Kinetic properties of particle-in-cell simulations compromised by Monte Carlo collisions

M. M. Turner*

School of Physical Sciences and National Centre for Plasma Science and Technology, Dublin City University, Dublin 9, Ireland

Abstract

The particle-in-cell method with Monte Carlo collisions is frequently employed when a detailed kinetic simulation of a weakly collisional plasma is required. In such cases, one usually desires, *inter alia*, an accurate calculation of the particle distribution functions in velocity space. However, velocity space diffusion affects most, perhaps all, kinetic simulations to some degree, leading to numerical thermalization, and consequently distortion of the velocity distribution functions, among other undesirable effects. The rate of such thermalization can be considered a figure of merit for kinetic simulations. This paper shows that, contrary to previous assumption, the addition of Monte Carlo collisions to a one-dimensional particle-in-cell simulation seriously degrades certain properties of the simulation. In particular, the thermalization time can be reduced by as much as three orders of magnitude. This effect makes obtaining a strictly converged simulation results difficult in many cases of practical interest.

*Electronic address: `miles.turner@dcu.ie`

I. INTRODUCTION

A wide range of methods is available for simulating the behavior of plasmas using computers. These methods vary in the level of detail that they represent, and there is usually some relationship between the level of detail and the computational cost: A more detailed simulation will normally be more costly. Consequently, there is inevitably some compromise to be made between physical fidelity and computational expense. Some such compromises are more critical than others. For example, in simulations of low-temperature collisional plasmas, such as those used in many technological applications, it is generally agreed that the Boltzmann equation is acceptable as the fundamental description of the plasma, so the central problem is to compute an accurate solution of that equation. However, the Boltzmann equation is usually solved by methods involving approximations. In most cases, and especially in fluid models or hybrid models with fluid elements, critical compromises are made in the velocity space representation and in the coupling between the field equations and the Boltzmann equation. Such compromises are critical in the sense that they entail approximations that are difficult to verify either *a priori* (by reference to the physical parameters of the case) or *a posteriori* (by reference to internal consistency conditions of the simulation). As a consequence, it is difficult to be clear about the expected level of accuracy. A widely used approach to this difficulty is to employ a hierarchy of simulations. For example, one might carry out a reference calculation using a direct simulation of the Boltzmann equation that is presumed to be highly accurate, and then use the results of such a calculation as a benchmark for other more economical simulation procedures. At the present time, this type of benchmarking is probably one of the more important applications of direct Boltzmann solvers, because the computational cost of such solvers remains excessive for many routine calculations, especially calculations involving the complicated geometrical configurations that are necessary in many engineering studies. There are other reasons for adopting direct Boltzmann solvers, but a common feature of all applications of such solvers is a desire for an accurate representation of the physics. Consequently, one wants to be able to make strong statements about the accuracy to be expected in any given circumstances, and ideally such statements should be related to internal consistency conditions of the simulation that can be readily verified. Such conditions are usually derived from some kind of theory of the simulation algorithm, and are stated in terms of numerical parameters, such

as discretization parameters in space and time, and macroscopic plasma parameters, such as density and temperature. Note that we do not refer here to simulation errors caused by inaccurate or unavailable input data, such as cross sections or rate constants. This category of error is often important, especially when comparison with experiment is the objective, but they are not the subject of the present discussion. Concern about the issues discussed above is not new. Previous responses have included detailed analysis of the well-posedness of fluid models [1] (leading to some criticism of certain earlier works), suggestions for validation procedures [2] and a benchmark comparison of a number of different models [3]. It might be concluded that this is a fairly sparse literature, and that these issues of convergence and correctness have not received the attention that they perhaps deserve, especially in the case of fully kinetic models that are intended to be highly accurate.

In the remainder of this paper, we discuss problems that arise in Monte Carlo solutions of the Boltzmann equation, specifically, the particle-in-cell method. We refer in particular to particle-in-cell simulations in which an explicit Monte Carlo method is used to represent collisional interactions. Such simulations employ three numerical parameters, namely the cell size, Δx, the time step Δt, and the number of particles per Debye length, N_D. Confidence in the accuracy of simulations using this method requires that the simulation results should be insensitive to the choice of the numerical parameters, for proper choices of those parameters. We will show that the previously articulated theory of particle in cell simulations does not hold for this modification of the classical particle-in-cell procedure, in that the kinetic properties of the simulation are changed in an unwelcome way by the introduction of the Monte Carlo collision operator. In particular, insensitivity of the results with respect to N_D then occurs only with many more particles that is suggested by commonly accepted rules of thumb. This result has relevance to anyone using particle-in-cell simulation to model collisional plasmas, especially in the benchmark context discussed above.

II. BACKGROUND

The behavior of a typical weakly-coupled plasma is well approximated by a set of classical point particles, with motion subject to Newton's equations and with forces derived from Maxwell's field equations. For most analytical and computational purposes, this representation is intractable, and simplifications must be sought immediately. A usual procedure is to

approximate the set of particles by a continuum distribution function. Formally, this step involves some kind of averaging procedure to remove certain spatial and temporal frequencies that are associated with the graininess of the particle description. The details of this averaging are not relevant here, but it is important that an element of approximation is introduced in passing from the particle representation to the continuum distribution function. The theoretical procedure involves a small parameter, the so-called "plasma parameter," which is the inverse of the number of particles contained in a Debye sphere. The collisionless Boltzmann equation, also known as the Vlasov equation, emerges from the averaging procedure as the approximation at zero order in the plasma parameter. In a spatially uniform plasma, an arbitrary function of the velocity coordinates satisfies this collisionless Boltzmann equation, but this is not a correct description of a real plasma, where there is some finite Coulomb collision frequency, as a consequence of which the velocity space distribution relaxes toward a Maxwellian [4–6]. These collisional effects are discarded at the zeroth order approximation in the plasma parameter. In the first order approximation, there appears the so-called Lennard-Guernsey-Balescu collision operator [5, 7], which is a representation of the Coulomb collision effects. In addition to these collisional phenomena involving charged particles, a description of collisions between charged and neutral particles is also needed, and for this purpose the usual Boltzmann collision operator is appropriate [5, 8]. For most practical purposes, then, the fundamental description of the charged particles in a low temperature plasma is the Boltzmann equation, with separate collision operators to describe Coulomb collisions between charged particles and collisions between charged particles and neutrals. As the preceding discussion indicates, this is a mesoscopic and not a microscopic description, because of the spatio-temporal averaging involved in deriving the Boltzmann equation.

The main idea of the particle-in-cell method is to model the plasma using a set of so-called superparticles, which are subject to Newtonian or relativistic equations of motion, where the fields are calculated by solving Maxwell's equations using source terms derived from the particle positions and velocities. The particles move with a finite time step Δt, and the fields are calculated at the boundaries of a set of finite sized cells of width Δx. In the simplest case of a one-dimensional electrostatic plasma with non-relativistic particles, the central equations are those for the particle position, x, and velocity, v:

$$\frac{v_{i,n+1/2} - v_{i,n-1/2}}{\Delta t} = \frac{q}{m} E(x_{i,n}) \tag{1}$$

$$\frac{x_{i,n+1} - x_{i,n}}{\Delta t} \;=\; v_{i,n+1/2} \tag{2}$$

where the i is a superparticle index, n is a time level, q is the particle charge, m is the particle mass and E is the electric field. The corresponding field equation is

$$\frac{\phi_{j+1,n} - 2\phi_{j,n} + \phi_{j-1,n}}{\Delta x^2} = \frac{\rho_{j,n}}{\epsilon_0}, \tag{3}$$

where j is an index for spatial cell boundaries, ϕ is the electrostatic potential, and ρ is the charge density. The algorithm is completed by a prescription for determining the charge density from the set of particle positions, and the electric field at the particle positions from the potential defined at the cell boundaries.

