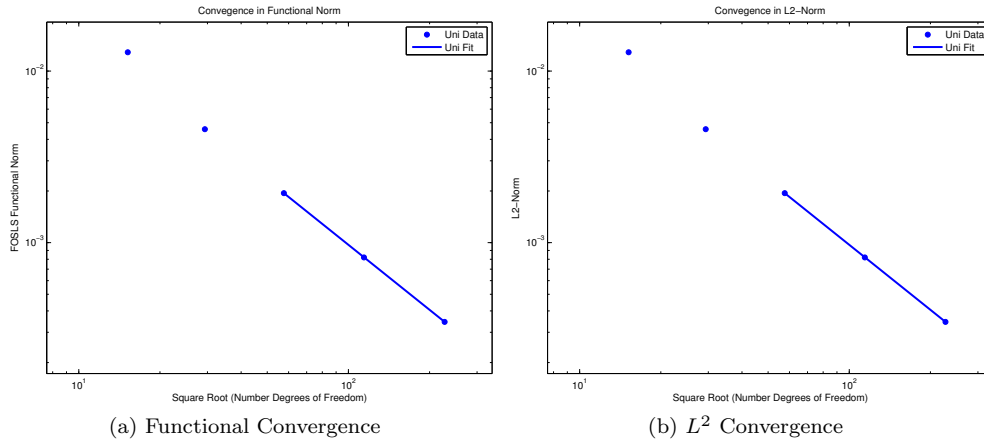


Fig. 6.2: Test 1 - Finite Element Convergence

**Test 2:** Figure 6.3a shows the approximate angular flux. Note the sharp change of slope along the material interface as the characteristics leave the inner sphere. The large absorption results in a rapid decrease toward the outflow boundary. Figure

6.3b shows the approximate scalar flux computed using adaptive mesh refinement compared to the exact solution. Note the lack of flux dip at the sphere center and unphysical oscillations near the material interface.

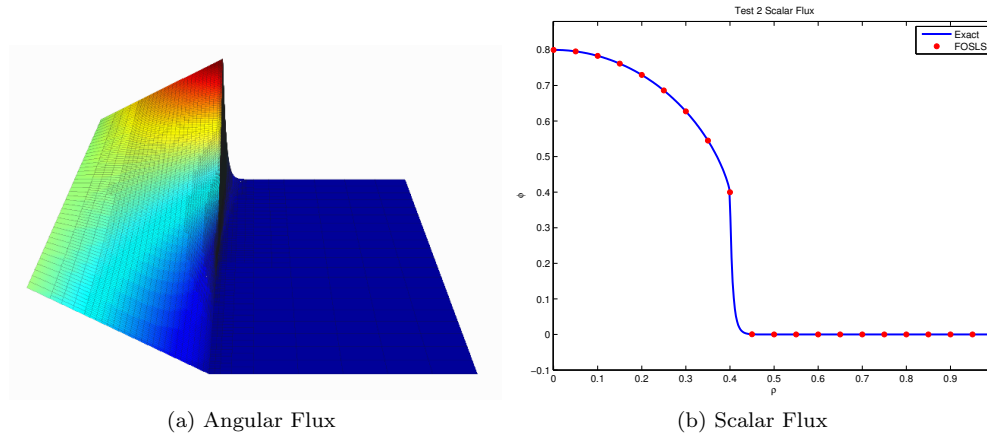


Fig. 6.3: Test 2 - Solutions




Fig. 6.4: Test 2 - Refined Mesh

Figure 6.4 shows the adaptively refined grid on the finest level. Notice the concentration of elements along the material interface, where the slope changes dramatically. The finite element space is constructed to have element boundaries along this interface, mitigating some of the difficulty of approximating the solution in this region.

Figures 6.5a and 6.5b show the convergence of the method to the true solution in the functional norm and  $L^2$ -norm, respectively. Convergence results are shown for both uniform and adaptive mesh refinement. In each measure it is clear that adaptive mesh refinement can be extremely beneficial in terms of accuracy per computational cost, with approximately 2% of the total degrees of freedom required with the adaptively refined mesh to achieve the same accuracy as on the uniform mesh. This is because it is only necessary to refine elements along the material interface.

