

Molecular Physics

An International Journal at the Interface Between Chemistry and Physics

ISSN: 0026-8976 (Print) 1362-3028 (Online) Journal homepage: <https://www.tandfonline.com/loi/tmph20>

Beyond direct simulation Monte Carlo (DSMC) modelling of collision environments

O. Schullian & B. R. Heazlewood

To cite this article: O. Schullian & B. R. Heazlewood (2019) Beyond direct simulation Monte Carlo (DSMC) modelling of collision environments, Molecular Physics, 117:21, 3076-3087, DOI: 10.1080/00268976.2019.1602740

To link to this article: <https://doi.org/10.1080/00268976.2019.1602740>



© 2019 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 25 Apr 2019.



Submit your article to this journal [↗](#)



Article views: 562



View related articles [↗](#)



View Crossmark data [↗](#)

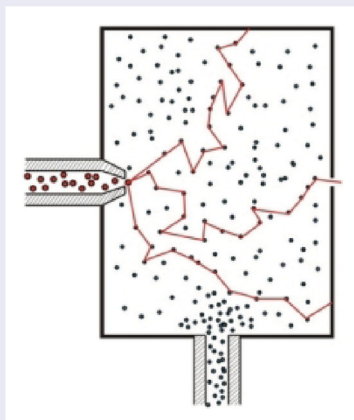
Beyond direct simulation Monte Carlo (DSMC) modelling of collision environments

O. Schullian^a and B. R. Heazlewood^b

^aDepartment of Theory and Biosystems, Max Planck Institute of Colloids and Interfaces, Potsdam, Germany; ^bDepartment of Chemistry, Physical and Theoretical Chemistry Laboratory, University of Oxford, Oxford, UK

ABSTRACT

Direct simulation Monte Carlo (DSMC) models have been successfully adopted and adapted to describe gas flows in a wide range of environments since the method was first introduced by Bird in the 1960s. We propose a new approach to modelling collisions between gas-phase particles in this work – operating in a similar way to the DSMC model, but with one key difference. Particles move in a mean field, generated by all previously propagated particles, which removes the requirement that all particles be propagated simultaneously. This yields a significant reduction in computation effort and lends itself to applications for which DSMC becomes intractable, such as when a species of interest is only a minor component of a large gas mixture.



ARTICLE HISTORY

Received 20 December 2018
Accepted 26 March 2019

KEYWORDS

DSMC; gas-phase collisions;
rarefied gas flows; buffer gas
cell

1. Introduction

The direct simulation Monte Carlo (DSMC) method has successfully been applied to describe a plethora of gas-phase systems in the decades since it was first introduced [1,2]. Indeed, DSMC is the typical method adopted for simulating rarefied gas flows that are not accessible to molecular dynamics simulations (due to the high number of particles in the system) or the Navier–Stokes equation (because of the break down of the transport term [2]). With DSMC, a collection of simulated particles is taken to represent the properties of a large number of physical particles. All simulated particles are propagated simultaneously. The simulation volume is divided into a series of subcells, and the probability of a collision between a

pair of simulated particles during interval δt is calculated by establishing the ratio of the volume swept out by the cross sections of the particles to the volume of the subcell. The position of each simulated particle is subsequently updated, taking energy transfer into account if a collision is deemed to have occurred, and the next timestep of the calculation is performed. There are several key assumptions that underlie the DSMC method that have to be satisfied for the approach to yield meaningful results. In particular, a relatively small number of simulated particles (each representing f_n physical particles) must be able to represent the behaviour of the whole system. These simulated particles are propagated ballistically and collisions are performed stochastically, with the collision

CONTACT B. R. Heazlewood  brianna.heazlewood@chem.ox.ac.uk  Department of Chemistry, Physical and Theoretical Chemistry Laboratory, University of Oxford, South Parks Road, Oxford OX1 3QZ, UK

© 2019 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

probability adjusted by f_n to reflect the properties of the physical system.

We have previously developed a DSMC model incorporating rotational state-changing cross sections, and this model has been shown to accurately reproduce the rotational and translational cooling of ammonia molecules seeded in a helium supersonic expansion [3], in addition to successfully describing the thermalisation of ammonia molecules colliding with helium atoms in a buffer gas cell [4]. While we have demonstrated that a DSMC model can represent these important collision environments very well, it can be computationally expensive to implement; a DSMC model becomes unfeasible for systems with high numbers of particles, or where there are mixtures of gases with a wide concentration range. As such, there are limitations on the collision environments that DSMC methods can simulate. For example, our DSMC model cannot be easily applied to buffer gas cells longer than approximately 0.5 cm, whereas many experimental buffer gas cells are 1–2 cm in length [5].

In this paper, we propose a self-consistent mean field DSMC (SCMFD) method as an alternative approach to Bird's DSMC method, for modelling collisions between particles in certain gas-phase environments. This new SCMFD method relies on two key assumptions. Firstly, the trajectory of a single particle is determined by the mean field, but does not influence the mean field – it is one particle in an ensemble of trillions (or more) particles. Secondly, the combination of a limited number of trajectories is sufficient to reconstruct the mean field, which consists of the density and velocity fields. The probability of a collision between two particles is calculated *a priori*, in a similar way to the DSMC approach – but without the need to convert the simulated particles back to the real particles each simulated particle represents (i.e. removing the need for the f_n correction factor). There are three main benefits to this SCMFD approach:

- (i) There is no need to distinguish between simulated particles and physical particles.
- (ii) It is not necessary to recalculate the collision probability at timesteps where no collisions occur (as the system has not changed and thus does not need to be updated). This significantly improves the speed of the calculation.
- (iii) Individual particles can be followed – meaning that one need not explicitly consider the same relative number of particles of each type as exist in the physical system. This extends the applicability of SCMFD to systems where the species of interest is a minor component of the gas mixture, which is one of the reasons that the DSMC method can become intractable. It removes the requirement that all

particles be propagated simultaneously, enabling resources to be focused on the collisions of interest, as there is no need to explicitly consider collisions between particles that are not of interest.

While the SCMFD method is currently only valid for stationary flows (i.e. time-independent mean fields), it reduces the number of interactions that need to be calculated between particles (which scales as N^2), leading to a significant decrease in the memory and time required to perform the simulations. Collisions within a gas mixture with many orders of magnitude difference in the concentration of each species can be easily treated. This would be especially advantageous for understanding the collision dynamics where there are low concentration 'seed' gases, or where there are complex cooling mechanisms at play, as frequently occurs in supersonic expansions and buffer gas cooling cells.

