

SANDIA REPORT

SAND96-2485

Unlimited Release

Printed November 1996

RECEIVED
OCT 19 1998
OSTI

Spatial Domain-Based Parallelism in Large Scale, Participating-Media, Radiative Transport Applications

Shawn P. Burns, Mark A. Christon

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: 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, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent 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, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01



DISCLAIMER

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

SPATIAL DOMAIN-BASED PARALLELISM IN LARGE SCALE, PARTICIPATING-MEDIA, RADIATIVE TRANSPORT APPLICATIONS

Shawn P. Burns, Engineering Sciences Center
Mark A. Christon, Computational/Computer Sciences and Math Center
Sandia National Laboratories
Albuquerque, NM 87185-0835

ABSTRACT

Parallelism for gray participating media radiation heat transfer may be placed in two primary categories: spatial and angular domain-based parallelism. Angular, e.g., ray based, decomposition has received the greatest attention in the open literature for moderate sized applications where the entire geometry may be placed on processor. Angular based decomposition is limited, however, for large scale applications ($O(10^6)$ to $O(10^8)$ computational cells) given the memory required to store computational grids of this size on each processor. Therefore, the objective of this work is to examine the application of spatial domain-based parallelism to large scale, three-dimensional, participating-media radiation transport calculations using a massively parallel supercomputer architecture. Both scaled and fixed problem size efficiencies are presented for an application of the Discrete Ordinate method to a three dimensional, non-scattering radiative transport application with nonuniform absorptivity. The data presented shows that the spatial domain-based decomposition paradigm results in some degradation in the parallel efficiency but provides useful speedup for large computational grids.

ACKNOWLEDGEMENTS

This work was performed in part using the Intel PARAGON parallel supercomputer maintained by the Massively Parallel Computing Research Laboratory at Sandia National Laboratories in Albuquerque, New Mexico. Overall support was provided by the U.S. Department of Energy contract DE-AC04-94AL85000.

NOMENCLATURE

A	Coefficients of the discrete transport equation.
$e_{\phi,i}$	Relative solution error at node i for quantity ϕ .
h	Mesh spacing.
I	Radiative intensity.
$\hat{i}, \hat{j}, \hat{k}$	Unit Cartesian direction vectors.
L	Edge length of the computational domain.
n	Order of the angular quadrature set.
\hat{n}	Outward directed unit surface normal.
N	Number of nodes in the global (distributed) computational grid.
N_{flop}	Number of floating point operations performed by the local processor for the global solution.
N_L	Number of nodes in the local subdomain.
$N_{L,x}$	Number of grid lines in the local subdomain in the x direction, $N_{L,x} = N_{L,y} = N_{L,z}$.
N_p	Number of processors (subdomains).
n_p	Number of processors in the x, y, and z directions.
n_n	Number of nodes in the x, y, and z directions.
Q	Radiative source term.
q_r	Radiative heat flux vector.
r, θ , ϕ	Spherical coordinates.
\mathbf{r}	Location vector.
s	Distance along the ordinate direction \hat{s} .
\hat{s}	Unit direction vector.
T	Temperature.
t_c	Communication time.
t_g	Grind time including communication.
x,y,z	Cartesian coordinates.

Greek

Γ	Global boundary.
Γ_s	Portion of the global boundary over which Dirichlet conditions are applied.
ϵ	Emissivity.
η_f	Absolute fixed problem size efficiency.
$\eta_{f,i}$	Incremental fixed problem size efficiency.
η_s	Scaled problem size efficiency.
η^*	Theoretical efficiency.
μ	Extinction cross section.
μ_a	Absorption cross section.
ν	Radiation frequency.
σ	Stefan-Boltzmann constant.

INTRODUCTION

The motivation for this study derives from the need for a efficient, parallel radiation transport algorithm to be used for high temperature, transient, reacting flow simulations where complex three-dimensional geometries are involved and high-resolution grids are required. Multiphysics simulations in which the radiative transport is tightly coupled with a simultaneous solution of the conservation equations for mass, momentum, energy and species transport place extreme demands on the computational efficiency of the participating media radiation calculation.

For the complex three dimensional geometries of ultimate interest, and the associated complex flow fields, high resolution computational grids are required to resolve the important length scales in the flow, i.e., grids with on the order of 10^6 to 10^8 cells. In general, combined radiation-flow problems may require grids with a different resolution for each transport mode. However, here the radiation grid is assumed to conform to the flow grid as a first step in developing a parallel algorithm. The spatial and temporal resolution requirements for this class of multiphysics problem can only be considered tractable with the most powerful supercomputer technology, specifically, massively parallel supercomputers.

The potential of parallel computing for heat transfer in general and radiation heat transport specifically has been recognized by several authors in recent years [1-5]. Shih, et al. [1] provides a description of parallel and vector architectures as well as a conceptual description of how parallel processing techniques might be applied to heat transfer. The review articles by Howell [2, 3] focus more narrowly on radiation heat transport and discuss how parallel processing techniques may effect future directions in radiation modeling research. Howell [3] also discusses a National Science Foundation workshop held in 1993 [4], the purpose of which was to improve the communication between the radiation heat transfer and supercomputing communities. Finally, Azmy [5] reviews parallel processing techniques employed in neutron transport and provides a useful taxonomy of parallel decomposition paradigms.

Azmy describes parallelism in terms of the portion of the solution space which is decomposed. In general, the radiation heat transfer solution domain is defined by three spatial coordinates (x, y, z), two angular coordinates (θ, ϕ), and a radiation frequency (ν). Using Azmy's nomenclature, spatial domain-based parallelism decomposes the spatial domain (x, y, z) and assigns a processor to each sub-domain. Conversely, angular and frequency domain-based parallelism subdivides the angular (θ, ϕ) or frequency (ν) domain and assigns processors to each of these subdomains. In each case, the portions of the solution space which are not decomposed are stored on each processor in their entirety.

The work of Burns, et al. [6, 7] demonstrated angular domain-based decomposition strategies for Monte Carlo (MC) surface-based neutron, photon and electron transport on single-instruction multiple-data (SIMD) supercomputers. In effect, this work focused upon the obvious and inherent fine-grained parallelism in ray tracing to achieve a high degree of vectorization, i.e., SIMD parallelism. Haferman, et al. [8] also employed angular domain-based decomposition with the Discrete Ordinate (S_n) method to model three dimensional atmospheric radiative transport. Finally, Howell [9] suggested the use of angular decomposition to improve the performance of the MC method for multiphysics applications.