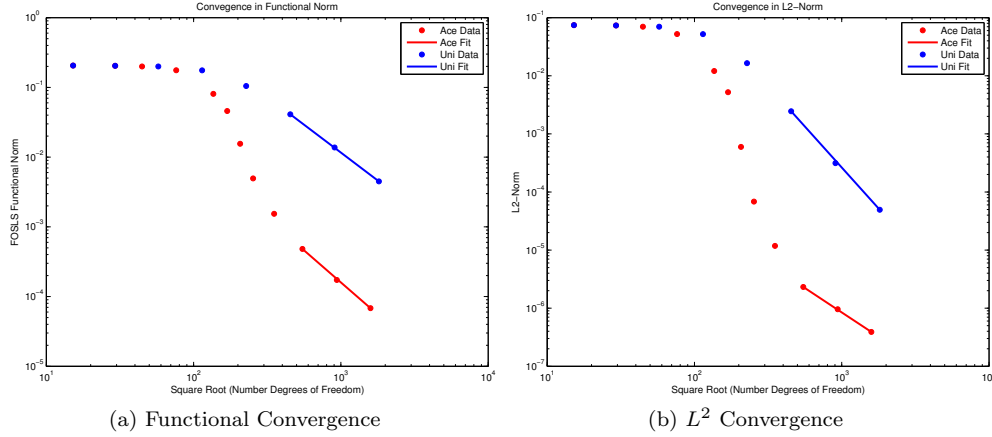


Fig. 6.5: Test 2 - Finite Element Convergence

**Test 3:** Figure 6.6a shows the approximate angular flux. Note that the square root behavior is carried along the osculating characteristic until the outflow boundary. Figure 6.6b shows the approximate scalar flux for an adaptively refined mesh as well as convergence results for both uniform and adaptively refined meshes. Again note the lack of presentation of the flux dip and spurious oscillations near the material interface.

Figure 6.7 shows the adaptively refined mesh. Note the concentration of elements along the osculating characteristic. The mesh interface is not so highly refined because the finite element space is chosen to have element boundaries along the interface and the change in absorption, thus change in slope, is not as dramatic as in the Test 3 Figures.

Figures 6.8a and 6.8b show the convergence of the method to the true solution in the functional norm and  $L^2$ -norm, respectively. Convergence results are shown for both uniform and adaptive mesh refinement. Again, it is clear that adaptive mesh refinement can be extremely beneficial in terms of accuracy per computational cost, with approximately 2% of the total degrees of freedom required with the adaptively refined mesh to achieve the same accuracy as on the uniform mesh. This is because it is only necessary to refine elements along the osculating characteristic. Convergence for this test is slower than for Test 2 because of the square root behavior of the solution along a curving characteristic.

**Smoothed Aggregation Results.** In this section, we report performance results for the SA method discussed in Section 5. F(2,2)-cycles are used to precondition CG to within a relative residual tolerance of  $10^{-8}$ . The relaxation scheme is symmetric Gauss-Seidel, i.e. a forward and then a backward sweep in both (5.1a) and (5.1c). F-cycles are chosen over V-cycles because they offer more scalable work estimates for the solver.

We report the iteration count, average convergence rate over all iterations ( $\gamma$ ), operator complexity (Op. Comp.), the number of degrees of freedom to the nearest thousand (ndof (K)), the cost of setup phase in work units (WU), the total cost of the solve phase in work units, and the solve phase WU per digit of accuracy (WU/digit).




Fig. 6.6: Test 3 - Solutions




Fig. 6.7: Test 3 - Refined Mesh

Operator complexity represents the total number of nonzeros in the hierarchy divided by the fine-grid number of nonzeros and is a relative measure of the storage cost of the method. The setup phase is measured in total WU, where all the operations, such as the sparse matrix-matrix multiplications during the energy minimization process that computes  $P$ , are measured. The solve phase is also measured in WU, showing both the total WU required to solve for 8 digits of accuracy and for the WU per digit of accuracy. The latter is the amount of work required to reduce the residual by one order of magnitude. This quantity is computed with  $c/|\log_{10}(\gamma)|$ , where  $c$  is the cycle complexity. The cycle complexity  $c$  is the cost in floating-point operations (relative to one fine grid matrix-vector multiply) to apply one F-cycle. This quantity is computed by weighting the number of nonzeros in the matrix at each multigrid level by the number of times a relaxation sweep is done at that level during one F-cycle. This total amount is then normalized by the number of nonzeros present in the original fine grid matrix to obtain  $c$ . The setup WU are shown to better understand the cost