This particle-in-cell method [9, 10] has certain obvious structural similarities to a direct simulation of the microscopic model of the plasma outlined above, but this is essentially misleading: Particle-in-cell simulation is properly understood to be a Monte Carlo solution of the Boltzmann equation (or the Vlasov equation). The physical content of a particle-in-cell simulation is in principle the same as that of solution of the Boltzmann equation by finite differences or any other method that might be proposed, despite the fact that the formulation of the particle-in-cell algorithm need contain no explicit reference to the Boltzmann equation. Consequently, and again in principle, the only consideration entering into a choice between the particle-in-cell method and other competing algorithms is numerical efficiency. The main alternatives to the particle-in-cell method are various finite difference methods, such as the well-known "convective scheme" [11, 12]. Both the convective scheme and the particle-in-cell method require that finite intervals in space and time be specified as part of the algorithm. In the particle-in-cell method, these intervals must resolve the Debye length and the period of the plasma frequency, subject to well-known criteria. No such criteria seem to be generally accepted for the convective scheme, but it seems reasonable to suppose that similar limitations exist, and that these will point to numerical parameters of the same order of magnitude as in the particle-in-cell simulation. In addition, the number of particles is a parameter in a particle-in-cell simulation, and a finite interval in velocity space is a parameter in the convective scheme. There are no generally accepted criteria for choosing these parameters. It is consequently not easy to evaluate the relative efficiency of these methods on paper, and there are at the present time no detailed practical comparisons available. When the number of real space dimensions is greater than one, it is usually

conceded that Monte Carlo methods enjoy an advantage over finite difference methods, essentially because a finite difference method entails an explicit computational mesh spanning all of the real and velocity space dimensions that are represented, and this leads to difficulties when the number of dimensions is large. Adaptive meshing techniques have produced improvements in the performance of finite difference solutions of the Vlasov equation [13], but it is presently not clear that these methods can be extended to the collisional case. So it seems that particle methods are likely to remain the direct Boltzmann solver of choice for most problems in the immediate future, unless there are important innovations in some other type of direct solver.

The classical particle-in-cell procedure applies to a collisionless plasma. When collisional effects are to be included, they must be added via an explicit collision operator, which is commonly some kind of Monte Carlo procedure [14]. In particular, Coulomb collisions are not included unless they are explicitly represented in the collision operator [15, 16]. However, as the discussion below shows, there are spurious numerical phenomena with effects resembling those of Coulomb collisions.

III. LIMITATIONS OF THE PARTICLE-IN-CELL METHOD

The best known limitations of the explicit particle-in-cell procedure applied to plasmas that are electrostatic and unmagnetized are the accuracy conditions, usually expressed as:

$$\frac{\lambda_D}{\Delta x} \gtrsim \mathrm{O}(1) \tag{4}$$

$$\omega_p \Delta t \lesssim O(1). \tag{5}$$

These well-known results constrain the time step Δt and the cell size Δx, which must resolve the highest temporal and spatial frequencies of the plasma, respectively associated with the plasma frequency ω_p and the Debye length λ_D. These results are consequences of analytical and empirical investigations of the particle-in-cell algorithm [9, 10], with the optimum values of the numerical constants depending somewhat on the aims in view. A third numerical parameter is the number of superparticles, which is normally characterized by the number of particles per cell, N_C, or the number of particles per Debye length, N_D, which is analagous the inverse plasma parameter. The latter is the more fundamental, and therefore the more useful, metric. Apart from these explicit numerical parameters, a

particle-in-cell simulation also entails a choice of superparticle shape, or weighting scheme, which is used to define,*e.g.* the charge density in terms of the superparticle positions. Most of the simulations in this paper are carried out using bilinear weighting, sometimes called the cloud-in-cell scheme. Although this is the most commonly used particle shape, other choices are possible, and the kinetic properties of the simulation depend appreciably on the shape factor [9, 10].

Apart from outright numerical instability, particle-in-cell simulations are subject essentially to two kinds of error. One of these is self-heating (or occasionally cooling), *i.e.*, in a simulation of a closed physical system, the total energy of the simulation will change with time. In a simulation of an open system, such as a discharge, one should ensure that this spurious heating rate is small compared to the physical heating rate. The dependence of the self-heating rate on the numerical parameters has been extensively discussed in earlier works [9, 10, 17]. The second kind of error is spurious relaxation of the distribution function toward a Maxwellian, an effect sometimes called "collisions," although to minimize confusion we will generally avoid this term in the present discussion. This also has been discussed in considerable detail [18–23]. However, the all these works discuss simulation of collisionless plasmas. The aim of the present work is to discuss the effect of explicitly introduced collisions on the rate of relaxation of the velocity distribution function.

As indicated above, in the absence of explicit collisions included by a Monte Carlo procedure, a particle-in-cell algorithm is formally a solution of the Vlasov equation. In the absence of spatial gradients and of externally imposed fields, the Vlasov equation for a quasi-neutral plasma is solved by an arbitrary function of velocity. Clearly, an ideal simulation should reproduce this feature, but in practice most, if not all, simulations will exhibit some spurious velocity space transport, meaning that any initial velocity distribution function will relax to a Maxwellian in a finite amount of time. In a particle-in-cell simulation, such relaxation is associated with the graininess of the particles, which introduces a spurious stochastic component into the electric field. This stochastic part of the field induces diffusion in velocity space, which in turn causes the particle velocity distribution to relax to a Maxwellian. This process is analogous to the relaxation of the velocity distribution functions of the particles in a real plasma brought about by Coulomb collisions, which is also caused by the underlying graininess of the real plasma. A crucial insight in the early development of particle-in-cell simulation occurred when it was realized that the use of finite

size particles could dramatically suppress fluctuations, and therefore inhibit spurious relaxation of the velocity distribution function [18]. This is why a particle-in-cell simulation with $N_D \sim 100$ may be a tolerable representation of a real plasma with $N_D \sim 10^6$ or more. This does not, of course, mean that spurious relaxation is never important, and kinetic properties of collisionless particle-in-cell simulations have been extensively studied [19–21, 24–27]. In the context of one-dimensional simulations, the most salient result of these enquiries is that the rate of relaxation of the velocity distribution is proportional to $1/N_D^2$, that is, there is no relaxation at first order in the plasma parameter. This result has been established both theoretically [20] and by observation of computer simulations. Despite the clarity of the theoretical result, most authors have been surprised by this behavior. It has been noted that all binary collisions in a strictly one-dimensional system are co-linear, and that a co-linear collision between two particles of equal mass either produces no change in the particle velocities or an exchange of velocities between the particles; no other outcome is consistent with the conservation laws, and neither outcome changes the distribution function in any way. However, the relevance of these facts to a system where interactions between particles do not appear to be restricted to binary collisions is not altogether transparent [9, 21], but fortunately this is an issue we need not pursue here.

Several procedures have been employed for quantifying the rate of relaxation of the velocity distribution function [21, 24, 27], and in general these lead to a relation of the form:

$$\omega_p \tau_R \propto N_D^2 \tag{6}$$

where τ_R is a characteristic relaxation time. The numerical constant in this equation, and the definition of τ_R, depend on the procedure that is adopted. A convenient and precise approach is the one suggested by Virtamo and Tuomisto [27]. We will discuss this approach in some detail, because we will adopt it below to study collisional plasmas. They assume that the thermalization of the velocity distribution is governed by an operator, O, with the Fokker-Planck form:

$$O \equiv \frac{1}{\tau_R} \frac{\partial}{\partial u} w \frac{\partial}{\partial u} \frac{1}{w} \tag{7}$$

where u is a velocity coordinate normalized to the thermal velocity $v_{\rm th}^2 = k_B T_e / m_e$, and

$$w = \frac{\exp(-u^2/2)}{\sqrt{2\pi}}. \tag{8}$$

This linear operator has eigenvalues and eigenfunctions given by

$$\lambda_l = -l/\tau_R \tag{9}$$

$$\psi_l(u) = w(u)\mathrm{He}_l(u), \tag{10}$$

where l is a integer and the functions He_l are Hermite polynomials [28]. These eigenfunctions satisfy the orthogonality relation:

$$\int_{-\infty}^{+\infty} du\, w^{-1}\psi_l(u)\psi_m(u) = l!\delta_{l,m}. \tag{11}$$

It is possible to express any distribution function $f(u,t)$ in terms of an expansion in these eigenfunctions so that

$$f(u,t) = \sum_{l=0}^{\infty} a_l(t)\psi_l(u). \tag{12}$$

Moreover, if the processes in the simulation are correctly described by the operator in Eqn. (7), then the coefficients in this expansion evolve in time as:

$$a_l(t) = a_l(0)\exp(-lt/\tau_R). \tag{13}$$

Note that ψ_0 is the Maxwellian distribution, and this eigenfunction does not decay in time, so that in this formulation all $f(u,0)$ relax to a Maxwellian, as they should. Additional details of this procedure are in [27], but, essentially, further use of the orthogonality relations enables the coefficients $a_l(t)$ to be extracted from a simulation. In practice, it is found that these coefficients indeed decay approximately exponentially, suggesting that Eqn. (7) adequately approximates the real relaxation mechanism. In [27], the numerical coefficient in Eqn. (6) is determined to be 28.6; other procedures lead to values of the same order of magnitude [21, 24].

