Journal of Computational Physics 228 (2009) 8766–8778

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

A combined Event-Driven/Time-Driven molecular dynamics algorithm for the simulation of shock waves in rarefied gases

Paolo Valentini *, Thomas E. Schwartzentruber

Department of Aerospace Engineering and Mechanics, University of Minnesota, Minneapolis, MN 55455, USA

ARTICLE INFO

Article history:
Received 23 April 2009
Received in revised form 14 August 2009
Accepted 26 August 2009
Available online 2 September 2009

PACS:
47.11.Mn
47.40.Ki
47.45.−n
47.61.Cb

Keywords:
Molecular Dynamics
Event-Driven MD
Shock Waves
Non-continuum effects

ABSTRACT

A novel combined Event-Driven/Time-Driven (ED/TD) algorithm to speed-up the Molecular Dynamics simulation of rarefied gases using realistic spherically symmetric soft potentials is presented. Due to the low density regime, the proposed method correctly identifies the time that must elapse before the next interaction occurs, similarly to Event-Driven Molecular Dynamics. However, each interaction is treated using Time-Driven Molecular Dynamics, thereby integrating Newton's Second Law using the sufficiently small time step needed to correctly resolve the atomic motion. Although infrequent, many-body interactions are also accounted for with a small approximation. The combined ED/TD method is shown to correctly reproduce translational relaxation in argon, described using the Lennard-Jones potential. For densities between $\rho = 10^{-4}$ kg/m^3 and $\rho = 10^{-1}$ kg/m^3, comparisons with kinetic theory, Direct Simulation Monte Carlo, and pure Time-Driven Molecular Dynamics demonstrate that the ED/TD algorithm correctly reproduces the proper collision rates and the evolution toward thermal equilibrium. Finally, the combined ED/TD algorithm is applied to the simulation of a Mach 9 shock wave in rarefied argon. Density and temperature profiles as well as molecular velocity distributions accurately match DSMC results, and the shock thickness is within the experimental uncertainty. For the problems considered, the ED/TD algorithm ranged from several hundred to several thousand times faster than conventional Time-Driven MD. Moreover, the force calculation to integrate the molecular trajectories is found to contribute a negligible amount to the overall ED/TD simulation time. Therefore, this method could pave the way for the application of much more refined and expensive interatomic potentials, either classical or first-principles, to Molecular Dynamics simulations of shock waves in rarefied gases, involving vibrational nonequilibrium and chemical reactivity.

1. Introduction

The highly nonequilibrium structure of a shock wave in a gas has been thoroughly investigated for decades. Studies have been devoted to determine density and temperature profiles, shock thickness, distributions of molecular velocities, and their dependence on the free stream conditions.

Attempts to solve the Boltzmann equation produced the first theoretical results that significantly improved the description based on the Navier–Stokes model, which failed to describe shocks for Mach numbers (M_1) greater than 2. The bimodal approximation of Mott-Smith [1] is one of the most successful approaches, as it provides very good predictions of the shock

* Corresponding author. Tel.: +1 612 626 7792; fax: +1 612 626 1558.
E-mail addresses: vale@aem.umn.edu (P. Valentini), schwartz@aem.umn.edu (T.E. Schwartzentruber).

0021-9991/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2009.08.026

thickness, despite determining symmetric profiles unlike those measured experimentally, very asymmetric for strong shock waves [2]. The reader is referred to Cercignani et al. [3] for a more complete survey of theoretical methods.

The need for validation of analytical solutions, obtained based on some assumptions and approximations, prompted a large number of experimental studies. Some of these works include Robben and Talbot [4], Schmidt [2], Alsmeyer [5], Muntz and Harnett [6], Holtz and Muntz [7], and Pham-Van-Diep et al. [8]. In all experiments, a beam of electrons is shone on the shock to measure its density profile and, in some cases, velocity distributions (using the electron beam fluorescence technique). Therefore, because electrons have to travel through a chamber at a fairly low pressure, the free stream density is usually set to very low values, i.e., $\sim 10^{-4}$ kg/m^3.

Introduced in 1963 for a hard sphere (HS) gas [9], the DSMC method of Bird [10] represents a particularly valuable and efficient tool to investigate the nonequilibrium structure of the normal shock wave. The method is valid for dilute gases, and full-scale simulations of normal shock waves matching the experimental conditions are very easily run on a modern computer workstation (e.g., Ref. [11]). Since its birth, the DSMC method has been equipped with collision models of increasing sophistication and realism, in order to reproduce real gas viscosity laws, to model rotational and vibrational nonequilibrium, and more recently to account for chemical reactivity. Because experiments can provide only limited information on the shock structure, i.e., mainly density profiles, experimental data are not sufficient to construct or inform DSMC collision models. For example, it is known [10] that a comparison with experimental density profiles alone is not enough to assess the validity of a particular DSMC collision scheme. This is demonstrated computationally by Valentini and Schwartzentruber [12] with full MD simulations of normal shock waves in argon, using the realistic Lennard-Jones potential. Temperature data, being a higher order moment of the molecular velocity distribution functions, are found to be more relevant to judge the applicability of the Variable Hard Sphere (VHS) model for a monatomic gas. Indeed, distribution functions, as solutions of Boltzmann's equation, are the most susceptible to the energetic details of atomic interactions. Since such details are either not available experimentally or affected by a large uncertainty, most of DSMC models have been designed to reproduce the empirical laws used in continuum simulations, which assume near-equilibrium conditions, with the rationale that they will be transferable enough to also describe well extreme nonequilibrium.

Similar to DSMC, Molecular Dynamics is a particle method designed to study dynamical properties of many-body systems. The solution of the classical equations of motion for a set of hard spheres was first accomplished in the late 1950s by Alder and Wainwright [13,14]. Atoms were modeled as billiard balls, moving at constant velocity between perfectly elastic collisions. With the recent developments in supercomputing, MD has rapidly gained large popularity to describe solids and liquids at the atomic level, both for equilibrium and nonequilibrium systems. The progress in computer hardware has also been accompanied by the development of highly sophisticated and transferable interatomic potentials, which represent the only modeling effort required for MD.

An accurate inter-particle potential is essential to describe the energy transfer among the various molecular degrees of freedom, as stated by Tokumasu and Matsumoto [15], who use MD to statistically characterize binary collisions between diatomic nitrogen molecules to compile a database of collision cross sections for use in a DSMC code. While this approach is shown to well-reproduce the rotational relaxation through low and high Mach number shock waves, it is not so appealing as the simulation of the full shock wave using MD directly. The construction of a database can become a daunting task as the sampling space becomes of higher order, and approximations become necessary. Hence, simulating the full shock wave with MD equipped with sophisticated potentials may provide all the necessary information to build and validate a DSMC collision model, to the finest details unavailable experimentally.

Of course, MD solutions using such potential models also require experimental validation. Although it is shown in Ref. [12] that the shock thickness results, obtained at a free stream density $\rho_1 = 1$ kg/m^3, compare very well with the data of Alsmeyer [5] at $\rho_1 \sim 10^{-4}$ kg/m^3, thus indicating that the translational nonequilibrium behavior of argon is essentially the same for $\rho_1 \lesssim 1$ kg/m^3, it has not been established if that is true for rotational or vibrational nonequilibrium, and their coupling. Thus, it would be desirable to compare the numerical results to experiments, *under the same low density conditions*.

However, as known from kinetic theory, the mean collision time τ scales as $\frac{1}{\rho\sqrt{T}}$, where ρ and T are the gas density and temperature, respectively. Therefore, at rarefied conditions $(\rho \sim 10^{-4}\ \text{kg/m}^3), \tau$ becomes of the order of microseconds. And because the time step required in MD simulations to integrate the equations of atomic motion must be of the order of 10^{-15} s to adequately resolve the interactions, MD becomes extremely inefficient, and prohibitively expensive at very low densities, where DSMC is the preferred choice.