2. Method

The principles underlying the SCMFD method are described – and demonstrated, in the subsequent sections – using two well-known scenarios. A homogeneous gas and the 'one-dimensional' systems introduced by Bird are considered, with the latter referring to systems which are only confined in the x direction, and are otherwise translationally invariant [2]. See the Appendix for a derivation of the underlying principles of SCMFD and several of the key equations in this section.

There are three main scenarios where the SCMFD method can be readily applied:

- (1) **A closed system**, where the number of particles is finite and cannot change. In such cases, a particle is tracked for a fixed number of steps and the number density – the quantity of interest – is given by

$$n_i = \frac{1}{V_i} \frac{NN_i}{N_s g}. \quad (1)$$

The total volume of the system, V , is divided into subcells – as occurs in the DSMC method [2] – and V_i is the volume of subcell i . N is the number of particles in the system, N_s is the number of simulated particles, N_i is the number of times a particle is located in subcell i and g is the number of steps a particle is propagated for.

- (2) **A steady flow**. Given an inflow flux J (expressed as the number of particles entering a subcell per unit of time), N will increase with time. As the increase is linear, the density is given by

$$n_i = \frac{1}{T} J \cdot T \langle dt_i \rangle = \frac{1}{V_i} J \frac{1}{N_s} N_i dt, \quad (2)$$

where T is the temperature of the system.

- (3) **When N goes to infinity**, for example an ideal gas in infinite space. In this case, the total volume must also go to infinity. Assuming there is an average density ρ_{gas} and the total volume $V = \infty$, then

$$n_i = \frac{1}{T} \cdot \rho_{\text{gas}} \cdot \frac{V}{V_i} \langle dt_i \rangle. \quad (3)$$

The value of $(V/V_i) \langle dt_i \rangle$ must be finite, and this can be achieved by either taking $V_i \rightarrow \infty$ (as we do here) or $\langle dt_i \rangle \rightarrow 0$.

The above-mentioned cases do not provide an exhaustive list, but simply outline a few of the possible implementations of our SCMFD approach. As each particle is treated independently, it is possible to calculate as many trajectories as one wishes for each particle type once the mean field conditions have been satisfied.

2.1. Calculation of the collision probability

To illustrate how collision probabilities are calculated we consider two particles, 1 and 2, with velocities v_1 and v_2 . The particles are transformed into a frame of reference where particle 2 is stationary, $v_r = v'_1 = v_1 - v_2$ and $v'_2 = 0$. The probability that particle 1 does not collide with any of the particles in volume V is given by

$$P_n = \prod_i (1 - P(v_{r,i})) \approx \exp \left(- \sum_i P(v_{r,i}) \right), \quad (4)$$

where the index i runs over all particles and P is defined in appendix A3. We can evaluate the sum in the exponential as

$$\begin{aligned} \sum_i P(v_{r,i}) &= \frac{1}{V} \sum_i |v_{r,i}| \delta t \\ &= \frac{N}{V} \frac{1}{N} \sum_i |v_i| \delta t = n \delta t \langle v_{r,i} \sigma \rangle. \end{aligned} \quad (5)$$

2.2. Variable hard sphere collisions

From the law of conservation of momentum, it can be established that $V_i = V_f$. Conservation of energy dictates that the absolute value of the relative velocity of two colliding particles must also be constant. This means that the relative velocity can only change in direction following a collision. Using a variable hard sphere model, we define an impact parameter b and consider the angle the impact parameter makes with an arbitrarily chosen 'origin' plane,

$$b = (r_1 + r_2) \cos(\chi/2), \quad (6)$$

where χ is the angle of the final relative velocity with respect to the initial relative velocity. In the simulation

procedure, a random impact parameter is chosen. In the frame where the initial relative velocity is in the z direction, the relative velocity is rotated around the y axis by the angle χ and then rotated around the z axis by the angle ϕ . (The angle ϕ is conserved due to angular momentum conservation.) This gives the final relative velocity, from which one can calculate the final velocity of both particles.

The diameter of each species in the variable hard sphere collision model is established using Bird's approach [2],

$$d = d_{\text{ref}} \sqrt{\frac{[2kT_{\text{ref}}/(m_r v_r^2)]^{\omega-1/2}}{\Gamma(5/2 - \omega)}}, \quad (7)$$

where d_{ref} is the reference diameter, T_{ref} is the reference temperature, ω is the reference viscosity index, v_r is the relative velocity, m_r is the reduced mass, and k is the Boltzmann constant.

2.3. Trajectory calculations

In order to calculate the trajectory of a single particle through the mean field, we follow a conceptually similar approach to that taken by the DSMC method [2,3]. A single particle is generated with position x and velocity v . This particle is propagated for a timestep of δt , after which the probability a collision occurred is calculated using the procedure outlined above. If and only if a collision is accepted, then the DSMC-like collision probabilities (see the Appendix) between the single particle of interest and the mean field particles are calculated. The probabilities are normalised and a collision partner is chosen. The collision is executed by adjusting the relative velocity of the particles, as described above using a random impact parameter and angle ϕ .

3. Simulation algorithm

3.1. Initialisation

The SCMFD method introduced in this work relies on the accurate description of the mean field, or the properties of the background gas, through which selected particles are propagated and tracked. The background gas is described using two arrays: one containing the number density of each species within each subcell, and one containing a distribution of velocities for each species. An initial configuration is then chosen – which, in the first example simulating argon in a one-dimensional box, is an ideal gas at temperature $T_i = T_1 \cdot (\text{dist} - x)/\text{dist} + T_2 \cdot x/\text{dist}$ (i.e. a linear approximation of the wall temperatures, T_1 and T_2).

3.2. Propagation

A particle is initially placed at one of the walls. The chosen particle is propagated by timestep dt (or αdt for the first timestep, with $\alpha \in [0, 1]$ uniformly distributed and different for every particle). After the propagation step, it is determined whether the particle is still within the confines of the simulated volume. If it has exited, then the particle is diffusely reflected from the corresponding wall. Then, using the protocol outlined above, it is established if a collision occurs and Equation (4) is evaluated against all entries in the background gas array. If no collision occurs, the particle is propagated for another timestep and the loop is repeated.

