Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Computational Physics 189 (2003) 539–556
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

A residual Monte Carlo method for discrete thermal radiative diffusion

T.M. Evans *, T.J. Urbatsch, H. Lichtenstein, J.E. Morel

CCS–4, MS D409, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

Received 30 October 2002; received in revised form 16 April 2003; accepted 22 April 2003

Abstract

Residual Monte Carlo methods reduce statistical error at a rate of $\exp(-bN)$, where b is a positive constant and N is the number of particle histories. Contrast this convergence rate with $1/\sqrt{N}$, which is the rate of statistical error reduction for conventional Monte Carlo methods. Thus, residual Monte Carlo methods hold great promise for increased efficiency relative to conventional Monte Carlo methods. Previous research has shown that the application of residual Monte Carlo methods to the solution of continuum equations, such as the radiation transport equation, is problematic for all but the simplest of cases. However, the residual method readily applies to discrete systems as long as those systems are monotone, i.e., they produce positive solutions given positive sources. We develop a residual Monte Carlo method for solving a discrete 1D non-linear thermal radiative equilibrium diffusion equation, and we compare its performance with that of the discrete conventional Monte Carlo method upon which it is based. We find that the residual method provides efficiency gains of many orders of magnitude. Part of the residual gain is due to the fact that we begin each timestep with an initial guess equal to the solution from the previous timestep. Moreover, fully consistent non-linear solutions can be obtained in a reasonable amount of time because of the effective lack of statistical noise. We conclude that the residual approach has great potential and that further research into such methods should be pursued for more general discrete and continuum systems.

1. Introduction

Monte Carlo methods are commonly used to simulate radiation transport phenomena. The overwhelming majority of these methods are constrained by the Central Limit Theorem. Thus, convergence in these methods is limited to $1/\sqrt{N}$, where N is the number of particles.

Halton [1] has shown that a Monte Carlo method based upon a residual approach could be used to solve certain matrix equations. The exciting aspect of this work was that exponential convergence was obtained; the statistical error was proportional to $\exp(-bN)$. This work has not been disseminated widely in the transport community.

* Corresponding author. Tel.: +1-505-665-3677; fax: +1-505-665-5538.
E-mail address: tme@lanl.gov (T.M. Evans).

0021-9991/$ - see front matter © 2003 Elsevier Science B.V. All rights reserved.
doi:10.1016/S0021-9991(03)00233-X

Over the last several years, researchers at Los Alamos National Laboratory have applied residual Monte Carlo methods to radiation transport calculations with the intent of demonstrating exponential convergence for continuum systems [2–6]. While they were able to demonstrate exponential convergence for simple, homogeneous, one-speed, 1D problems, significant difficulties were encountered for simple multidimensional problems. As we will explain later, such difficulties can be expected to arise for a broad class of problems. Presently, the true potential of residual Monte Carlo methods relative to conventional Monte Carlo methods for solving complex problems of physical interest is not well established.

The purpose of the present study is to demonstrate that enormous speedups, relative to conventional Monte Carlo, can be obtained using a residual Monte Carlo method to solve a discrete system that *approximates* a continuum system. Most importantly, we show that statistical errors in the solution can be efficiently reduced to negligible levels making it possible to solve a strongly nonlinear, discrete system with a fully implicit treatment of the nonlinearities. It is essentially impossible to do this with conventional Monte Carlo techniques.

2. The residual Monte Carlo method

2.1. The central theme of residual Monte Carlo methods

We wish to solve the following linear equation:

$$\hat{A}x = b, \tag{1}$$

where $\hat{A}$ is a linear operator, x is the solution, and b is the driving function. Suppose that we have an approximate solution, $\tilde{x}$, to Eq. (1). The error associated with $\tilde{x}$ is defined as follows:

$$\delta\tilde{x} = x - \tilde{x}. \tag{2}$$

The error term, $\delta\tilde{x}$, can also be considered a correction term for the approximate solution, $x = \tilde{x} + \delta\tilde{x}$.

Substituting Eq. (2) into Eq. (1) yields

$$\hat{A}\delta\tilde{x} = r, \tag{3}$$

where r is the residual and is defined

$$r = b - \hat{A}\tilde{x}. \tag{4}$$

Clearly, any method that solves Eq. (1) can be used to solve Eq. (3). Furthermore, if $\tilde{x} = 0$ then Eq. (3) becomes identical to Eq. (1).

Based upon the previous discussion, we can postulate a residual Monte Carlo algorithm:

1. Start with a solution estimate, $\tilde{x}$.
2. Do a *stage*
   2.1. perform a conventional Monte Carlo calculation to solve Eq. (3) for the correction, $\delta\tilde{x}$.
   2.2. add $\delta\tilde{x}$ to $\tilde{x}$ to obtain a new solution estimate.
3. Run stages until $\tilde{x}$ converges.

This algorithm constitutes the theme of residual Monte Carlo methods.

2.2. Difficulties for continuum equations

Previous attempts at residual Monte Carlo methods for the radiation transport equation have used global polynomial expansions for the continuum transport solution. Some of these expansions have in-

volved high-order, highly oscillatory Legendre polynomials [4,3]. Unfortunately, these attempts applied the residual method to the expansion terms such that the residual method's domain of effectiveness was a priori ill defined. Indeed, the cited studies [4,3] showed that a prohibitively large number of histories was required during each stage to achieve convergence of the iterative process.

An alternative to high-order, global basis functions is local, low-order, finite-element basis functions that would be used in conjunction with a weak formulation of the transport equation. However, it is unclear that such an approach would be more effective than previous attempts at continuum solution residual methods.

Instead, in this paper we have focused on residual Monte Carlo methods for discrete systems of equations. The performance of the residual method on discrete systems represents an upper bound on the possible performance for continuum systems. In the following section, we will derive a scheme for solving a discrete approximation of the nonlinear, thermal radiative equilibrium diffusion equation. Afterward, we will apply a residual scheme to the discrete system of equations.

3. A discrete scheme for nonlinear thermal radiative diffusion

3.1. Equilibrium diffusion

Consider the gray, equilibrium diffusion equation, in the absence of external sources or scattering, that is derived from an asymptotic expansion of the radiation transport equation [7]

$$(C_v + 4aT^3)\frac{\partial T}{\partial t} - \frac{\partial}{\partial x}\frac{4acT^3}{3\sigma_R}\frac{\partial T}{\partial x} = 0, \tag{5}$$

where $C_v \equiv C_v(x, T(x,t))$ (jerks cm^{-3} keV^{-1}) is the specific heat capacity of the material, $a = 0.01372$ (jerks cm^{-3} keV^{-4}) is the radiation constant, $c = 299.79$ (cm sh^{-1}) is the vacuum light speed, and $T \equiv T(x,t)$ (keV) is the temperature that characterizes both the radiation and the material. The opacity, $\sigma_R \equiv \sigma_R(x, T(t))$ (cm^{-1}), is the Rosseland mean opacity. Additionally, we can define a diffusion coefficient, $D \equiv D(x, T(x,t))$ (cm), as follows:

$$D = \frac{1}{3\sigma_R}. \tag{6}$$

The left-most term in Eq. (5) represents an effective heat capacity of the coupled radiation and material. The equation is nonlinear in temperature, both from terms containing T explicitly and implicitly through σ_R and C_v.

