International Conference on Mathematics, Computational Methods & Reactor Physics (M&C 2009)
Saratoga Springs, New York, May 3-7, 2009, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2009)

PARALLEL DISCRETE ORDINATES METHODS IN THE SCEPTRE PROJECT

Shawn Pautz, Bill Bohnhoff, Clif Drumm and Wesley Fan
Sandia National Laboratories*
Albuquerque, NM 87185-1179
sdpautz@sandia.gov, wjbohnh@sandia.gov, crdrumm@sandia.gov, wcfan@sandia.gov

ABSTRACT

Although the transport of photons and electrons/positrons can be described by the same Boltzmann transport equation, the cross sections are very different, resulting in very different iterative convergence properties for the two particle types. The SCEPTRE project contains a suite of deterministic codes for solving the linear steady-state Boltzmann transport equation, including two very different solution algorithms: a sweeps-based approach for solving the first-order transport equation that is efficient for photon transport, and a conjugate-gradient algorithm for solving the second-order transport equation that is efficient for electron/positron transport. These two solution algorithms are coupled through a common set of interfaces to handle the finite-element approximation in space, the discrete-ordinates method in angle, the scattering source terms, input parsing, pre- and post-processing capabilities. We evaluate the two solution algorithms by comparing the run times for photon- and electron-transport problems, investigating parallel performance and effects from finite-element basis function types, preconditioning, and scaling with angular quadrature order and Legendre cross section expansion order.

Key Words: coupled photon-electron transport, finite elements, parallel

1. INTRODUCTION

SCEPTRE (**S**andia's **C**omputational **E**ngine for **P**article **T**ransport for **R**adiation **E**ffects) is a suite of deterministic codes for solving the linear steady-state Boltzmann transport equation. It solves the first- and/or second-order forms of the transport equation by means of the multigroup energy discretization, the discrete-ordinates angular discretization, and finite-element spatial discretization on unstructured meshes [1,2]. Currently, SCEPTRE is primarily applied to predict the effects of X-rays and secondary electrons on cables and other electronic components. The high resolution needed for the accurate modeling of electron transport near conductor-dielectric boundary layers requires the use of large meshes and massively parallel computations.

The first-order solver in SCEPTRE uses discontinuous finite-element differencing. The solution method employs conventional source iteration and a parallel wavefront (sweeping) algorithm on distributed meshes. The sweeping algorithm mostly respects the streaming dependency graph; some dependencies may be ignored during the sweep in order to avoid cyclic dependencies and to improve performance. This approach typically preserves the iteration count needed for

* Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DEAC04-94AL85000

solution convergence of a serial computation; there is little or no degradation in iterative effectiveness as the processor count increases. However, this approach also causes degradation in the scaling of each iteration as the number of processors increases, since it becomes increasingly difficult to keep each processor occupied with useful tasks. Although SCEPTRE can in principle be used with any mesh partitioning, to date we have used only conventional partitioning approaches, which further limits the parallel scaling of the first-order method.

The second-order solver in SCEPTRE uses continuous finite-element differencing. This discretization yields a symmetric positive definite (SPD) matrix that couples all angular and spatial variables, permitting the use of a parallel conjugate-gradient (PCG) solver and eliminating the need for source iteration. This approach, unlike that for the first-order form, yields good scalability of each iteration as the processor count increases. SCEPTRE includes three different second-order formulations: even-odd parity flux (EOPF) [3], self-adjoint angular flux (SAAF) [4] and least-squares finite element (LSFE) [5]. The three formulations have different strengths and weaknesses, e.g. the LSFE method can be used for problems containing internal void regions without special treatment, and the problem size for EOPF method is half that of the other methods due to symmetry. SCEPTRE includes both S_N (discrete ordinates) and P_N (spherical harmonics) second-order solver algorithms, but only the S_N algorithm will be considered here since the P_N algorithm has not yet been extensively tested.