If a collision occurs, a routine randomly selects a number of particles from the background gas array and evaluates $P(v_r)$. A cumulative distribution function is subsequently calculated (and normalised, if necessary) and a collision partner is selected. A random impact parameter is chosen along with a random angle in $[0, 2\pi]$, with the resulting velocity rotated accordingly. A check on momentum and energy ensures that both are conserved. The particle is then assigned a new post-collision velocity. Only in the case where the particle being propagated collides with another particle of the same species (in the background gas) is the background gas velocity also adjusted. The reason for this is twofold; it improves convergence and enables us to preferentially follow only

the particles of interest. If we were to adjust the properties of the background gas following collisions between two different species, then we would need to reproduce the physical number ratio of all species in the system in the trajectory calculations – else we would overestimate the energy exchange between the species of interest and the background gas. Finally, when propagating a trajectory,

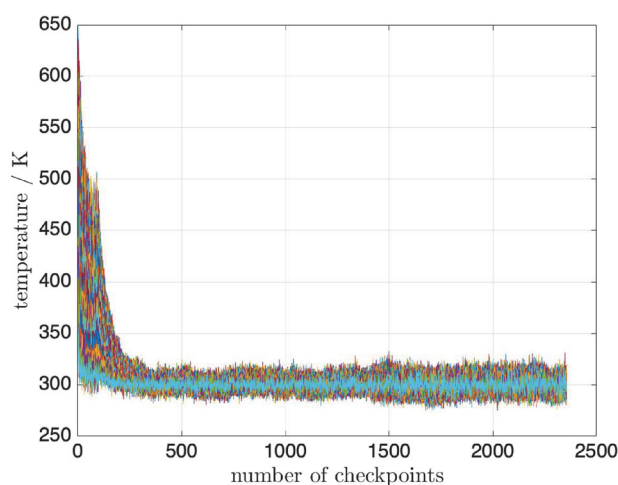


Figure 1. The temperature of the mean field gas is plotted as a function of iteration steps. Argon gas at an initial temperature of 600 K is cooled down through collisions with the 300 K walls of a box, with convergence achieved after approximately 400 checkpoints.

Table 1. Parameters adopted for the simulation of argon gas in a 1 m box.

Parameter	Value	Description
n_p	1	Number of different particles
$n_{\text{background}}$	1000	Number of velocities used to describe the background gas per subcell
	100,000	Total number of trajectories
	100,000,000	Number of timesteps per trajectory
δt	5×10^{-8} s	Timestep
a	5	After every a trajectories a checkpoint is written and the background gas is updated
	0.1	Relative amount of the background gas information (density and velocity distribution) that is updated
<i>Wall specifications</i>		
dist	1 m	Distance between walls (x direction)
n_{cell}	100	Number of subcells
v_y^1	0 m s^{-1}	Velocity in y direction of left wall
v_z^1	0 m s^{-1}	Velocity in z direction of left wall
T_1	300 K	Temperature of left wall
v_y^2	0 m s^{-1}	Velocity in y direction of right wall
v_z^2	0 m s^{-1}	Velocity in z direction of right wall
T_2	300 K	Temperature of right wall
<i>Particle parameters</i>		
	1	Number of particles per trajectory of this particle type
d_{ref}	4.17×10^{-10}	Reference diameter
T_{ref}	273 K	Reference temperature
ω	0.8	Reference viscosity index
m	39.948 u	Mass of particle
n_p	10^{20} m^{-3}	Number density of particle

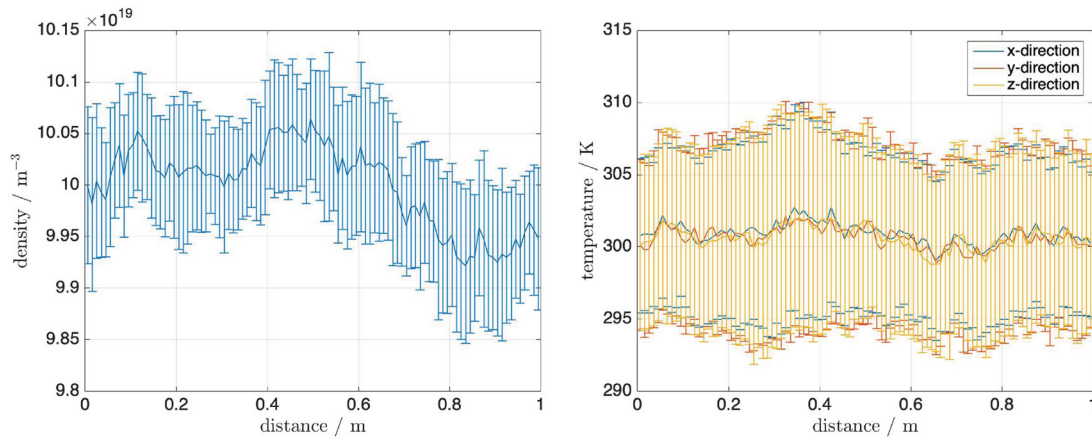


Figure 2. Argon gas, initially at 600 K, is cooled down via collisions with the 300 K walls of a 1 m box. The average local density and temperature of the gas established from the SCMF model are plotted (averaged over the last 200 checkpoints). The error bars indicate one standard deviation. It can be seen that argon is uniformly distributed throughout the box, and that the gas has successfully been cooled to 300 K through collisions with the walls, as expected.

Table 2. Parameters adopted for the simulation of a gas mixture in a 1 cm box.

Parameter	Value	Description
n_p	3	Number of different particles
$n_{\text{background}}$	1000	Number of velocities used to describe the background gas per subcell
	100,000	Total number of trajectories
	5,000,000	Number of timesteps per trajectory
δt	1×10^{-8} s	Timestep
a	5	After every a trajectories a checkpoint is written and the background gas is updated
	0.1	Relative amount of the background gas information (density and velocity distribution) that is updated
<i>Wall specifications</i>		
dist	0.01 m	Distance between walls (x direction)
n_{cell}	100	Number of subcells
v_y^1	0 m s^{-1}	Velocity in y direction of left wall
v_z^1	0 m s^{-1}	Velocity in z direction of left wall
T_1	300 K	Temperature of left wall
v_y^2	0 m s^{-1}	Velocity in y direction of right wall
v_z^2	0 m s^{-1}	Velocity in z direction of right wall
T_2	300 K	Temperature of right wall
<i>Particle parameters</i>		
	1	Number of particles per trajectory of this particle type
d_{ref}	4.17×10^{-10}	Reference diameter
T_{ref}	273 K	Reference temperature
ω	0.5	Reference viscosity index
m	39.948 u	Mass of Ar
n_p	10^{20} m^{-3}	Number density of Ar
	1	Number of particles per trajectory of this particle type
d_{ref}	4.17×10^{-10}	Reference diameter
T_{ref}	273 K	Reference temperature
ω	0.5	Reference viscosity index
m	4 u	Mass of He
n_p	10^{18} m^{-3}	Number density of He
	1	Number of particles per trajectory of this particle type
d_{ref}	4.17×10^{-10}	Reference diameter
T_{ref}	273 K	Reference temperature
ω	0.5	Reference viscosity index
m	28 u	Mass of N_2
n_p	10^{14} m^{-3}	Number density of N_2