In what follows, we shall define Eq. (5) in terms of T^4 instead of T,

$$\left(\frac{C_v}{4acT^3} + \frac{1}{c}\right)\frac{\partial \phi}{\partial t} - \frac{\partial}{\partial x} D \frac{\partial \phi}{\partial x} = 0, \tag{7}$$

where $\phi \equiv \phi(x,t)$ (jerks cm^{-2} sh^{-1}) is the angle-energy integrated radiation intensity. The intensity is defined, in the local thermodynamic equilibrium (LTE) limit,

$$\phi(x,t) = \int_\Omega d\Omega \int_\nu d\nu \psi(x,\nu,\Omega,t) = \int_\nu d\nu 4\pi B(\nu,T) = acT^4, \tag{8}$$

where ψ is the angular radiation intensity. The Planck function (or Planckian) is defined as

$$B(\nu,T) = \frac{2h\nu^3}{c^2}\left(e^{h\nu/kT} - 1\right)^{-1}, \tag{9}$$

where h is Planck's constant and k is the Boltzmann constant.

3.2. Discrete equilibrium diffusion

We wish to formulate a set of discrete probability equations in terms of the face-centered partial fluxes that we can solve using Monte Carlo methods. Casting the equations in terms of the face-centered partial fluxes provides a convenient form for a Monte Carlo scheme because particle flow provides the inter-cell coupling. We begin by applying a backward-Euler method to discretize Eq. (7) in time

$$\left(\frac{C_v^n}{4acT^{n\,3}}+\frac{1}{c}\right)\frac{\phi^{n+1}-\phi^n}{\Delta t}-\frac{\mathrm{d}}{\mathrm{d}x}D^n\frac{\mathrm{d}\phi^{n+1}}{\mathrm{d}x}=0. \tag{10}$$

We begin the spatial treatment of the time-discretized equation by considering Fig. 1. Operating on Eq. (10) with $\int_{x_{i-1/2}}^{x_{i+1/2}}(\cdot)\,\mathrm{d}x$ and applying Fick's Law,

$$F(x,t)=-D\frac{\partial\phi(x,t)}{\partial x}, \tag{11}$$

yields

$$\tilde{\sigma}_i^n\phi_i^{n+1}\Delta x+F_{i+1/2}^{n+1}-F_{i-1/2}^{n+1}=\tilde{\sigma}_i^n\phi_i^n\Delta x, \tag{12}$$

where

$$\tilde{\sigma}_i^n=\left(\frac{C_v^n}{4acT_c^{n\,3}\Delta t}+\frac{1}{c\Delta t}\right)\quad(\mathrm{cm}^{-1}), \tag{13}$$

$$\phi_i^n=acT_i^{n\,4}\quad(\text{jerks cm}^{-2}\ \text{sh}^{-1}), \tag{14}$$

$$D^n=\frac{1}{3\sigma_{\mathrm{R}}^n}\quad(\mathrm{cm}). \tag{15}$$

Here, $\Delta x=x_{i+1/2}-x_{i-1/2}$, which we will consider constant throughout the mesh for simplicity. Also, F is the face-centered flux in units of (jerks cm^{-2} sh^{-1}), defined as the net difference of the right-going and left-going partial fluxes

$$F^{n+1}=f^{(+)\,n+1}-f^{(-)\,n+1}. \tag{16}$$

Eq. (12) is a photon-energy balance equation that states that the photon absorption at time t^{n+1} and the sum of the net outgoing fluxes on the cell-edges is equal to the material radiation emission in the cell at time t^n. To complete the discrete diffusion scheme we need to define expressions for the outgoing partial fluxes in terms of the cell-centered intensities, sources, and incoming partial fluxes.

We can apply forward-differencing to Fick's Law, Eq. (11), to achieve expressions for the face-centered fluxes

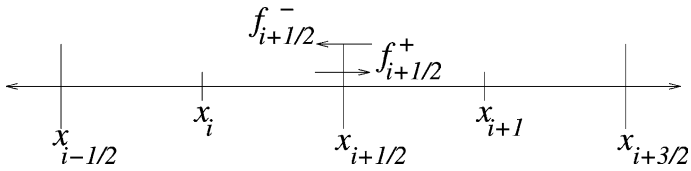

Fig. 1. 1D discretized mesh. We integrate over the cell from $x_{i-1/2}$ to $x_{i+1/2}$.

$$F_{i-1/2}^{n+1} = \frac{-2D^n}{\Delta x}\left(\phi_i^{n+1} - \phi_{i-1/2}^{n+1}\right), \tag{17a}$$

$$F_{i+1/2}^{n+1} = \frac{-2D^n}{\Delta x}\left(\phi_{i+1/2}^{n+1} - \phi_i^{n+1}\right). \tag{17b}$$

For reasons of symmetry, D^n is evaluated at the cell center. However, the diffusion coefficient for a particular cell must be chosen properly to allow correct wave propagation. Evaluating at the cell-centered temperature results in stagnated waves. Therefore, we first obtain face-centered temperatures with a harmonic average of the adjoining cell temperatures according to the following formula:

$$T_{j\pm1/2}^m = \frac{1}{3}\left(T_j^m + \sqrt{T_j^m T_{j\pm1}^m} + T_{j\pm1}^m\right), \tag{18}$$

where m is chosen to best match the underlying data [8,9]. We choose $m = 4$. The maximum of a cell's two face-centered temperatures is used to evaluate D^n.

We can use the Marshak boundary condition

$$4f_{\text{bnd}}^{(\pm)} = \phi(x) + 2D(x)\frac{\partial\phi(x)}{\partial x}\mathbf{n}\cdot\mathbf{i}, \tag{19}$$

where $\mathbf{i}$ is the particle direction and $\mathbf{n}$ is the outward normal, to write equations for the incoming partial fluxes. Substituting Fick's Law into Eq. (19), we write

$$f_{i-1/2}^{(+)\,n+1} = \frac{1}{4}\phi_{i-1/2}^{n+1} + \frac{1}{2}F_{i-1/2}^{n+1}, \tag{20a}$$

$$f_{i+1/2}^{(-)\,n+1} = \frac{1}{4}\phi_{i+1/2}^{n+1} - \frac{1}{2}F_{i+1/2}^{n+1}. \tag{20b}$$

Now, eliminating the cell-edge intensities between Eqs. (17a), (17b) and (20a), (20b) and using the definition of the net flux, we can write expressions for the outgoing, face-centered, partial fluxes,

$$f_{i+1/2}^{(+)\,n+1} = \frac{2D^n}{\Delta x + 4D^n}\phi_i^{n+1} + \frac{\Delta x - 4D^n}{\Delta x + 4D^n}f_{i+1/2}^{(-)\,n+1}, \tag{21a}$$

$$f_{i-1/2}^{(-)\,n+1} = \frac{2D^n}{\Delta x + 4D^n}\phi_i^{n+1} + \frac{\Delta x - 4D^n}{\Delta x + 4D^n}f_{i-1/2}^{(+)\,n+1}. \tag{21b}$$

As implied earlier, these equations are symmetric because we use a cell-centered diffusion coefficient.