The work of McGhee and Morel [10] has also focused upon primarily SIMD implementations of the discrete ordinates (S_n) and simplified spherical harmonics (SP_n) methods on the Thinking Machines CM-200 and CM-5 by exploiting the parallelism available in multi-frequency (multi-

group) radiation transport. The work of Benmalek, et al. [11] followed this approach choosing to distribute wavelength dependent radiation calculations across clusters of workstations, i.e., each workstation performed a complete radiation calculation for a given wavelength.

de Oliveira, et al. [12] employed spatial domain-based decomposition in a spherical harmonics (P_n) formulation to model neutron transport in complex geometries using a Cray T3D distributed memory MIMD supercomputer. de Oliveira, et al. asserted that the spatial domain-based decomposition paradigm yielded poor parallel performance for problems with localized sources but performed better for problems with more uniform sources. de Oliveira, et al. also suggested the use of dynamic processor re-mapping to improve the parallel performance of their model.

Novo, et al. [13] considered the application of both spatial and angular domain-based parallelism to two dimensional radiative heat transport using the Discrete Transfer (DT) method. Novo observed poor fixed problem size efficiency (see Gustafson, et al. [14] for definitions of fixed and scaled problem size efficiency) using spatial domain-based parallelism relative to angular domain-based parallelism. Novo, et al. suggested that the relatively poor performance of the spatial domain-based decomposition paradigm was due to an increase in the number of global solution iterations resulting from the non-local character of radiative transport. The behavior observed by both de Oliveira [12] and Novo, et al. will be discussed in greater detail in terms of the S_n method later in this work.

Why Spatial Domain-Based Parallelism?

The principal justification for the use of spatial domain-based parallelism in this work is the need for fine scale spatial resolution. On a distributed memory MIMD architecture, spatial domain-based decomposition allows for large distributed computational grids. In contrast, the spatial resolution of angular and frequency based decomposition is severely limited by the available on-processor memory.

For example, Sandia's Intel PARAGON parallel supercomputer has approximately 16 Mbytes of available memory on each processor. Assuming that $O(10)$ floating point values are required to fully describe each element in an unstructured, three dimensional finite element grid (8 bytes per float), then a maximum of $O(10^5)$ elements may be stored on each processor without considering the memory required for material properties, element connectivities, etc. Since angular and frequency based parallelism require that the entire spatial grid be stored on each processor, the maximum spatial resolution obtainable using angular based parallelism would be several orders of magnitude below the level required by the application of interest in this work.

Azmy [5] points out that angular decomposition with the S_n method has a practical limit in the number of angular directions that may be used for parallelism. In fact, the degree of parallelism falls far short of the massive parallelism provided by current supercomputer architectures. For example, the S_{16} level weighted quadrature set of Lathrop and Carlson [15] discretizes the unit sphere into 288 ordinate directions which is considerably less than the number of processors available on evolving parallel supercomputer architectures such as Sandia's TeraFlop machine with over 4500 computational nodes.

Novo, et al. [13] employed spatial domain-based decomposition because there is no analog for angular or frequency based parallelism in advection/diffusion transport. This is an important consideration for coupled, multiphysics applications since the use of different decomposition

paradigms for the advection/diffusion transport and the radiation transport would significantly complicate the overall parallel algorithm and increase the communication overhead.

The principal disadvantage of the spatial domain-based decomposition paradigm is that it is not a good match for the global nature of radiative transport. As will be demonstrated in this work, the global nature of radiative transport for applications which are not optically thick results in a serial character in the global solution sweeps when spatial domain-based parallelism is employed. This serial character is manifested by a decrease in the parallel efficiency of the algorithm as observed by Novo, et al. [13]. One of the objectives of this work is to examine how serious this effect is in terms of the usable parallel speedup provided by spatial-domain based parallelism.

Why Discrete Ordinates?

The texts by Modest [16] and Siegel and Howell [17], as well as the review article by Viskanta and Mengüç [18] describe a number of common radiative transport models including the YIX, finite element (FE), discrete transfer (DT), Monte Carlo (MC), discrete ordinate (Sn), and spherical harmonics (Pn) methods. An order of magnitude analysis of the operation count and communication overhead for each of these models suggests that the most promising candidates for providing an efficient radiation heat transport algorithm include the differential (Siegel and Howell [17]) Sn and Pn methods as well as the ray tracing DT method (Lockwood and Shah [19]). For brevity, this computational complexity analysis will not be reproduced here although details relating to the Sn method will be provided later in this work.

The Pn method was deemed unacceptable for this work given the complexity of the method relative to its accuracy [3]. For example, the Pn method evaluates a coupled system of partial differential equations for the directional radiative intensity distribution for both scattering and nonscattering applications. In contrast, the system of directional radiative intensity equations evaluated by the Sn method is uncoupled for nonscattering applications. This coupling causes the computational expense of the Pn method to increase rapidly with the degree of angular resolution.

Some preliminary work by the authors with the DT [19] method showed that optimization of the asynchronous interprocessor communication required by this method was problematic and resulted in a large amount of wait time during interprocessor communication. This wait time seriously degraded the parallel performance of the method. The DT algorithm described by Novo, et al. [13] provided better synchronization between processors at the expense of the directional continuity of rays communicated across processor boundaries. Application of the Sn model using spatial domain-based parallelism did not require a compromise of this sort.

Over the past decade, the Sn method has also been applied to a number of industrial furnace and boiler applications [20, 21]. These industrial applications involve the same coupled, multiphysics character as the application of interest in this work. Given the similarities, as well as the limitations in the other radiative transport models, the Sn method was chosen for this work.

Outline

The succeeding sections of this work describe the application of a spatial domain-based parallel S_n algorithm to three dimensional radiative transport. A demonstration problem selected from the literature is defined and a benchmark solution provided. Following the demonstration problem description, the parallel S_n algorithm is outlined and a computational complexity analysis conducted. Numerical results for the demonstration problem obtained using the parallel S_n algorithm are presented and the spatial and angular convergence rates examined. The performance of the parallel S_n algorithm is then examined in terms of the fixed and scaled problem size efficiencies as well as communication overhead and cycle time (e.g., computation time per node per ordinate direction). Finally, several conclusions are drawn from the data and recommendations for additional work made.

DEMONSTRATION PROBLEM

The demonstration problem used in this work was first introduced by Hsu and Farmer [22] with additional data provided by Burns, et al. [23] and Tan and Hsu [24]. The demonstration problem consists of an isothermal unit cube centered on the origin and oriented so that the sides of the cube are orthogonal to the principal Cartesian axes. The walls of the cube are cold and black and the interior of the cube consists of a gray, nonscattering, absorbing/emitting material with a spatially varying absorptivity. The absorptivity varies according to the trilinear function:

$$\mu_a L = 0.9 \left(1 - 2 \left| \frac{x}{L} \right| \right) \left(1 - 2 \left| \frac{y}{L} \right| \right) \left(1 - 2 \left| \frac{z}{L} \right| \right) + 0.1, \quad (1)$$

The domain is discretized using a uniform, orthogonal finite difference grid of $n_n \times n_n \times n_n$ computational nodes where the total number of nodes in the global mesh is $N = n_n^3$. The computational domain is decomposed into $n_p \times n_p \times n_p$ nonoverlapping spatial subdomains with one subdomain per processor where the total number of processors is $N_p = n_p^3$. Since nodes are shared on the interprocessor boundaries, the number of nodes on each subdomain is given by:

$$N_L = \left[\frac{n_n - 1}{n_p} + 1 \right]^3, \quad (2)$$

Figure 1 provides an example of the sub-domain mapping for $n_p = 4$ where the subdomains are separated for clarity.