the subcell and velocity of the particle are saved periodically. Every a trajectories, a selected percentage of the velocity entries in the background gas array are replaced (the 'oldest' k entries are replaced by k 'new' entries). Finally, a checkpoint file is created, enabling one to restart the simulation from this point and generating a complete description of the background gas.

4. Results

To evaluate the performance of the SCMFD method, we investigate the properties of different gas phase environments and compare the output of SCMFD simulations with data available in the literature or obtained from theoretical models. Each of the systems considered is either found in Bird's book [2], facilitating a direct comparison of our SCMFD approach with the results obtained from Bird's DSMC method, or lend themselves to direct

quantitative assessment from other well-established theoretical models.

4.1. Single component homogeneous gas

The first example selected is intentionally unrealistic – in terms of both size and initial conditions – in order to demonstrate the stability, convergence, accuracy and efficiency of the SCMFD method. Argon gas is simulated in a box of length 1 m. The mean-field gas, which fills the box initially, has a temperature of 600 K and the temperature of the walls is 300 K. The parameters utilised in the simulation are presented in Table 1.

Adopting the simulation procedure outlined above, we can follow the trajectories of gas particles as they move through the box and establish an approximate temperature from the variance at each checkpoint. While convergence is slow – as expected, given the large size of the system and the low frequency with which particles

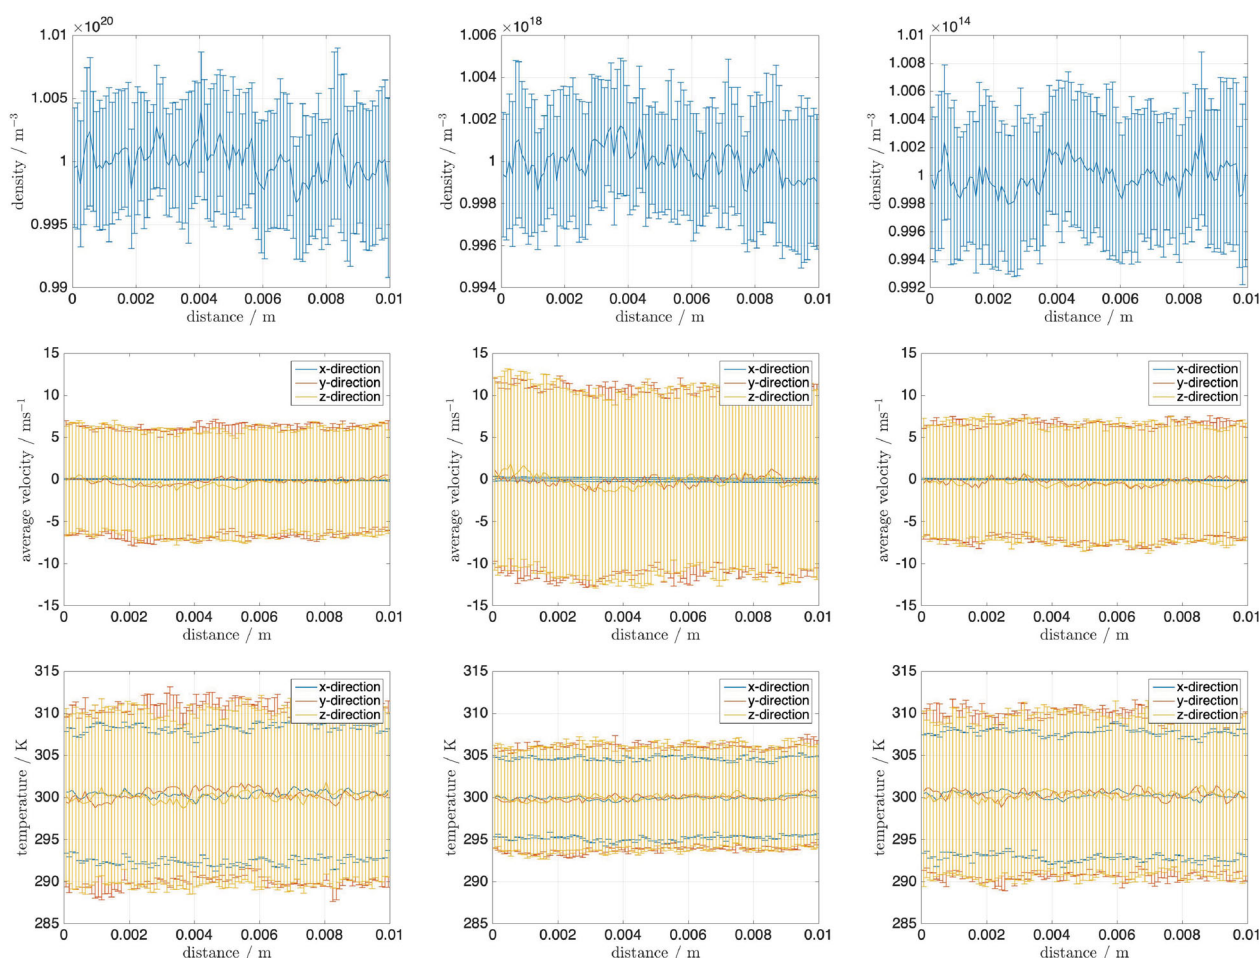


Figure 3. The density (top panel), average velocity (middle panel) and temperature (bottom panel) of each of the three gases in the mixture (Ar, He and N₂, respectively) are plotted across the length of the 1 cm box (averaged over the last 300 checkpoints). The error bars indicate one standard deviation.

collide with the walls – the gas does indeed cool and equilibrate with the walls (see Figure 1). The density of the gas in the box is found to be uniform when accounting for statistical scatter (which is below 1%), and the temperature of the gas also reaches the expected value of 300 K (see Figure 2). While it was not deemed appropriate for a test system of this size, one could increase the number of trajectories propagated to improve the sampling of the background gas and decrease the uncertainty in the resulting output.