Having defined equations for the outgoing partial fluxes, we now require an expression for the cell-centered intensity, ϕ_i^{n+1}, in terms of the incoming partial fluxes. We can write expressions for the ace-centered fluxes by considering Eqs. (17a), (17b) and (20a), (20b)

$$F_{i-1/2}^{n+1} = \frac{8D^n}{\Delta x + 4D^n}f_{i-1/2}^{(+)\,n+1} - \frac{2D^n}{\Delta x + 4D^n}\phi_i^{n+1}, \tag{22a}$$

$$F_{i+1/2}^{n+1} = \frac{2D^n}{\Delta x + 4D^n}\phi_i^{n+1} - \frac{8D^n}{\Delta x + 4D^n}f_{i+1/2}^{(-)\,n+1}. \tag{22b}$$

Plugging these equations into the balance equation, Eq. (12), yields the following expression for the cell-centered intensity:

$$\phi_i^{n+1} = \frac{\tilde{\sigma}_i^n \phi_i^n \Delta x + \frac{8D^n}{\Delta x + 4D^n}\left(f_{i-1/2}^{(+)\,n+1} + f_{i+1/2}^{(-)\,n+1}\right)}{\tilde{\sigma}_i^n \Delta x + \frac{4D^n}{\Delta x + 4D^n}}. \tag{23}$$

We can substitute Eq. (23) into Eqs. (21a), (21b); this gives us the following equations for the outgoing, partial fluxes:

$$f_{i+1/2}^{(+)\,n+1} = P_{\mathrm{T},i} f_{i-1/2}^{(+)\,n+1} + P_{\mathrm{R},i} f_{i+1/2}^{(-)\,n+1} + q_{i+1/2}^{(+)\,n}, \tag{24a}$$

$$f_{i-1/2}^{(-)\,n+1} = P_{\mathrm{R},i} f_{i-1/2}^{(+)\,n+1} + P_{\mathrm{T},i} f_{i+1/2}^{(-)\,n+1} + q_{i-1/2}^{(-)\,n}. \tag{24b}$$

Here, $P_{\mathrm{T},i}$ is the probability of transmission through cell i, $P_{\mathrm{R},i}$ is the probability of reflection from cell i, and $P_{\mathrm{A},i} = 1 - P_{\mathrm{T},i} - P_{\mathrm{R},i}$ is the probability of absorption in cell i. Because of the symmetry resulting from having a cell-centered diffusion coefficient, the probabilities are the same in both equations. The form for the probability of transmission is

$$P_{\mathrm{T},i} = \frac{4\left(\frac{2D^n}{\Delta x + 4D^n}\right)^2}{\tilde{\sigma}_i^n \Delta x + \frac{4D^n}{\Delta x + 4D^n}} = \frac{4R^2}{\tilde{\sigma}_i^n \Delta x + 2R}, \tag{25}$$

and the probability of reflection is

$$P_{\mathrm{R},i} = \frac{4\left(\frac{2D^n}{\Delta x + 4D^n}\right)^2}{\tilde{\sigma}_i^n \Delta x + \frac{4D^n}{\Delta x + 4D^n}} + \frac{\Delta x - 4D^n}{\Delta x + 4D^n} = 1 - 4R\left(\frac{\tilde{\sigma}_i^n \Delta x + R}{\tilde{\sigma}_i^n \Delta x + 2R}\right), \tag{26}$$

where

$$R = \frac{2D^n}{\Delta x + 4D^n}. \tag{27}$$

It follows that the probability of absorption, though not explicitly used in the method, is

$$P_{\mathrm{A},i} = 1 - P_{\mathrm{T},i} - P_{\mathrm{R},i} = \frac{4R\tilde{\sigma}_i^n \Delta x}{\tilde{\sigma}_i^n \Delta x + 2R}. \tag{28}$$

The outgoing face-centered sources from each cell are equal in magnitude in this case because of symmetry. The sources are expressed as

$$q_{i+1/2}^{(+)\,n} = q_{i-1/2}^{(-)\,n} = \frac{2D^n}{\Delta x + 4D^n}\left(\frac{\tilde{\sigma}_i^n \phi_i^n \Delta x}{\tilde{\sigma}_i^n \Delta x + \frac{4D^n}{\Delta x + 4D^n}}\right). \tag{29}$$

Eqs. (24a), (24b) are the fundamental equations of the discrete equilibrium diffusion Monte Carlo (EqDDMC) method [10].

3.3. Equilibrium discrete diffusion Monte Carlo

Having defined a set of probability equations for the partial fluxes at each face, we can now describe a conventional Monte Carlo method for solving these equations.

1. Calculate the $P_{\mathrm{R},i}$ and $P_{\mathrm{T},i}$ for each cell.
2. Sample source particles in each cell i with weights

$$w_i^{(+)} = \frac{q_{i+1/2}^{(+)n} + \frac{1}{2} Q_i^n \Delta x}{N_i^{(+)}}, \tag{30a}$$

$$w_i^{(-)} = \frac{q_{i-1/2}^{(-)n} + \frac{1}{2} Q_i^n \Delta x}{N_i^{(-)}}, \tag{30b}$$

where the Q_i^n are cell-centered external sources and the N_i are the numbers of plus/minus particles per cell. The source numbers are calculated

$$N_i^{(+)} = \text{int}\left[\frac{q_{i+1/2}^{(+)n} + \frac{1}{2} Q_i^n \Delta x}{\bar{Q}} * N \right], \tag{31a}$$

$$N_i^{(-)} = \text{int}\left[\frac{q_{i-1/2}^{(-)n} + \frac{1}{2} Q_i^n \Delta x}{\bar{Q}} * N \right], \tag{31b}$$

where N is the total number of source particles and $\bar{Q}$ is the total source energy in the problem

$$\bar{Q} = \sum_i \left(q_{i+1/2}^{(+)n} + q_{i-1/2}^{(-)n} + Q_i^n \Delta x \right) + q_{1/2}^{(+)n} + q_{I+1/2}^{(-)n}. \tag{32}$$

3. Stochastically transport particles from cell to cell based on the pre-calculated probabilities for transmission, reflection, and absorption. Tally the $f^{(+)n+1}$ and $f^{(-)n+1}$ on each face. A particle history is terminated when a particle escapes the system or is absorbed.
4. After transporting all particles, use Eq. (23) to calculate ϕ_i^{n+1}. Eq. (8) can be used to calculate T^{n+1}.
5. Update material properties and run next timestep until the problem end time is reached.

These steps represent the application of a conventional Monte Carlo algorithm to solve Eqs. (24a), (24b). In Section 4 we will derive the residual form of these equations that will follow the same basic solution procedure illuminated here.

3.4. Negative probabilities

The probabilities of absorption, transmission, and reflection always sum to one, and the probability of absorption and transmission are always non-negative. However, the equations are such that the probability of reflection may be negative. The probability of reflection is

$$1 - 4R\left(\frac{\tilde{\sigma}_i^n \Delta x + R}{\tilde{\sigma}_i^n \Delta x + 2R} \right), \tag{33}$$

where R is defined in Eq. (27). Substituting, we find that the condition for a positive reflection probability is

$$\frac{4D^n}{\Delta x + 4D^n} > \frac{4D^n - \Delta x}{c\Delta t} \left[1 + \frac{C_v^n}{4aT^{n\,3}} \right]. \tag{34}$$

Whenever $\Delta x > 4D^n$, the probability of reflection is always positive. In the limit of zero cell size, a positive probability of reflection requires that

$$\Delta t > \frac{4D^n}{c} \left[1 + \frac{C_v^n}{4aT^{n\,3}} \right]. \tag{35}$$

We can easily and correctly handle negative probabilities, but the cost of doing so is an increased variance [11]. For this work, we choose to remain in the "diffuse-cell" regime by requiring that $\Delta x > 4D^n$.