The benchmark results used for comparison to the S_n results were obtained from direct numerical integration of the radiative heat flux equation [16]. Highly accurate benchmark results were obtained for this work using zeroth order (piecewise constant) Newton-Cotes integration in

spherical coordinates. The spherical integration provided faster convergence under spatial refinement than the Cartesian formulation described in Burns, et al [23]. Sensitivity studies were performed to assess the influence of the spatial resolution on the benchmark solution accuracy.

The final spatial resolution used to obtain the benchmark results employed 320 intervals in the radial, azimuthal, and zenith (r, θ, ϕ) directions respectively. Furthermore, optical distances between points in the cube were calculated using Simpson's 1/3 rule and 10 integration intervals. The maximum change (absolute) in the solution using 160 (r, θ, ϕ) intervals and 320 (r, θ, ϕ) intervals was less than 5×10^{-5} for the radiative flux divergence and less than 10^{-5} for the components of the radiative heat flux. Increasing the number of Simpson's integration intervals from 10 to 50 for the optical distance calculation resulted in a change of less than 10^{-5} in both the flux and flux divergence values. Table 1 summarizes the benchmark values for the radiative heat flux and flux divergence at various locations within the cube. Values obtained by Tan and Hsu [24] are also included in Table 1 for comparison.

ALGORITHM DESCRIPTION

This section provides a detailed description of the parallel S_n algorithm used to obtain the numerical and parallel performance results for the demonstration problem described above. This section begins with a review of the continuous governing equations and boundary conditions employed by the S_n method as described in the texts by Modest [16] and Siegel and Howell [17]. The discrete forms of the continuous equations are then derived and the solution logic is described. Finally, a theoretical estimate of the on-processor operation count is made.

Governing Equations

The S_n method evaluates a system of Boltzmann transport equations corresponding to a number of ordinated directions, \hat{s} . Each transport equation describes the variation of the directional radiative intensity, $I(\mathbf{r}, \hat{s})$, throughout the domain. For gray radiative transport, the Boltzmann transport equation may be written in the form:

$$\hat{s} \cdot \nabla I(\mathbf{r}, \hat{s}) + \mu(\mathbf{r})I(\mathbf{r}, \hat{s}) = Q, \quad (3)$$

where, μ is the total extinction cross section (i.e., scattering + absorption cross section), and Q is the radiative source term which includes volume emission and in-scattering. Each Boltzmann equation in the set corresponds to an ordinate direction from an angular quadrature set that discretizes the unit sphere. For nonscattering applications, the set of equations is uncoupled since the in-scattering term is zero.

Boundary conditions for Eq. (3) are only required along those portions of the global boundary, Γ , over which the inner product of the unit direction vector, \hat{s} , and the outward directed unit normal vector, \hat{n} , is less than zero:

$$\Gamma_s = \{\mathbf{r}: (\mathbf{r} \in \Gamma) \wedge (\hat{\mathbf{s}} \cdot \hat{\mathbf{n}}(\mathbf{r}) < 0)\} \subset \Gamma, \quad (4)$$

All global boundaries are assumed to be gray, diffuse boundaries where the boundary intensity is given by:

$$I(\mathbf{r}, \hat{\mathbf{s}}) = \varepsilon \sigma T^4(\mathbf{r}) + (1 - \varepsilon)[\mathbf{q}_r(\mathbf{r}) \cdot \hat{\mathbf{n}}(\mathbf{r})], \quad \forall \mathbf{r} \in \Gamma_s, \quad (5)$$

where, ε is the emissivity, and \mathbf{q}_r is the radiative heat flux vector. For the demonstration problem described above, all global boundaries are assumed to be cold and black so

$$I(\mathbf{r}, \hat{\mathbf{s}}) = 0, \quad \forall \mathbf{r} \in \Gamma_s.$$

The radiative heat flux and flux divergence at any point in the domain are obtained by integrating the directional intensities over all solid angles about location \mathbf{r} [16]:

$$\mathbf{q}_r(\mathbf{r}) = \int_0^{2\pi} \int_0^\pi I(\mathbf{r}, \hat{\mathbf{s}}) \hat{\mathbf{s}} \sin \theta d\theta d\phi, \quad (6)$$

$$\nabla \cdot \mathbf{q}_r(\mathbf{r}) = \mu_a(\mathbf{r}) \left\{ 4\sigma T^4(\mathbf{r}) - \int_0^{2\pi} \int_0^\pi I(\mathbf{r}, \hat{\mathbf{s}}) \sin \theta d\theta d\phi \right\}, \quad (7)$$

where, T is the temperature, μ_a is the absorptivity, σ is the Stefan-Boltzmann constant, and θ and ϕ are the zenith and azimuthal angles respectively, i.e., $\hat{\mathbf{s}} = -\sin \theta \sin \phi \hat{\mathbf{i}} - \cos \theta \hat{\mathbf{j}} - \sin \theta \cos \phi \hat{\mathbf{k}}$.

For this work, Eq. (3) is discretized using a one sided finite difference stencil on an orthogonal computational grid, cf. Figure 1. The resulting discrete representation of Eq. (3) for the intensity at node p may be written:

$$A_p I_p - A_e I_e - A_w I_w - A_n I_n - A_s I_s - A_t I_t - A_b I_b = Q(\mathbf{r}_p), \quad (8)$$

$$A_p = A_w + A_e + A_s + A_n + A_b + A_t + \mu(\mathbf{r}_p), \quad (8a)$$

$$A_w = \left(\frac{1}{h}\right) [\hat{\mathbf{s}} \cdot \hat{\mathbf{i}}, 0], \quad (8b)$$

$$A_e = \left(\frac{1}{h}\right) [-\hat{\mathbf{s}} \cdot \hat{\mathbf{i}}, 0], \quad (8c)$$

$$A_s = \left(\frac{1}{h}\right) [\hat{\mathbf{s}} \cdot \hat{\mathbf{j}}, 0], \quad (8d)$$

$$A_n = \left(\frac{1}{h}\right) \lceil -\hat{s} \cdot \hat{j}, 0 \rceil, \quad (8e)$$

$$A_b = \left(\frac{1}{h}\right) \lceil \hat{s} \cdot \hat{k}, 0 \rceil, \quad (8f)$$

$$A_t = \left(\frac{1}{h}\right) \lceil -\hat{s} \cdot \hat{k}, 0 \rceil, \quad (8g)$$

where, $\lceil a, b \rceil$ is the maximum value of a and b , r_p is the location of node p , and h is the node spacing. One sided differencing is used in the present work to preserve the diagonal dominance of the differencing operator as the extinction cross section approaches zero.