None of the works so far mentioned has discussed the effects of Monte Carlo collisions on the kinetic properties of the simulation. Indeed, it seems to be tacitly and universally assumed that there are no such effects. In the next section we will show that this is an overly optimistic view, and that in fact the properties of the simulation are seriously degraded when Monte Carlo collisions are introduced.

IV. KINETIC PROPERTIES OF COLLISIONAL PARTICLE SIMULATIONS

In this section we examine the effect on Monte Carlo collisions on the kinetic properties of a one-dimensional particle-in-cell simulation. The simulation procedure is conventional,

and uses a first order weighting scheme, the so-called cloud-in-cell method [9, 10]. We have employed periodic boundary conditions, and no external forcing fields are applied. Only the electrons are treated in the simulations; a neutralizing positive background charge is assumed. The initial conditions are uniform in real space, and the particles are distributed in velocity space as

$$f(u,0) = \begin{cases} 1/2\sqrt{3}, \text{when } |u| < 1 \\ 0, \text{ otherwise} \end{cases} \tag{14}$$

where $u = v/v_{\mathrm{v_t}}$ and $v_{\mathrm{th}}^2 = k_B T_e/m_e$ as above. This distribution has the same mean energy as a Maxwellian with temperature T_e. Collisions occur at a frequency ν independent of particle speed, with partners that are assumed to be of infinite mass, so that there is no change in kinetic energy caused by collisions. In order to preserve the strictly one-dimensional character of the simulation, all scattering is backward, so that no energy is coupled into velocity components normal to the spatially resolved axis. This could be regarded as isotropic scattering with frequency 2ν, since there is no change of energy or momentum in a collision with forward scattering in this model. A conventional Monte Carlo procedure was adopted [14], in which each superparticle is tested for a collision on each time step. The numerical parameters for the simulation were chosen to marginally satisfy the inequalities of Eqs. (5) and (4), while the length of the system was approximately $160\lambda_D$ and N_D was 102, unless otherwise stated. For these numerical parameters, the self-heating time, τ_H, is generally considerably larger than τ_R, so that the total energy of the system changes relatively little during the time of interest. Consequently, in a change of procedure relative to [27], we did not find it necessary to impose energy conservation artificially by rescaling the particle velocities during the course of each calculation. The relaxation time for the velocity distribution function is defined in the manner discussed above, by examining the decay of the amplitudes of the eigenmodes of the assumed collision operator, Eq. (7). As in [27], we examine the eigenmode amplitudes for $l = 4, 6, 8$ and 10. It can be shown that insofar as the conservation laws for momentum and energy hold, a_1 and a_2 must vanish. Moreover, the symmetry of our initial $f(u)$ with respect to u means that all the other anti-symmetric eigenmodes have zero amplitude, or at least that they are excited only by fluctuations, so that their decay rates are difficult to measure.

Fig. 6 shows that the decay of the mode amplitudes remains approximately exponential, even in the presence of Monte Carlo collisions. That figure also shows that the decay rate is

appreciably increased when Monte Carlo collisions take place. A more systematic overview of this effect is in fig. 2, which shows the relaxation time of the distribution function as a function of the normalized collision frequency, ν/ω_p. The values of τ_R shown here are an unweighted average of the values inferred from the decay rates of the four eigenmode amplitudes shown in fig. 6; these decay rates were found by a least squares fit of an exponential function, as shown in the examples in the figure. Normally, an upper limit on the time step in a particle-in-cell simulation with Monte Carlo collisions is suggested to be $\nu\Delta t < 0.1$ [14, 29]. Since one normally also seeks to satisfy the inequality given by Eq. (5), this implies $\nu/\omega_p < \mathrm{O}(1)$. This limit may be closely approached in practical situations. Fig. 2 shows that in such a case, the thermalization time may be almost three orders of magnitude less than one would infer from standard references [9, 10]. A further point of interest is the effect of the parameter N_D on the thermalization time when Monte Carlo collisions are present. Fig. 3 shows that τ_R scales linearly with N_D in this case, contrary to the collisionless situation where the scaling is with N_D^2. All these data are well described by the relation:

$$\omega_p \tau_R = \frac{34.4}{N_D^{-2} + 28.0 N_D^{-1}(\nu/\omega_p)}, \tag{15}$$

which is shown as the solid line in figs. 2 and 3. We note that the numerical factor of 28.6 determined in [27] is here found to be 34.4, suggesting that the present simulation has a slightly longer collisionless thermalization time. No reason for this difference has been identified. We have also examined the effect of the cell size, Δx, on the thermalization time. The results of these calculations, which appear in fig. 4, show that the thermalization time is insensitive to the cell size, as long as the accepted accuracy criterion given by Eq. (4) is respected. These changes in the thermalization time, τ_R, are accompanied by corresponding changes in the self-heating time. Since the simulation energy grows approximately linearly with time due to self-heating, τ_H is defined as the time taken for the kinetic energy of the simulation particles to double [9, 10]. Fig. 5 shows the change in τ_H as a function of collision frequency. There is a marked degradation of the quality of the simulation in this respect also as the Monte Carlo collision frequency increases.

Some of the calculations discussed above were repeated using a lower order interpolation scheme, namely, nearest grid point weighting [9, 10]. This method is distinctly inferior for particle simulations employed as Vlasov solvers [9, 10]. In particular, the heating time is increased in a manner not compensated by the slightly more economical computational

procedure. However, for highly collisional plasmas, $\nu/\omega_p > 10^{-2}$, we found that this zero order method has a thermalization time essentially indistinguishable from that of the first order method used in most of our calculations.

As the opening discussion suggested, the occurrence of thermalisation in a real or simulated plasma is intimately connected with the stochastic part of the electric field, which is described by the fluctuation spectrum. Consequently, one expects that a large change in the rate of thermalization will be accompanied by a corresponding change in the fluctuation spectrum. Measurements of the fluctuation spectra in the simulations discussed above indeed show large variations. Relative to a collection of uncharged particles, a plasma ordinarily exhibits reduced density fluctuations at frequencies below the plasma frequency and wavelengths longer than the Debye length. In the presence of collisions, this effect is less marked, presumably because collisions destroy the correlations between particles that collective effects create. This is presumably how Monte Carlo collisions produce an enhancement of the fluctuation spectrum and consequently an increase in the thermalization rate.

V. EXAMPLES

In the examples discussed in this section we examine the implications of the results of section IV for four practical situations, specifically, a radio-frequency discharge in argon under two different conditions, a direct current discharge in helium and a positive column discharge in oxygen. We show that there are difficulties with convergence in the first three of these cases. The simulation method was again conventional [14, 29], with the only unusual feature the inclusion of Coulomb collisions in the Monte Carlo collision operator using the procedure of [16]. As an economy, we considered only electron-electron collisions. Electron-ion and ion-ion collisions are unlikely to be of significance under the conditions considered below [30].