4. A residual method for nonlinear thermal radiative diffusion

As described in Section 2, we can derive a residual method from the conventional Monte Carlo method defined in Section 3. Furthermore, the Monte Carlo part of the solution algorithm remains unchanged between the conventional and residual schemes. Eqs. (24a), (24b) are solved with discrete Monte Carlo particle transport to get estimates for the partial fluxes, $\tilde{f}^{(+)\,n+1}$ and $\tilde{f}^{(-)\,n+1}$, on each face. Recalling Eq. (4),

$$r = b - \hat{A}\tilde{x}. \tag{4}$$

we can use these partial flux estimates to calculate residuals on each face via an application of Eq. (4) to Eqs. (24a), (24b)

$$r^{(+)}_{i+1/2} = P_{\mathrm{T},i}\tilde{f}^{(+)\,n+1}_{i-1/2} + P_{\mathrm{R},i}\tilde{f}^{(-)\,n+1}_{i+1/2} + q^{(+)\,n}_{i+1/2} - \tilde{f}^{(+)\,n+1}_{i+1/2}, \tag{36a}$$

$$r^{(-)}_{i-1/2} = P_{\mathrm{R},i}\tilde{f}^{(+)\,n+1}_{i-1/2} + P_{\mathrm{T},i}\tilde{f}^{(-)\,n+1}_{i+1/2} + q^{(-)\,n}_{i-1/2} - \tilde{f}^{(-)\,n+1}_{i-1/2}. \tag{36b}$$

In a 1D problem, the boundary conditions are sampled exactly; thus, the residuals on all problem boundaries are identically zero,

$$r^{(+)}_{1/2} = r^{(-)}_{I+1/2} = 0. \tag{37}$$

Now, if Eqs. (36a), (36b) are subtracted from Eqs. (24a), (24b), we obtain the analog of Eq. (3) for corrections to the partial flux estimates

$$\delta\tilde{f}^{(+)\,n+1}_{i+1/2} = P_{\mathrm{T},i}\delta\tilde{f}^{(+)\,n+1}_{i-1/2} + P_{\mathrm{R},i}\delta\tilde{f}^{(-)\,n+1}_{i+1/2} + r^{(+)}_{i+1/2}, \tag{38a}$$

$$\delta\tilde{f}^{(-)\,n+1}_{i-1/2} = P_{\mathrm{R},i}\delta\tilde{f}^{(+)\,n+1}_{i-1/2} + P_{\mathrm{T},i}\delta\tilde{f}^{(-)\,n+1}_{i+1/2} + r^{(-)}_{i-1/2}. \tag{38b}$$

Here, $\delta\tilde{f}$ is the correction to the original estimates of the partial fluxes. We can solve for these corrections in the same manner that we solved Eqs. (24a), (24b). In fact, Eqs. (38a), (38b) are identical to Eqs. (24a), (24b) with the exception that the residuals replace the source terms, and we solve for partial flux corrections instead of the partial fluxes.

Eqs. (24a), (24b) and (38a), (38b) imply a residual scheme that is run in stages. Each stage, k, provides a correction to the initial estimates of the partial fluxes

$$\tilde{f}^{k+1} = \tilde{f}^{k} + \delta\tilde{f}^{k}. \tag{39}$$

We now can postulate a residual scheme (REqDDMC) based on these equations:

(1) Do a standard, Section 3, discrete iteration to solve for $\tilde{f}^0$ in the first timestep.
(2) Use Eqs. (36a), (36b) to solve the residuals.
(3) Solve Eqs. (38a), (38b) to get corrections to the partial fluxes, which is equivalent to the EqDDMC transport scheme [Eqs. (24a), (24b)] with a modified source.
(4) Use Eq. (39) to update the flux estimates and test for convergence on the partial fluxes,

$$\frac{\left|\tilde{f}^{k+1}_{i+1/2} - \tilde{f}^{k}_{i+1/2}\right|}{\tilde{f}^{k+1}_{i+1/2}} < \epsilon \quad \text{for } i = 0, \ldots, I. \tag{40}$$

(5) If not converged return to step (2).
(6) If converged, calculate T^{n+1} using Eq. (23) and update material properties. Use $\tilde{f}^{(-)\,n+1}$ and $\tilde{f}^{(+)\,n+1}$ to calculate the first stage residual for the t^{n+2} timestep. Continue running timesteps until end time is reached.

As we shall show, the great advantage of the residual scheme is that work is only performed where the residuals are high. There is no source in Eqs. (38a), (38b) where $r \sim 0$. Additionally, using $\tilde{f}^{(-)\,n+1}$ and $\tilde{f}^{(+)\,n+1}$ to estimate the residuals in the following timestep means that the costly EqDDMC transport step need only be run in the first stage of the first problem cycle.

5. Linear error analysis

Quantifying the gain of the REqDDMC method over the conventional EqDDMC method is not straightforward because each method converges the statistical error, σ, at a different rate. Also, the amount of work in each method varies differently with the number of particles and cells. The usual generic metric for comparing linear Monte Carlo methods, the Figure-of-Merit [12] (FOM) will not suffice here, except for snapshot comparisons. The FOM,

$$\mathrm{FOM} = \frac{1}{\sigma^2 T_{\mathrm{CPU}}}, \tag{41}$$

where σ is the standard deviation and T_{CPU} is the computer (CPU) time, assumes that the methods converge the statistical error according to the Central Limit Theorem,

$$\sigma = \frac{c}{\sqrt{N}}, \tag{42}$$

where N is the number of particles. The FOM is used typically to measure the effectiveness of variance reduction techniques in minimizing the constant c in Eq. (42).

To compare REqDDMC with EqDDMC, we will compare their errors as a function of computer time for the linear case of a single timestep. For each method, we have analytical expressions for both the convergence rate and the computer time. The coefficients of these expressions are unknown, but we can determine them heuristically. Once the coefficients are determined, we can parameterize in numbers of particles, particles per stage, stages, and residual convergence criterion to functionally determine the residual gain as a function of computer time,

$$\mathrm{Gain(t)} = \frac{\epsilon_{\mathrm{EqDDMC}}(t)}{\epsilon_{\mathrm{REqDDMC}}(t)}, \tag{43}$$

where the error, ϵ, for each method is evaluated for the same computer time. After obtaining these coefficients, we will plot the residual gain as a function of computer time for two cases, a 100-cell problem and a 10^6-cell problem.

For this linear error analysis, we consider one timestep of a steady-state, homogeneous, infinite medium problem, which has an analytic solution equal to its initial condition. In our simulation, the nominal cell size is 0.005 cm, the nominal time step is 10^{-4} shakes, the density is 3.0 g cc^{-1}, the absorption coefficient is a constant 100.0 cm^2 g^{-1}, and the specific heat is 0.1 jerks g^{-1} keV^{-1}. The problem has reflecting boundary conditions and an initial temperature of 1 keV. We obtain the general coefficients for the gain function using 100 cells.