As stated above, the boundary intensity is set to zero at all nodes on the global boundary. The specification of the intensity at nodes on interprocessor boundaries is described in the next section.

Angular discretization is provided by the level weighted, even moment quadrature set of Lathrop and Carlson [15]. The directional intensity distribution is evaluated for each direction in the quadrature set and the radiative flux and flux divergence are approximated by:

$$\mathbf{q}_r(\mathbf{r}_p) \approx \sum_{i=1}^{n(n+2)} I(\mathbf{r}_p, \hat{s}_i) \hat{s}_i w_i, \quad (9)$$

$$\nabla \cdot \mathbf{q}_r(\mathbf{r}_p) = \mu_a(\mathbf{r}_p) \left\{ 4\sigma T^4(\mathbf{r}_p) - \sum_{i=1}^{n(n+2)} I(\mathbf{r}_p, \hat{s}_i) w_i \right\}, \quad (10)$$

where, m is the number of ordinate directions in the quadrature set, and w_i are the angular quadrature weights.

Logic Flow

The logic flow diagram for the parallel Sn algorithm developed for this work is given in Figure 2. At the start of the algorithm, each processor performs preprocessing tasks including initializing the message passing interface, reading the problem description file, and generating the angular direction (ordinate) set and the local (on-processor) computational grid. Once preprocessing is completed, all of the processors are synchronized and the total grind time counter is started before the radiative transport is calculated for each of the ordinate directions in the quadrature set.

For each ordinate direction, each processor calculates the coefficients of Eq. (8) for each node in the interior of the local mesh and sets the intensity to zero for nodes on Γ_s . During the first global solution sweep, the intensities at nodes on the interprocessor boundaries along which

$\hat{s} \cdot \hat{n}(r_p) < 0$ are set to zero. At subsequent sweeps, the intensities at these shared nodes are set to the value communicated from the neighboring processor/subdomain during the previous global solution sweep.

Once the linear system and boundary conditions are set, each processor conducts two Gauss-Seidel sweeps through the local system of equations. To accelerate convergence, the Gauss-Seidel sweeps are ordered relative to the ordinate direction. That is, the value of $r_p \cdot \hat{s}$ is evaluated at each node and the corresponding nodal equations are evaluated in order from the minimum to maximum value of $r_p \cdot \hat{s}$. In this way the boundary condition along Γ_s propagates naturally through the local grid in the \hat{s} direction and local, e.g., on processor, convergence may be obtained in a small number of Gauss-Seidel solution sweeps.

After the local Gauss-Seidel solution sweeps, a global reduction step is then executed to calculate the global residual norm from the local residual values and the global residual value is broadcast back to all processors. If the global residual value is less than some convergence criteria (10^{-6}) then the solution proceeds to the next ordinate direction. If the global solution is *not* converged, then each processor communicates the intensities at nodes on the interprocessor boundaries where $\hat{s} \cdot \hat{n}(r_p) > 0$ to the corresponding neighboring processor and the algorithm returns to the Gauss-Seidel solution step.

Computational Complexity Analysis

The computational complexity analysis presented here provides a theoretical estimate for the on-processor operation count required by the algorithm outlined above. The operation count is estimated for each global solution sweep and then the total number of solution sweeps is estimated to provide the total operation count. An expression is sought in terms of the number of nodes in the global domain and the number of subdomains.

Within an order of magnitude, the operation count per local Gauss-Seidel solution sweep, per ordinate direction is proportional to the number of nodes in the local mesh (Eq. (2)). Since $n_n \gg 1$ and $n_n/n_p \gg 1$, the average on-processor operation count per solution sweep, per ordinate direction is proportional to $(n_n/n_p)^3 = N/N_p$. Additionally, for a three dimensional angular quadrature set of order n , there are a total of $n(n+2)$ ordinate directions which discretize the unit sphere [17]. Thus, the on processor operation count per local Gauss-Seidel solution sweep is $O[n(n+2)N/N_p]$.

A consequence of the global nature of radiation transport for applications which are not optically thick is that the number of global solution sweeps required to obtain convergence using spatial domain-based parallelism increases as the number of processors. The cause of this behavior lays in the directional (i.e., hyperbolic) nature of the Boltzmann transport equation (Eq. (3)). Figure 3 provides an illustration of this behavior in one dimension.

Figure 3 shows a one dimensional domain decomposed into N_p nonoverlapping subdomains. During the first global solution sweep, only processor zero obtains the correct solution since this subdomain contains the global boundary at $s=0$. All other subdomains converge to an incorrect solution since these local solutions are based upon an unknown "upstream" boundary condition. During the second global solution sweep, the intensity at $s=s_0$ is communicated to subdomain 1

and now both subdomains 0 and 1 obtain the correct local solution. The global solution proceeds in this fashion until the boundary condition at $s=0$ has propagated through all the subdomains and the correct global solution is obtained.

As a result of the serial algorithmic effect described above, *the number of global iterations required to obtain a converged global solution increases with the number of processors.* The situation is improved somewhat for three dimensions since the solution proceeds along "planes" of subdomains orthogonal to the ordinate direction so there are more processors performing useful work. Also, for optically thick problems, the solution on each subdomain will depend only on the solution in neighboring subdomains and the number of global solution sweeps will be less dependent on the number of subdomains. For moderate optical thickness, the subdomain mapping summarized in Figure 1, $O(n_p)$ global solution sweeps will be required for each ordinate direction.

By combining the expressions described above, the operation count for a single processor for the entire computation is given by:

$$N_{flop} = Cn(n+2) \frac{N}{N_p^{2/3}}, \quad (11)$$

where, C is a constant of $O(1)$.

RESULTS

All of the results presented in this report were obtained on Sandia National Laboratories Intel PARAGON massively parallel supercomputer using the native NX message passing library [25]. Some development work was also performed using the MPI message passing library [26] on a network of SUN SPARC workstations as well as an IBM SP distributed memory parallel machine. The Intel PARAGON architecture consists of 1840 compute nodes each consisting of two Intel i860XP processors with 16 Mbytes of memory per compute node. Each processor runs under the SUNMOS operating system developed jointly by Sandia National Laboratories and the University of New Mexico in Albuquerque, New Mexico. The compute nodes are connected through a two dimensional mesh communication network with links in each direction capable of communicating at up to 200 Mbytes/second. Intel rates the peak speed of the PARAGON at 140 GFLOPs.