4.2. Gas mixtures

One powerful and important feature of the SCMFD method is the ability to efficiently model systems where there are several species present in very different concentrations. A mixture of three gases, Ar, He and N₂, is chosen and these gases are simulated at vastly different densities – spanning six orders of magnitude. A 1 cm box is chosen to contain the gases, to obtain fast convergence. All gases are initially at 300 K, with all other input parameters as set out in Table 2. This model gas mixture system is especially well suited for benchmarking how accurately SCMFD can reproduce quantities such as mean free path,

mean free time and collision rate, as the resulting values can be directly compared to those derived analytically (see [2] for details on the derivation).

The mean free path for particle p is established from

$$\lambda_p = \left(\sum_{q=1}^{n_p} n_q \frac{\langle \sigma_{pq} v_r^{pq} \rangle}{\langle |v| \rangle} \right)^{-1} = \left(\sum_{q=1}^{n_p} n_q \sigma_{pq} \sqrt{\frac{m_r}{m_p}} \right)^{-1}, \quad (8)$$

with the mean free time given by

$$t_{\text{MFT}} = \frac{\lambda_p}{\langle |v_p| \rangle} = \frac{\lambda_p}{\sqrt{\frac{8kT}{m_p \pi}}}. \quad (9)$$

The number of collisions that occur per unit time and unit volume is dependent on the type of collision partner; when the collision partners are of the same species (i.e. both particle type p),

$$v_{pp} = \frac{1}{2} n_p^2 \sigma_{pp} \langle |v_r^{pp}| \rangle. \quad (10)$$

When the collision involves two different species (i.e. between particle types p and q), then

$$v_{pq} = n_p n_q \sigma_{pq} \langle |v_r^{pq}| \rangle. \quad (11)$$

The density of all three gases in the box is found to be uniform, with the temperature and velocity of each

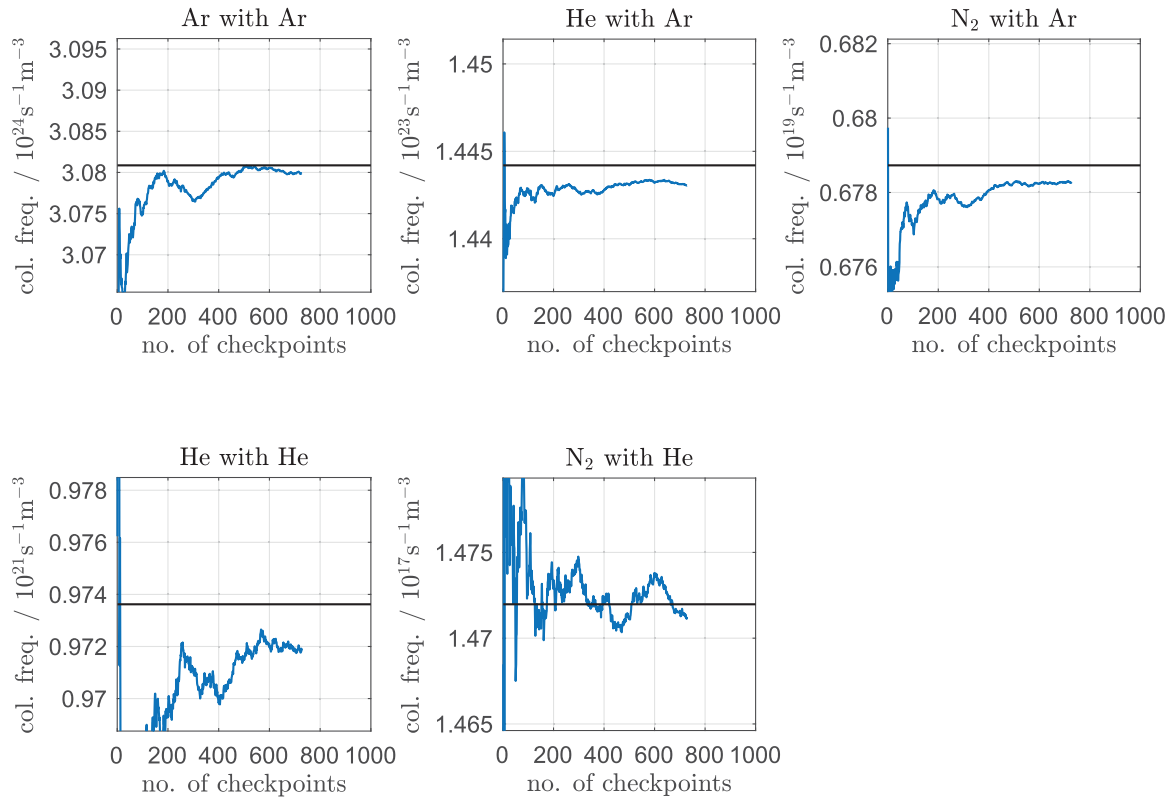


Figure 4. The calculated frequency of collisions between the different collision partner combinations is plotted. The horizontal line in each plot indicates the collision frequency predicted by theory.

species converged and stable (see Figure 3). A numerical estimate of how well the simulation reproduces the properties expected from a quantitative theoretical treatment of the system can be derived from considering the mean free path, mean free time and collision frequency per unit volume. This information can be straightforwardly obtained from the simulation, although it is worth commenting on the collision frequency calculations. The frequency of collisions between two different species can be obtained in two ways – by following the trajectories of species 1 and monitoring interactions with the mean field gas of species 2, or by following species 2 and tracking its interactions with the mean field gas of species 1. While either approach will yield the same result, it is much more efficient to follow the lower density particle, species 2 (He), as it propagates through the high-density field of species 1 (Ar), as far more collisions of interest occur per trajectory than would take place in the reverse situation. The collision frequency between each of the gases

is shown in Figure 4, with the exception of N_2-N_2 collisions (species 3 colliding with species 3). This is because no N_2-N_2 collisions took place over the length of the simulation – as expected, given that the mean free time for such collisions is approximately 13 s. The simulated collision frequencies between all other collision partners, Ar–Ar, Ar–He, Ar– N_2 , He–He, and He– N_2 , are in excellent agreement with theoretical values (within 0.5%). As Table 3 shows, the mean free path and mean free time established from the simulations are also in line with the expected theoretical values.