The absolute errors for one cycle of both EqDDMC and REqDDMC are shown in Fig. 2. The errors are measured as the l_∞ norm of the temperature with respect to the analytic solution of unity. The REqDDMC

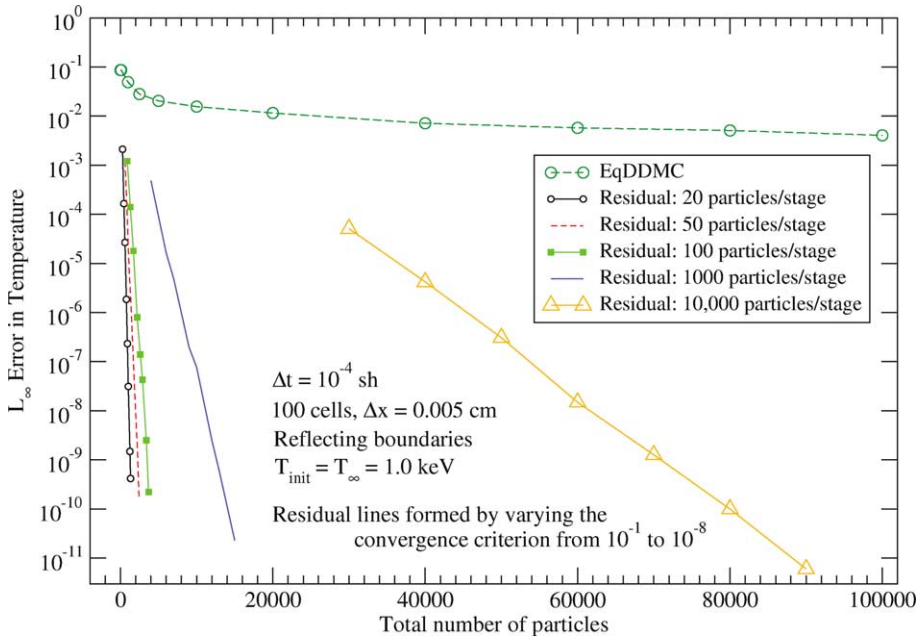

Fig. 2. The exponential error convergence of the REqDDMC method is shown compared to the inverse-root-N convergence of the stock EqDDMC for the linear case of one timestep.

curve for each selection of particles per stage was formed by running REqDDMC eight times, varying the convergence criterion from 10^{-1} to 10^{-8} in powers of ten.

The residual gains, which approach 10 orders of magnitude, are shown in Fig. 3. From the 100-cell residual gain curves, we see the effects of the competing functions of the overall exponential convergence and the first stage's inefficient $1/\sqrt{N}$ convergence. The 10^4 particles in just one stage takes longer than the 20 particles/stage case for any convergence criterion. The deterministic overhead work for each stage is insignificant.

From the 10^6-cell residual gain curves, we see that the deterministic overhead work for each stage dominates the cost of transporting particles. In fact, with too few particles per stage, more stages are required, and the overall gain diminishes. Note that the dependence on number of particles per stage is reversed: more particles per stage is better than fewer for the larger number of cells.

Our linear error analysis shows rapidly increasing residual gains of up to 10 orders of magnitude. This residual gain represents what can be achieved in a linear Monte Carlo calculation or in one linearized timestep of a nonlinear Monte Carlo calculation. The actual gain in a REqDDMC calculation depends on the linear, residual gain each timestep and the propagation of deterministic and stochastic errors over multiple timesteps. The balance between transport work, deterministic overhead work per stage, stability, and adequate sampling will become more important when we extend the REqDDMC method to higher spatial dimensions.

6. Bias

Nonlinear mathematical operations on stochastic quantities produce biased statistical estimates [13]. Biased estimates do not possess the correct expected value. The estimates of the probabilities in the Eq-

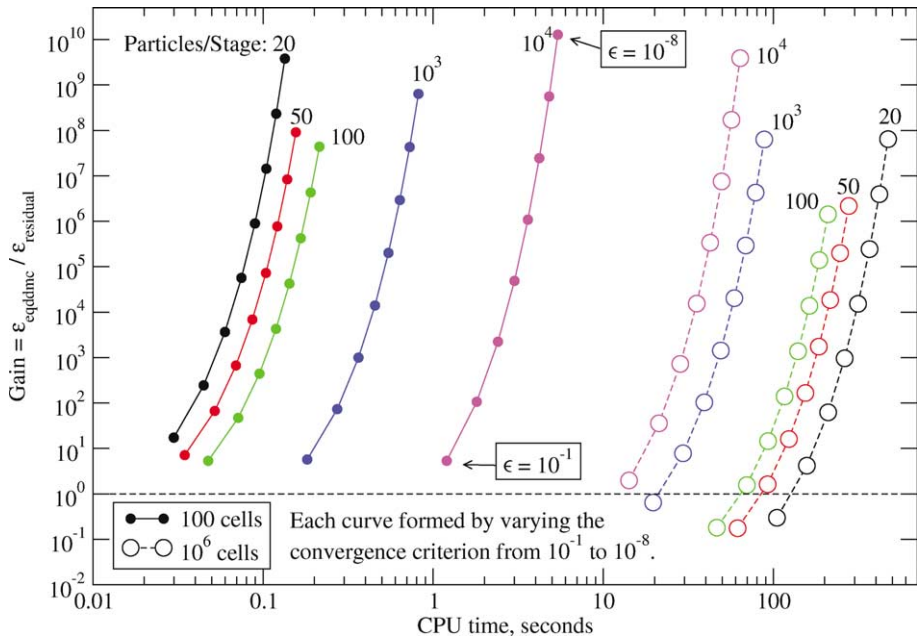

Fig. 3. Residual gain as a function of CPU time for 100 cells and 10^6 cells and varying numbers of particles per stage.

DDMC method are biased because they depend nonlinearly on stochastic estimates of the temperature. The reflection, transmission, and absorption probabilities always sum to unity, but the bias affects the individual values of the probabilities. The overall EqDDMC bias is negative because a symmetric temperature distribution produces distributions of transmission and reflection probabilities that are skewed toward smaller values. The bias is positive for the absorption probability but the EqDDMC method never explicitly uses the absorption probability. The negative bias accumulates each timestep because statistically lower partial flux estimates result in lower temperature estimates, which result in lower transmission and reflection probabilities, which result in statistically lower partial flux estimates. Thus, statistically smaller values are preferentially retained. The EqDDMC bias goes to zero as the number of particles increases, but unfortunately, it goes to zero more slowly than the standard deviation of the temperature.

The same bias also exists in each stage of the REqDDMC method, but the staging and absence of statistical noise virtually eliminate any bias in the REqDDMC results.

We will demonstrate the bias in the EqDDMC method and the near absence of bias in the REqDDMC method with a simple problem. The problem is the steady-state, homogeneous, infinite medium that was used in Section 5. The system begins with a temperature of unity that should stay constant over time. As shown in Fig. 4, the EqDDMC temperature decreases over time, whereas the REqDDMC temperature holds steady at unity. Incidentally, the residual calculation took 25 seconds, and the EqDDMC calculations took about 0.23 seconds per particle on two processors, which translates to about 128 CPU-hours for the 10^6-particle run.

The magnitude of the EqDDMC bias depends on the physical, computational, and runtime parameters. For instance, if the demonstration problem above had an opacity with a $1/T^3$ dependence, the bias would be present but somewhat mitigated. If, instead, the demonstration problem had a specific heat proportional to $(T^n)^3$, then the equations become linear, and there would be no bias because the $(T^n)^{-3}$ in the effective opacity $\tilde{\sigma}_t^n$, Eq. (13), would cancel out.