For these reasons, relatively few MD studies have investigated normal shock waves. Most have been restricted to condensed systems, such as liquid argon [16], or very dense argon and nitrogen [17]. At dilute conditions, the HS model is often used [19,18], being cheaper from the computational viewpoint, but less accurate physically. More recently, large-scale MD simulations of normal shock waves in dilute argon have been performed using the realistic Lennard-Jones potential, but at densities of $\mathcal{O}(1)$ kg/m^3 [12].

If atoms or molecules essentially evolve independently, except for discrete asynchronous instantaneous interactions, Event-Driven Molecular Dynamics (EDMD) was soon recognized to be the best choice to solve the dynamical problem exactly, within the limits of machine accuracy. Indeed, in their seminal paper [14], Alder and Wainwright essentially describe the first Event-Driven approach to Molecular Dynamics simulation. The algorithm computes the time at which the first binary collision occurs. All molecules are then moved at constant velocity to that time, when only one pair undergoes a collision, and thus new velocities must be determined depending on the nature of the interaction (e.g., elastic or inelastic).

For a rarefied gas of point-like atoms interacting through continuous at-a-distance forces, many of the assumptions for EDMD are verified, as (i) atoms or molecules move independently of each other for an $\mathcal{O}(\tau)$ time (this is true so long as the potential has a cut-off), (ii) interactions are essentially binary and asynchronous, and (iii) the duration of the interactions is usually small compared to τ. Clearly, this is not true for a condensed phase, such as a liquid, where molecules are constantly interacting with their very close neighbors.

The ED approach to the dynamical simulation of many-body systems composed of rigid objects is much more efficient and accurate if compared to time-stepping methods, in which Newton's equations are solved for a series of small time slices. The requirement of a small time step derives from the singular nature of the interaction, which exhibits a discontinuity at some inter-particle separation. Thus, depending on the density, a very large number of small steps must be taken before a collision is detected.

Alder and Wainwright already recognize that a further improvement of the algorithmic efficiency is obtained if a list of future collision times $\theta^{(\alpha)}$ ($\alpha = 1, 2, \ldots, M \in \mathbb{N}$ and $\theta^{(\alpha-1)} < \theta^{(\alpha)}$) is retained. That is because the difference between the shortest $\theta^{(1)}$ just elapsed and the next larger $\theta^{(2)}$ is the time required for the next collision. This is true if and only if none of the particles involved in the first collision taking place at $\theta^{(1)}$ is immediately involved in a successive collision. Starting from their basic scheme, numerous more sophisticated and efficient algorithms have been developed, particularly focusing on the following four aspects:

1. Improvement of the efficiency to determine, for each particle, the time of the next collision. This is accomplished by dividing the systems into cells or constructing *so-called* neighbor lists for each molecule [20]. With such strategies, the time taken for the calculation of the first collision event for a given particle can be made independent of N, the total number of atoms in the system [21];
2. Efficient management of a *priority queue* to process, in the correct order, the sequence of collision events among all particles. The article by Paul [22] offers a comprehensive survey of priority queue implementations. While the oldest algorithms usually have an $\mathcal{O}(\log N)$ scaling (e.g., implicit heaps binary trees, leftist trees, binomial queues, etc.), more recently an approach based on bounded priority queues is proposed, with all operations requiring $\mathcal{O}(1)$ time [22];
3. Parallelization aspects. Recently, Miller and Luding [23] have presented a parallel EDMD algorithm which includes dynamic load-balancing to achieve good scalability;
4. Development of methods capable of handling hard particles of various shapes, for example very non-spherical or rod-like shapes [21].

EDMD has been shown to be more efficient and accurate than Time-Driven MD (TDMD) for the simulation of hard spheres, however the physical validity of HS potentials to model real fluids is limited. Therefore, our objective is to extend the Event-Driven methodology to more realistic spherically symmetric soft potentials. In this article, we present a combined ED/TD Molecular Dynamics algorithm designed to simulate gases under rarefied conditions ($\rho \leqslant 0.1\ \mathrm{kg/m^3}$). While we do not focus on aspects (1) through (4), our belief is that they could be adopted to further improve the efficiency of our approach, particularly (1) and (2).

The ED/TD MD algorithm is applied to study translational relaxation in argon, for densities ranging from $\rho = 10^{-4}\ \mathrm{kg/m^3}$ to $\rho = 10^{-1}\ \mathrm{kg/m^3}$, using the Lennard-Jones potential. A normal shock wave in Ar is also simulated, setting the free stream conditions to $\rho_1 = 10^{-3}\ \mathrm{kg/m^3}, T_1 = 300\ \mathrm{K}$, at a Mach number of 9. Detailed comparisons between TDMD, ED/TD MD, and DSMC are made. The paper is organized as follows. Section 2 outlines the combined ED/TD Molecular Dynamics algorithm. Section 3 contains the results of the relaxation calculations. In Section 4, the results of the normal shock wave are presented. Finally, the summary and the conclusions are contained in Section 5.

2. Numerical methods

2.1. A simple EDMD algorithm for rigid spheres

Hard particles exert forces on one another only when they collide. Between collisions, provided no external force field is present (e.g., gravity), they travel at constant velocity along straight lines. Hence, a typical EDMD algorithm just computes the times of collisions, instead of trajectories. If collisions are assumed to be perfectly elastic, the calculation is purely algebraic (particles are not allowed to deform, and no energy is transferred to internal degrees of freedom).

If these particles are rigid spheres, the calculation of the collision time for any two spheres is straight-forward. Assuming for simplicity that all spheres have the same diameter r_c, at the collision instant their centers will be separated by a distance r_c:

$$\|\mathbf{r}_i(\theta_{ij} + t_0) - \mathbf{r}_j(\theta_{ij} + t_0)\| = r_c, \tag{1}$$

where θ_{ij} is the time elapsed from t_0, and $\mathbf{r}$ is the position vector. Because between collisions sphere i moves at constant velocity, given its position at t_0, it follows that at a later time $t_0 + \theta$,

$$\mathbf{r}_i(\theta + t_0) = \mathbf{r}_i(t_0) + \theta \mathbf{v}_i(t_0), \tag{2}$$

where $\mathbf{v}$ is the velocity vector. Combining Eq. (2) with Eq. (1) leads to the well-known quadratic equation for θ_{ij} [24], whose solution is

$$\theta_{ij} = \frac{(-\mathbf{v}_{ij}\cdot\mathbf{r}_{ij}) \pm \sqrt{(\mathbf{v}_{ij}\cdot\mathbf{r}_{ij})^2 - \mathbf{v}_{ij}\cdot\mathbf{v}_{ij}(\mathbf{r}_{ij}\cdot\mathbf{r}_{ij} - r_c^2)}}{(\mathbf{v}_{ij}\cdot\mathbf{v}_{ij})}, \tag{3}$$

where $\mathbf{v}_{ij} = \mathbf{v}_i - \mathbf{v}_j, \mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$, and $(\cdot)$ denotes the standard scalar product in $\mathbb{R}^3$. Depending on the discriminant in Eq. (3), $\theta_{ij} \in \mathbb{R}$ is determined.

The calculation of each θ_{ij} for all ij colliding partners would be easily achieved by checking all $\frac{1}{2}N(N-1)$ pairs. Evidently, such an $\mathcal{O}(N^2)$ approach is not efficient. Therefore, the standard $\mathcal{O}(N)$ linked-list cell method [20] is typically employed. Finally, if $\theta = \min_{\{\forall ij\}}\{\theta_{ij}\}$, then $t_0 + \theta$ is the time of the *next* collision. At $t_0 + \theta$, the interaction is processed, i.e., the post-collision velocities are computed, and the algorithm searches for the next θ. Within machine accuracy, this procedure is exact for a HS gas.

At a given density, if N increases, θ decreases, as the likelihood that two particles are close and immediately collide becomes larger. Hence, re-calculating θ (usually an $\mathcal{O}(N)$ operation), after each single collision is processed, starts to be inefficient. However, it is exactly on this aspect that methods based on priority queues have been developed in order to avoid the calculation of the next θ after each event is processed, taking advantage of the asynchronous nature of the interactions [22].