4.3. Couette flow

The flow of a gas between two parallel surfaces, where one surface is stationary and the other is moving with a fixed velocity, can be described as a Couette flow. To demonstrate the versatility and robustness of the SCMFD model, we simulate the Couette flow of Ar gas and compare

Table 3. The mean free path and mean free time between collisions for all three types of particles in the gas mixture.

Particle type	Theoretical value	SCMFD value	Relative deviation
<i>Mean free path</i>			
1. Ar	6.323×10^{-3} m	$6.33 \pm 0.06 \times 10^{-3}$ m	-4.67×10^{-4}
2. He	8.609×10^{-3} m	$8.62 \pm 0.05 \times 10^{-3}$ m	-6.82×10^{-4}
3. N_2	6.868×10^{-3} m	$6.87 \pm 0.06 \times 10^{-3}$ m	-8.24×10^{-4}
<i>Mean free time</i>			
1. Ar	1.5857×10^{-5} s	$1.586 \pm 0.014 \times 10^{-5}$ s	-4.66×10^{-4}
2. He	6.832×10^{-6} s	$6.84 \pm 0.05 \times 10^{-6}$ s	-8.54×10^{-4}
3. N_2	1.442×10^{-5} s	$1.443 \pm 0.012 \times 10^{-5}$ s	-7.25×10^{-4}

Table 4. Input parameters adopted to simulate the Couette flow of ar gas.

Parameter	Value	Description
n_p	1	Number of different particles
$n_{\text{background}}$	2000	Number of velocities used to describe the background gas per subcell
	100,000	Total number of trajectories
	200,000,000	Number of timesteps per trajectory
δt	5×10^{-8} s	Timestep
a	10	After every a trajectories a checkpoint is written and the background gas is updated
	0.1	Relative amount of the background gas information (density and velocity distribution) that is updated
<i>Wall specifications</i>		
dist	1 m	Distance between walls (x direction)
n_{cell}	200	Number of subcells
v_y^1	0 m s^{-1}	Velocity in y direction of left wall
v_z^1	0 m s^{-1}	Velocity in z direction of left wall
T_1	273 K	Temperature of left wall
v_y^2	300 m s^{-1}	Velocity in y direction of right wall
v_z^2	0 m s^{-1}	Velocity in z direction of right wall
T_2	273 K	Temperature of right wall
<i>Particle parameters</i>		
	1	Number of particles per trajectory of this particle type
d_{ref}	4.17×10^{-10}	Reference diameter
T_{ref}	273 K	Reference temperature
ω	0.81	Reference viscosity index
m	39.948 u	Mass of particle
n_p	10^{20} m^{-3}	Number density of particle

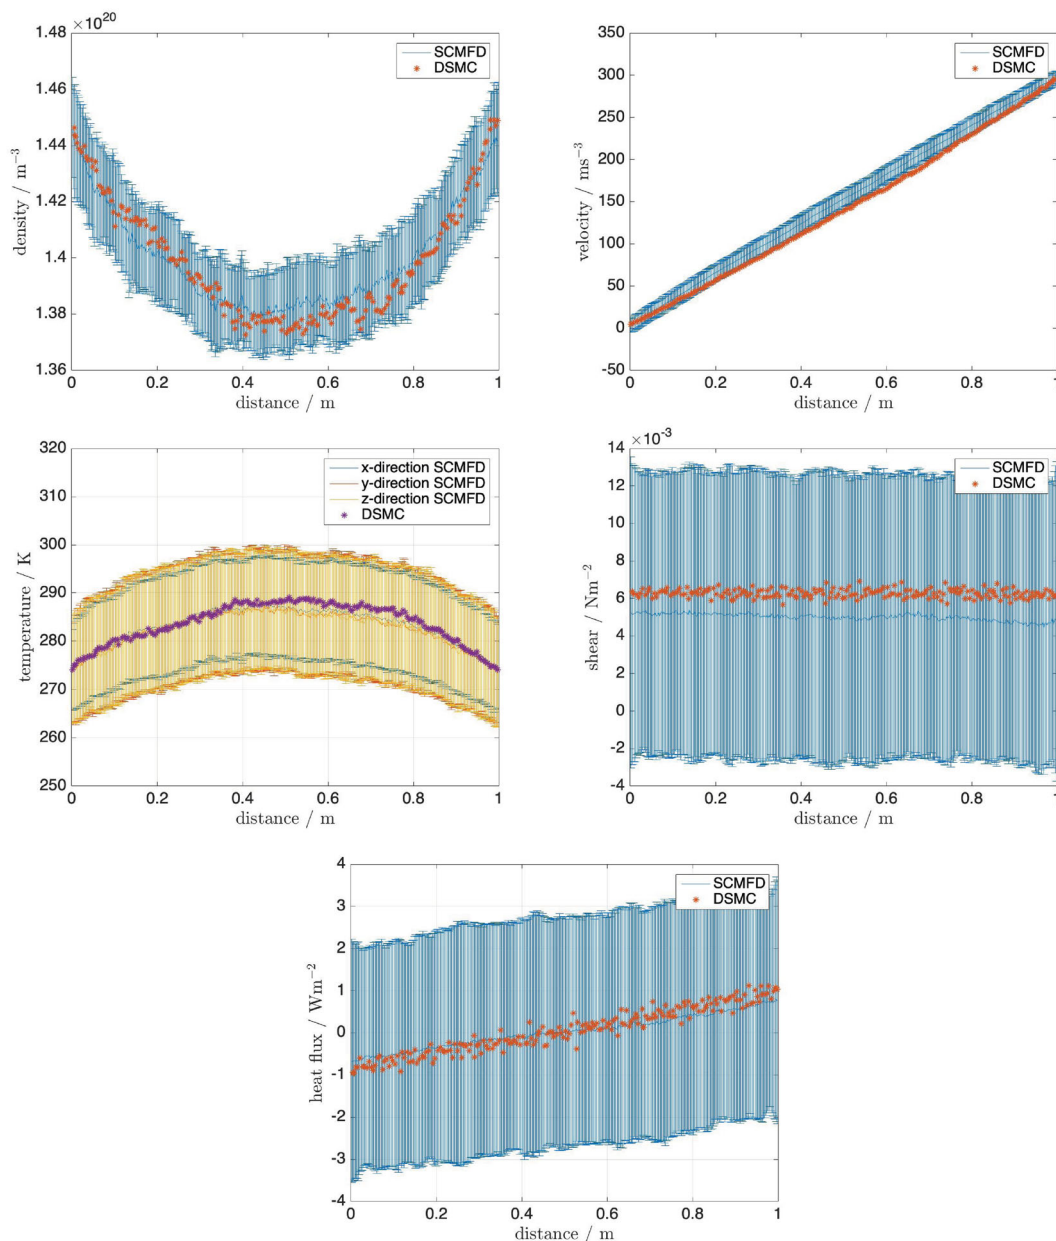


Figure 5. The density, velocity, temperature, shear stress and heat flux profiles calculated for a Couette flow of Ar gas, using the SCMFD model introduced in this work and Bird's DSMC model [2]. The error bars represent one standard deviation.

the results to existing measurements for this well-studied system. The input parameters and properties of the simulated system are set out in Table 4.