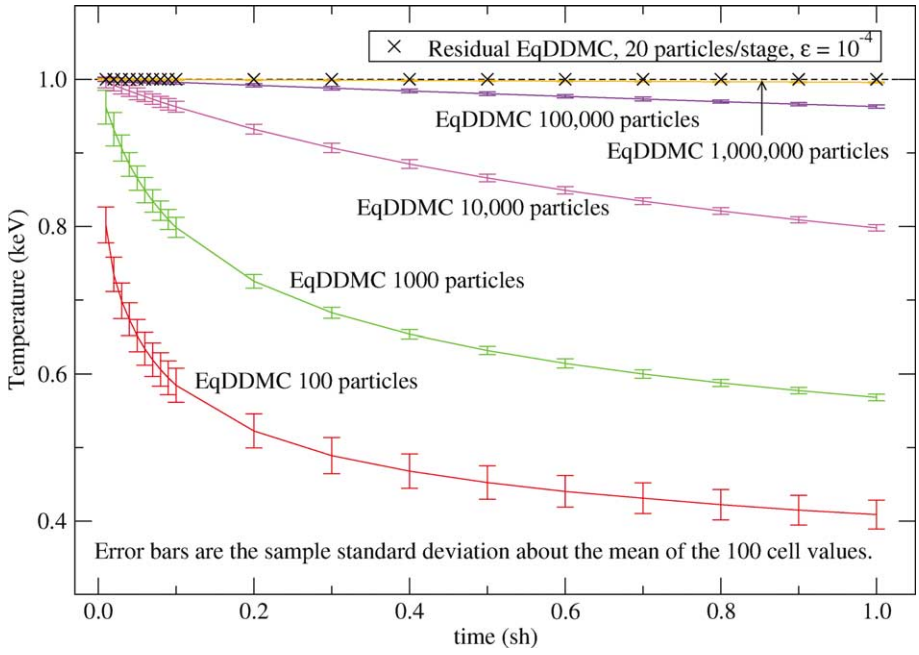

Fig. 4. A steady-state, homogeneous, infinite medium problem demonstrates the bias in the EqDDMC method and the essentially nonexistent bias in the REqDDMC method.

7. Results

We have tested the residual EqDDMC method on the nonlinear Marshak Wave problem [14]. This 1D equilibrium problem has an isotropic intensity impinging on a cold, pure-absorbing slab. The problem is solved analytically by assuming that the heat capacity is constant, the opacity follows a power law, and the material energy density dominates the radiative energy density, i.e., $4aT^3 \ll C_v$. The Marshak wave problems solved below have the following parameters:

$$C_v = 0.1 \quad (\text{jerks g}^{-1}\ \text{keV}^{-1}),$$

$$\sigma = \frac{100}{T^3} \quad (\text{cm}^{-1}),$$

$$f_{\text{left}}^{(+)} = \frac{ac}{4} \quad (\text{jerks cm}^{-2}\ \text{sh}^{-1}).$$

The timestep begins at 1.25×10^{-10} sh and increases 10% each timestep until it reaches the specified maximum.

Fig. 5 compares results of the residual Monte Carlo method with the analytic solution to the Marshak problem at 1, 5, and 10 shakes. The residual method results agree well with the analytic solution. Errors between the analytic result and residual result arise because the problem is not spatially resolved enough to account for gradients in the cell-centered diffusion coefficients. The spatial resolution is limited by the restriction on positive probabilities that requires $\Delta x > 4D^n$.

Because of the improved efficiency of the residual method, both in terms of noise reduction and runtime, we can afford to iterate on the nonlinearities in the material coefficients. We use the Picard iteration scheme

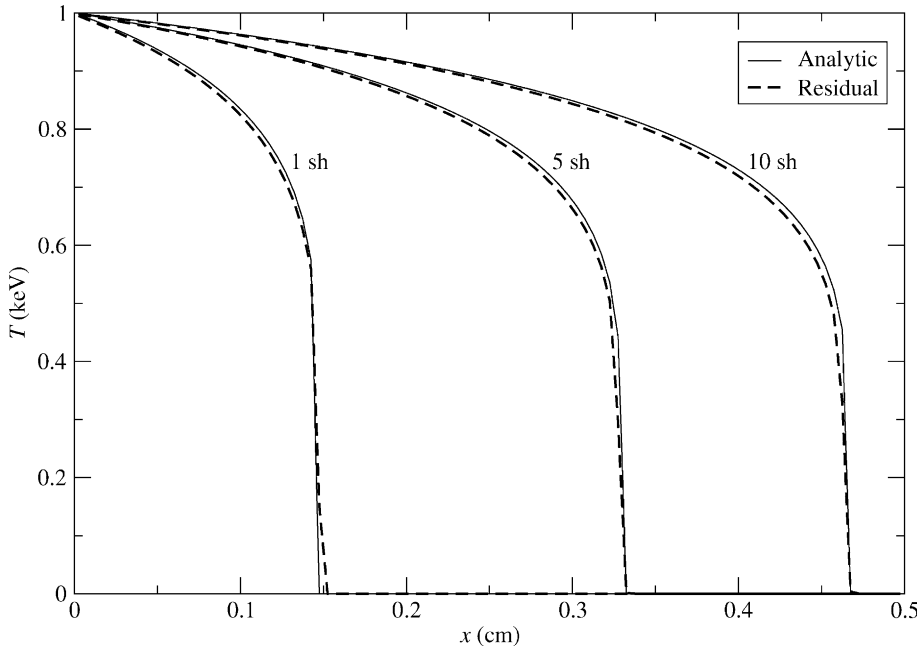

Fig. 5. The REqDDMC solution to the Marshak wave problem at 1.0, 5.0, and 10.0 shakes.

[15] to estimate the material properties at $T^{n+1/2}$ as we convergence on T^{n+1}. The Picard iterations converge on the following condition:

$$\frac{\left|T_i^{n+1,l+1} - T_i^{n+1,l}\right|}{T_i^{n+1,l+1}} < \epsilon \quad \text{for } i = 0, \ldots, I, \tag{44}$$

where l is the Picard iteration index. The solutions in Fig. 5 were iterated to convergence using this scheme and a convergence criterion of 0.001. Iterating on material nonlinearities is a fully consistent treatment of the nonlinearities and, in principle, allows larger timesteps.

As seen in Section 4 there are several degrees of freedom in the residual method. Primary among these are the number of particles per stage and timestep. When the timestep is large, the fluxes from the previous timestep may not be good estimates of the new fluxes, and more particles are required to get a good estimate of the residual in the first stage of a timestep. Fig. 6 shows a convergence study of timestep and number of particles per stage for the Marshak wave problem. This analysis for relatively small meshes shows that running fewer particles per stage with a smaller timestep can be as efficient as running larger timesteps that require more particles per stage. However, running too few particles per stage will result in an under-sampling of phase space and possibly instabilities and bias, the latter of which is demonstrated in Section 6 and in Fig. 7. As shown in Fig. 7, to get good answers at 10 shakes, when there is a lot of energy dispersed over most of the system, the number of particles per stage must be increased to cover the increased amount of discrete phase space.

Fig. 7 also shows the advantages of the Picard iterations over using one iteration with material coefficients evaluated at temperatures extrapolated to mid-timestep. This extrapolation is feasible because of the negligible noise in the residual method, but it is not as accurate as Picard iterations. Making this one-pass approach more accurate may involve simply running smaller timesteps. Then the overall efficiency of the residual method is a balance between the cost of Picard iterations and running additional cycles at a smaller timestep.