A. Radio Frequency Discharge in Argon

In this example we discuss the problem of simulating a radio frequency discharge in argon at low pressure. The experimental context for this example is the extensively discussed experimental demonstration that, at sufficiently low pressure, the electron energy distribution

function has a bi-Maxwellian character [31]. As the pressure is increased, this remarkable distribution function vanishes and is replaced by a less surprising form, qualitatively resembling a Druyvesteyn. There are many points of experimental and theoretical interest relating to this work, but the issue to which we wish to draw attention here is the suggestion made in the original work that electron-electron collisions are the primary mechanism controlling the temperature of the colder group of electrons. More specifically, it was suggested that at sufficiently low pressure, a group of colder electrons becomes trapped in a potential well at the centre of the discharge, where they are heated inefficiently by the radio frequency field. This effect is particularly marked in argon because of the Ramsauer-Townsend minimum in the elastic cross section for collisions between electrons and argon atoms. The proposal made in [31] was that electron-electron collisions between the cold and hot electrons is the dominant heating mechanism for low-energy electrons. However, there are a number of particle simulation studies [32–34], including one by the present author, that show reasonable agreement with this or similar experiments without including electron-electron collisions, and a larger number of studies of related problems, *e.g.*, [35–42], also neglecting electron-electron collisions. In this section, we re-examine these simulation results, with special attention to convergence with respect to N_D. In the cases where the two-temperature electron energy distribution function is exhibited, we show that electron-electron collisions are indeed important. Therefore, the good agreement obtained previously was essentially fortuitous, and arose because the numerical relaxation effects discussed above chanced to produce an approximately correct rate of relaxation.

The measurements referred to above were carried out in a discharge formed between a pair of parallel plates separated by 2 cm, with an applied voltage oscillating at 13.56 MHz. For the purposes of the present paper, we have carried out fresh simulations of this experiment, for two experimental conditions. One of these conditions is at 100 mTorr, which is in the pressure regime where a bi-Maxwellian electron energy distribution is observed experimentally. The other is at 500 mTorr, where a Druyvestyn-like electron energy distribution is observed. The simulations were carried out for a range of values of N_D, and in the case of the lower pressure, both with and without electron-electron collisions included explicitly through the Monte Carlo collision operator. Collisions between electrons and neutrals were modelling using essentially the cross sections of Tachibana [43], with adjustments to the low-energy part of the momentum transfer cross section from Pack *et al* [44], and with isotropic scattering

in the centre of mass frame. Ion-neutral collisions followed the recommendation of Phelps [45], with the anisotropic scattering cross section approximated by an isotropic scattering component and a backward scattering component.

The results of these calculations are shown in figs. 6 to 12. At 100 mTorr, experimental measurements show that the electron energy distribution function is bi-Maxwellian, with the majority of the electrons in a bulk group with a temperature of less than 1 eV, and a minority group of electrons, perhaps 10 % of the total, in a hot tail with a temperature of about 3 eV. These tail electrons are responsible for essentially all of the excitation and ionization that occurs, and they also absorb most of the power that is coupled into the discharge. As we indicated above, the bulk group of electrons are confined in a potential well around the mid-plane of the discharge, where they oscillate almost collisionlessly in a weak radio-frequency field. These conditions are moderately collisional, in that $\nu/\omega_p \sim 0.02$. The simulation results appearing in figs. 6 and 7 show that the density and temperature of the plasma at the discharge mid-plane vary appreciably as a function of the number of superparticles per Debye length, N_D, and that for large N_D, the results also depend on whether or not Coulomb collisions are included. In the absence of Coulomb collisions, the plasma is denser and colder. For similar N_D, these differences are almost entirely due to the cold group of electrons. Fig. 9 compares the electron energy distribution function for two cases that have approximately the same N_D, and otherwise differ only in the inclusion or otherwise of Coulomb collisions. The density and temperature of the hot electrons are essentially identical in both cases, but there is a large difference in these parameters for the bulk group. Qualitatively, these observations are easy to understand. The density and temperature of the bulk electrons depends on a subtle balance of cooling and heating processes, where the cooling processes are evaporation and elastic recoil, and the heating processes are an interaction with the radio-frequency field and Coulomb collisions with the hot group of electrons. Since Coulomb collisions are not the only heating mechanism, we expect to observe convergence as a function of N_D whether such collisions are included or not, and figs. 6 and 7 show at least evidence that this convergence takes place. Coulomb collisions only shift the equilibrium values by a substantial amount. Other discharge parameters are much less affected by these convergence issues. For instance, the data in fig. 8 show that the time-averaged ion flux arriving at an electrode is almost independent of N_D and is not appreciably influenced by Coulomb collisions.

At 500 mTorr, the collision frequency is much larger, and $\nu/\omega_p \sim 0.3$. Coulomb collisions are not expected to be important, but the effects of the superparticle density can still be seen. In this case, the plasma density is almost independent of N_D, but there is a substantial change in the temperature, as can be seen in figs 11 and 10. Fig. 12 shows that this is attributable to a distortion of the electron energy distribution function in the energy range below the inelastic threshold, where spurious Maxwellianization takes place. Clearly, the gross properties of the plasma are insensitive to the shape of the distribution function in this region.

B. Direct Current Discharge in Helium

The direct current discharge is a classic problem of low-temperature plasma physics, which has been attracting interest for more than a century [46]. Such a discharge has a complicated spatial structure, especially in the vicinity of the cathode, where a number of regions can be distinguished and have been named. For the present discussion, the important point is that there is a region of positive space charge with large electric fields adjacent to the cathode. These fields accelerate ions toward the cathode, and electrons in the opposite direction. Electrons enter the cathode region after being emitted from the cathode surface by any of several mechanisms, including positive ion impact, fast neutral atom impact, and the photoelectric effect. Since electron impact ionization is rapid under the conditions in the cathode region, the electron density rises rapidly as one moves away from the cathode surface, and at some point a transition occurs to a region where there is a plasma, with an approximately equal number of electrons and positive ions. This is known as the negative glow. Relatively recent work [47, 48] has shown that the negative glow contains a remarkably high density ($\sim 10^{17}\mathrm{m}^{-3}$) of cold ($\sim 0.1$ eV) electrons, which persist because of the presence of a potential well, with an associated reversal of the electric field. The temperature of these electrons is established by a balance between cooling by elastic collisions with neutral gas atoms, and heating by Coulomb collisions with untrapped electrons. Experimentally, these phenomena can be studied in an approximately one-dimensional context if the distance between the anode and the cathode is small compared to the other dimensions of the electrodes, and such a scenario is readily reproduced in a one-dimensional simulation. In the present context, the point of interest is the dependence of the density and temperature

of the trapped electrons on the number of super-particles.

Our simulation example is a discharge formed between a pair of plane electrodes separated by 1 cm. Helium gas with a density of 3.2×10^{22} m^{-3} and a temperature of 300 K fills the space between the electrodes, corresponding to a pressure of about 1 Torr. All the electrons emitted from the cathode are assumed to be produced by ion impacts, and the emission coefficient has the constant value of 0.2. Collisions of helium ions with the background gas are described by two elastic processes, with isotropic and backward scattering in the centre of mass frame, as in the case of argon discussed above [45, 49]. Electron impact processes are elastic scattering, lumped excitation, and ionization [49]. Under these conditions, $\nu\Delta t \sim 0.03$ and $\nu/\omega_p \sim 0.1$, so this is a collisional plasma, and in view of the results discussed above, convergence with respect to the superparticle number is expected to be difficult. This is the case, as figs. 13 and 14 show. Indeed, it is not clear that satisfactory convergence has been achieved even at the right hand side of these figures, where approximately 10^7 superparticles were employed. (Convergence is not expected for the case without Coulomb collisions, since it is not clear that there is any other physical untrapping mechanism.) Not all of the discharge parameters are affected, however. Fig. 15 shows that the ion flux collected at the cathode is highly insensitive to either the superparticle density or the inclusion or otherwise of electron-electron collisions.