The ability of our SCMFD model to describe the Couette flow of Ar can be assessed in a number of ways. In the first instance, the density, velocity, temperature, shear stress and heat flux profiles are directly compared to those generated by Bird's DSMC model [2] in Figure 5. In all cases, the distribution calculated by SCMFD is in close agreement with the reference data generated by Bird's DSMC approach. There is a quantitative agreement in the heat flux profiles generated by the two models, and this is a property which DSMC is acknowledged to describe well

[2]. There is a very close agreement in the density, velocity and temperature profiles generated by the two methods. Furthermore, the SCMFD model yields profiles that are more symmetric across the distance between the two surfaces – and thus, it could be argued, achieves a more physically accurate description of the system. The trend in the shear stress profile for the DSMC and SCMFD models is in agreement, but the SCMFD value is consistently approximately 15% lower. The DSMC shear stress distribution corresponds to a coefficient of viscosity for Ar of $2.13 \times 10^{-5} \text{ N s m}^{-2}$, which is somewhat higher than the nominal value of $2.117 \times 10^{-5} \text{ N s m}^{-2}$ (at 273 K) but lower than the $2.18 \times 10^{-5} \text{ N s m}^{-2}$ value expected under

these conditions (as the average temperature across the flow is 283 K). Bird himself notes that the discrepancy between the DSMC model and the observed viscosity of Ar is ‘rather larger than would be expected’ [2].

To evaluate the performance of the SCMFD model in calculating the viscosity coefficient of argon, we must first consider the conditions of the system in more detail. Chapman–Enskog theory allows the viscosity of a gas to be derived from the Boltzmann equation, providing a continuum model of the system as a perturbed ideal flow [6]. Chapman–Enskog theory is best applied to ideal gas systems at high densities; the further a system is from such a state, the less reliable the predicted viscosity becomes. The Knudsen number for the conditions in the simulated Couette flow is 0.00925 [2], corresponding to a density where Chapman’s continuum model begins to break down. Lowering the density of the gas increases the Knudsen number, and the system transitions into a regime where the principles of free molecular flow apply: collisions occur far more frequently with the wall than with other gas particles, leading to the limiting case where the properties of the system are fully determined by the properties of the wall. Under the conditions of molecular free flow, the viscosity coefficient is proportional to the number density. As such, we expect the viscosity of argon in the Couette flow simulation to be directly proportional to the number of gas particles in the system at low densities. When the density of the system becomes sufficiently high, we expect the viscosity to transition from the molecular free flow behaviour to follow the Chapman–Enskog trend line. As Figure 6 shows, the viscosity coefficient of argon as calculated from the SCMFD simulation follows the molecular free flow prediction at low densities, and exhibits Chapman–Enskog behaviour at high densities. In the transition region, corresponding to densities of approximately 10^{20} – 10^{21} particles per m^{-3} in this system, the viscosity calculated by SCMFD is lower than that predicted by both molecular free flow and Chapman–Enskog. However, it is not immediately clear whether either model (molecular free flow or Chapman–Enskog) can be accurately applied in this regime. It is possible that SCMFD might underestimate correlations between velocity components, because these correlations are calculated by following the collisions of single particles. Alternatively, the traditional DSMC approach might overcorrelate these quantities, as the collisions of relatively few particles are taken as being representative of the whole system. Further work is needed to establish what the expected behaviour is in this transition region before we can quantitatively assess the performance of SCMFD for simulating Couette flow over a vast range of densities. For the moment, we can confidently assert that SCMFD accurately simulates the

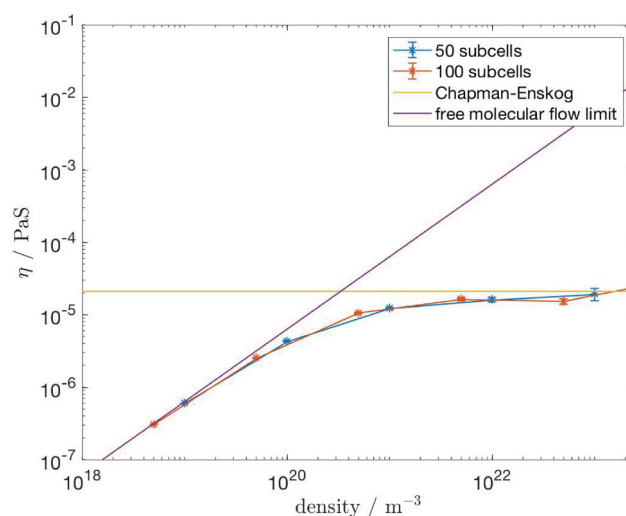


Figure 6. The viscosity coefficient of argon is plotted against the density of the system, calculated using the SCMFD approach with a 1 cm box divided into 50 (blue) and 100 (red) subcells. The behaviour expected from the molecular free flow (purple) and Chapman–Enskog (yellow) theories is also shown. The error bars indicate the standard error of the mean (colour online).

Couette flow of argon at both low densities and high densities.

5. Discussion and conclusion

The benefits and potential applications of an SCMFD model have been clearly demonstrated for several important systems. In particular, the use of a mean field approximation means that gas systems where there are vast differences in the densities of the components can be efficiently – and accurately – modelled. This will facilitate the simulation of, for example, the collisions of a seed gas in a buffer gas cell, where treatment with traditional DSMC methods is unfeasible for all but the smallest cell sizes.