If inter-particle forces are derived from a soft potential, such as the Lennard-Jones potential, the standard EDMD algorithm cannot be used, because (i) collisions are now ill-defined, (ii) post-collision states are not the result of simple elastic re-bouncing, and (iii) at higher densities, many particles may interact at once, thus violating the assumption of binary interactions. Moreover, while undetected overlaps do not affect the stability of a HS gas calculation (where the system energy is purely kinetic), they will produce unphysically high potential energies for a realistic pair potential, causing the simulation to crash.

2.2. A combined Event-Driven/Time-Driven algorithm for spherically symmetric soft potentials

As a gas becomes rarefied, the mean free path $\lambda \sim 1/\rho$ gets very large compared to the potential cut-off distance r_c. For instance, at $\rho \sim 10^{-4}$ kg/m^3, $\lambda \sim 1.15$ mm. For typical soft-sphere potentials not including the Coulomb interaction, r_c is usually set between 2 or 3 times the $\mathcal{O}(10^{-10})$ m molecular diameter [20]. Therefore, at these conditions, atoms do not interact until two or more of them are separated by a distance smaller than r_c. If a standard TDMD approach is used, Newton's equations are integrated for many thousands of time steps, during which atoms or molecules move along straight lines, because the resultant force acting on them is zero. This is very inefficient.

The combined Event-Driven/Time-Driven approach that we propose is designed to speed-up the MD simulation of a rarefied gas. This is achieved by correctly detecting the next interaction, like in the Event-Driven approach, thus avoiding integrating the trajectory for all atoms. At the same time, each interaction is correctly resolved using the necessarily small $\mathcal{O}(10^{-15})$ s time step of a Time-Driven simulation. Furthermore, many-body interactions, although rare, are detected and simulated, with a small approximation.

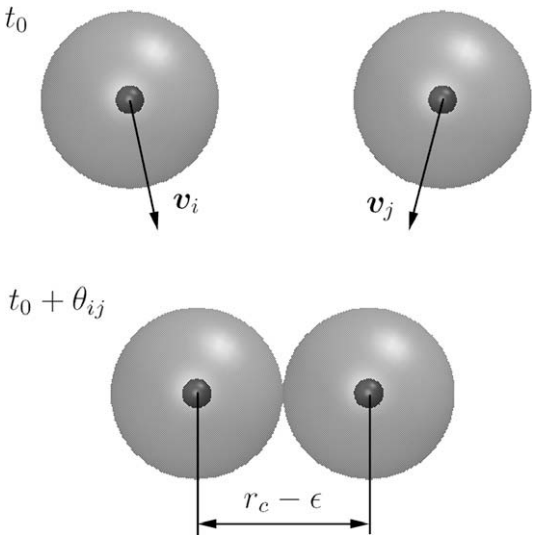

Fig. 1. Geometry of the impending *collision* between two atoms surrounded by their interaction sphere of diameter r_c. The point-like atoms are represented with the smaller spheres for clarity.

To develop the algorithm, we imagine each point-like atom surrounded by its *interaction sphere* of diameter r_c, as shown in Fig. 1. Analogously to EDMD, for a binary molecular encounter, we can express the time of interaction between any two atoms as

$$\|\mathbf{r}_i(t_0+\theta_{ij})-\mathbf{r}_j(t_0+\theta_{ij})\| = r_c-\epsilon, \tag{4}$$

where $\epsilon \ll r_c$. If at $t=t_0, i$ and j are already separated by a distance smaller than r_c, then θ_{ij} is set to zero, and therefore $\theta = \min_{\{\forall ij\}}\{\theta_{ij}\} = 0$. To compute each θ_{ij} more efficiently, we use a standard linked-list cell method [20]. The system is binned into cubic cells of side $\psi\lambda$, with $\psi \sim 1/100$.

If $\theta \neq 0$, the system is advanced in time to $t_0+\theta$, using Eq. (2). At this time, a pair of molecules, which we denote with α and β, will be separated by a distance $r_c-\epsilon$, and thus they will exert a force on each other. Because of the low density regime, it is reasonable to assume that α and β will evolve independently of the rest of the system, and that the duration of such encounter will be very short compared to τ. Therefore, we integrate in time the sub-system $\Gamma=\{\alpha,\beta\}$, while keeping the rest of the atoms still. Any interacting molecules at $t=t_0$ ($\theta=0$) are added to Γ, and processed analogously.

Although infrequent, the possibility that α and β start interacting with a third body during their evolution cannot be excluded *a priori*. Hence, while computing the trajectory of α and β, if a frozen atom ϕ gets within $r_c-\epsilon$ of either α or β or both, it is added to Γ, and its trajectory starts to be integrated too. When all atoms in Γ are separated by a distance larger than r_c, the next θ is determined, and the procedure repeated. We will denote with Θ_Γ the duration of such an interaction. The dynamics of the particles in Γ is solved using a standard second-order accurate symplectic velocity Verlet algorithm [25]. The time step Δt is set to 10^{-15} s, to adequately resolve the interactions. Depending on the simulation conditions, $\Theta_\Gamma/\Delta t$ can be of the order of a thousand, but $\Theta_\Gamma \ll \tau$ must hold.

The task to efficiently determine whether a third (or more) body needs to be added to Γ is not trivial. In a typical TDMD simulation, for each atom i, a neighbor list is constructed, usually every time step. Such list contains all the molecules located within a radius r_c from i, and it is used to compute the resultant force on i. Here, we construct two such neighbor lists, denoted with $\Lambda_i^{(1)}$ and $\Lambda_i^{(2)}$. $\Lambda_i^{(1)}$ contains i's neighboring atoms within a distance r_c and is used to compute the resultant force on i, as in any standard MD code. $\Lambda_i^{(2)}$ contains atoms that are distant from i between r_c and Qr_c, with $Q>1$. Therefore, in a typical situation $\Lambda_\alpha^{(1)}=\{\beta\}$ and $\Lambda_\beta^{(1)}=\{\alpha\}$. Also typical is the case that $\Lambda_\alpha^{(2)}=\{\emptyset\}$ and $\Lambda_\beta^{(2)}=\{\emptyset\}$. However, it may happen that $\Lambda_\alpha^{(2)}\neq\{\emptyset\}$ or $\Lambda_\beta^{(2)}\neq\{\emptyset\}$, particularly if Q is increased. For the sake of the argument, we assume that $\Lambda_\alpha^{(2)}=\{\phi\}$. So while Γ is evolved in time, the algorithm checks whether either α or β or both get within a distance smaller than r_c with their neighbors contained in $\Lambda_{(\alpha,\beta)}^{(2)}$. If $\|\mathbf{r}_\alpha-\mathbf{r}_\phi\|<r_c$ at some instant, then ϕ is added to $\Lambda_\alpha^{(1)}$ and included in the force calculation. Because the Δt is $\mathcal{O}(10^{-15})$ s, the overlap $\|\mathbf{r}_\alpha-\mathbf{r}_\phi\|-r_c$ is very small as well. Moreover, to further guarantee that any additional molecule is correctly included in the interaction (if, say, Q is chosen too small), $\Lambda_i^{(2)}$ is refreshed at some predefined intervals during the time evolution of Γ. Although this is a costly operation ($\mathcal{O}(N)$), it is carried out only a few times during Θ_Γ (in all our simulations, every 100 steps during the time integration of the trajectories of the Γ atoms), whereas in a standard MD code $\Lambda_i^{(1)}$ are usually regenerated at every time step. For the conditions of interest here, because $\theta \gg \Delta t$, $\Lambda_i^{(1,2)}$ are refreshed much fewer times per physical time unit than conventional TDMD. While we found that $Q=3,4$ worked for all our simulations, we did not find a simple way to determine the optimal Q *a priori*, but we relied on numerical trials that provided stable calculations.