In SCEPTRE the user can choose whether to use the first-order solver or one of the second-order solvers on a group-by-group basis. The solver that is best suited for a particle type at a specific energy can be chosen to provide the most efficient solution to the entire problem. This is extremely beneficial for coupled photon-electron transport problems, since the cross section and scattering ratio can dictate the convergence properties for different particle types and different energy regimes. For example, the transport mean free path of 200-keV photons is 4-5 orders of magnitude larger than that of 200-keV electrons, and the scattering ratio of 200-keV photons is about 0.1, while that of 200-keV electrons is typically greater than 0.9.

In general, the first-order solver works well for streaming-dominated problems, such as photon transport in electronic devices, and the second-order solvers work well for problems with forward-peaked scattering and higher scattering ratios, such as electron transport. The convergence rate of the first-order solver is strongly influenced by the scattering ratio, while the convergence rate of the second-order solvers depends upon the condition number of the system matrix resulting from the space-angle discretization of the problem. For the second-order solvers, effective preconditioning can greatly reduce the condition number of the matrix, resulting in drastic reduction in solver times, but further work remains to be done to develop an optimal preconditioning. The convergence properties of the first- and second-order algorithms in many ways can complement each other.

The first- and second-order transport algorithms share many code components in common, including finite-element tools, cross section handling, angular quadrature, angular moments, distributed source data, boundary conditions and transport field containers. The transport field containers and distributed source data are based on the discontinuous finite-element representation. Since the second-order solvers use continuous finite elements, a mapping is required from discontinuous data structures to continuous, and then back again to map the results

to the discontinuous data structures for compatibility with the first-order solver. These mappings are performed seamlessly and efficiently with the Trilinos tools [6].

2. PROBLEM DESCRIPTION

The problem considered for performance evaluation is a two-dimensional cross section of a miniature, coaxial cable, with dimensions and materials shown in Table I. The cable is small and relatively transparent to photons but not to electrons, as can be seen from the cross sections shown in Tables IIa and IIb. These cross sections were computed by the CEPXS code [7], which is a physics code to generate multigroup, Legendre cross sections for use in deterministic and multigroup Monte Carlo radiation transport codes.

Table I. Coaxial cable specifications

Layer	Material	Outer Radius (cm)	Thickness (cm)
Center Conductor	Iron	0.0478	0.0956
	Copper	0.0594	0.0116
	Silver	0.0606	0.0012
Dielectric	Teflon	0.15113	0.09053
Outer Conductor	Copper	0.17907	0.02794

Table IIa. Photon cross sections at 200 keV (cm^{-1})

	Iron	Copper	Silver	Teflon
σ_t	1.17004×10^{0}	1.42261×10^{0}	3.16203×10^{0}	2.62917×10^{-1}
σ_{s0}	1.13938×10^{-1}	1.45778×10^{-1}	3.08826×10^{-1}	1.30645×10^{-2}
σ_{s1}	1.08554×10^{-1}	1.38100×10^{-1}	2.83913×10^{-1}	1.28180×10^{-2}
σ_{s2}	1.01109×10^{-1}	1.28283×10^{-1}	2.62529×10^{-1}	1.23410×10^{-2}
σ_{s3}	9.31516×10^{-2}	1.18137×10^{-1}	2.40030×10^{-1}	1.16613×10^{-2}
Scattering Ratio	0.0974	0.102	0.0977	0.0497

As stated earlier, the primary application of SCEPTRE is to predict the effect of X-rays and secondary electrons on cables and other electronic components. The key quantity of interest in these types of calculations is the electron emission from the conductors and penetration into the dielectric materials. Because of the large magnitude of the electron cross sections and the large magnitude of the production of electrons in high atomic number materials, the gradient of the electron fluence near material interfaces can be very steep, so that a fine spatial resolution (typically sub micron) is needed near the interfaces to achieve acceptable accuracy. Ideally, it would be desirable to use a different mesh for different particle types and even for different energy regimes, with the spatial dependence of the mesh refinement driven by the spatial gradients in the solution. However, the mapping of source terms in parallel from coarse to fine meshes can be a difficult bookkeeping problem and is not currently implemented in SCEPTRE.