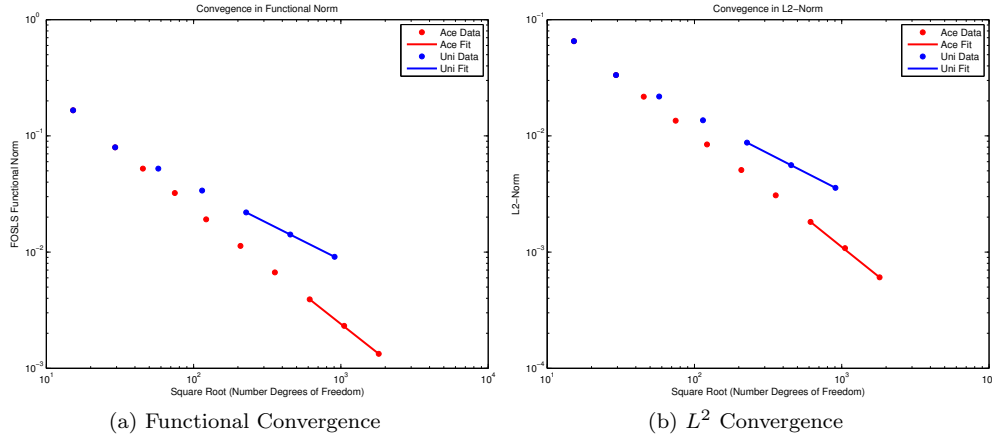


Fig. 6.8: Test 3 - Finite Element Convergence

of the nonstandard setup phase by allowing a comparison to a standard F-cycle solve phase with a moderate operator complexity. In other words, the solve phase work is typical for an efficient multigrid scheme and is a useful yardstick.

Table 6.1 depicts the results for a sequence of uniformly refined grids for Tests 2 and 3. The results for both problems show a slow growth in iterations and in WU required to solve the problem. Overall, the approach is significantly more scalable than classical SA, and more efficient at large grid sizes. For instance, consider a classical SA method, using the same F(2,2) cycle and sophisticated strength-of-connection method [20], but with standard filtered Jacobi prolongation smoothing. Here, the iterations for the last three Test 3 grids are 82, 125 and 194 leading to total solve phase WU of 493, 740 and 1144. The setup is much cheaper, with 82 WU required for each of the last three Test 3 grids. Whereas for the chosen approach, the WU required to solve the problem to 8 digits of accuracy grows slowly each uniform grid refinement. The corresponding iterations for the last three Test 2 grids are 101, 163 and 249, leading to even less scalable total solve phase WU of 586, 950 and 1457. The setup cost is again 82 WU for each of the last three Test 2 grids. Thus, the chosen approach is even more efficient relative to classic SA for Test 2. For both problems as the problem size increases, the chosen approach eventually becomes more efficient by costing less total WU.

Last, we examine the solver for a sequence of adaptively refined grids for Tests 2 and 3 in Tables 6.2 and 6.3, respectively. The grid numbering reflects the number of adaptive mesh refinements from an initial uniform mesh. The results for Test 2 show nearly scalable behavior, while Test 3 again exhibits a slow growth in iterations<sup>2</sup>. It should be noted that these matrices based on AMR are extremely difficult for classical multigrid methods. For instance when considering the same classical SA method as above, the iterations for the last three adaptive Test 3 grids are 482, 871

<sup>2</sup>This growth in iterations can be slowed by increasing the setup phase WU. If extremely wide interpolation stencils are used, leading to operator complexities around 3.25, then the iteration counts for the last three adaptive Test 3 grids are 19, 23 and 25. However, the setup phase now costs approximately 4,250 WU. Overall, care was taken to choose solver parameters that balance the setup and solve phase costs.

Grid	Test 2					Test 3				
	1	2	3	4	5	1	2	3	4	5
ndof (K)	13	51	206	821	3,281	3	13	51	206	821
Iterations	30	28	37	42	38	26	30	28	32	34
$\gamma$	0.54	0.51	0.60	0.64	0.61	0.49	0.53	0.51	0.56	0.58
Op. Comp.	1.78	1.80	1.59	1.55	1.55	1.78	1.77	1.67	1.60	1.57
Setup WU	1039	1222	799	771	775	804	878	849	810	801
Solve WU	314	308	340	370	335	241	315	270	301	310
Solve $\frac{\text{WU}}{\text{digit}}$	39	39	43	46	42	30	40	34	38	39

Table 6.1: SA results, uniform refinement, Tests 2 and 3

and 1501, leading to total solve phase WU of 2,964, 5,334 and 9,154. The setup is much cheaper, with 84 WU required for each of the last three adaptive Test 3 grids. The corresponding numbers for the last three adaptive Test 2 grids are 165, 251, and 442 iterations, leading to total solve phase WU of 988, 1492 and 2649. The setup is again 84 WU for each of the last three adaptive Test 2 grids. Comparing these numbers to the chosen approach, where the overall WU's required is much smaller and grows slowly, it is clear that we have made an appropriate choice for solving these problems.