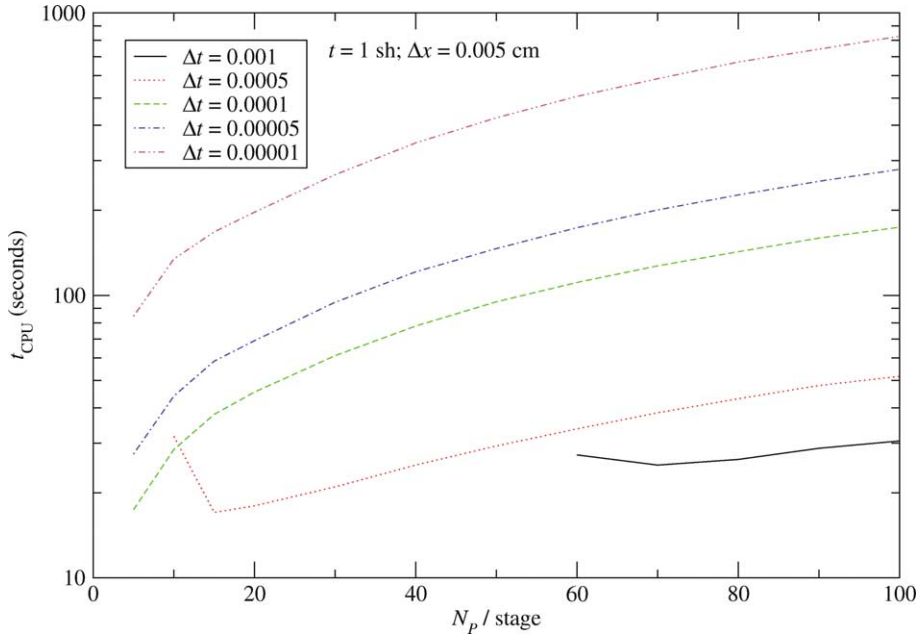

Fig. 6. The computer time versus number of particles per stage for five different maximum timesteps of the REqDDMC method. The Marshak problem was run to a maximum time of 1.0 shakes. The timesteps are in shakes.

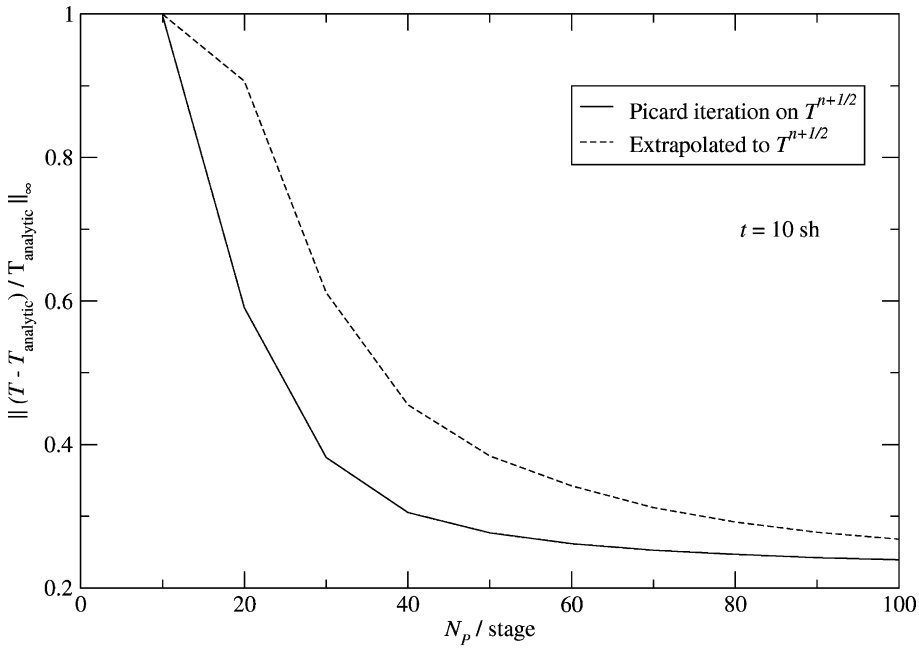

Fig. 7. The l_∞ errors between the analytic and REqDDMC solutions to the Marshak wave problem at 10 shakes for Picard iterations and extrapolated temperatures. Both options were run with a maximum timestep of 0.0002 shakes.

As stated in Section 5, the traditional figure of merit (FOM) [12] is insufficient for a general error analysis of the residual method. However, we can use the FOM to evaluate a snapshot comparison between the EqDDMC and the REqDDMC methods at a particular time. That is to say, that, if we were to imagine that the residual method converged at a rate of $c/\sqrt{N}$, this evaluation gives us an effective gain of REqDDMC over EqDDMC in the constant c. Fig. 8 shows the sample standard deviation of the temperature taken from 30 replicated calculations (each with a different random number seed) of the Marshak wave problem at 1.0 shakes. The sample error is more than an order of magnitude less then the EqDDMC sample error in the wavefront, and it is two to three orders of magnitude smaller behind the wave. For each of the residual and analog methods, Table 1 shows runtimes, l_∞ norms of the standard deviation of the temperature, and effective FOMs. Clearly, when error and runtime are factored together, the benefits from the residual method are apparent.

Finally, given the dramatically improved efficiency exhibited by the residual method over conventional Monte Carlo, we want to compare the REqDDMC method with a deterministic calculation. Table 2 shows runtimes for the Marshak wave problem run to 1.0, 5.0, and 10.0 shakes for the residual and deterministic diffusion solvers. The maximum timesteps were chosen to yield the smallest runtimes, without adversely affecting accuracy, for each method. Both methods iterate on the nonlinearities in the material coefficients using the Picard iteration scheme with the convergence criterion given in Eq. (44). The results are essentially a comparison between the residual method and Gaussian elimination on a tri-diagonal mesh. Gaussian elimination on a tri-diagonal mesh is a $O(N)$ algorithm; thus, it is very efficient. Fig. 9 shows the results of both the residual and deterministic calculations. The residual method results compare very favorably to the deterministic results because there is negligible noise. Although the deterministic calculations are considerably faster in 1D, the residual method could be a legitimate competitor to deterministic diffusion in multi-dimensional calculations that require iterative linear solvers. The multi-dimensional results are the focal point of a future study.

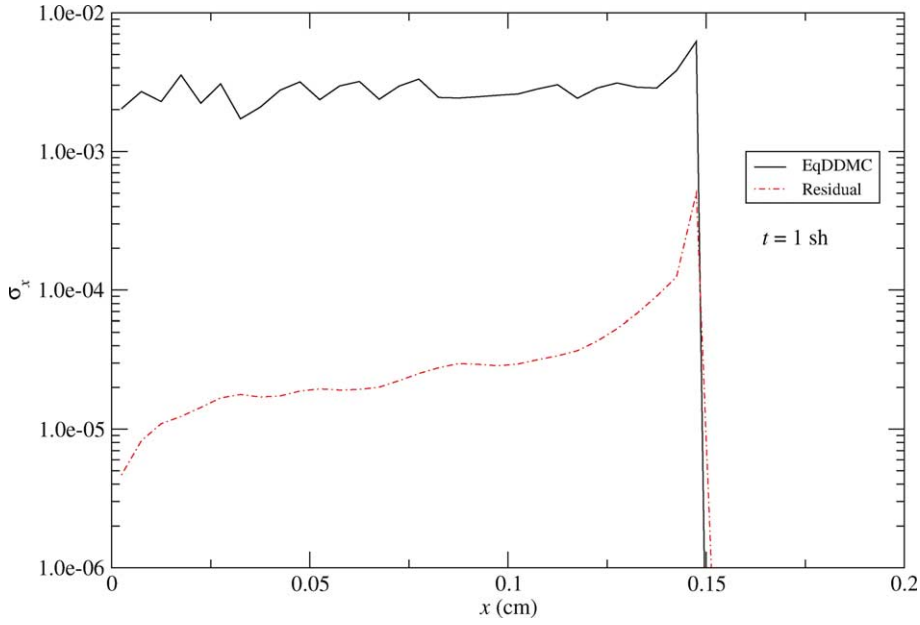

Fig. 8. Plot of the standard deviation of the temperature in each cell for the REqDDMC and EqDDMC methods on the Marshak wave problem. The Marshak wave was run to 1.0 shakes. Both the EqDDMC and REqDDMC problems were run with a maximum timestep of 0.0001 shakes.