As we will show later, photon transport on an overly refined mesh does not pose any problem for the first-order sweeps-based algorithm, but for the second-order algorithms, this increases the condition number of the resulting matrix considerably and reults in slow convergence. However, effective preconditioning of the matrix can substantially mitigate this effect.

Table IIb. Electron cross sections at 200-keV (cm^{-1})

	Iron	Copper	Silver	Teflon
σ_t	6.97961×10^4	8.34048×10^4	1.18091×10^5	1.04219×10^4
σ_{s0}	6.44678×10^4	7.75741×10^4	1.12025×10^5	8.58885×10^3
σ_{s1}	6.39921×10^4	7.69881×10^4	1.10984×10^5	8.54100×10^3
σ_{s2}	6.32079×10^4	7.60309×10^4	1.09359×10^5	8.45897×10^3
σ_{s3}	6.21969×10^4	7.48034×10^4	1.07343×10^5	8.34969×10^3
σ_{s4}	6.10023×10^4	7.33574×10^4	1.05020×10^5	8.21724×10^3
σ_{s5}	5.96529×10^4	7.17264×10^4	1.02443×10^5	8.06450×10^3
σ_{s6}	5.81699×10^4	6.99346×10^4	9.96485×10^4	7.89366×10^3
σ_{s7}	5.65695×10^4	6.80013×10^4	9.66656×10^4	7.70649×10^3
Scattering Ratio	0.924	0.930	0.949	0.824

Two finite-element meshes of the cross section of a coaxial cable are shown in Figures 1a and 1b. Figure 1a shows a quadrilateral mesh of the coaxial cable, refined to 2 μm near the conductor-dielectric interfaces. A triangular mesh is shown in Figure 1b, which was obtained by splitting each quadrilateral in half. It is noted that the triangular mesh created in this way do possess large element aspect ratios, which can be decreased with the penalty of increasing the element count.

3. RESULTS

In this section, we evaluate the performance of the second-order SAAF algorithm and the first-order sweeps-based algorithm by comparing convergence iteration count and runtimes for a variety of properties, including the finite-element basis functions, mesh refinement, preconditioning, parallel performance, and scaling with angular quadrature order and Legendre expansion order of the scattering cross section. These calculations were performed on the Sandia's Thunderbird cluster, with each compute node consisting of dual 3.6 GHz Intel EM64T processors and 6 GB RAM. Thunderbird's high-speed message passing fabric is Infiniband. The Infiniband fabric is a two level CLOS topology with eight top-level core switches and 280 leaf switches (24 ports per leaf switch) that the compute nodes connect to. Each leaf switch has 16 downlinks (16 compute nodes per leaf switch) and 8 uplinks. Thus, the network is 2-to-1 oversubscribed in terms of raw number of links. SCEPTRE was compiled with optimization with GCC version 3.4.6 and Open MPI version 1.2.7.