Grid	2	3	4	5	6	7	8	9	10	11
ndof (K)	1	6	19	29	43	64	124	300	881	2,524
Iterations	14	25	28	27	31	33	35	36	36	41
$\gamma$	0.26	0.47	0.51	0.50	0.55	0.57	0.59	0.60	0.59	0.64
Op. Comp.	1.52	1.78	1.79	1.83	1.85	1.83	1.87	1.86	1.83	1.76
Setup WU	285	700	872	948	996	1041	1,205	1,293	1,305	1,183
Solve WU	63	229	298	299	353	375	422	444	424	464
Solve $\frac{\text{WU}}{\text{digit}}$	8	29	37	37	44	47	53	56	53	58

Table 6.2: SA results, adaptive refinement, Test 2.

**7. Conclusion** We have demonstrated that the FOSLS discretization coupled with adaptive mesh refinement provides an accurate discrete approximation of the spherically symmetric 1D transport equation in spherical coordinates. The method has good convergence properties and avoids troublesome unphysical numerical artifacts such as flux dip at the sphere center and oscillations near the material interface. We have also demonstrated the resulting linear system can be solved scalably using an algebraic multigrid method.

Future work will involve extending the FOSLS discretization methodology to more difficult problems related to neutron transport. This includes problems where discretization in spherical coordinates is advantageous but the domain, source, and material coefficients are not necessarily spherically symmetric. The method should also prove useful for problems involving general hyperbolic conservation laws. Future

Grid	2	3	4	5	6	7	8	9	10
ndof (K)	1	2	6	15	43	127	377	1,108	3,261
Iterations	22	24	26	25	28	30	36	40	53
$\gamma$	0.43	0.45	0.48	0.47	0.51	0.54	0.59	0.63	0.70
Op. Comp.	1.65	1.83	1.90	1.91	1.83	1.79	1.73	1.68	1.65
Setup WU	385	874	930	1053	1034	1025	951	899	851
Solve WU	100	233	254	298	314	337	380	407	523
Solve $\frac{\text{WU}}{\text{digit}}$	13	29	32	37	39	42	48	51	65

Table 6.3: SA results, adaptive refinement, Test 3.

work will also involve improving the performance and efficiency of the smoothed aggregation solver. Since the nature of the anisotropy in the resulting linear system is entirely defined by the geometry of the domain, it may be beneficial to define aggregation based on geometric data instead of (or in addition to) the matrix coefficients as is typically done in standard smoothed aggregation multigrid.