At low density, for a finite size system, it may happen that no θ is found. Hence, if that is the case, similarly to Alder and Wainwright [14], we prescribe a θ^m to advance the system in time using Eq. (2). θ^m should be set large enough not to compromise the speed-up of the method (if $\theta^m=\Delta t$, it would be a traditional TDMD), but not too large to produce unphysical results (θ^m should be a fraction of τ). By numerical experiments, we found that setting $\theta^m = \min\{B_x/2, B_y/2, B_z/2\}/\max_{\{\forall i\}}\{v_i^{x,y,z}\}$ provided the correct relaxation rates (see Section 3), where B_x, B_y, B_z are the simulation super-cell dimensions.

Summarizing, the basic steps of the algorithm are the following:

```
do t=1, Nsteps
   call linked_list_cell(ψλ, Λ_i^(0))
   call collision_time(Λ_i^(0), θ)
   call advance(θ)
   call linked_list_cell(r_c, Λ_i^(1))
   call compute_interactions(Qr_c, Λ_i^(1))
end do.
```

Initially, by subdividing the domain into cubes of side $\psi\lambda$, the subroutine `linked_list_cell`$\left(\psi\lambda, \Lambda_i^{(0)}\right)$ constructs neighbor lists $\Lambda_i^{(0)}$ for each molecule i, containing particles within a distance $\psi\lambda$ from i. With the $\Lambda_i^{(0)}$ lists, subroutine `collision_time()` calculates the time θ of the next interaction. Using Eq. (2), a simple routine `call advance`(θ) fast-forwards the system to $t+\theta$, when at least one pair of atoms is interacting. The routine `linked_list_cell`$\left(r_c, \Lambda_i^{(1)}\right)$ identifies such pair(s) (separated by a distance smaller than the cut-off radius r_c) by constructing neighbor lists $\Lambda_i^{(1)}$, once

again using the linked-list cell method. The interacting atoms are then evolved in time by `compute_interactions`$\left(\Lambda_i^{(1)}\right)$, using the velocity-Verlet integration scheme with $\Delta t = 1$ fs. Periodically (in our simulations, every 100 time steps), subroutine `compute_interactions`$\left(\Lambda_i^{(1)}\right)$ calls `linked_list_cell`$\left(Qr_c, \Lambda_i^{(1)}\right)$, which returns the $\Lambda_i^{(2)}$ lists used to check for many-body interactions and to update the $\Lambda_i^{(1)}$ lists.

2.3. The Direct Simulation Monte Carlo method and collision model

The Direct Simulation Monte Carlo (DSMC) solutions presented in this article are obtained using the MONACO code [26]. The DSMC method takes advantage of the dilute nature of gases and decouples movement and collision processes. A representative number of simulated particles are translated without interacting for time steps that are on the order of τ. Then, nearby particles (located within a cell with dimensions on the order of λ) are paired and undergo collisions in a statistical manner. The collision probability for a pair of particles is determined through the use of collision cross-sections. The simplest DSMC collision model is the hard-sphere model in which the collision cross-sections are HS diameters. In general, the collision cross-section of an interaction is dependent on the inter-molecular potential and the relative speed of colliding particles, among other factors. More accurate than the HS model and widely used in DSMC simulation is the variable-hard-sphere (VHS) model of Bird [10]. In fact, VHS collision cross-sections $Q^{(2)}$ are based on a reference HS diameter that is allowed to vary as a power-law function of the relative speed g:

$$Q^{(2)} \sim g^{-2\zeta}. \tag{5}$$

Upon integration of $Q^{(2)}$ over a Maxwellian velocity distribution function, the usual viscosity law is obtained:

$$\mu \sim T^{\omega}, \tag{6}$$

where μ is the gas shear viscosity, and $\omega = 1/2 + \zeta$. Thus, ζ is chosen to obtain ω of the order of 0.75, which is characteristic of real gases. It is known that this power law assumption is very accurate for collisions with high relative velocities (high temperatures). In fact, for hot monatomic gases, the VHS Chapman–Enskog collision integrals (and, therefore, the transport properties) match very well those obtained with the Lennard-Jones potential [12]. Hence, the VHS model is expected to perform well for simulating shocks characterized by a high free stream temperature. This was recently demonstrated by Valentini and Schwartzentruber [12] by comparing the VHS–DSMC solutions with large-scale Lennard-Jones MD simulations for normal shock waves in argon. Near-perfect agreement is found for high temperatures (with $\omega = 0.7$), but discrepancies arise at low temperatures, where the power law assumption is less valid.

All DSMC simulations presented in this article employ the VHS collision model.

3. Translational relaxation calculations

For the relaxation simulations, a total of $N = 7538$ argon atoms were used. Argon is very well modeled by the Lennard-Jones potential [27]:

$$\psi(r_{ij}) = 4\epsilon\left[\left(\frac{\sigma}{r_{ij}}\right)^{12} - \left(\frac{\sigma}{r_{ij}}\right)^{6}\right], \tag{7}$$

where $\epsilon/k = 119.18$ K (k is Boltzmann's constant), $\sigma = 3.42$ Å, and $r_{ij} = \|\mathbf{r}_{ij}\|$. m is set to 6.624×10^{-26} kg, and the cut-off $r_c = 3\sigma$. At $t = 0, N/2$ particles were assigned $v_{x,i} = v_0$ and the other half $v_{x,i} = -v_0$, with $v_0 = 10\sqrt{2kT_e/m}$. $v_{y,i}$ and $v_{z,i}$ were distributed according to a Gaussian distribution. The total kinetic energy $\mathcal{K} = 1/2\sum_{i=1}^{N} m_i \mathbf{v}_i \cdot \mathbf{v}_i$ was then rescaled to obtain a global temperature $T_e = 300$ K. The macroscopic temperature is computed with the following relation:

$$T = \frac{m}{3k}\left[\langle v_x^2 + v_y^2 + v_z^2\rangle - (\langle v_x\rangle^2 + \langle v_y\rangle^2 + \langle v_z\rangle^2)\right], \tag{8}$$

where we define

$$T_{x,y,z} = \frac{m}{k}\left(\langle v_{x,y,z}^2\rangle - \langle v_{x,y,z}\rangle^2\right). \tag{9}$$

If such calculations are carried out with a standard TDMD approach using a symplectic time integration scheme, the NVE ensemble is the most straight-forward choice, and the total energy E is conserved within numerical accuracy. Because of the low density regime, the total energy of the system is mainly kinetic. Therefore, provided that E remains constant, $T_e = 300$ K is the equilibrium temperature at $t \gg \tau_r$, when all molecular velocity components are distributed according to a Maxwell–Boltzmann distribution with variance set by T_e. Periodicity was enforced in all directions and the box dimensions were scaled to obtain the desired density. A time step of 1 fs was employed to evolve the Γ sub-sets in time.

If the one-particle velocity distribution function (vdf) $f(v,t)$ is close to its equilibrium value and decays with a uniform time constant τ_r, the collision operator in Boltzmann's equation can be simplified to

$$\frac{\partial f}{\partial t} = -\frac{f - f_0}{\tau_r}, \tag{10}$$

using the well-known Bhatnagar–Gross–Krook (BGK) approximation [28]. f_0 denotes the equilibrium vdf, i.e., the Maxwell–Boltzmann distribution.

Initially, $T_y = T_z \neq T_x$, but under the BGK approximation (Eq. (10)), the difference between any two should decay exponentially:

$$\Delta T(t) = \Delta T(0)\exp(-t/\tau_r) = \Delta T(0)\exp(-\nu_r t), \tag{11}$$

where $\nu_r = 1/\tau_r$. At equilibrium, $T_x = T_y = T_z = 300$ K.