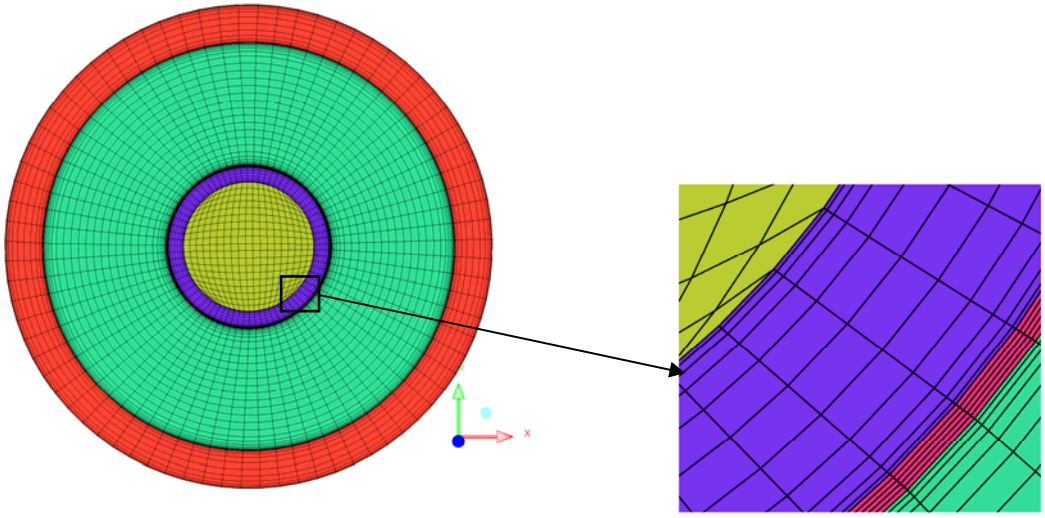

Figure 1a. A quadrilateral mesh of coaxial cable, showing close up near inner conductor-dielectric layer.

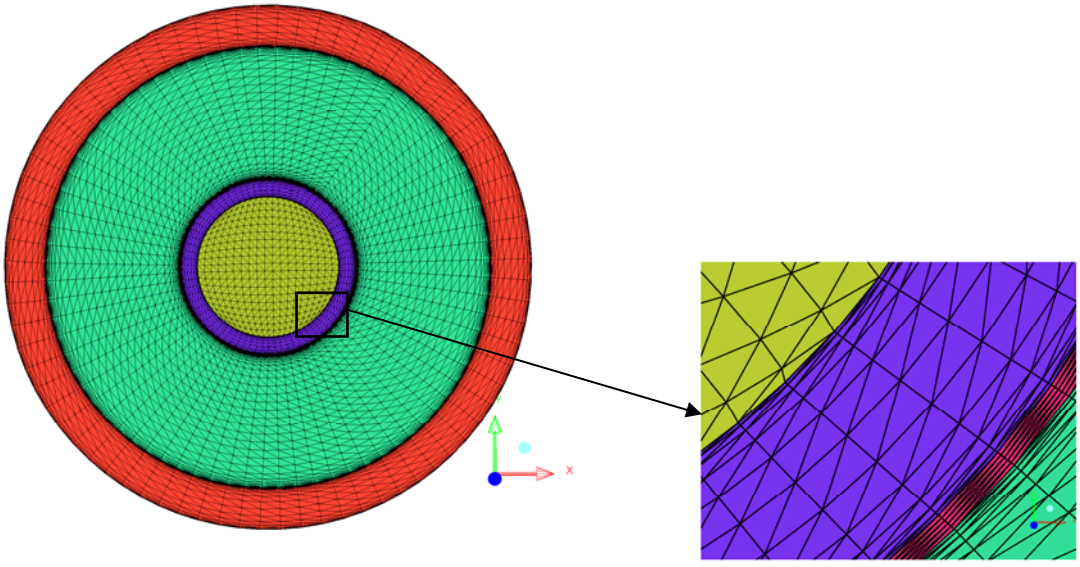

Figure 1b. A triangular mesh of coaxial cable, showing close-up near inner conductor-dielectric layer.

3.1 Effects of Finite-Element Basis Functions and Mesh Refinement

Table IIIa shows results for the transport of 200-keV photons on the coaxial-cable meshes shown in Figs. 1a and 1b. Linear (quad4 and tri3) and quadratic (quad8 and tri6) finite-element basis functions were compared. S_8 angular quadrature and P_3 Legendre cross section expansion were used in the calculations, with a convergence tolerance of 10^{-8} and four processors. The first- and second-order algorithms use different expressions of convergence tolerance: the SAAF solver uses the 2-norm of the pointwise relative residual from the PCG iterations, and the first-order solver uses the infinity-norm of the solution difference between successive iterates. More work is needed to evaluate the effect of the convergence tolerance on the results of interest in the calculations. The SAAF calculations used a block-diagonal preconditioner, with KLU sparse LU factorization of the preconditioner system, which is described in Sec. 3.2.

