Title: **Parallel 3-D Spherical-Harmonics Transport Methods**

Author(s): J. E. Morel, CIC-19
J. M. McGhee, CIC-19
T. Manteuffel, Dept. Math., U. Colorado

Submitted to: DOE Office of Scientific and Technical Information (OSTI)

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Los Alamos

NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Form 836 (10/96)
ST 2629
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Parallel 3-D Spherical-Harmonics Transport Methods

Jim E. Morel* and John M. McGhee

Computing, Information, and Communications Division, Los Alamos National Laboratory

Thomas Manteuffel

Department of Mathematics, University of Colorado

Abstract

This is the final report of a three-year, Laboratory-Directed Research and Development (LDRD) project at the Los Alamos National Laboratory (LANL). We have developed massively parallel algorithms and codes for solving the radiation transport equation on 3-D unstructured spatial meshes consisting of arbitrary combinations of hexahedra, wedges, pyramids, and tetrahedra. Three self-adjoint forms of the transport equation are solved: the even-parity form, the odd-parity form, and the self-adjoint angular flux form. We developed this latter form, which offers several significant advantages relative to the traditional forms. The transport equation is discretized in space using a trilinear finite-element approximation, in direction using a spherical-harmonic approximation, and in energy using the multigroup approximation. Our discrete equations are solved using a parallel conjugate-gradient. All of our parallel algorithms were implemented on the CM-5 computer at LANL. Calculations are presented which demonstrate that our solution technique is both highly parallel and efficient.

Background and Research Objectives

The Boltzmann transport equation describes the transport of particles (photons, neutrons, electrons, ions, etc.) through space and matter.[1] The numerical solution of this important equation is an essential modeling tool in a wide variety of fields. Among these are nuclear reactor design and analysis, inertial confinement fusion, criticality safety, radiation therapy, medical imaging, nuclear weapons safety, and oil-well logging. All existing numerical methods for solving the transport equation can be broadly characterized as either stochastic or deterministic. Stochastic (Monte Carlo) methods are easily applied in complex 3-D geometries, but severe statistical errors often persist if differential rather than integral quantities are being calculated. Deterministic methods do not suffer from statistical error, but rather from discretization error. The most popular deterministic method in use today is the discrete-ordinates or S_n method.[1] This method has proven to be extremely useful for nuclear reactor analysis, but for a wide class of problems it exhibits a serious lack of rotational invariance commonly known as ray effects. Thus, even though the S_n method is widely used in neutron transport, it is unsuitable for many other applications. Indeed, there are many

* Principal Investigator, E-mail: jim@lanl.gov

important applications for which both S_n and Monte Carlo methods are poorly suited. For instance, the calculation of radiation symmetrization effects is of critical importance in inertial-confinement fusion (ICF) analysis. S_n methods are considered to be inappropriate for such calculations because of their lack of rotational invariance. At the present time, the only true transport techniques available for ICF analysis are Monte Carlo techniques. Unfortunately, the statistical errors associated with such techniques make calculations for irradiation symmetry analysis prohibitively expensive in most instances.

The calculation of electron and hole transport under the influence of electromagnetic fields is another example. This problem is of primary importance in semiconductor modeling. The S_n method is not suitable for such calculations because the electromagnetic fields can take on arbitrary orientations as a function of space and time, and particle motion along electric field lines cannot be properly modeled unless each field line is coincident with a discrete S_n characteristic. In general, one cannot expect such coincidence. Although S_n -like deterministic transport methods have been used in semiconductor modeling, they have been restricted to 1-D geometries where their lack of rotational invariance is of no consequence. All major multi-dimensional semiconductor modeling codes currently use Monte Carlo transport methods, and the execution times are often excessive. The development of a rotationally invariant, deterministic neutral-particle transport solver capable of efficiently modeling complex 3-D geometries would provide a significant advance over existing transport methods. It would also provide an essential building block for the eventual development of more general transport capabilities such as those for thermal radiation transport and charged-particle transport with electromagnetic fields.