Numerical Performance

The overall numerical accuracy of the Sn solution was gauged by the accuracy of the radiative flux and flux divergence values along the lines $(x, 0.5, 0)$ and $(x, 0, 0)$ respectively. The nodal Sn values were compared to the benchmark solution (cf. Table 1) by defining a relative error at each node:

$$e_{\phi, i} = 100 \frac{(\phi_{Sn} - \phi_{b.m.})_i}{\|\phi_{b.m.}\|_{max}}, \quad (12)$$

where, e is the relative error, ϕ refers to the radiative flux or flux divergence, i refers to a specific nodal location, Sn refers to the Sn solution value, $b.m.$ refers to the benchmark value, and $\|\phi_{b.m.}\|_{max}$ refers to the maximum benchmark value along the line.

Figures 4 and 5 show the variation of the nodal relative errors for the normalized [23] radiative flux and flux divergence respectively for angular quadrature order $n=10$ and node spacings of $h=1/10, 1/20, 1/40, 1/80,$ and $1/160$. The data in Figures 4 and 5 show that the errors decrease in general with increasing spatial refinement. The maximum errors in the radiative flux divergence occur near the center of the cube where the absorptivity is maximum (cf. Eq. (1)). The maximum errors in the radiative heat flux, by comparison, occur in the corners of the cube.

An RMS error norm was also defined to express the overall solution accuracy in a compact form:

$$\|e\|_{RMS} = \left\{ \frac{\sum_{i=1}^{n_n} e_{\phi, i}^2}{n_n} \right\}^{1/2} \quad (13)$$

Figures 6 and 7 show the RMS error norm for the radiative heat flux and flux divergence for angular quadrature orders $n=4, 6, 8, 10, 12, 14, 16, 18,$ and mesh spacings of $h=1/10, 1/20,$ and $1/40$. Additionally, Figures 6 and 7 show RMS error values using a rotated angular quadrature set which was obtained by rotating the Lathrop and Carlson [15] quadrature set by 45° relative to the principal Cartesian axes. This data set was obtained to evaluate the effect of the structure of the quadrature set on the solution accuracy.

The data shown in Figure 6 shows that the spatial discretization level has a much stronger effect on the radiative heat flux at the wall of the cube than does the angular discretization level. For $h=1/10$, increasing the angular quadrature set order from $n=4$ to $n=10$ decreases the RMS error by approximately 1%. Further increase in the angular quadrature order does not significantly decrease the RMS error. For $h=1/20$ and $h=1/40$, the angular quadrature set order has a negligible effect on the RMS error. Rotating the quadrature set also does not significantly improve the accuracy of the Sn solution. In contrast, increasing the spatial resolution from $h=1/10$ to $h=1/40$ reduces the RMS radiative heat flux error by a factor of approximately 3 from 5-6% to 2%.

Figure 7 shows that both angular and spatial resolution have a strong effect on the accuracy of the radiative flux divergence. The RMS flux divergence error decreases significantly for all three spatial resolution levels when the angular quadrature set order is increased from $n=4$ to $n=10$. The improvement in the solution accuracy is much less pronounced for $n>10$ although the RMS error does continue to decrease for $h=1/20$ and $h=1/40$. Rotating the quadrature set decreases the RMS flux divergence error significantly over the unrotated set for all quadrature set orders. The

rotated quadrature set RMS error is also nearly constant relative to the quadrature set order.

Figure 8 shows the spatial convergence rate of the S_n results for an unrotated angular quadrature set of order $n=10$ and spatial resolution of $h=1/10, 1/20, 1/40, 1/80, 1/160$. The data in Figure 8 shows that the solution accuracy for both the radiative flux and flux divergence continues to improve as the spatial resolution is increased but that the overall convergence rate for the radiative flux and flux divergence is sublinear. This relatively slow convergence rate may be a result of the first order differencing used to obtain the directional intensity combined with the angular integration of the intensity required to evaluate the flux and flux divergence (cf. Eqs. (6) and (7)). The slope of the convergence data shown in Figure 8 appears to be fairly constant for $h=1/10$ to $1/80$ but appears to decrease to some extent for $h=1/160$. This may indicate that additional angular refinement ($n>10$) is necessary to maintain the same spatial convergence rate.

Parallel Performance

Gustafson, et al. [14] define two measures of parallel efficiency: fixed problem size and scaled problem size. Fixed problem size implies that the size of the global (distributed) mesh is constant as the number of processors is increased. For ideal parallelism, the ratio of the single processor grind time to the grind time with N_p processors (parallel speedup) should be equal to the number of processors. The fixed problem size efficiency may be defined as:

$$\eta_f = \frac{t_g(1)}{N_p t_g(N_p)}, \quad (14)$$

where, $t_g(N_p)$ is the grind or wall clock time for the entire calculation using N_p processors.

For large computational grids, it may not be possible to fit the entire grid on a single processor. In this case an incremental fixed problem size efficiency may be defined by extending Eq. (14):

$$\eta_{f,i} = \frac{N_{p,i} t_g(N_{p,i})}{N_p t_g(N_p)}, \quad (15)$$

where, $N_{p,i}$ is the smallest number of processors required to store the problem of interest. The incremental fixed size efficiency reduces to the absolute efficiency for $N_{p,i}=1$, $\eta_{f,1}=\eta_f$.

Scaled problem size implies that the size of the local on-processor grid is fixed and, therefore, the size of the distributed global grid increases as the number of processors increases. Theoretically, since each processor has the same amount of work to do, the grind time should not increase as the number of processors increases. Thus, the scaled problem size efficiency may be defined:

$$\eta_s = \frac{t_g(1)}{t_g(N_p)} \quad (16)$$

Theoretical expressions for each of the efficiencies defined above may be obtained by substituting Eq. (11) for the grind time in each expression. The resulting theoretical efficiencies, η^* , may be written:

$$\eta^*_f = \left(\frac{1}{N_p}\right)^{1/3}, \quad (17)$$

$$\eta^*_{f,i} = \left(\frac{N_{p,i}}{N_p}\right)^{1/3}, \quad (18)$$

$$\eta^*_s = \frac{N_L N_p^{2/3}}{N}. \quad (19)$$

From Eqs. (17)-(19), it is expected that the efficiency of the parallel Sn algorithm to be relatively poor. For example, for $N_p=8$, the fixed size efficiency should be on the order of 50% and drop to as low as 10% for $N_p=1000$. As discussed above, this rapid decrease in the parallel efficiency is very similar to the performance observed by de Oliveira, et al. [12] and Novo, et al. [13].