Using the nominal collision frequency ν_{nom} derived from viscosity [29]:

$$\nu_{nom} \equiv \frac{4}{\pi}\frac{\rho R T_e}{\mu(T_e)}, \tag{12}$$

we can express $\nu_r = \nu_{nom}/Z$, where we introduce the collision number Z to characterize the relaxation process. Hence, with the above definition, Eq. (11) can be rewritten as

$$\Delta T(t) = \Delta T_0 \exp(-\nu_{nom} t/Z). \tag{13}$$

For a simple gas of spherical non-interacting molecules, under no external force field and close to equilibrium, τ_r is given by [28]:

$$\frac{1}{\tau_r} = \nu_r = \frac{\rho R T_e}{\mu(T_e)}, \tag{14}$$

and, therefore, according to Eq. (12):

$$\nu_{nom} = \frac{4}{\pi}\nu_r, \tag{15}$$

giving $Z_{HS} = 4/\pi = 1.273$. In the calculations, $\mu(T_e) = 2.27 \times 10^{-5}$ kg m^{-1} s^{-1}. Because in the BGK approximation $Z = \text{const}$ contains the information on the microscopic interactions between the gas particles, it will differ depending on the particle method used to describe the relaxation process. However, at dilute or rarefied conditions and moderately high temperatures, we do not expect Z to deviate significantly from Z_{HS}.

The DSMC simulations were performed for a two-dimensional box with specularly reflecting walls. The simulation employs 1.6 million particles. Although not relevant for a uniform gas relaxation problem, the box was subdivided into square cells of sides $\lambda/4$, resulting in approximately 250 particles/DSMC cell. Initially, particles were randomly positioned within each cell to obtain a density of $\rho = 10^{-4}$ kg/m^3 throughout the box and were given velocities according to the non-equilibrium velocity distribution described above. The simulation time step was set to $\Delta t = 10^{-10}$ s. This is significantly less than the mean collision time, but is necessary since the initial velocity distribution is highly nonequilibrium and the collision time experienced by a particle varies widely. Temperatures and velocity distribution functions were computed by sampling all particles in the entire box at a specific time step during the DSMC simulation. The VHS power-law exponent ζ was set to obtain $\omega = 1/2 + \zeta = 0.8$, and the VHS reference parameters were set such that $\mu_{VHS}(300\text{ K}) = 2.27 \times 10^{-5}$ kg m^{-1} s^{-1} [10].

As Fig. 2 shows, the combined ED/TD method well reproduces the approach to equilibrium for densities between $\rho = 10^{-4}$ kg/m^3 to $\rho = 10^{-1}$ kg/m^3. This demonstrates that the ED/TD MD method predicts the correct collision rates for

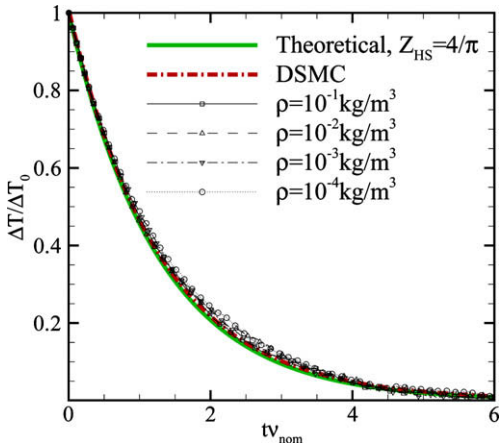

Fig. 2. Relaxation history at various densities.

a wide range of densities. We estimated an average $Z \simeq 1.36$, with a least-square fit of the time history of $T_y - T_x$ and $T_z - T_x$ for all cases, close both to the theoretical value $Z_{HS} = 1.273$ and $Z_{DSMC} = 1.31$. We also ran a further relaxation calculation using standard TDMD, but at $\rho = 1$ kg/m^3, at which the TD approach is still practical but computationally expensive. Once again, using the average of $T_y - T_x$ and $T_z - T_x$ and a least-square fitting procedure, we computed $Z_{MD} \simeq 1.4$, which is within 3% of the translational collision number estimated from the ED/TD relaxation curves.

To further investigate the relaxation process, we also plotted the normalized speed distribution function (sdf) at various times for the $\rho = 10^{-4}$ kg/m^3 case. In order to compute the sdfs, each simulation was repeated 10 times, each time stopped at the tv_{nom} of interest, and the results averaged. The sdfs obtained with DSMC and ED/TD MD are shown in Fig. 3(a)–(c). Speeds are rescaled using $c_e = \sqrt{2kT_e/m}$, and the bin sizes are $\Delta c/c_e = 0.02$ $(tv_{nom} = 2), \Delta c/c_e = 0.04 (tv_{nom} = 4)$, and $\Delta c/c_e = 0.06$ $(tv_{nom} = 6)$.

The essential qualitative and quantitative features of the sdfs are well captured by the ED/TD method, particularly the highly spiked shape at $tv_{nom} = 2$, suggesting that most atoms still preserve memory of their initial configuration. However, the marked spike exhibited by the DSMC sdfs at both $tv_{nom} = 2$ and $tv_{nom} = 4$ is not reproduced by the ED/TD MD method, whose sdfs are slightly more spread out. The sdf at $tv_{nom} = 2$ computed with the ED/TD algorithm is compared with that obtained from a standard TDMD relaxation simulation at $\rho = 1$ kg/m^3. As Fig. 3(d) shows, the agreement is excellent.

Therefore, the difference between ED/TD MD and DSMC is attributed to assumptions in the DSMC collision model, which differ from the collision dynamics simulated by MD. Specifically, via the nature of the DSMC method, in which particle pairs are collided probabilistically (i.e., either they collide or they do not), at $tv_{nom} = 2$ or $tv_{nom} = 4$, there is still a significant number of particles which have not yet undergone a collision. These particles maintain their exact initial velocities and compose the spike in the sdfs. Furthermore, individual collision probabilities are computed using the VHS collision cross-section which assumes a specific power-law dependence $(\omega = 0.8)$. The relaxation process presented here is a stringent test case to compare DSMC with MD, since the initial velocity distribution is highly nonequilibrium (involving a wide range of collision relative speeds). Despite this, Figs. 2 and 3(a)–(c) demonstrate that the relaxation process is very accurately simulated by the ED/TD method at $\rho = 10^{-4}$ kg/m^3.

The average duration $\langle \Theta_\Gamma \rangle$ of the interaction among the atoms in Γ is shown in Fig. 4, where it is plotted relative to the mean collision time. $\langle \Theta_\Gamma \rangle$ denotes the average of Θ_Γ over the course of each simulation. τ is estimated using the hard-sphere

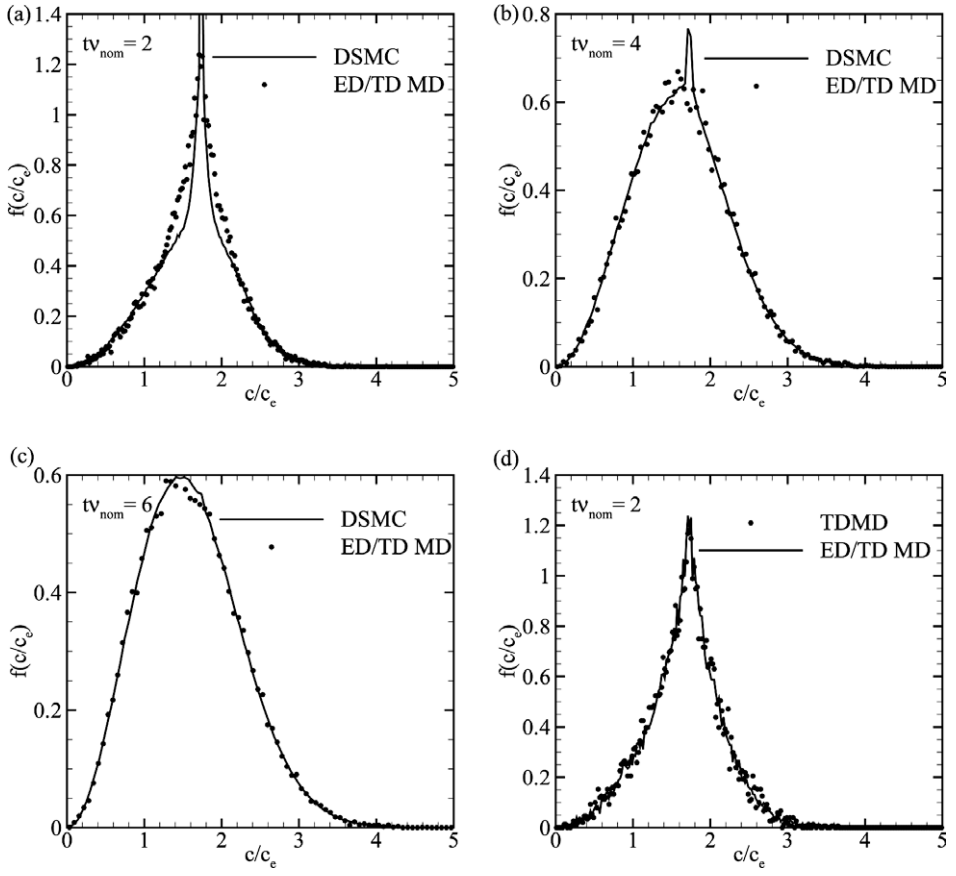

Fig. 3. Speed distribution functions obtained with the combined ED/TD algorithm: comparison with (a)–(c) DSMC at $tv_{nom} = 2, 4, 6$; (d) with TDMD at $tv_{nom} = 2$.