The only deterministic transport method that rigorously preserves the rotational invariance of the transport operator is the spherical-harmonics or P_n method. This is one of the oldest methods for solving the transport equation. It was quite popular in the early days of the reactor physics community, but the P_n method was quickly displaced by the S_n method. This occurred because the first-order equations associated with the S_n method are much easier to solve than the second-order equations associated with the classical P_n method, and most importantly, because ray effects do not appear in reactor physics calculations. Since the reactor physics community has historically funded essentially all transport code development, it is not surprising that production P_n transport codes were never developed in the United States. The British have developed production P_n codes, but they have focused almost entirely upon the development of accurate finite-element and variational spatial-discretization schemes, and neglected the development of sophisticated techniques for solving the associated discretized equations. The absence of methods for performing rotationally invariant deterministic transport calculations in complex 3-D geometries represents a serious deficiency in the U.S. technology

base. The research that we have performed represents a first step towards the elimination of this deficiency.

The P_n approximation is generally applied to two traditional second-order self-adjoint forms of the transport equation: the even-parity (EP) and odd-parity (OP) transport equations. This is done in order to obtain symmetric positive-definite matrix equations. Highly sophisticated and effective numerical methods such as the conjugate-gradient method exist for solving such matrix equations. Unfortunately, calculations with void regions are highly problematic with these traditional self-adjoint forms of the transport equation.

Our research objectives are as follows:

1. To develop a parallel algorithm for solving the neutral-particle radiation transport equation with a spherical-harmonics angular discretization on unstructured 3-D meshes consisting of arbitrary combinations of hexahedra, wedges, pyramids, and tetrahedra.
2. To investigate new self-adjoint formulations of the transport equation that can be applied to problems containing void regions.

Importance to LANL's Science and Technology Base and National R&D Needs

The algorithms and the associated computer code that we have developed can be applied to a wide variety of problems of interest to the DOE: reactor design and analysis, reactor safety research, Inertial Confinement Fusion (ICF) research, health-physics, criticality safety, nuclear weapons safety, controlled fusion research, neutron radiation therapy, and oil-well-logging tool development. Logging tool development is of particular interest to the oil industry, and a collaborative research and development agreement (CRADA) already exists between LANL and a consortium of oil companies for applying Monte Carlo techniques to this problem. As previously noted, there are currently no production neutral-particle transport codes in the United States that offer a spherical-harmonics angular discretization. The standard S_n method lacks rotational invariance and thus yields non-physical ray effects in certain types of problems. Thus it is clearly important to have a rotationally invariant alternative to the S_n method. The code that we developed provides such a capability and is unique in doing so. In addition, there are also no production transport codes that use any type of 3-D unstructured mesh, much less a hexahedral mesh. The use of unstructured hexahedral meshes results in a capability for efficiently modeling complex 3-D geometries. The production 3-D neutral-particle transport codes presently available in the United States use only rectangular meshes. Modeling complex 3-D geometries with rectangular meshes is a very inefficient process. Thus the computer code that we have developed represents a significant advance in our ability to perform radiation transport calculations for complex 3-D geometries.

Perhaps most importantly, the potential long-term impact of our algorithms and code is far-reaching in that they represent a critical building block for the eventual development of more general transport methods such as those that include the modeling of charged particles with electromagnetic fields. Deterministic charged-particle transport methods are needed for a wide range of plasma physics and solid state applications. Transport tools for such applications are currently based only on Monte Carlo methods and are usually very expensive to use.

Although we have designed our algorithms on the CM-5, our code can easily be ported to workstations, thus making our methods widely available to a large number of researchers. As previously noted, the almost total absence of tools for performing rotationally invariant deterministic transport calculations represents a serious deficiency in our technology base. The algorithms and associated code that we have developed provide the technical foundation upon which to partially remedy this deficiency. In particular, our algorithms and associated code can be used as a starting point for the development of production neutron and thermal radiation transport codes that can provide rotationally-invariant solutions.