There are several improvements that we will seek to address in the future. At present, the initial configuration and parameters of the system must be guessed such that convergence is achieved; if the initial inputs are too far from the physical properties of the system, convergence can be slow. It is also possible for the system to get ‘trapped’ in a stationary solution that is not, in fact, an equilibrium configuration, owing to the cancellation of certain terms. (For example, if there is an equal and opposite transfer of energy from collisions between the particles and the walls during the updating procedure, this can mimic the zero change in energy expected from a converged system.) A more involved check for convergence is being developed to avoid such situations. One key requirement of the current SCMFD model is that the

system is stationary. This means considering only closed, equilibrated systems or systems with an equal influx and outflux of particles. While it should be possible to generalise our approach to treat time-dependent systems, a sophisticated procedure will be required to manage the large amount of data generated at every timestep (on the order of 100 MB).

Acknowledgments

The authors wish to dedicate this paper to Professor Tim Softley, to thank him for his ongoing support and guidance. O.S. thanks H. Antila and M. Miettinen for their helpful discussions. Supporting data can be obtained from the Oxford Research Archive.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

B.R.H. acknowledges funding provided by the Engineering and Physical Sciences Research Council (EPSRC) (projects EP/N004647/1 and EP/N032950/1) and the Christ Church Research Fund.

References

- [1] G.A. Bird, *Phys. Fluids* **6**, 1518 (1963).
- [2] G.A. Bird, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows* (Clarendon Press, Oxford, 1994).
- [3] O. Schullian, J. Loreau, N. Vaeck, A. van der Avoird, B.R. Heazlewood, C.J. Rennick and T.P. Softley, *Mol. Phys.* **113**, 3972 (2015).
- [4] M.J. Doppelbauer, O. Schullian, J. Loreau, N. Vaeck, A. van der Avoird, C.J. Rennick, T.P. Softley and B.R. Heazlewood, *J. Chem. Phys.* **146**, 044302 (2017).
- [5] N.R. Hutzler, H.-I Lu and J.M. Doyle, *Chem. Rev.* **112**, 4803 (2012).
- [6] S. Chapman and T.G. Cowling, *The Mathematical Theory of Non-Uniform Gases: An Account of the Kinetic Theory of Viscosity, Thermal Conduction and Diffusion in Gases*, 3rd ed. (Cambridge University Press, Cambridge, 1970).

Appendix

A.1 Basic principles of self-consistent mean field DSMC (SCMFD)

Let $\rho(x, v)$ be the density distribution function of a stationary system. An initial distribution is chosen, and trajectories are subsequently calculated, from which a new density arises. This process is repeated as long as the density changes – that is, until a stationary state is reached.

A system (propagated in time) is the set of trajectories $\{q_i(t)\}$. From a set of initial parameters, (x_i^0, v_i^0) , a particle is propagated through the mean field, giving rise to a new set of trajectories. Taking the average of an observable, such as the density, with respect to these trajectories yields a set of

distribution functions,

$$\rho'(x, v) = \frac{1}{T} \int dt \sum_i \delta(q - \mathcal{P}_t^\rho(x_i^0, v_i^0)) \delta\left(v - \frac{d}{dt} \mathcal{P}_t^\rho(x_i^0, v_i^0)\right). \quad (\text{A1})$$

When a sufficient number of trajectories have been calculated, and the system reaches a stationary state, then $\rho' = \rho$. By calculating a series of fields, where ρ^i is established from ρ^{i-1} , one can arrive at a stationary point such that the series converges towards ρ . As with the DSMC approach, achieving convergence with the SCMFD method implies that a solution to the Boltzmann equation has been found.

A.2 Calculation of the density

The calculation of ρ using Equation (A1) can be tedious. A more straightforward approach can be taken if one assumes that we can express the distribution function as the product of the number density $n(x)$ and a velocity distribution function,

$$\rho(x, v) = n(x) \cdot \rho^v(v; x), \quad (\text{A2})$$

normalised such that $\int_V \rho^v dv \rho(v; x) = 1$. In a given simulation, ρ^v is easily obtained by saving the velocity of a trajectory at regular intervals. (One could also use a functional form for ρ^v instead of a long list of velocities. We opted not to take this approach, as we did not wish to make assumptions about what functional form would best reproduce the velocity distribution under a given set of conditions. However, using a functional form to describe ρ^v would depend on fewer parameters; it would save a lot of memory space and potentially be beneficial for higher dimensional systems.)

The estimation of $n(x)$ is more challenging, because it depends both on the way trajectories are calculated and on the boundary conditions. As the number of particles in a given volume is a time-average, we introduce the function

$$\theta_{ij}(t) = \begin{cases} 1 & \text{if particle } j \text{ is in subcell } i \text{ at time } t, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A3})$$

In this way,

$$V_i n_i(t) = \sum_j \theta_{ij}(t) \quad (\text{A4})$$

gives the instantaneous density. We are interested in the time-averaged density given by

$$V_i n_i = \frac{1}{T} \sum_j \int_0^T dt \theta_{ij}(t), \quad (\text{A5})$$

where the summands are the average time a particle spends in a given subcell. Given N particles, and defining the average time a particle spends in subcell i as $\langle dt_i \rangle$, then $V_i n_i$ can be expressed as

$$V_i n_i = \frac{1}{T} N \langle dt_i \rangle. \quad (\text{A6})$$

Finally, the value of $\langle dt_i \rangle$ is approximated as

$$\langle dt_i \rangle = \frac{1}{N_s} N_i dt. \quad (\text{A7})$$

Recalling that N_s is the number of particles that must be simulated to accurately reproduce the properties of the system, with

N_i is the number of times a particle is located in subcell i , one arrives at an expression for the number density

$$n_i = \frac{NN_i dt}{V_i TN_s}. \quad (\text{A8})$$

A.3 Establishing the collision probability

Two particles, labelled 1 and 2, are transformed into the frame where $v_r = v'_1 = v_1 - v_2$ and $v'_2 = 0$ (i.e. particle 2 is stationary). In this frame of reference, after timestep dt particle 1 flies out a volume

$$V(v_r) = dt|v_r|\sigma(v_r). \quad (\text{A9})$$

A collision is deemed to have occurred if particle 2 (which is stationary in this frame) was located in the flown out volume. The ratio of the flown out volume to the total volume gives the overall probability of a collision having occurred. For small values of $V(v_r)$ this can be approximated by

$$P_c(v_{r,i}) = V(v_r)/V, \quad (\text{A10})$$

or, conversely, the probability of no collision occurring can be expressed as

$$P(v_{r,i}) = 1 - P_c(v_{r,i}). \quad (\text{A11})$$