Table 1
Runtimes, l_∞ norms of the standard deviation of the temperature, and FOMs for the REqDDMC and EqDDMC methods

Method	CPU time (s)	$\|\sigma_x\|_\infty$	FOM
EqDDMC	4541.9	6.2×10^{-3}	5.7
REqDDMC	17.0	5.0×10^{-4}	2.4×10^{5}

The sample standard deviations are generated from 30 replicate calculations of the Marshak wave problem to 1.0 shake. Calculations were performed on a LINUX-based Pentium IV processor.

Table 2
Runtimes of the REqDDMC and deterministic solvers on the Marshak wave problem run to 1.0, 5.0, and 10.0 shakes

Method	Δt (shakes)	CPU time (s)		
		$t = 1$ sh	$t = 5$ sh	$t = 10$ sh
REqDDMC	0.0002	19.44 (5)	269.09 (35)	724.11 (50)
Deterministic	0.001	0.75	2.32	4.16

Both methods utilized identical meshes of 100 cells with $\Delta x = 0.005$ cm. The numbers in parenthesis for the REqDDMC runtimes are the number of particles per stage.

Each residual entry in Table 2 is a separate run with a constant number of particles per stage. As we noted earlier, the Marshak Wave problem requires more particles per stage at late times. Therefore, the overall efficiency of the REqDDMC method could be improved by making the number of particles per stage a variable that starts out small and increases with t. We will analyze methods for performing incremental particle populations in a future study.

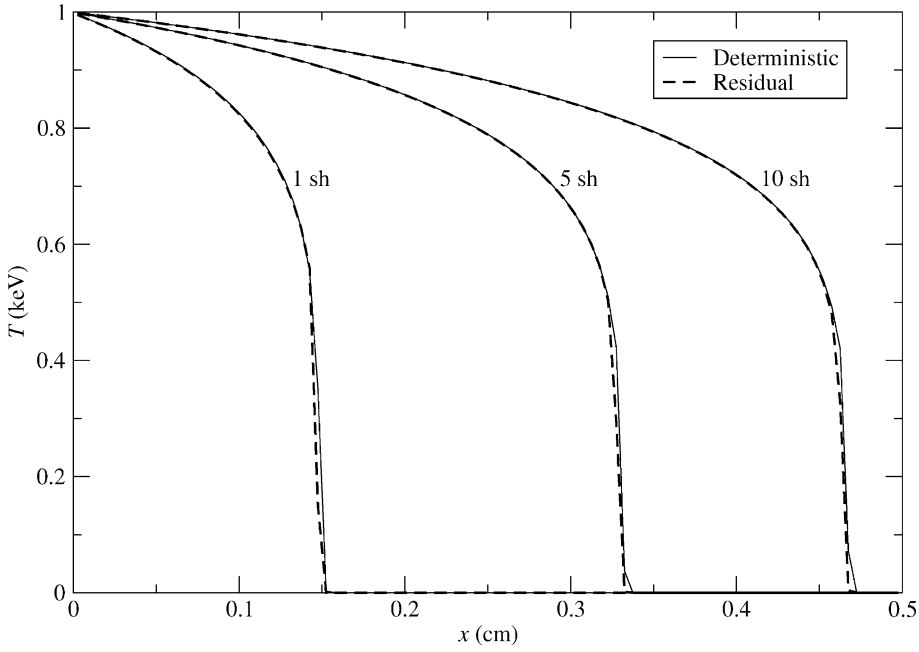

Fig. 9. REqDDMC and deterministic solutions to the Marshak wave problem at 1.0, 5.0, and 10.0 shakes.

8. Conclusions

We have demonstrated the application of the residual Monte Carlo method to discrete, nonlinear, thermal radiative equilibrium diffusion in one dimension. The residual method shows efficiency gains of several orders of magnitude over the discrete conventional Monte Carlo method on which it is based. The residual Monte Carlo method is successful here because it is applied to a simple discrete system that approximates the continuum diffusion equations instead of to the continuum equations themselves. The observed efficiency gains warrant further research into more sophisticated discrete systems that better approximate the continuum equations.

The residual Monte Carlo method reduces statistical errors to such small levels that it is possible to solve a strongly nonlinear, discrete system with a fully implicit treatment of the nonlinearities. Although the residual Monte Carlo method is not as efficient as a deterministic method in 1D, the results presented here are promising enough to warrant study in multi-dimensional problems. With increasing dimension, Monte Carlo methods generally become more competitive with deterministic methods. The REqDDMC method in multi-dimensions will be the topic of future investigations.

Acknowledgements

The authors thank Dr. Gordon Olson, Transport Methods Group, Los Alamos National Laboratory, for aid in completing this study. This work was performed under US Government contract W-7405-ENG-36 for Los Alamos National Laboratory.

References

[1] J. Halton, Sequential Monte Carlo techniques for the solution of linear systems, Journal of Scientific Computing 9 (2) (1994) 213–257.
[2] T. Booth, Exponential convergence on a continuous Monte Carlo transport problem, Nuclear Science and Engineering 127 (3) (1997) 338–345.
[3] J. Favorite, H. Lichtenstein, Exponential Monte Carlo convergence of a three-dimensional discrete ordinates solution, Transactions of the American Nuclear Society 81 (1999) 147–148.
[4] H. Lichtenstein, Exponential convergence rates for reduced-source Monte Carlo transport in $[x, \mu]$ geometry, Nuclear Science and Engineering 133 (1999) 11.
[5] T. Booth, Adaptive importance sampling with a rapidly varying importance function, Nuclear Science and Engineering 136 (3) (2000) 399–408.
[6] T. Booth, An approximate Monte Carlo adaptive importance sampling method, Nuclear Science and Engineering 138 (1) (2001) 96–103.
[7] J.E. Morel, T.A. Wareing, K. Smith, A linear-discontinuous spatial differencing scheme for S_N radiative transfer calculations, Journal of Computational Physics 128 (1996) 445–462.
[8] C.E. Knapp, C.W. Cranfill, Comparison of numeric to analytic solutions for a class of nonlinear diffusion equations, Laboratory Report LA-UR-92-3538, Los Alamos National Laboratory (27 October 1992).
[9] R. Szilard, G. Pomraning, Numerical transport and diffusion methods in radiative transfer, Nuclear Science and Engineering 112 (3) (1992) 256–269.
[10] T.M. Evans, T.J. Urbatsch, H. Lichtenstein, 1-D equilibrium discrete diffusion Monte Carlo, in: Proceedings of the MC2000 – International Conference, Lisbon, Portugal, 2000.
[11] T. Urbatsch, J. Morel, J. Gulick, Monte Carlo solution of spatially-discrete transport equation, Part I: Transport, in: J. Aragonés (Ed.), Mathematics and Computation, Reactor Physics and Environmental Analysis in Nuclear Applications, vol. 1, Institute of Nuclear Fusion (DENIM), Senda Editorial, S.A., Madrid, Spain, 1999, pp. 251–261.
[12] J.F. Briesmeister, MCNP-A General Monte Carlo N-Particle Transport Code, No. LA-12625-M, Version 4B, Los Alamos National Laboratory, 1997.
[13] M.H. Kalos, P.A. Whitlock, Monte Carlo Methods, vol. I, Wiley, New York, 1986.

[14] A. Petschek, R. Williamson, Penetration of radiation with constant driving temperature, Technical Report LAMS-2421, Los Alamos Scientific Laboratory, May 1960.
[15] C.R. Doering, J.D. Gibbon, Applied Analysis of the Navier–Stokes Equations, first ed., Cambridge Texts in Applied Mathematics, Cambridge University Press, Cambridge, MA, 1995.