C. Positive Column in Oxygen

This final example differs from the previous ones in dealing with a discharge in a molecular gas, namely oxygen. In this case, there is no significant range of electron energies where elastic collisions are the dominant energy loss mechanism, because vibrationally and rotationally inelastic collisions can occur with threshold energies as low as 0.02 eV. Moreover, the example we discuss is a model positive column where the sustaining electric field penetrates the entire volume of the discharge, so that electrons cannot become trapped in spatial regions where no heating mechanism is effective. In particular, we consider a one-dimensional Cartesian model, where the discharge is formed between a pair of plane electrodes separated by 5 cm. Oxygen molecules with a density of 3.2×10^{21} m^{-3} and are assumed to fill the space between the electrodes. Both electrodes are grounded, and a discharge is sustained by an oscillating electric field applied in the plane of the electrodes. This electric field

oscillates at a frequency of 13.56 MHz, and its amplitude is determined by requiring that the line-averaged current density amplitude parallel to the electric field be 20 A m^{-2}. The charged species followed in the particle simulation are electrons, O_2^+ ions and O^- ions. We makes the probably extreme assumption that half of the neutral oxygen molecules occupy the metastable $O(a\,^1\Delta_g)^*$ state, so that the dominant destruction mechanism for negative ions is detachment in collisions with molecules in these metastable states. The cross sections and rate constants used in these calculations are relatively voluminous, and the details are not especially germane to the present discussion. We therefore refer interested readers to other recent work dealing with modelling of oxygen discharges for further particulars (see for example [50] and references therein). The results of these calculations are summarized in figs. 16 to 18. Fig. 16 shows the line-averaged densities of electrons and negative ions as a function of N_D, the number of electrons per Debye length, while fig. 17 shows corresponding data for the effective temperature of electrons and negative ions. The final figure shows the electron energy distribution functions from the extreme cases shown in figs. 16 and 17. Inspection of these figures shows that the present case is unlike the three examples discussed above, in that a high degree of convergence is achieved when $N_D \sim 100$. This is possible because the physical processes that shape the electron energy distribution function are relatively fast, owing to the prevalance of inelastic processes. Consequently, a relatively small number of particles is needed to ensure that these physical processes dominate the spurious numerical effects discussed earlier in this paper.

VI. DISCUSSION

In this section we comment on the practical conclusions that can be drawn from the previous sections. Chief among these is of course the observation that, for collisional plasmas, the results of particle simulations may be much more sensitive to the number of superparticles than has hitherto been supposed. In particular, it is commonly assumed that an adequately large value of N_D might be 10–100, but the results presented above show this to be wildly misleading. In order to determine an appropriate value of N_D, some insight is needed into the physical processes that shape the electron energy distribution function. Essentially, one needs to associate some physical relaxation time with the electron energy distribution function, and then require that this be smaller than the numerical relaxation

time, τ_R. The physical relaxation time may be different for electrons in different energy ranges. For instance, in the examples discussed above, there is no case where the electron energy distribution function above the inelastic collision threshold is affected by numerical relaxation. We note that the problem of estimating this physical relaxation time arises also in the context of deciding whether Coulomb collisions should be considered or not, and discussions on that topic are useful [8, 30, 51, 52]. Various formulae for estimating the rate of relaxation of a distribution function due to Coulomb collisions are also available [4]. Whatever physical considerations may arise, τ_R should be adequately estimated from Eq. (15), and so one has in principle a method of choosing a suitable value of N_D. In some cases, the result of such a calculation will be unwelcome, because the conclusion will be that N_D needs to be more than 1000, and this is an unwieldy number, even for modern desktop computers.

These considerations show that under some circumstances it will be difficult to obtain particle-in-cell simulation results that can be demonstrated to be converged with respect to all three numerical parameters $(\Delta t, \Delta x, N_D)$. This might not be important. For example, many discharge parameters are practically unaffected by the density and temperature of trapped electrons, or the exact shape of the electron energy distribution function. For instance, in the examples discussed, the impedance of the discharge and power deposition depend on phenomena occuring predominantly in sheath regions that are barely affected by these issues. Consequently, the issues examined in this paper may be of slight importance in some practical cases. Moreover, for some purposes it might be satisfactory to select the superparticle weight using Eq. (15) such that the numerical thermalization rate approximates the expected physical thermalization rate by Coulomb collisions, as appears to have happened fortuitously in some previous works. However, there will remain some cases where a strictly converged solution is desired, and then the question arises of whether some method other than the classical particle-in-cell procedure is preferable. In the present state of knowledge, it is difficult to give a clear answer to this question, not because other methods do not exist, but because few of them have been widely adopted, and even less have been characterized in a helpful way. For example, the convective scheme is a relatively widely-used method, but there are no detailed studies of convergence properties or thermalization behaviour. Variant particle-in-cell schemes may also have promise, *e.g.* [53], but again, detailed investigations of their properties are not yet forthcoming.

VII. CONCLUSION

In the main section of this paper, we studied the kinetic behavior of a one-dimensional particle-in-cell simulation in the presence of Monte Carlo collisions. We found that the kinetic properties of the simulation are appreciably degraded when the $\nu/\omega_p \gtrsim 10^{-4}$, and very seriously degraded when $\nu/\omega_p \sim 1$. If we take as figures of merit for the kinetic behavior of the simulation the thermalization time, τ_R, and the heating time, τ_H, then the amount of degradation may be as much as a factor of a thousand in the former case and a hundred in the latter. By reference to four examples, we showed that these effects may sometimes be responsible for serious difficulties in achieving simulation results that can be shown to be properly converged as a function of all the relevant numerical parameters. Not all of the physical parameters are affected by these issues, and there may be many practical cases where serious issues are not raised by these results, such as simulations involving predominantly molecular gases. Nevertheless, the function of fully kinetic simulations as benchmarks for other models is clearly compromised if there are cases where full convergence cannot practically be achieved. It is not clear whether some other kinetic model should now be preferred to the particle-in-cell method for collisional plasmas, because the kinetic properties of other methods are essentially unknown.

---

[1] M. H. Wilcoxson and V. I. Manousiouthakis, IEEE Trans. Plasma Sci. **21**, 213 (1993).

[2] W. N. G. Hitchon, G. J. Parker, and J. E. Lawler, IEEE Trans. Plasma Sci. **21**, 228 (1993).

[3] M. Surendra, Plasma Sources Sci. Technol. **4**, 56 (1995).

[4] B. A. Trubnikov, in *Reviews of Plasma Physics* (Consultants Bureau, New York, 1965), vol. 1, pp. 105–204.

[5] I. P. Shkarofsky, T. W. Johnston, and M. P. Bachynski, *The Particle Kinetics of Plasmas* (Addison-Wesley, Reading, Massachusetts, 1966).

[6] S. Ichimaru, *Statistical Plasma Physics, Volume I: Basic Principles*, vol. 87 of *Frontiers in Physics* (Addison-Wesley, Reading, Massachusetts, 1992).

[7] R. Balescu, *Statistical Mechanics of Charged Particles*, vol. 4 of *Monographs in Statistical Physics and Thermodynamics* (Wiley, New York, 1963).

[8] B. E. Cherrington, *Gaseous Electronics and Gas Lasers* (Pergamon, Oxford, 1979).

[9] C. K. Birdsall and A. B. Langdon, *Plasma Physics via Computer Simulation* (McGraw-Hill, New York, 1985).

[10] R. W. Hockney and J. W. Eastwood, *Computer Simulation using Particles* (Adam Hilger, Bristol, 1988).

[11] W. N. G. Hitchon, D. J. Koch, and J. B. Adams, J. Comp. Phys. **83**, 79 (1989).

[12] G. J. Parker, W. N. G. Hitchon, and J. E. Lawler, J. Comp. Phys. **106**, 147 (1993).

[13] M. Gutnic, A. Haefele, I. Paun, and E. Sonnendrücker, Comput. Phys. Commun. **194** (2004).

[14] C. K. Birdsall, IEEE Trans. Plasma Sci. **19**, 65 (1991).

[15] T. Takizuka and H. Abe, J. Comp. Phys. **25**, 205 (1977).