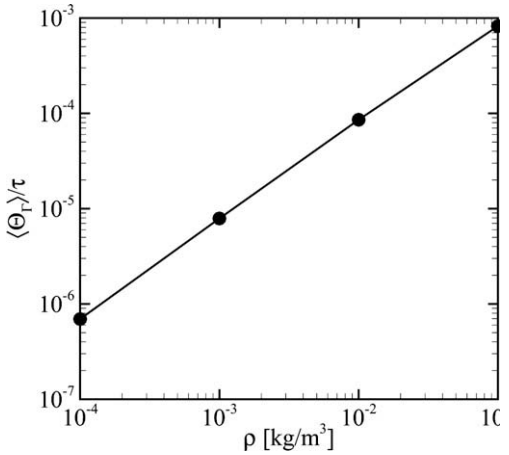

Fig. 4. Average interaction duration compared to the mean collision time.

approximation $\tau/\tau_{LJ} \simeq 0.14/\rho^*\sqrt{T^*}$, where $\rho^* \simeq \rho\sigma^3/m, T^* = kT_e/\epsilon \simeq 2.52$, and $\tau_{LJ} = \sigma\sqrt{m/\epsilon} \simeq 2.2 \times 10^{-12}$ s. Even at the highest density, $\langle\Theta_\Gamma\rangle/\tau$ is very small ($\simeq 1/1000$). Hence, molecules essentially travel unperturbed for most of their time evolution, because, on average, two of them interact for a very short time compared to the time that must elapse before they interact with any of the other atoms again.

Fig. 5 shows $\langle\theta\rangle/\Delta t$ at the various densities. We call this ratio the nominal speed-up as it is independent of the particular implementation of the method. In the density range considered, $\langle\theta\rangle/\Delta t$ varies from roughly 20,000 to 100, for the lowest and highest density, respectively. Also in Fig. 5, we plot a measure of the computational overhead of the ED/TD method, i.e., the ratio between the wall clock times to perform 1000 steps (on a 2.2 MHz Dual Core AMD Opteron™ CPU), with and without the use of the ED/TD algorithm. Except at the highest density, our implementation of the ED/TD method is roughly 4 times more computationally expensive per time step than the conventional TDMD. The reason for an approximately 20-fold increase at the highest ρ may be due to the larger probability that additional particles start interacting with the Γ atoms during their time evolution. The presence of more atoms in Γ means that more time is spent integrating their trajectories and more checks are performed to update the $\Lambda_i^{(1,2)}$ lists. However, as shown in Fig. 2, the method still correctly captures the relaxation process at $\rho = 10^{-1}$ kg/m^3, signifying that the main assumptions are still verified at this density. However, even though each ED/TD step (in our implementation) is up to 20 times more costly, each is equivalent to hundreds or even thousands of TDMD steps, thus making MD simulations of rarefied gases computationally tractable while maintaining accuracy.

Because very few atoms interact at once, the computational time spent to integrate the trajectories of the Γ atoms is completely negligible compared to the determination of θ and the construction of the $\Lambda_i^{(1,2)}$ lists and their refresh. The results of a code profiling indicate that the routine `linked_list_cell()` takes up roughly 80% of the computational time, whereas only 0.5% is spent to integrate the trajectories of the Γ atoms in the subroutine `compute_interactions()`. This suggests that algorithmic improvements discussed as aspects (1) and (2) in Section 1 could potentially greatly increase the efficiency

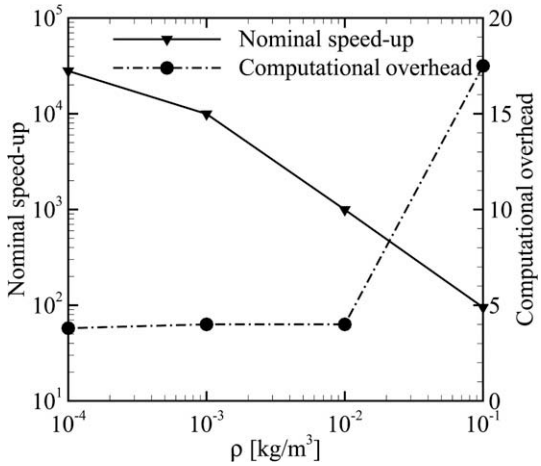

Fig. 5. Speed-up of the ED/TD algorithm.

of our combined ED/TD approach. Moreover, `compute_interactions()` can be easily made as independent as possible from the rest of the simulation code. Hence, we claim that the use of extremely accurate potentials, either empirical or more *ab initio*, should be feasible. This would be particularly rewarding when, for example, vibrational relaxation and chemical reactions need be investigated.

Finally, the total energy is conserved to better than $0.001\%/10^5$ time steps, for all cases. Hence, the NVE ensemble is reproduced, and the ED/TD algorithm ensures that no spurious overlaps compromise the stability of the calculation.

4. Mach 9 normal shock wave in rarefied argon

The second test case is a Mach 9 normal shock wave in argon, and unlike in the previous relaxation simulations, this is characterized by a strong density gradient.

The free stream conditions are $\rho_1 = 10^{-3}$ kg/m^3 and $T_1 = 300$ K. All variables are rescaled as

$$\rho_n = \frac{\rho - \rho_1}{\rho_2 - \rho_1}, \quad T_n = \frac{T - T_1}{T_2 - T_1}, \tag{16}$$

where the subscripts 1 and 2 denote far upstream and far downstream quantities, respectively. For all DSMC and MD results, the length scale is normalized with the upstream mean free path λ_1 estimated using:

$$\lambda_1 \simeq \frac{16}{5}\sqrt{\frac{\gamma}{2\pi}}\frac{\mu_1}{\rho_1 a_1} = 116\ \mu\text{m}, \tag{17}$$

where $\gamma = 5/3$ is the specific heat ratio, $\mu_1 \simeq 2.27 \times 10^{-5}$ kg m^{-1} s^{-1} the shear viscosity, a_1 the speed of sound.

The total number of simulated MD argon atoms was roughly N = 18,000. Although this gas dynamics problem is essentially one-dimensional, MD simulations are three-dimensional. The simulation box extended for roughly $12\lambda_1$ upstream and $12\lambda_1$ downstream of the shock. The other two dimensions were set to $B_y = 200$ Å and $B_z = 100$ Å.