Scientific Approach and Accomplishments

During the first year of our project we developed a tetrahedral-mesh algorithm for solving the self-adjoint even-parity form of the neutral-particle transport equation. A spherical-harmonics discretization was used in angle, a linear-continuous finite-element discretization was used in space, and the standard multigroup discretization was used in energy. A parallel solution technique was developed for the CM-5 computer. This solution algorithm is based upon a parallel conjugate-gradient method used in conjunction with the standard scattering source iteration method. The work expended in the conjugate-gradient method arises almost totally from matrix-vector multiplies. A matrix-vector multiply is a process that contains considerable inherent parallelism. For instance, both multiplications and sums are performed in a matrix-vector multiply. All of the multiplications can be performed simultaneously, and each row sum can be performed independent of the others. A simple analysis indicates that a least upper bound on the parallel speed-up associated with this matrix-vector algorithm is proportional to $2N$ for a "sufficiently large" $N \times M$ matrix. Thus this algorithm does indeed have a great deal of inherent parallelism associated with it.

We now give a brief description of our CM-5 algorithm. Let us assume that we wish to multiply the vector "x" by the matrix "A" to obtain the vector "y". All of the non-zero matrix elements are placed in a 1-D array, "B." In particular, each row of the matrix is sequentially packed into this array. The vector "x" is stored in a 1-D array "X." Since each matrix element must multiply a particular component of the vector "x," these x-components are gathered from the array "X" into a 1-D array "C," which is ordered in accordance with the matrix element

ordering in array "B." All of the multiplications are then performed in parallel using the CM-FORTRAN statement "D=B*C." The final steps are to first perform the sums associated with each row of the matrix "A," and then to scatter the resulting components into a 1-D array "Y," which contains the components of the "y" vector. The sums associated with each given row are carried out serially within that row, but the sums are carried out in parallel across rows. These two steps are accomplished with one CM-FORTRAN statement known as a "SCATTER-WITH-ADD." We later give computational results which demonstrate the effectiveness of our parallel conjugate-gradient algorithm. We note that our algorithm is parallelized over space, but is serial in angle and energy.

During the second year, we added a capability for solving the self-adjoint odd-parity form of the neutral-particle transport equation, and began investigating a new self-adjoint form of the transport equation known as the self-adjoint angular flux (SAAF) form. This equation can be derived in a variety of ways. For instance, by suitably defining an inner product, it can be derived as a normal form of the standard first-order form of the transport equation. It also can be derived by algebraic manipulation of the standard first-order form of the transport equation. Finally, it can be derived simply by summing the even-parity and odd-parity equations. We discovered that the SAAF equation offers many advantages relative to the traditional (EP and OP) self-adjoint equations:

1. After standard vertex-based finite-element discretization in space, the traditional self-adjoint equations have one component of the flux (either even or odd) located with full accuracy at each vertex, and the other component (either odd or even) located with reduced accuracy (one less order) within each cell. This property makes it impossible to construct the full angular flux at any one location with full accuracy. However, this problem does not arise with the SAAF equation because it yields the *complete* angular flux with full accuracy at each vertex and the *complete* angular flux with reduced accuracy within each cell.
2. Because the complete angular flux is present at the vertices, it is very simple to implement reflective and reflective-like boundary conditions with the SAAF equation. However, such boundary conditions are much more difficult to implement with the traditional self-adjoint equations.
3. Using the same spatial discretization technique and the same order P_n approximation for the three self-adjoint equations, one finds that the formal order of accuracy is identical for all of the equations, but for any given spatial mesh, the SAAF equation yields a solution that is much more accurate than the traditional equations. In particular, it is six times more accurate for the problems we considered.
4. Calculations in voids cannot be carried out with the traditional self-adjoint equations because the even-parity and odd-parity components completely decouple. However, because

the SAAF equation has the full angular flux as its unknown rather than a single component (even-parity or odd-parity) of the angular flux, voids do not present a problem.

The only disadvantage of the SAAF equation relative to the traditional equations is that it has twice as many angular unknowns using the same order P_n approximation. However, the enhanced accuracy of the SAAF equation more than compensates for the increased cost associated with these additional unknowns.

During the third year we modified our algorithms to accommodate standard finite-element meshes, i.e., unstructured meshes consisting of arbitrary combinations of hexahedra, wedges, pyramids, and tetrahedra. This modification was an important step from a practical point of view because it gave our transport methods compatibility with standard finite-element fluid-dynamics codes. Infrared radiation transport plays an important role in many types of combustion problems. We are unaware of any codes other than our own that can perform full infrared radiation transport calculations in participating media on 3-D finite-element meshes.