Table 2 summarizes the actual fixed problem size performance for global grid sizes of $N=5.06 \times 10^4$, 2.27×10^5 , and 1.77×10^6 nodes for a fixed angular quadrature order of $n=10$. The table shows the grind time, communication time, communication overhead, cycle time, and fixed size efficiencies as a function of the number of processors. The communication overhead is defined as the ratio of the communication time to the grind time and the cycle time is defined as the grind time per ordinate direction per node.

Both the incremental and absolute (i.e., $N_{p,i}=1$) fixed size efficiencies are compared to the theoretical expressions of Eqs. (17) and (18) in Figure 9. It is difficult to compare the incremental and absolute values directly since the incremental values do not take into account the dramatic decrease in the efficiency at $N_p=1$. As a result, the incremental efficiencies tend to be higher than the absolute values as in the $N=121^3$ case in Table 2. Nevertheless, the actual values compare well with the theoretical predictions.

Figure 9 also shows the fixed size efficiency values obtained by de Oliveira, et al. [12] and Novo, et al. [13]. de Oliveira, et al. asserted that higher fixed size efficiencies were obtained for applications with distributed rather than localized sources. This conclusion, however, was based on a comparison between incremental efficiencies, used for the distributed source case, and absolute efficiencies, used for the localized source case. The data in Figure 9 suggests that both the distributed and localized source efficiencies observed by de Oliveira, et al. are comparable to the data presented by Novo, et al. (uniform source) as well as the performance obtained in the present work.

Table 3 summarizes the actual performance data for two scaled problem cases using $N_L = 21^3$ and 39^3 for a fixed angular quadrature order of $n = 10$. The largest distributed grid considered in this report was obtained using 512 processors and $N_L = 39^3$. The resulting distributed grid for this case had 2.84×10^7 computational nodes.

Figure 10 compares the actual scaled problem size efficiency values in Table 3 with the theoretical expression given in Eq. (19). As with the fixed problem size cases, the scaled efficiency agrees well with the theoretical prediction and shows a dramatic decrease in the efficiency as the number of processors increases. Nevertheless, the scaled problem size cases showed that the size of the distributed global mesh could be increased 400 fold with only a 10 fold increase in the computation time using the spatial domain-based decomposition paradigm.

Figure 11 compares the actual parallel grind time to the equivalent sequential grind time for the fixed problem size cases. The sequential grind time was obtained by extrapolating the single processor cases assuming that the sequential grind time was directly proportional to the product of the number of nodes in the global mesh and the number of ordinate directions (cf. Eq. (11)). As shown in the figure, the grind time obtained with the spatial domain-based parallel Sn algorithm is more than an order of magnitude less than the equivalent sequential grind time for the largest grid.

The data in both Tables 2 and 3 show that the communication overhead is very low for the parallel algorithm outlined in Figure 2. The communication time required for the scaled problem cases was less than 5% of the total grind time for $N_L = 21^3$ and less than 3% for $N_L = 39^3$. Similar performance was obtained for the fixed problem size cases provided the number of nodes in each subdomain was greater than approximately 1000 nodes. As the number of nodes in each subdomain drops below this level (e.g., $N = 37^3$, $N_p = 216$) the communication time becomes significant relative to the number of operations performed by each processor.

The cycle times listed in Tables 3 and 4 decrease as the number of processors increases for both the fixed and scaled problem size cases. Approximately 100 microseconds are required per node per ordinate direction for a single processor. The cycle time decreases to several microseconds for $N_p > 100$. In contrast, the cycle times observed by de Oliveira, et al. [12] were on the order of several hundred microseconds. It is unclear from the data presented by de Oliveira, et al. whether these relatively high cycle times resulted from the complexity of the Pn model or from the unstructured finite element grid used for the spatial discretization.

CONCLUSIONS AND RECOMMENDATIONS

This study demonstrates the application of spatial domain based parallelism to three dimensional, participating media radiation heat transfer calculations using the Sn method. The objective was to provide data to support the development of an efficient parallel algorithm for use in large scale, transient, coupled, multi-mode, high temperature, reacting flow applications. A demonstration calculation was conducted using a well known benchmark solution from the open literature. Computational grids ranging from 9×10^3 to 28×10^6 nodes were considered with total computation times of several seconds to 1.7 hours depending on the number of processors employed. The following general observations may be made from the data presented:

- 1) The spatial domain based decomposition paradigm permits the use of very large computational grids on the order of 10^7 to 10^8 computational cells. Angular and frequency domain based decomposition paradigms are limited to moderate sized computational grids since these paradigms require that the entire computational grid be stored in the local processor memory.
- 2) Of the Sn, Pn, YIX, FE/FV, DT, and MC radiative transport models, the Sn, Pn, and DT models appeared to be the best candidates for spatial domain based parallelism based on operation count and communication overhead. Additional advantages of the Sn method include the simplicity of the formulation and the synchronous nature of the communication requirements.
- 3) The global character of radiative transport calculations in general gives rise to a corresponding sequential nature in the global solution sweeps which degraded the parallel efficiency of the Sn algorithm considered. Similar effects have been shown by other authors for the DT and Pn methods and are expected to arise in some form in all of the common radiative transport models currently in use.
- 4) The parallel efficiency should improve for optically thick problems as the local solution on each subdomain becomes less dependent of the solution in other subdomains.
- 5) Although the algorithmic efficiency is relatively low, the actual parallel efficiency is very close to the theoretical limit and significant parallel speedup was obtained using the parallel Sn algorithm. Scaled speedup tests showed that the mesh size could be increased more than 400 fold with only a 10 fold increase in the grid time.

It is anticipated that parallel supercomputer hardware with a processor speed approximately 10 times faster than the Intel PARAGON will become available in the near future. The new hardware will also have a greater memory capacity per processor and a larger number of processors than the Intel PARAGON. Additionally, the timing results presented in this report did not use the full capacity of the Intel PARAGON (1840 compute nodes) and employed a relatively high order angular quadrature set ($n=10$). With these considerations in mind, it is estimated that a 2-5 fold reduction in the current cycle time will be necessary to obtain reasonable turnaround time ($t_g < 10$ minutes) with $O(10^8)$ elements in an unstructured finite element grid and $n=6-10$ on the new hardware.

The results presented here were obtained using a structured finite difference grid. Work is currently underway to develop an unstructured, finite element algorithm based on the results and lessons learned in this analysis. The unstructured algorithm will serve as a prototype for a parallel, radiation heat transport library which may be called by a larger high temperature, reacting flow modeling program.