[16] K. Nanbu, Phys. Rev. E **55**, 4642 (1997).

[17] H. Ueda, Y. Omura, H. Matsumoto, and T. Okuzawa, Comput. Phys. Commun. **79**, 249 (1994).

[18] J. Dawson, Phys. Fluids **5**, 445 (1962).

[19] O. C. Eldridge and M. Feix, Phys. Fluids **5**, 1076 (1962).

[20] O. C. Eldridge and M. Feix, Phys. Fluids **6**, 398 (1963).

[21] J. M. Dawson, Phys. Fluids **7**, 419 (1964).

[22] H. Okuda and C. K. Birdsall, Phys. Fluids **13**, 2123 (1970).

[23] Y. Matsuda and H. Okuda, Phys. Fluids **18**, 1740 (1975).

[24] D. Montgomery and C. W. Nielson, Phys. Fluids **13**, 1405 (1970).

[25] H. Okuda, Phys. Fluids **15**, 1268 (1973).

[26] A. B. Langdon, Phys. Fluids **22**, 163 (1978).

[27] J. Virtamo and H. Tuomisto, Phys. Fluids **22**, 172 (1979).

[28] M. Abramowitz and I. A. Stegun, *Handbook of mathematical functions* (Dover, New York, 1965).

[29] V. Vahedi and M. Surendra, Comput. Phys. Commun. **87**, 179 (1995).

[30] S. D. Rockwood, Phys. Rev. A **8**, 2348 (1973).

[31] V. A. Godyak and R. B. Piejak, Phys. Rev. Lett. **65**, 996 (1990).

[32] M. Surendra and D. B. Graves, Phys. Rev. Lett. **66**, 1469 (1991).

[33] M. M. Turner, R. A. Doyle, and M. B. Hopkins, Appl. Phys. Lett. **62**, 3247 (1993).

[34] V. Vahedi, C. K. Birdsall, M. A. Lieberman, G. DiPeso, and T. D. Rognlien, Plasma Sources

Sci. Technol. **2**, 273 (1993).

[35] T. H. Chung, H. S. Yoon, and J. K. Lee, J. Appl. Phys. **78**, 6441 (1995).

[36] K. Nagayama, B. Farouk, and Y. H. Lee, Plasma Sources Sci. Technol. **5**, 685 (1996).

[37] M. Yan, A. Bogaerts, W. J. Goedheer, and R. Gijbels, Plasma Sources Sci. Technol. **9**, 583 (2000).

[38] V. Georgieva, A. Bogaerts, and R. Gijbels, J. Appl. Phys. **94**, 3748 (2003).

[39] J. K. Lee, N. Babaeva, H. C. Kim, O. V. Manuilenko, and J. W. Shon, IEEE Trans. Plasma Sci. **32**, 47 (2004).

[40] H. C. Kim and J. K. Lee, Phys. Rev. Lett. **93**, 085003 (2004).

[41] K. Matyash and R. Schneider, Contrib. Plasma Phys. **44**, 589 (2004).

[42] N. Y. Babaeva, J. K. Lee, and J. W. Shon, J. Phys. D: Appl. Phys. **38**, 287 (2005).

[43] K. Tachibana, Phys. Rev. A **34**, 1007 (1986).

[44] J. L. Pack, R. E. Voshall, A. V. Phelps, and L. E. Kline, J. Appl. Phys. **71**, 5363 (1992).

[45] A. V. Phelps, J. Appl. Phys. **76**, 747 (1994).

[46] G. Francis, *Gas Discharges II* (Springer, Berlin, 1956), vol. XXI of *Handbuch der Physik*, chap. The Glow Discharge at Low Pressure, pp. 53–208.

[47] E. A. Den Hartog, D. A. Doughty, and J. E. Lawler, Phys. Rev. A **38**, 2471 (1988).

[48] J. E. Lawler, E. A. Den Hartog, and W. N. G. Hitchon, Phys. Rev. A **43**, 4427 (1991).

[49] A. V. Phelps, URL `ftp://jila.colorado.edu`.

[50] M. W. Kiehlbauch and D. B. Graves, J. Vac. Sci. Tech. A **21**, 660 (2003).

[51] S. D. Rockwood, J. Appl. Phys. **45**, 5229 (1974).

[52] C. J. Elliott and A. E. Greene, J. Appl. Phys. **47**, 2946 (1976).

[53] I. V. Schweigert and V. A. Schweigert, Plasma Sources Sci. Technol. **13**, 315 (2004).

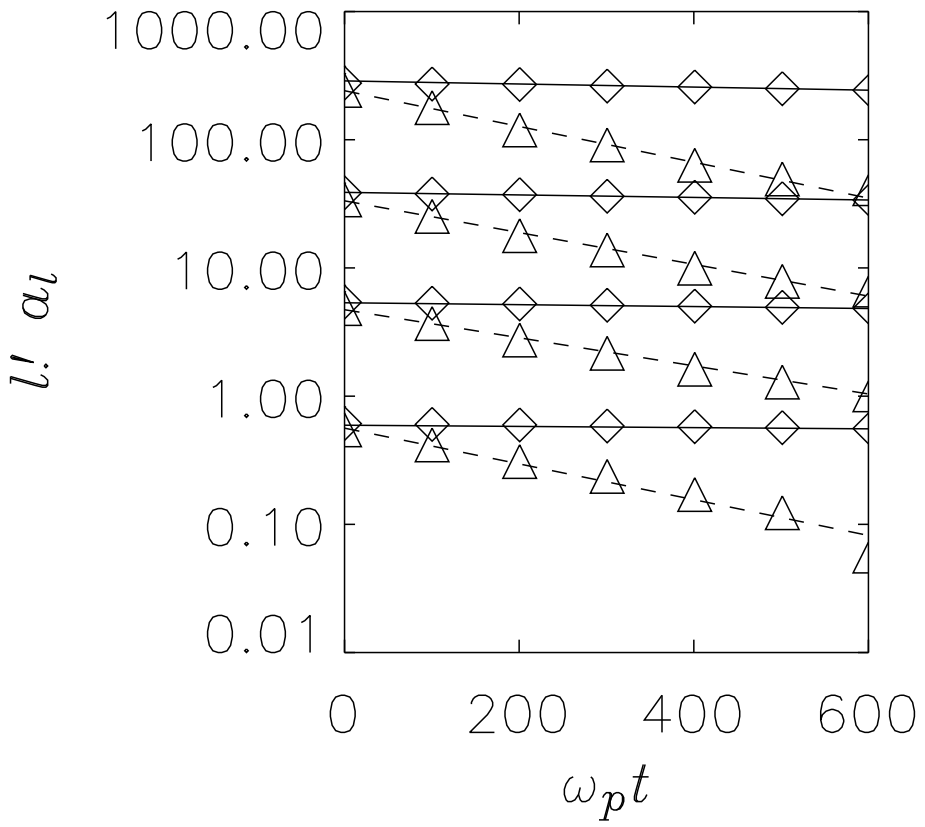

FIG. 1: Temporal decay of the eigenmodes of the collision operator Eq. (7) as observed in a particle-in-cell simulation with and without collisions. The $\diamond$ symbols are simulation data without collisions, and the solid lines are a least squares fit of an exponential decay to these points. The $\triangle$ symbols and dashed lines are corresponding results for $\nu/\omega_p \simeq 5 \times 10^{-3}$. Data are shown for the eigenmodes a_4, a_6, a_8 and a_{10}, in order beginning at the bottom of the figure. The factorial weighting factor arises naturally in the computation of the mode amplitudes and its retention conveniently compresses the range of the graph; it is not otherwise significant.