A number of interesting observations can be made from these results. The first-order algorithm scales roughly with the number of elements in the mesh, while the SAAF algorithm scales more directly with the number of nodes. The SAAF runtimes are much longer for the quadratic basis functions (quad8 and tri6) as compared with the linear basis functions, as there are more unknowns per element.

Table IIIa. Comparison of transport algorithms for various element types and mesh refinement for 200-keV photon transport

S_8P_3, tol=10^{-8}, 4 processors			SAAF (preconditioned)		First Order	
Element Type	N_{elem}	N_{node}	N_{iter}	Run Time (s)	N_{iter}	Run Time (s)
quad4	4845	4884	6	8.61	6	7.45
quad8	4845	14612	6	47.5	6	12.2
tri3	9690	4884	6	8.37	6	14.1
tri6	9690	19457	6	55.4	6	17.8
Refined Meshes:						
quad4	19380	19457	6	67.1	6	35.0
tri3	38760	19457	6	53.5	6	59.0

The results for 200-keV electron transport are shown in Table IIIb. Linear (quad4 and tri3) and quadratic (quad8 and tri6) finite-element basis functions were compared. S_{12} angular quadrature and P_7 Legendre cross section expansion were used in these calculations, with a convergence tolerance of 10^{-8} and 32 processors. A higher-order Legendre cross section expansion is used for the electron transport calculations, due to the higher degree of anisotropy in the electron scattering physics. No preconditioning was used for the SAAF results, as an effective preconditioner for the electron transport has not yet been found, so that preconditioning has little effect or even a negative effect on convergence of electron groups.

The first-order algorithm iteration counts are much higher for electron transport than those for photon transport due to the higher scattering ratio of the electron cross sections. Consequently,

the run times are quite a bit higher for electron transport than for photon transport, even with using 32 processors instead of 4.

Like the photon-transport results, the first-order algorithm scales roughly with the number of elements in the mesh, while the SAAF algorithm scales more directly with the number of nodes. As observed previously, the SAAF runtimes are much longer for the quadratic basis functions (quad8 and tri6) as compared with the linear basis functions.

Unlike the photon results, the SAAF algorithm performs better for a refined mesh with linear basis functions, rather than a coarser mesh with quadratic basis function. Further work remains to be done to determine computational efficiency as a function of accuracy of refining a mesh vs. using a higher basis function order.

Table IIIb. Comparison of transport algorithms for various element types and mesh refinement for 200-keV electron transport

$S_{12}P_7$, tol=10^{-8}, 32 processors			SAAF (unpreconditioned)		First Order	
Element Type	N_{elem}	N_{node}	N_{iter}	**Run Time (s)**	N_{iter}	**Run Time (s)**
quad4	4845	4884	29	3.5	55	36.9
quad8	4845	14612	133	49.6	55	71.1
tri3	9690	4884	29	3.1	55	53.2
tri6	9690	19457	78	32.9	55	98.5
Refined Meshes						
quad4	19380	19457	56	19.1	55	155
tri3	38760	19457	57	16.5	55	234

3.2 Preconditioning the Second-Order Transport Algorithm

As mentioned previously, modeling photon transport on an over-refined mesh, such as a mesh refined to accurately model electron transport, results in an ill-conditioned matrix and slow convergence for the second-order transport algorithms. Effective preconditioning can dramatically improve the convergence. We have found that for the photon transport, a block-diagonal preconditioner is very effective [8]. Rather than solving the unpreconditioned linear system

$$A\boldsymbol{x} = \boldsymbol{b}, \tag{1a}$$

a different linear system, $M^{-1}A$, with a lower condition number than the original linear system, A, is solved instead:

$$M^{-1}A\boldsymbol{x} = M^{-1}\boldsymbol{b}. \tag{1b}$$

An effective preconditioner, *M*, should be relatively easy to invert and a good approximation to the original linear system.