Because of the non-periodic nature of the problem, the correct far field boundary conditions must be imposed in the flow direction x, while periodicity can be prescribed along y and z. To this purpose, atoms whose $x_i \leqslant \lambda_1$ and $x_i \geqslant 23\lambda_1$ were periodically removed and regenerated using the far field primitive variables $(\rho_{(1,2)}, u_{(1,2)}, v_{(1,2)} = 0, w_{(1,2)} = 0, T_{(1,2)}$, where u, v, w denote the bulk velocity components), related through the perfect gas inviscid one-dimensional Rankine–Hugoniot jump conditions corresponding to $M_1 = 9$. In both the upstream and downstream regions of size $\lambda_1 \times B_y \times B_z$, the number of particles was determined to match the pre- and post-shock densities:

$$n_{(1,2)} = \texttt{floor}\left(\frac{\rho_{(1,2)}\lambda_1 B_y B_z}{m}\right). \tag{18}$$

Because these boundary atoms were placed at random spatial locations, overlaps might occur. Therefore, similarly to Valentini and Schwartzentruber [12], we minimized their potential energy, using a Steepest Descent routine, to correctly remove overlaps, while keeping the rest of the system frozen. Each atom velocity was then drawn from a Maxwell–Boltzmann distribution whose variance is determined by $T_{(1,2)}$. Such a procedure was repeated roughly 200 times per τ_1, to ensure that atoms in the $[0, \lambda_1]$ region did not drift more than $5\%\lambda_1$ before being replaced. In this way, particles in the region where the sampling occurred, i.e., $x \in [\lambda_1, 23\lambda_1]$, were effectively in contact with infinite reservoirs, thus not feeling the *free surfaces* at $x = 0$ and $x = 24\lambda_1$.

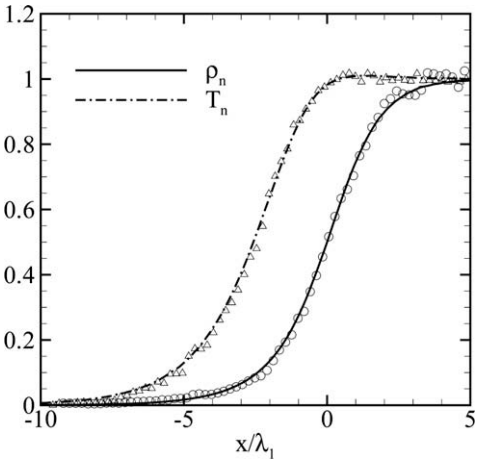

Fig. 6. Comparison between ED/TD MD and DSMC temperature and density profiles. Symbols are ED/TD MD, lines are DSMC.

In the domain $\mathcal{D} = \{x \in (\lambda_1, 23\lambda_1)\}$, the sampling occurred. To this purpose, $\mathcal{D}$ was subdivided into 100 bins in the x direction. Density was computed by time averaging the number of atoms contained in each bin. The macroscopic temperature was computed using Eq. (8).

For the DSMC simulations, the cell sizes and time steps were verified to be $\leqslant \lambda_1/2$ and $\leqslant \tau_1/4$, respectively. A two-dimensional computational domain was used, with only 10 cells in the y-direction where symmetry boundary conditions are prescribed. The total number of simulation particles was set to roughly 8 million, to produce smooth profiles and distribution functions, and to prevent the shock wave from drifting upstream or downstream. The VHS model was used with $\omega = 0.7$, as recommended in Ref. [12].

Fig. 6 shows a remarkable agreement between the temperature and density profiles obtained with the ED/TD method and those computed using DSMC. The origin of the x-axis was set at the point where $\rho_n = 0.5$. Although the density profile shown in Fig. 6 appears fairly smooth, its slope is not so, particularly around $\rho_n = 0.5$. Hence, we evaluated the inverse shock thickness δ by approximating the derivative $d\rho_n/d(x/\lambda_1)$ with a finite difference scheme on a 5-point stencil, as suggested by Macrossan and Lilley [29]. We have estimated $\delta \simeq 0.239$, which is within the $\pm 4\%$ uncertainty of the reported $\delta = 0.23$ in Alsmeyer's experiments [5].

In addition, the perpendicular velocity distribution functions were extracted at four locations throughout the nonequilibrium region of the shock, and they are shown in Fig. 7(a)–(d). The x-axis was rescaled using $c_1 = \sqrt{2kT_1/m}$, each vdf was sampled using a bin size $\Delta v_x/c_1 = 0.2$ and was normalized. The vdfs obtained with the ED/TD method agree very well with those obtained with DSMC. Small discrepancies may be due to the slight difference in the ρ_n values at which they are extracted (but always within 5%), and statistical scatter (DSMC vdfs are much smoother due to the much larger number of particles used). However, the essential features are very well captured, both qualitatively and quantitatively, namely their strongly bimodal shape, and the location and value of the local maxima and minimum.

In this simulation, $\langle\theta\rangle/\Delta t \simeq 300$. While in the relaxation calculation with $\rho = 10^{-3}$ kg/m^3, such value was roughly 10,000, a smaller value can be expected due to the much higher temperatures reached here, and thus, much larger relative velocities are expected. Moreover, the downstream density is approximately four times ρ_1, and the total number of particles ($N \simeq 18,000$) is more than twice that used in the relaxation calculations ($N = 7538$).

For this flow, one mean collision time $\tau_1 \simeq 3 \times 10^{-7}$ s is covered by approximately 900,000 ED/TD steps, the equivalent of 300 million TDMD time steps, and roughly $20\tau_1$ are necessary to reach steady state and obtain fairly smooth profiles. Given that $\langle\theta\rangle/\Delta t \simeq 300$ and a computational overhead of roughly 4 for our implementation (Fig. 5), compared to a standard TDMD

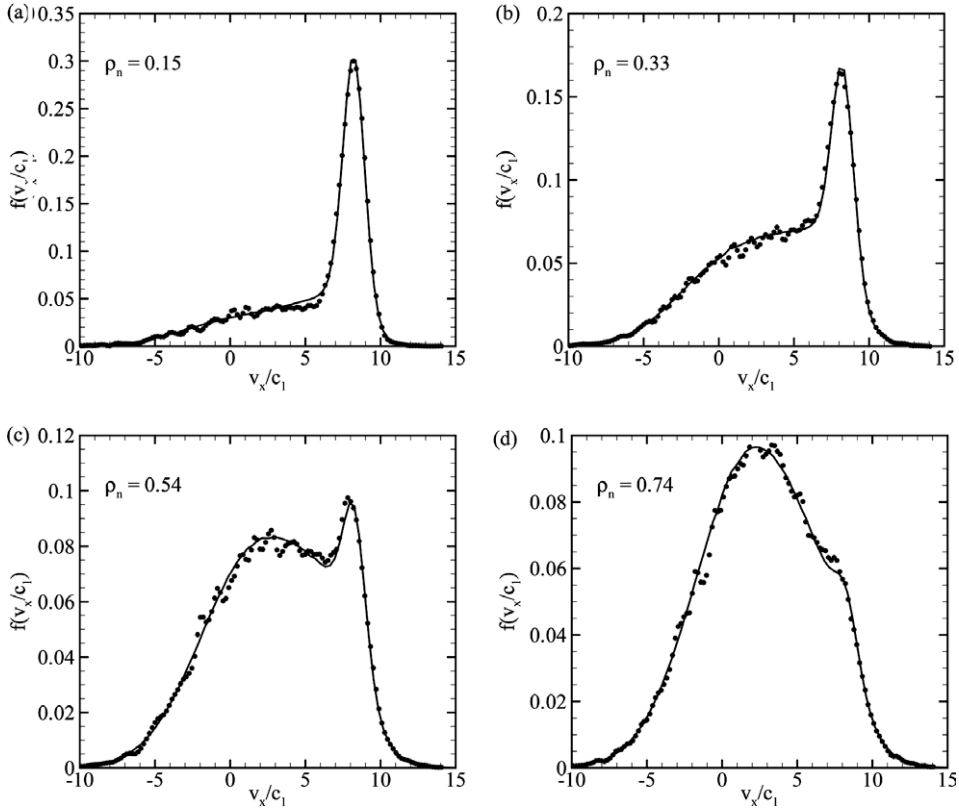

Fig. 7. Velocity distribution function at (a) $\rho_n = 0.15$, (b) $\rho_n = 0.33$, (c) $\rho_n = 0.54$, and (d) $\rho_n = 0.74$. Symbols are ED/TD MD, lines are DSMC.

calculation, the ED/TD algorithm accelerated such a simulation by a factor of approximately 75, thus making it feasible even on a single processor.