## REFERENCES

- [1] M. L. ADAMS AND E. W. LEWIS, *Fast Iterative Methods for Discrete-Ordinate Particle Transport Calculations*, Progress in Nuclear Energy, 40(1), 2002, pp. 3-159.
- [2] T. AUSTIN, AND T. MANTEUFFEL, *A Least-Squares Finite-Element Method for the Linear Boltzmann Equation with Anisotropic Scattering*, SIAM J. Numer. Anal., 44(2), (2006), pp. 540-560.
- [3] T. S. BAILEY AND R. D. FALGOUT, *Analysis of massively parallel discrete-ordinates transport sweep algorithms with collisions*, Proc. Int. Conf. on Math., Comp. Phys. and Reactor Phys., Saratoga Springs, NY, 2009.
- [4] M. BREZINA AND R. FALGOUT AND S. MACLACHLAN AND T. MANTEUFFEL AND S. MCCORMICK AND J. RUGE, *Adaptive smoothed aggregation ( $\alpha$ SA) multigrid*, SIAM Rev., (47) 2005, pp. 317-346.
- [5] M. BREZINA AND J. GARCIA AND T. MANTEUFFEL AND S. MCCORMICK AND J. RUGE AND L. TANG, *Parallel adaptive mesh refinement for first-order system least squares*, J. Num. Lin. Alg. Appl., (19) 2012, pp. 343-366.
- [6] A. J. CLEARY AND R. D. FALGOUT AND V. E. HENSON AND J. E. JONES AND T. A. MANTEUFFEL AND S. F. MCCORMICK AND G. N. MIRANDA AND J. W. RUGE, *Robustness and scalability of algebraic multigrid*, SIAM J. Sci. Comput., 21 (2000), pp. 1886-1908.
- [7] H. DE STERCK AND T. MANTEUFFEL AND S. MCCORMICK AND J. NOLTING AND J. RUGE AND L. TANG, *Efficiency-based h- and hp-refinement strategies for finite element methods*, J. Num. Lin. Alg. Appl., (15) 2008, pp. 249-271.
- [8] H. DE STERCK AND T. MANTEUFFEL AND S. MCCORMICK AND L. OLSEN, *Least-squares finite element methods and algebraic multigrid solvers for linear hyperbolic PDEs*, SIAM J. Sci. Comp., (26) 2004, pp. 31-54.
- [9] H. DE STERCK AND T. MANTEUFFEL AND S. MCCORMICK AND L. OLSEN, *Numerical Conservation Properties of  $H(\text{div})$ -Conforming Least-Squares Finite Element Methods for the Burgers Equation*, SIAM J. Sci. Comp., (26) 2005, pp. 1573-1597.
- [10] R. D. FALGOUT AND P. S. VASSILEVSKI, *On Generalizing the Algebraic Multigrid Framework*, SIAM J. Numer. Anal., (42) 2004, pp. 1669-1693.
- [11] J. HANSEN AND J. PETERSON AND J. MOREL AND J. RAGUSA AND Y. WANG, *A new least-squares transport equation compatible with voids*, Proc. 2013 Int. Conf. on Math. and Comp. Methods Applied to Nuclear Science and Engineering, 2013.
- [12] E. W. LARSEN, *Diffusion theory as an asymptotic limit of transport theory for nearly critical systems with small mean free path*, Ann. Nuclear Energy, (7) 1980, pp. 249-255.
- [13] K.D. LATHROP, *A comparison of angular difference schemes for one-dimensional spherical*

- geometry  $s_n$  equations*, Nucl. Sci. Eng., 134, (2000), pp. 239-264.
- [14] K. D. LATHROP, *Ray effects in discrete ordinate equations*, Nucl. Sci. Eng., (32) 1968, pp. 357-369.
  - [15] K. D. LATHROP, *Remedies for ray effects*, Nucl. Sci. Eng., (45) 1971, pp. 255-268.
  - [16] E. E. LEWIS AND W. F. MILLER, *Computational Methods of Neutron Transport*, John Wiley, New York, 1984.
  - [17] E. MACHORRO, *Discontinuous Galerkin finite element method applied to the 1-D spherical neutron transport equation*, J. Comp. Phys., 223(1), (2007), pp. 67-81.
  - [18] T. MANTEUFFEL, AND K. RESSEL, *Least-Squares Finite-Element Solution of the Neutron Transport Equation in Diffusive Regimes*, SIAM J. Numer. Anal., 35(2), (1998), pp. 806-835.
  - [19] M. M. MATHIS AND N. M. AMATO AND M. L. ADAMS, *A general performance model for parallel sweeps on orthogonal grids for particle transport calculations*, Proc. ACM Int. Conf. Supercomputing (ICS), Santa Fe, NM, 1980, pp. 255-263.
  - [20] L. N. OLSON AND J. B. SCHRODER AND R. S. TUMINARO, *A New Perspective on Strength Measures in Algebraic Multigrid*, Numer. Linear Algebra Appl., 17 (2010), pp. 713-733.
  - [21] L. N. OLSON AND J. B. SCHRODER AND R. S. TUMINARO, *A general interpolation strategy for algebraic multigrid using energy-minimization*, SIAM J. Sci. Comput., (33) 2011, pp. 966-991.
  - [22] J. B. SCHRODER, *Smoothed Aggregation Solvers for Anisotropic Diffusion*, Numer. Linear Algebra Appl., 19 (2012), pp. 296-312.
  - [23] L. R. SCOTT AND S. ZHANG, *Finite element interpolation of nonsmooth functions satisfying boundary conditions*, Math. Comp., 54 (1990), pp. 483-493.
  - [24] R. TUMINARO AND C. TONG, *Parallel smoothed aggregation multigrid: aggregation strategies on massively parallel machines*, Supercomputing 2000 Proc., 2000.
  - [25] P. VANĚK AND J. MANDEL AND M. BREZINA, *Algebraic Multigrid Based On Smoothed Aggregation for Second and Fourth Order Problems*, Computing, 56 (1996), pp. 179-196.