To demonstrate certain properties of our parallel algorithm, we now present a few computational results. All calculations were performed on either the CM-5 at the LANL Advanced Computing Laboratory or on a single node of the LANL IBM cluster. The CM-5 is a massively parallel computer. It has total of 1024 processing nodes, with each node composed of four floating-point vector processors. Each vector processor has 8 megabytes (MB) of memory and a peak computation rate of 32 megaflops (MFLOPS). The machine can operate with independent partitions of 32, 64, 128, or 512 processing nodes. The CM-5 has a fat-tree communications network with data transfer rates of 256 MB per second (MB/s) within a processing node, 20 MB/s between nodes that are 1 to 4 nodes apart, 10 MB/s between nodes that are 5 to 16 nodes apart, and 5 MB/s on nodes that are more than 16 nodes apart. The CM-5 can operate in both an SIMD (Single Instruction Multiple Data) mode and a MIMD (Multiple Instruction Multiple Data) mode. However, all of our algorithms were designed for the SIMD mode.

Our first set of calculations is intended to demonstrate the scaling of our algorithm as the number of processors is increased, and to give the reader an overall feel for the performance of our algorithm relative to a single node of the IBM cluster at Los Alamos. A single node of this cluster is basically an IBM-590 RS-6000 workstation, which has a peak computation rate of 256 MFLOPS. The simple problem we consider consists of a homogeneous cube of isotropically scattering material, 100 cm on a side, with a total macroscopic cross section of 0.05657 cm^{-1} , a scattering ratio of 0.5, a unit isotropic distributed source, and vacuum boundary conditions. Calculations were performed for each of four tetrahedral meshes. These meshes are described in Table 1. It can be seen from Table 1 that

the first mesh consists of a relatively small number of cells, while the fourth mesh consists of a very large number of cells. All of the calculations were performed with a P_3 spherical-harmonic approximation. The CPU times are given in Table 2 for a 32 processing node (PN) partition of the CM-5, a 64 PN partition, a 128 PN partition, and a single node of the IBM cluster. Several significant trends can be identified from Table 2. For instance, with a small mesh (few unknowns), the IBM actually gives a smaller CPU time than the CM-5, and as the number of processors is doubled, the CM-5 CPU time actually *increases* slightly. This is to be expected because the calculation has too few spatial unknowns to properly utilize the CM-5 architecture. We stress that our algorithm is parallel over space but not angle or energy. Conversely, for the larger spatial meshes, doubling the number of processors nearly halves the CPU time. For instance, using the largest spatial mesh it is reduced by a factor of 0.55 when increasing from 32 PNs to 64 PNs, and by 0.6 when increasing from 64 PNs to 128 PNs. Perfect scaling would reduce the CPU time by a factor of 0.5. Thus our algorithm scales well. Furthermore, using the largest mesh with 128 PNs, the CM-5 has a CPU time that is approximately 40 times faster than that of the IBM. Thus our algorithm obtains good overall CPU performance from the CM-5.

The second set of calculations demonstrates that although our algorithm is *not* parallel over angle, it scales in several desirable ways as the P_n order is increased. The same physical problem used in the first set of calculations was also used in the second set. All calculations were performed using the third grid (the second-largest) described in Table 1. Calculations were performed using P_1 through P_9 spherical-harmonic approximations on both the CM-5 with 64 PNs and a single node of the IBM cluster. The CPU times for both machines, and the memory usage on the CM-5 are given in Table 3. In addition, the number of spherical-harmonic moments (angular unknowns) are given for each P_n order of approximation, so as to indicate how the number of unknowns scales as the P_n order is increased. It can be seen from Table 3 that the CM-5 has CPU times that are roughly 12 times smaller than those of the IBM for all calculations of order greater than 1. This ratio of CPU times quickly saturates with P_n order because our algorithm is not parallel over angle. It can also be seen from Table 3 that both the CPU time and the memory usage grow linearly with the number of angular moments (angular unknowns). This is the best scaling that one can expect. It is not clear that the observed CPU scaling will occur in all problems. It is certainly conceivable that increasing the number of angular unknowns in certain problems will significantly increase the condition number of the coefficient matrix. This would cause the CPU time to increase at a superlinear rate. However, one might have to go to very high P_n