The second-order SCEPTRE algorithm contains several preconditioning options provided by Trilinos, including multilevel (ML) and incomplete factorization (IF) preconditioners, but need further investigation. In addition, we have found that an effective preconditioner for the streaming-dominated problems, such as photon transport in small geometries, by using the pure-absorption transport equation which is nearly equivalent to using the diagonals of the blocks in the discretized transport system, *A*, as a preconditioner. This effectively replaces the full system solve, which scales quadratically as the number of directions, with a set of preconditioner linear systems, which scale linearly with the number of directions.

Implementing the block-diagonal preconditioner requires a procedure for inverting the preconditioning operator, *M*. There are several ways to solve the preconditioner system and the results from two different methods are shown here. The first one is to use a PCG algorithm for each discrete direction. This method is fairly easy to implement and scales well with number of processors, but has the drawback that if *A* is ill-conditioned, *M* will also likely be ill-conditioned, resulting in slow convergence.

An alternative is to use a direct solver to invert the preconditioner linear system for each discrete direction, and then store it for use in subsequent iterations. The serial KLU sparse LU factorization algorithm [9] is available from the Trilinos package and has been implemented in SCEPTRE. This approach results in very efficient convergence of the SAAF algorithm despite the serial component of the implementation. Parallel direct solvers for the preconditioner system have not yet been implemented into SCEPTRE. Table IV compares the results of the SAAF algorithm with the first-order algorithm for several preconditioning options. With the KLU option, the iteration count is reduced to that of the first-order solver, with comparable run times, but the KLU approach does not scale well due to the serial portion of the implementation.

Table IV. Effect of preconditioning of SAAF algorithm on 200-keV photon transport

S_8 P_3, tol=10^{-8}, quadratic-triangular (tri6) mesh								
N_{proc}	SAAF						First Order Solver	
	No Preconditioning		Block-Diagonal Preconditioning				N_{iter}	Run Time (s)
			PCG		KLU			
	N_{iter}	Run Time (s)	N_{iter}	Run Time (s)	N_{iter}	Run Time (s)		
1	20268	42848	23	1821	6	113	6	56.9
4	20268	11340	23	866	6	55.4	6	17.8

3.3 Scattering Order, Extended-Transport Correction and Scattering Ratio

The convergence rate of the first-order algorithm is strongly affected by the scattering ratio, i.e. the ratio of the scattering cross section to the total cross section. The convergence rate of the

second-order algorithm, however, depends on the condition number of the system matrix and is insensitive to the scattering ratio. It is for this reason that the second-order algorithm performs well for electron transport problems, where the scattering ratio is large.

Table V shows the solver times for the first- and second-order algorithms for the model coaxial cable problem with the quadratic-triangular (tri6) mesh, S_{16} angular quadrature, 32 processors, and a convergence tolerance of 10^{-8}. Since the electron scattering is highly forward peaked, the Legendre-polynomial approximation of the angular dependence of the electron cross sections converge slowly. An extended transport correction, or a δ-function correction, of the cross sections can improve the accuracy of the truncated representation of the cross sections [10]. In this approach, a δ-function scattering term with magnitude equal to the scattering cross section moment with order one greater than the Legendre order of the cross section expansion, is subtracted from the total cross section and the scattering moments:

$$\sigma_t^{ETC} = \sigma_t - \sigma_{s,L+1} \tag{1a}$$

$$\sigma_{s,l}^{ETC} = \sigma_{s,l} - \sigma_{s,L+1} \tag{1b}$$

The extended transport correction reduces the effective scattering ratio, resulting in faster convergence of the first-order algorithm. The order of the Legendre expansion needed for a given application is problem and accuracy dependent. A P_7 expansion is not an unreasonable scattering order to use in coupled electron-photon transport. Furthermore, a high-order angular quadrature must be used in correspondence to a high-order Legendre expansion to assure sufficient accuracy to evaluate the angular moments