REFERENCES

1. T.M. Shih, L.J. Hayes, W.J. Minkowycz, K.-T. Yang, W. Aung, Parallel Computations in Heat Transfer, *Numerical Heat Transfer*, vol. 9, pp. 639-662, 1986.
2. J.R. Howell, Thermal Radiation in Participating Media: The Past, the Present, and Some Possible Futures, *ASME Journal of Heat Transfer*, vol. 110, November, pp. 1220-1229, 1988.
3. J.R. Howell, Radiative Heat Transfer: Opportunities and Challenges, *Proceedings of the ASME/JSME Thermal Engineering Conference*, Honolulu, Hawaii, pp. 23-33, 1995.
4. R.D. Skocypec, L.A. Gritz, T.W. Tong, Summary of the NSF Workshop on Using High-Performance Computing to Solve Participating Media Radiative Heat Transfer Problems, presented at the 6th AIAA/ASME Joint Thermophysics and Heat Transfer Conference, Colorado Springs, CO, June, 1994.
5. Y.Y. Azmy, Multiprocessing for Neutron Diffusion and Deterministic Transport Methods, Invited for Publication in *Progress in Nuclear Energy*, 1995.
6. P. J. Burns, M. Christon, R. Schweitzer, O. M. Lubeck, H. J. Wasserman, M. L. Simmons, and D. V. Pryor, Vectorization of Monte Carlo Particle Transport: An Architectural Study Using the LANL Benchmark "GAMTEB", *Proceedings Supercomputing '89*, pp. 10-20, 1989.
7. P. J. Burns, J. D. Maltby, and M. A. Christon, Large-Scale Surface to Surface Transport for Photons and Electrons via Monte Carlo, *Computing Systems in Engineering*, Vol. 1, No. 1, pp. 75-99, 1990.
8. J.L. Haferman, T.F. Smith, W.F. Krajewski, Multi-Dimensional Radiative Transfer Computations Using a Scalable Parallel Implementation of the Discrete-Ordinates Method, *Proceedings of the International Geosciences and Remote Sensing Symposium (IGARSS)*, Pasadena, CA, vol. 3, pp. 1623-1625, 1994.
9. J.R. Howell, Improving the Monte Carlo Method for Radiative Transfer by Use of Parallel Processors, *Heat Transfer in Thermal Systems Seminar-Phase II*, National Cheng Kung University, Taiwan, January 13-14, pp. 53-57, 1986.
10. J. M. McGhee and J. E. Morel, NIKE/ATHENA User's Manual, LA-UR-94-4312, Los Alamos National Laboratory, Los Alamos, NM, 1994.
11. A. Benmalek, W. Li, and T. W. Tong, *MYST User's Manual - Sequential and Parallel Versions*, Department of Mechanical and Aerospace Engineering, Arizona State University, Tempe, Arizona, 1995.
12. C.R.E. de Oliveira, C.C. Pain, A.J.H. Goddard, Parallel Domain Decomposition Methods for Large-Scale Finite Element Transport Modeling, pp. 490-498, 1995

13. P.J. Novo, P.J. Coelho, M.G. Carvalho, Parallelization of the Discrete Transfer Method: Two Different Approaches, Proceedings of the 31st National Heat Transfer Conference, August 3-6, Houston, TX, *ASME HTD-Vol. 325*, pp. 45-54, 1996.
14. J.L. Gustafson, G.R. Montry, R.E. Benner, Development of Parallel Methods for a 1024-Processor Hypercube, *SIAM Journal on Scientific and Statistical Computing*, vol. 9, no. 4, pp. 609-638, 1988.
15. K.D. Lathrop, B.G. Carlson, Discrete Ordinates Angular Quadrature of the Neutron Transport Equation, *Los Alamos Scientific Laboratory Report, LA-3186*, Los Alamos, New Mexico, 1965.
16. M.F. Modest, *Radiative Heat Transfer*, McGraw Hill, Inc., New York, NY, 1993
17. R. Siegel, J.R. Howell, *Thermal Radiation Heat Transfer, 3rd ed.*, Hemisphere Publishing Corp., Washington, D.C., 1992
18. R. Viskanta, M.P. Mengüç, Radiation Heat Transfer in Combustion Systems, *Progress Energy Combustion Science*, vol. 13, pp. 97-160, 1987
19. F.C. Lockwood, N.G. Shah, A New Radiation Solution Method for Incorporation in General Combustion Prediction Procedures, Eighteenth Symposium (International) on Combustion, The Combustion Institute, Pittsburg, PA, pp. 1405-1414, 1988
20. W.A. Fiveland, Three-Dimensional Radiative Heat-Transfer Solutions by the Discrete-Ordinates Method, *Journal of Thermophysics and Heat Transfer*, vol. 2, no. 4, pp. 309-316, 1988.
21. W.A. Fiveland, Combustion Modeling for Fossil-Fired Applications, Presented at the 31st National Heat Transfer Conference, August 3-6, Houston, TX, 1996.
22. P. Hsu, J.T. Farmer, Benchmark Solutions of Radiative Heat Transfer Within Nonhomogeneous Participating Media Using the Monte Carlo and YIX Methods, Proceedings of the 30th National Heat Transfer Conference, *ASME HTD-vol. 315*, Portland, OR, pp. 29-36, 1995.
23. S.P. Burns, J.R. Howell, D.E. Klein, Finite Element Solution for Radiative Heat Transfer With Nongray, Nonhomogeneous Radiative Properties, Proceedings of the 30th National Heat Transfer Conference, *ASME HTD-vol. 315*, Portland, OR, pp. 3-10, 1995.
24. Z. Tan, P. Hsu, High Order Interpolation Schemes in the YIX Method, Presented at the 30th National Heat Transfer Conference, Open Forum on Radiative Heat Transfer in Participating Media, Portland, OR, 1995.
25. *PARAGON Fortran System Calls Reference Manual*, doc. no. 312488-BETA, Intel

Supercomputer Systems Division, Beaverton, OR, March, 1994

26. W. Gropp, E. Lusk, A. Skjellum, *Using MPI, Portable Parallel Programming with the Message-Passing Interface*, MIT Press, Cambridge, MA, 1995

FIGURES




FIGURE 1 - Subdomain Mapping for 4x4x4 Subdomain Mesh.




FIGURE 2 - Spatial Domain Based Parallel Discrete Ordinate Algorithm Logic Flow Diagram.




FIGURE 3 - One Dimensional Spatial Domain Decomposition Example.




FIGURE 4 - Distributed Relative Error for the Scaled Radiative Heat flux along $(x, 0.5, 0)$ Using Angular Quadrature Order $n=10$.




FIGURE 5 - Distributed Relative Error for the Scaled Radiative Flux Divergence Along $(x, 0, 0)$ Using Angular Quadrature Order $n=10$.




FIGURE 6 - Spatial and Angular Convergence for the Radiative Heat Flux Along $(x, 0.5, 0)$.