5. Summary and conclusions

In this work, we have presented a novel algorithm to speed-up the simulation of rarefied gases using spherically symmetric soft potentials. To achieve this, we have combined ideas from both the Event-Driven and the Time-Driven approach to Molecular Dynamics. In particular, the proposed algorithm correctly identifies the time of the next interaction, fast-forwards the system to that time, and it processes each interaction using Time-Driven MD with the sufficiently small time step to correctly resolve the atomic motion. Many-body interactions are also detected and simulated, with a small approximation.

The ED/TD method was applied to translational relaxation in argon between $\rho = 10^{-4}$ kg/m^3 and $\rho = 10^{-1}$ kg/m^3. The realistic Lennard-Jones potential was used in the simulations. For all densities, comparisons with theory, DSMC, and full TDMD calculations have revealed that the ED/TD approach correctly describes the relaxation process, thus reproducing the proper collision rates. Within the framework of the BGK approximation, a translational collision number $Z \simeq 1.36$ was determined, within 3% the value obtained with pure TDMD. Moreover, the NVE ensemble is correctly reproduced, thus avoiding spurious overlaps. Compared to standard TDMD, a speed-up as high as 20,000 was achieved, with each ED/TD step at most 20 times more expensive than each TDMD time step.

We also tested our method to simulate a Mach 9 shock wave in a rarefied argon flow, at conditions much more similar to those in typical experiments. The highly nonequilibrium nature of the problem, characterized by very strong temperature and density gradients, provided a fairly stringent test for our accelerated method. However, the results we obtained are remarkably accurate, both in terms of bulk quantity profiles, i.e., temperature and density, and in terms of molecular velocity distributions throughout the nonequilibrium region of the shock front. Using our approach, we could reduce the total number of steps per free stream mean collision time by roughly 300 times, thus making such simulations using a realistic soft potential feasible even on a single CPU.

The current implementation of our combined ED/TD method is very simple, particularly with respect to the calculation of the next interaction time and management of the sequence of events, and its performance worsened at the highest density considered here ($\rho = 10^{-1}$ kg/m^3). However, we think that the more complex algorithms that have been widely applied and tested for EDMD of rigid particles could be readily implemented in our ED/TD scheme to further improve its efficiency. Moreover, such enhancements might also prove beneficial to use this approach at higher densities ($\rho \sim 1$ kg/m^3), for example for simulating micro-scale flows at STP conditions, at which the ED/TD assumptions may still be valid.

Finally, a code profiling also revealed that most of the computational effort is spent to determine the next collision time (80%), unlike the integration of the trajectory of the interacting particles (0.5%). Hence, we argue that the use of much more sophisticated and realistic potentials (either classical or more first-principles) could be feasible, since the force calculation (which typically represents the biggest burden in a TDMD calculation for a condensed phase) only involves a few pairs of particles per ED/TD step. If so, full MD simulations of shock waves could be achieved to study rotational or vibrational nonequilibrium, and even chemical reactivity, without relying on empirical models.

Acknowledgments

The authors would like to acknowledge many helpful discussions with Professors Graham V. Candler and Ellad B. Tadmor. The authors would also like to thank Prof. Iain D. Boyd for the use of the MONACO code.

The research is supported by Air Force Office of Scientific Research (AFOSR) under Grant No. FA9550-04-1-0341. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the AFOSR or the U.S. Government.

References

[1] H.M. Mott-Smith, The solution of the Boltzmann equation for a shock wave, Phys. Rev. 82 (6) (1951) 885–892.
[2] B. Schmidt, Electron beam density measurements in shock waves in argon, J. Fluid Mech. 39 (2) (1969) 361–373.
[3] C. Cercignani, A. Frezzotti, P. Grosfils, The structure of an infinitely strong shock wave, Phys. Fluids 11 (9) (1999) 2757–2764.
[4] F. Robben, L. Talbot, Measurement of shock wave thickness by the electron beam fluorescence method, Phys. Fluids 9 (4) (1966) 633–643. Apr..
[5] H. Alsmeyer, Density profiles in argon and nitrogen shock waves measured by the absorption of an electron beam, J. Fluid Mech. 74 (2) (1976) 497–513.
[6] E.P. Muntz, L.N. Harnett, Molecular velocity distribution function measurements in a normal shock wave, Phys. Fluids 12 (10) (1969) 2027–2031.
[7] T. Holtz, E.P. Muntz, Molecular velocity distribution functions in an argon normal shock wave at Mach number 7, Phys. Fluids 29 (9) (1983) 2425–2436.
[8] G. Pham-Van-Diep, D. Erwin, E.P. Muntz, Nonequilibrium molecular motion in a hypersonic shock wave, Science 245 (1989) 624–626.
[9] G.A. Bird, Approach to translational equilibrium in a rigid sphere gas, Phys. Fluids 6 (10) (1963) 1518–1519.
[10] G.A. Bird, Molecular Gas Dynamics and the Direct Simulation of Gas Flows, Clarendon, Oxford, 1994.
[11] D.A. Erwin, G.C. Pham-Van-Diep, E.P. Muntz, Nonequilibrium gas flows. I: A detailed validation of Monte Carlo direct simulation for monatomic gases, Phys. Fluids A 3 (4) (1991) 697–705.
[12] P. Valentini, T.E. Schwartzentruber, Large-scale molecular dynamics simulations of normal shock waves in dilute argon, Phys. Fluids 21 (6) (2009) 066101-1–066101-9.
[13] B.J. Alder, T.E. Wainwright, Phase transition for a hard sphere system, J. Chem. Phys. 27 (1957) 1208–1209.
[14] B.J. Alder, T.E. Wainwright, Studies in molecular dynamics. I. General method, J. Chem. Phys. 31 (2) (1959) 459–466.
[15] T. Tokumasu, Y. Matsumoto, Dynamic molecular collision (DMC) model for rarefied gas flow simulations by the DSMC method, Phys. Fluids 11 (7) (1999) 1907–1920.

[16] B.L. Holian, W.G. Hoover, B. Moran, G.K. Straub, Shock-wave structure via non-equilibrium molecular dynamics and Navier–Stokes continuum mechanics, Phys. Rev. A 22 (6) (1980) 2797–2808.
[17] S. Schlamp, B.C. Hathorn, Higher moments of the velocity distribution function in dense-gas shocks, J. Comput. Phys. 223 (2007) 305–315.
[18] P. Kowalczyk, A. Palczewski, G. Russo, Z. Walenta, Numerical solutions of the Boltzmann equation: comparison of different algorithms, Eur. J. Mech. B-Fluid 27 (2008) 62–74.
[19] E. Salomons, M. Mareschal, Usefulness of the Burnett description of strong shock waves, Phys. Rev. Lett. 69 (2) (1992) 269–272.
[20] M.P. Allen, D.J. Tildesley, Computer Simulation of Liquids, Oxford, 1987.
[21] A. Donev, S. Torquato, F.H. Stillinger, Neighbor list collision-driven molecular dynamics simulation for nonspherical hard particles. I. Algorithmic details, J. Comput. Phys. 202 (2005) 737–764.
[22] G. Paul, A complexity O(1) priority queue for event driven molecular dynamics simulations, J. Comput. Phys. 221 (2007) 615–625.
[23] S. Miller, S. Luding, Event-Driven molecular dynamics in parallel, J. Comput. Phys. 193 (2003) 307–316.
[24] J.M. Haile, Molecular Dynamics Simulation – Elementary Methods, Wiley, 1997.
[25] D. Frenkel, B. Smit, Understanding Molecular Simulation, Academic Press, 2002.
[26] S. Dietrich, I.D. Boyd, Scalar and parallel optimized implementation of the Direct Simulation Monte Carlo method, J. Comput. Phys. 126 (1996) 328–342.
[27] J.O. Hirschfelder, C.F. Curtiss, R.B. Bird, Molecular Theory of Gases and Liquids, John Wiley and Sons, 1954.
[28] S. Chapman, T.G. Cowling, The Mathematical Theory of Non-Uniform Gases, Cambridge, 1970.
[29] M.N. Macrossan, *v*-DSMC: a fast simulation method for rarefied flow, J. Comput. Phys. 173 (2) (2001) 600–619.