Table V. Effect of scattering ratio on solver time for 200-keV electron transport

32 processors, S_{16} angular quadrature, tol=10^{-8}, quadratic-triangular (tri6) mesh						
P_L order	**Scattering Ratio**				**Run Time (s)**	
	Iron	Copper	Silver	Teflon	SAAF	First Order Solver
With Extended Transport Correction						
1	0.191	0.209	0.305	0.0662	96.5	25.5
3	0.394	0.420	0.536	0.169	96.2	52.3
7	0.643	0.666	0.753	0.372	97.1	193.4
Without Extended Transport Correction						
1	0.924	0.930	0.949	0.824	139	398
3	0.924	0.930	0.949	0.824	131	530
7	0.924	0.930	0.949	0.824	117	973

3.2 Parallel Performance

Parallel performance results are shown in Tables VIa and VIb for photon and electron transport, respectively. S_8 angular quadrature and P_3 Legendre cross section expansion were used in the photon transport calculations, with a convergence tolerance of 10^{-8}. Two measures of parallel performance are given in the tables: 1) a relative performance, which is parallel efficiency relative to that of the previous number of processors, and 2) a cumulative performance, which is the parallel efficiency relative to the serial run. S_{12} angular quadrature and P_7 Legendre cross section expansion were used in the electron transport calculations, with a convergence tolerance of 10^{-8}.

Table VIa. Parallel performance for 200-keV photon transport

S_8 P_3, tol=10^{-8}, quadratic-triangular (tri6) mesh						
N_{proc}	SAAF (preconditioned)			First Order Solver		
	Run Time (s)	Parallel Performance		Run Time (s)	Parallel Performance	
		rel	cum		rel	cum
1	1846	-	1	56.9	-	1
2	1703	0.54	0.54	31.5	0.90	0.90
4	900	0.95	0.51	17.8	0.88	0.80
8	446	1.0	0.52	9.50	0.94	0.75
16	216	1.0	0.53	5.15	0.92	0.69
32	112	0.96	0.52	2.81	0.92	0.63

Table VIb. Parallel Performance for 200-keV Electron Transport

S_{12} P_7, tol=10^{-8}, linear-triangular (tri3) mesh						
N_{proc}	SAAF (unpreconditioned)			First Order Solver		
	Run Time (s)	Parallel Performance		Run Time (s)	Parallel Performance	
		rel	cum		rel	cum
1	72.3	-	1	1,073	-	1
2	40.4	0.89	0.89	660.	0.81	0.81
4	20.7	0.98	0.87	377.	0.88	0.71
8	10.7	0.97	0.84	197.	0.96	0.68
16	5.65	0.95	0.80	108	0.91	0.62
32	3.06	0.92	0.74	52.3	1.0	0.64

3.4 Effect of Angular Quadrature Order

The first-order solver time scales linearly with number of directions whereas the dependence of the second-order solver time on number of directions is more complicated. The full linear solver time scales quadratically as the number of directions while the preconditioner system solver time scales linearly with the number of directions, so the effective dependence will be somewhere

between linear and quadratic, in practice. The observed orders of the dependence (p) of the solver time on the number of directions are presented in Table VIIa and VIIb for photon and electron transport, respectively.

As expected, the observed p for the first-order algorithm is nearly 1. For the unpreconditioned second-order algorithm, the observed p is nearly 2, while for the photon calculation including block-diagonal preconditioning with the sparse direct solve of the preconditioner system, p is nearly 1, and for the block-diagonal preconditioner with PCG solve, p is between 1 and 2.