FIGURE 7 - Spatial and Angular Convergence for the Radiative Flux Divergence Along (x, 0, 0).




FIGURE 8 - Spatial Convergence Rate for Angular Quadrature Order n= 10.




FIGURE 9 - Fixed Problem Size Efficiency for Angular Quadrature Order $n=10$.




FIGURE 10 - Scaled Problem Size Efficiency for Angular Quadrature Order $n=10$.




FIGURE 11 - Raw Grind Time Data for Angular Quadrature Order n= 10.

TABLES

TABLE 1 - Benchmark Results for Isothermal, Unit Cube with Absorptivity Varying Trilinearly and No Scattering

$\pm x$	$\frac{\nabla \cdot q_r(x, 0, 0)}{\sigma T_o^4}$ a	$\frac{n \cdot q_r(x, 0.5, 0)}{\sigma T_o^4}$ b	$\pm x$	$\frac{\nabla \cdot q_r(x, 0, 0)}{\sigma T_o^4}$ a	$\frac{n \cdot q_r(x, 0.5, 0)}{\sigma T_o^4}$ b
0.000	3.07561	0.18802	0.275	1.69575	0.15181
	3.07558 ^c	0.18806 ^c			
0.025	2.94305	0.18762	0.300	1.56180	0.14607
0.050	2.81677	0.18646	0.325	1.42485	0.14006
0.075	2.69362	0.18464	0.333	1.37850 ^c	0.13801 ^c
0.100	2.57196	0.18217	0.350	1.28473	0.13380
0.111	2.51807 ^c	0.18091 ^c	0.375	1.14130	0.12731
0.125	2.45069	0.17916	0.400	0.99443	0.12056
0.150	2.32896	0.17562	0.425	0.84405	0.11354
0.175	2.20616	0.17163	0.444	0.72462 ^c	0.10781 ^c
0.200	2.08181	0.16721	0.450	0.69010	0.10614
0.222	1.96963 ^c	0.16297 ^c	0.475	0.53255	0.09815
0.225	1.95549	0.16241	0.500	0.37165	0.08821
0.250	1.82691	0.15727			

a. $\pm 5 \times 10^{-5}$

b. $\pm 10^{-5}$

c. Z. Tan, P. Hsu, High Order Interpolation Schemes in the YIX Method, Presented at the 30th National Heat Transfer Conference, Open Forum on Radiative Heat Transfer in Participating Media, Portland, OR, 1995

TABLE 2 - Fixed Sized Timing Study Results, Parallel Sn Algorithm, n= 10

N_L	N_p	t_g (sec)	t_c (sec)	t_c/t_g	t_{cycle} (μs)	$\eta_f, \eta_{f,i}$
$N= 37^3$						
5.1×10^4	1	509.3	0.0	0%	84	100%
6.9×10^3	8	135.0	3.5	2.6%	22	47%
2.2×10^3	27	57.1	3.7	6.5%	9.4	33%
1.0×10^3	64	31.5	4.0	12.7%	5.2	25%
3.4×10^2	216	16.3	5.7	35.0%	2.7	14%
$N= 61^3$						
3.0×10^4	8	638.1	10.3	1.6%	23	100%
9.3×10^3	27	352.7	11.6	3.3%	13	54%
4.1×10^3	64	147.2	8.1	5.5%	5.4	54%
2.2×10^3	125	87.0	7.7	8.9%	3.2	47%
1.3×10^3	216	59.0	7.6	12.9%	2.1	40%
$N= 121^3$						
3.0×10^4	64	1086.5	25.8	2.4%	5.1	100%
1.6×10^4	125	649.3	21.2	3.3%	3.1	86%
9.3×10^3	216	441.6	19.0	4.3%	2.1	73%
4.1×10^3	512	244.8	18.4	7.5%	1.2	55%

TABLE 3 - Scaled Timing Study Results, Parallel Sn Algorithm, n= 10

N	N_p	t_g (sec)	t_c (sec)	t_c/t_g	t_{cycle} (μs)	η_s
$N_L = 21^3$						
9.3×10^3	1	81.1	0.0	0%	73	100%
6.9×10^4	8	186.3	4.5	2.4%	23	44%
2.3×10^5	27	352.7	11.6	3.3%	16	23%
5.3×10^5	64	437.0	16.6	3.8%	6.9	19%
1.0×10^6	125	377.5	16.9	4.5%	3.1	21%
1.8×10^6	216	609.5	26.3	4.3%	2.9	13%
2.8×10^6	343	708.4	33.1	4.7%	2.1	11%
4.2×10^6	512	788.1	38.7	4.9%	1.6	10%
$N_L = 39^3$						
5.9×10^4	1	601.6	0.0	0%	85	100%
4.6×10^5	8	1326.9	17.6	1.3%	24	45%
1.5×10^6	27	2670.5	45.7	1.7%	15	23%
3.6×10^6	64	3517.2	66.6	1.9%	8.2	17%
7.0×10^6	125	4141.1	90.2	2.2%	5.0	15%
1.2×10^7	216	3030.7	75.8	2.5%	2.1	20%
1.9×10^7	343	5554.0	114.3	2.1%	2.4	11%
2.8×10^7	512	6237.8	160.0	2.6%	1.8	9%

DISTRIBUTION:

- 1 Professor Hohn Howell
The University of Texas at Austin
Office of the Dean of Engineering
Austin, TX 78712
- 1 Dr. Dale Klein, Associate Chancellor
The University of Texas System
210 E. 6th Street
Austin, TX 78701
- 1 Professor Patrick Burns
Colorado State University
Mechanical Engineering Department
Fort Collins, CO 80523
- 1 Professor Erik Thompson
Colorado State University
Civil Engineering Department
Fort Collins, CO 80523
- 1 Dr. James D. Maltby
Lawrence Livermore National Laboratory
Livermore, CA 94550
- 1 MS 0841 P.J.Hommert, 9100
1 MS 0826 M.W.Glass, 9111
1 MS 0835 T.Bickel, 9113
50 MS 0835 S.P.Burns, 9113
1 MS 0835 R.J.Cochran, 9113
1 MS 0835 D.M.Hensing, 9113
1 MS 0835 T.E.Voth, 9113
1 MS 0321 W.Camp, 9200
1 MS 1111 J.N.Shadid, 9221
1 MS 0819 J.N.Peery, 9231
50 MS 0819 M.A.Christon, 9231
1 MS 1179 J.R.Lee, 9341
1 MS 1179 G.D.Valdez, 9341
1 MS 1179 L.Lorence, 9341
1 MS 9018 Central Technical Files, 8523-2
5 MS 0899 Technical Library, 4414
2 MS 0619 Review & Approval Desk, 12630
For DOE/OSTI