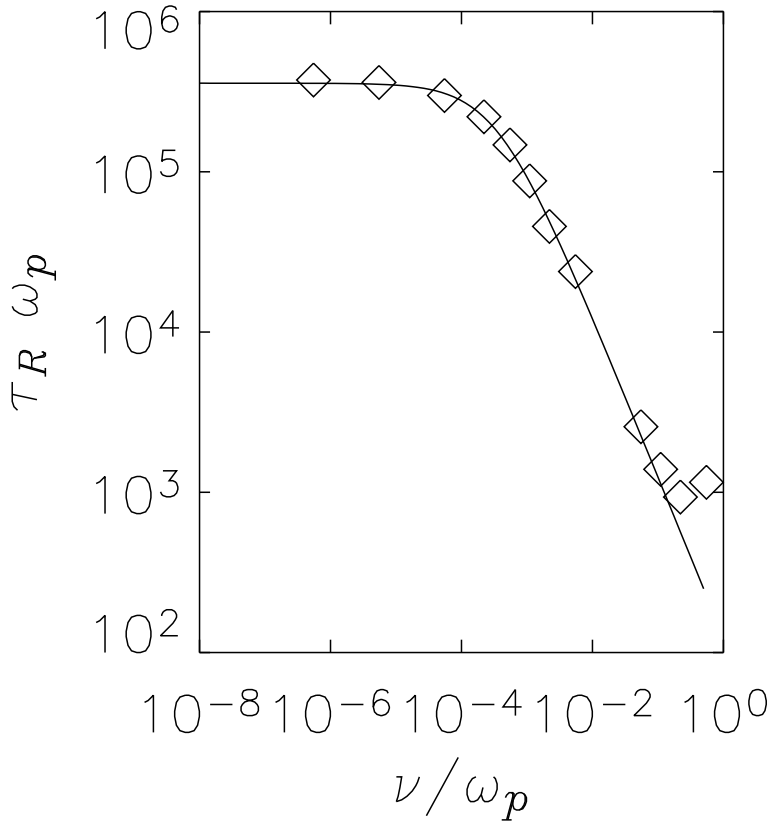

FIG. 2: The relaxation time of the electron velocity distribution function, τ_R, defined as discussed in the text, as a function of the normalized electron scattering frequency ν/ω_p. The points are simulation data and the solid line is the fit from Eqn. (15)

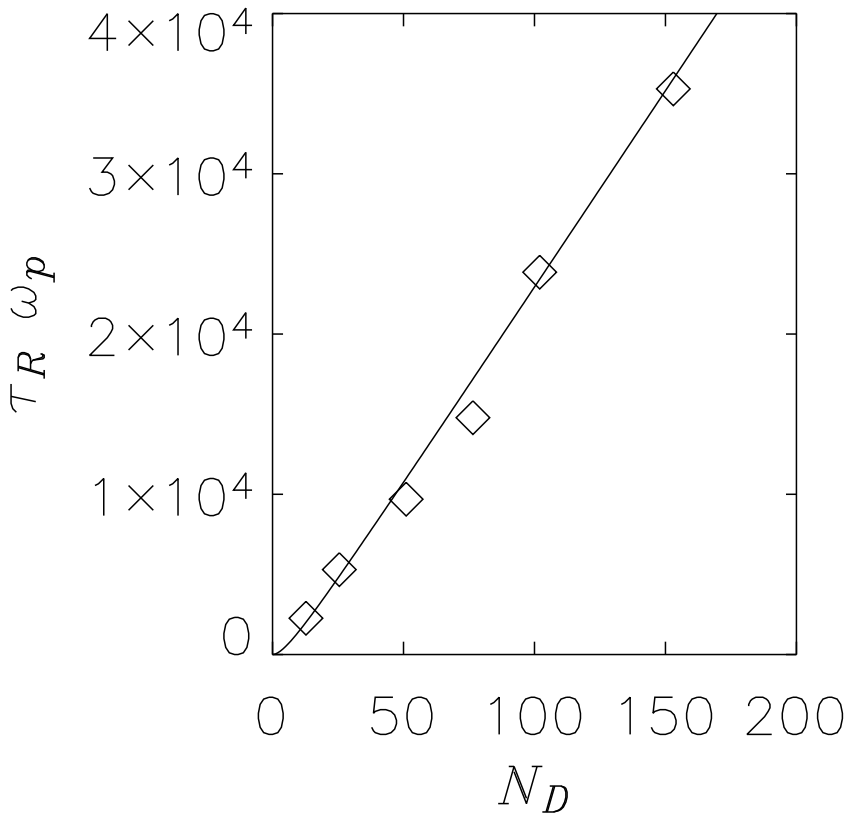

FIG. 3: The relaxation time of the electron velocity distribution function, τ_R, defined as discussed in the text, as a function of the number of particles per Debye length, N_D. The points are simulation data and the solid line is the fit from Eqn. (15)

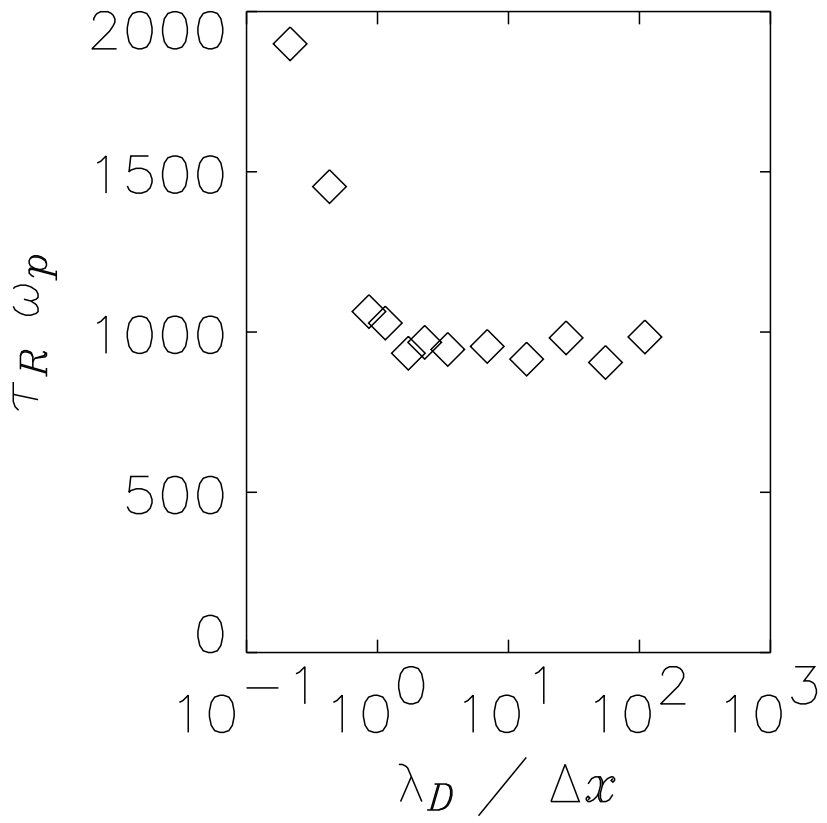

FIG. 4: The thermalization time, τ_R, of the simulation as a function of the ratio of the Debye length, λ_D, to the cell size Δx, for a normalized collision frequency $\nu/\omega_p = 0.044$

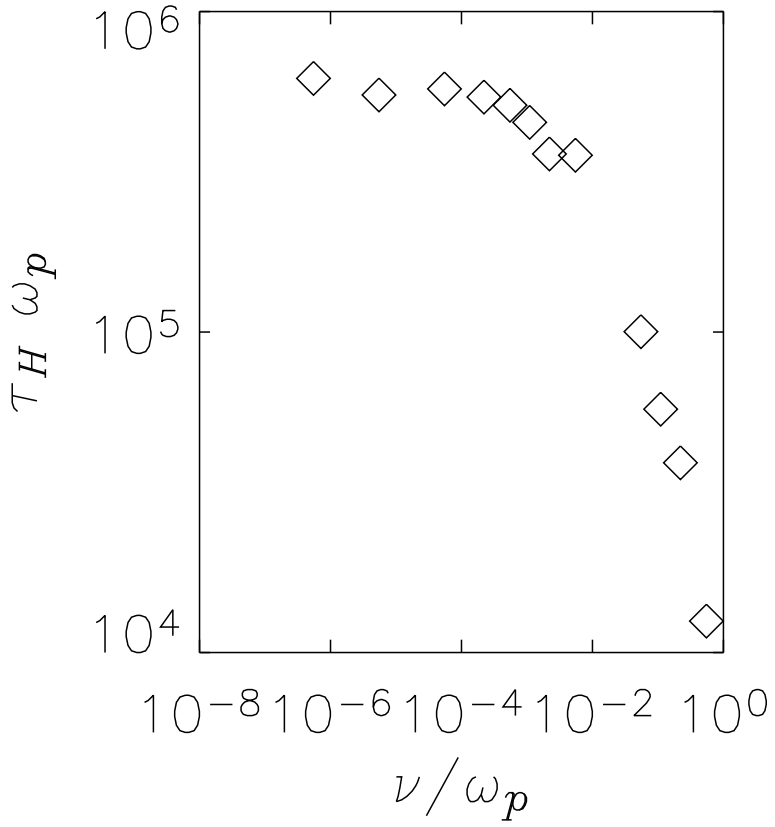

FIG. 5: The self-heating time, τ_H, of the simulation as a function of the normalized collision frequency.