Table 1: Mesh Parameters for P_3 Calculations

Calculation	Cell MFP	Vertices	Tets
1	2.0	71	192
2	1.0	429	1,536
3	0.5	2,969	12,228
4	0.25	22,065	98,304

Table 2: CPU Times for P_3 Calculations

CPU Times (seconds)				
Calculation	32 PN-CM	64 PN-CM	128 PN-CM	IBM
1	1.17	1.20	1.33	0.25
2	2.06	2.11	2.07	3.55
3	7.83	5.32	4.29	58.78
4	70.70	38.84	23.23	905.08

Table 3: CPU Times and Memory Usage for P_n Calculations

P_n Order	Number of Unknowns	CPU Time (seconds)		Memory (MB)
		64 PN-CM	IBM	
1	4	0.91	4.64	12.3
3	16	5.32	58.78	27.9
5	36	16.56	210.5	57.7
7	64	36.04	452.59	107.4
9	100	66.03	885.50	185.2

order to see this effect. For instance, we have seen the condition number of the coefficient matrix significantly increase as the number of spatial unknowns is increased in certain problems characterized by cells that are optically thin. The observed memory scaling results from the fact that the full coefficient matrix for our equations is never stored, but rather is generated as needed in blocks that do not grow in size as the number of angular unknowns is increased. This is an important component of our algorithm since the P_n matrix rapidly becomes huge with increasing P_n order.

We have included several graphics to demonstrate certain characteristics of the P_n approximation in general as well as our particular implementation of it. For instance, we previously discussed the fact that the P_n equations yield rotationally invariant solutions, while S_n solutions can suffer from ray effects due to a lack of such invariance. In Fig. 1 we show an iso-contour of the scalar flux calculated with the P_3 approximation for an isotropic point source of radiation in an optically thin medium. This solution exhibits the spherical symmetry of the exact physical solution. In Fig. 2 we show an S_8 iso-contour solution for the same problem. Spherical symmetry is grossly lacking due to the fact that particles effectively travel along discrete directions under the S_n approximation. The locations of these discrete directions can be clearly seen in the graphic.

In Fig. 3 we show the mesh and an associated plot of the scalar flux on a cut-plane slicing through the center of a sphere. The sphere is homogeneous and contains a constant isotropic distributed source of radiation. Colors are related to the magnitude of the scalar flux solution: red (center) corresponds to the highest magnitude and blue (outer edge) corresponds to the lowest magnitude. This particular mesh contains primarily hexahedra, but has wedges near the poles and pyramids near the center of the sphere. This demonstrates the general finite-element mesh capability that we implemented in our code during the third year of the project.

Publications

1. Morel, J. E., McGhee, J. M., and Roberts, R. M., "DANTE: A 3-D Unstructured Finite-Element Mesh Transport Code," to appear in the Proceedings of the Seminar on 3D Deterministic Radiation Transport Computer Programs: Features, Applications, and Perspectives, December 2-3, 1996, OECD Chateau de la Muette, Paris XVI.
2. Morel, J. E., and McGhee, J. M., "A Self-Adjoint Angular Flux Form of the Transport Equation," to be submitted to *Nuclear Science and Engineering*.
3. Morel, J. E., McGhee, J. M., and Roberts, R. M., "Parallel Algorithms for 3-D Unstructured-Mesh Self-Adjoint Transport Calculations," to be submitted to *Nuclear Science and Engineering*.

References

- [1] Lewis, E. E., and Miller, W. F., Jr., *Computational Methods of Neutron Transport*, American Nuclear Society, Inc., LaGrange Park, Illinois (1993).

Figure 1. Iso-contour scalar flux P_3 solution for a point source of radiation in an optically thin medium.

Figure 2. Iso-contour scalar flux S_g solution for a point source of radiation in an optically thin medium.

Figure 3. Mesh and cut-plane scalar flux plot for a constant isotropic radiation source in a homogeneous sphere.