Table VIIa. 200-keV photon transport

P_3 Scattering, tol=10^{-8}, quadratic-triangular (tri6) mesh, 32 processors										
S_N Order	N_{dir}	SAAF (Block-Diagonal Preconditioner)						First Order Solver		
		PCG			KLU					
		N_{iter}	Run Time (s)	p	N_{iter}	Run Time (s)	p	N_{iter}	Run Time (s)	p
4	12	23	20.8		6	13.2		6	0.92	
8	40	23	111	1.39	6	45.1	1.02	6	3.1	1.01
12	84	23	261	1.15	6	99.7	1.07	6	6.4	0.98
16	144	23	449	1.01	6	out of memory		6	10.8	0.97

Table VIIb. 200-keV electron transport

P_3 Scattering, tol=10^{-8}, quadratic-triangular (tri6) mesh, 32 processors							
S_N Order	N_{dir}	SAAF (unpreconditioned)			First Order Solver		
		N_{iter}	Run Time (s)	p	N_{iter}	Run Time (s)	p
4	12	60	0.59		26	9.45	
8	40	55	3.97	1.58	26	34.4	1.07
12	84	55	16.0	1.88	26	79.0	1.12
16	144	55	48.6	2.06	26	138	1.03

4. SUMMARY

This paper presents preliminary comparisons of the performance of the first-order sweeps-base algorithm and the second-order PCG algorithm in the SCEPTRE project. Generally, the first-order algorithm performs better for photon transport, due to the small scattering ratio of photon cross sections, and the second-order algorithm performs better for electron transport, due to the efficiency of the PCG algorithm. More work is needed to incorporate acceleration methods into the first-order algorithm, and to develop preconditioning methods that work well for second-order algorithm for electron transport. Some conclusions that may be drawn from this study include: 1) the first-order algorithm is efficient for higher-order finite-element basis functions, while the second-order performs better for linear basis functions, 2) both algorithms scale well in

parallel, 3) the first-order algorithm depends strongly on the order of Legendre expansion of the cross sections for electron transport, while the second-order algorithm is insensitive to it, 4) the first-order algorithm scales linearly with the number of S_N directions, while the unpreconditioned second-order algorithm scales with the number of directions squared.

REFERENCES

1. S. D. Pautz, "An Algorithm for Parallel Sn Sweeps on Unstructured Meshes," *Nucl. Sci. Eng.*, **140**, pp. 111-136 (2002).
2. C.R. Drumm and J. Lorenz, "Parallel FE Approximation of the Even/Odd-Parity Form of the Linear Boltzmann Equation," *Math. Comp. Modeling*, **31**, pp. 55 (2000).
3. J. J. Duderstadt and W. R. Martin, *Transport Theory*, Wiley-Interscience, New York, USA (1979).
4. J. E. Morel and J. M. McGhee, "A Self-Adjoint Angular Flux Equation," *Nucl. Sci. Eng*., **132**, pp. 312-325 (1999).
5. B. Jiang, *The Least-Squares Finite Element Method*, Springer-Verlag, Berlin, Germany (1998).
6. Michael Heroux, et al., "An Overview of Trilinos," SAND2003-2927, Sandia National Laboratories report, Albuquerque, NM (2003).
7. L. J. Lorence, Jr., J. E. Morel, and G. D. Valdez, "Physics Guide to CEPXS: A Multigroup Coupled Electron-Photon Cross-Section Generating Code Version 1.0," SAND89-1685, Sandia National Laboratories report, Albuquerque, NM (1989).
8. C. R. Drumm and W. C. Fan, "Uncollided-Flux Preconditioning of the Conjugate Gradients Solution of the Transport Equation," *Proc. Int. Conf. Nuclear Mathematical and Computational Sciences*, Gatlinburg, TN, April 6-10 (2003).
9. T. A. Davis, *Direct Methods for Sparse Linear Systems*, SIAM, Philadelphia, USA (2006).
10. J. E. Morel, "On the Validity of the Extended Transport Cross-Section Correction for Low-Energy Electron Transport," *Nucl. Sci. Eng*., **71**, pp. 64-71 (1979).