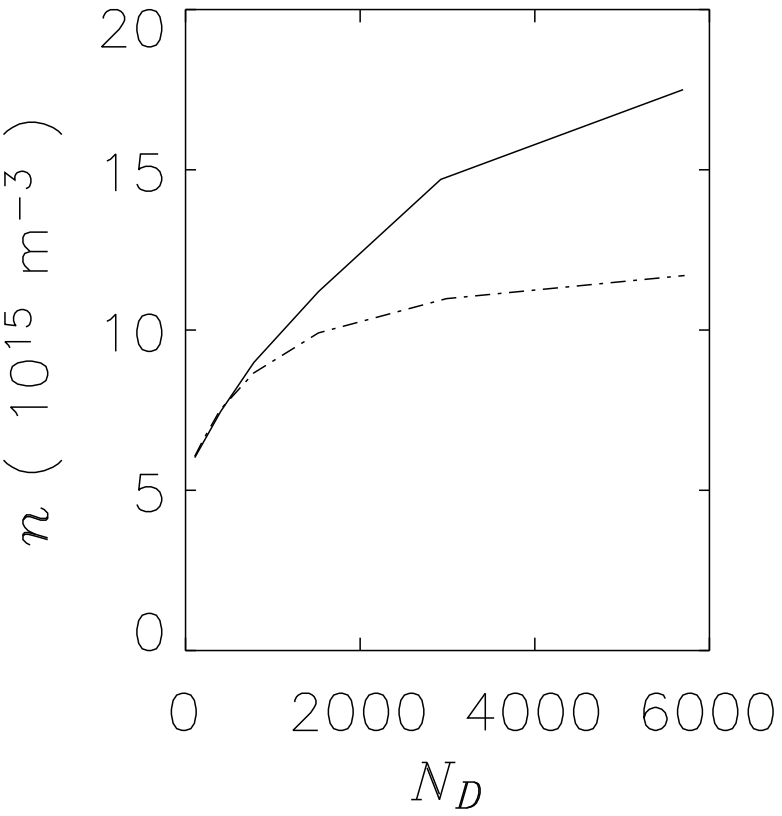

FIG. 6: Plasma density at the discharge mid-plane, as a function of the number of simulation particles per Debye length, for a radio frequency discharge in argon, with plane electrodes separated by 2 cm, an excitation frequency of 13.56 MHz, a pressure of 100 mTorr and a current density of 30 A m^{-2}. The solid line denotes the result without electron-electron collisions; the dashed line is the corresponding result with electron-electron collisions.

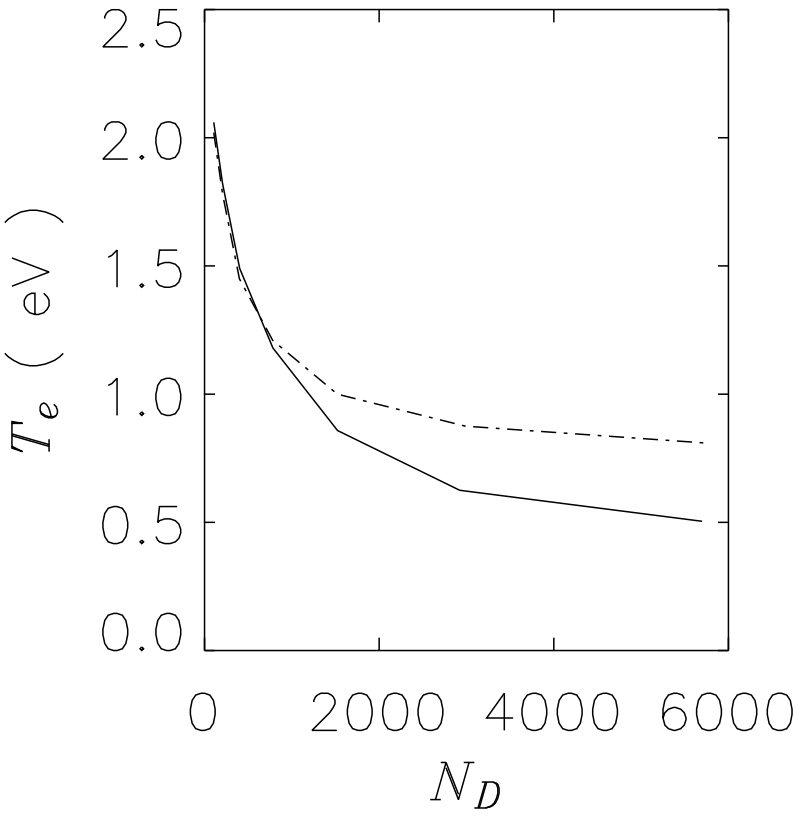

FIG. 7: Effective electron temperature at the discharge mid-plane, for a radio frequency discharge at the conditions and with the rubric of Fig. 6

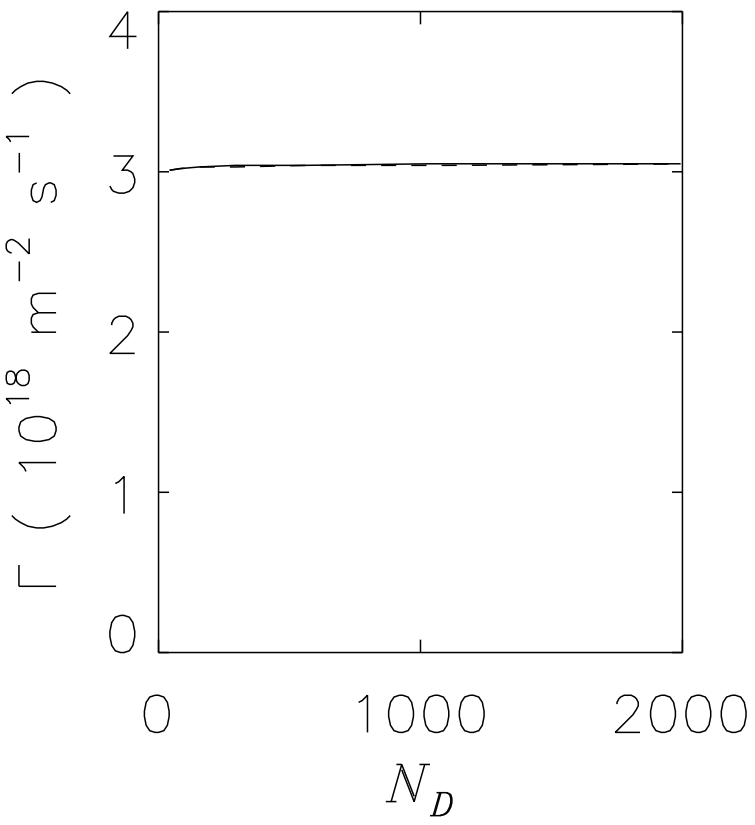

FIG. 8: Time averaged ion flux at an electrode, for a radio frequency discharge at the conditions and with the rubric of Fig. 6

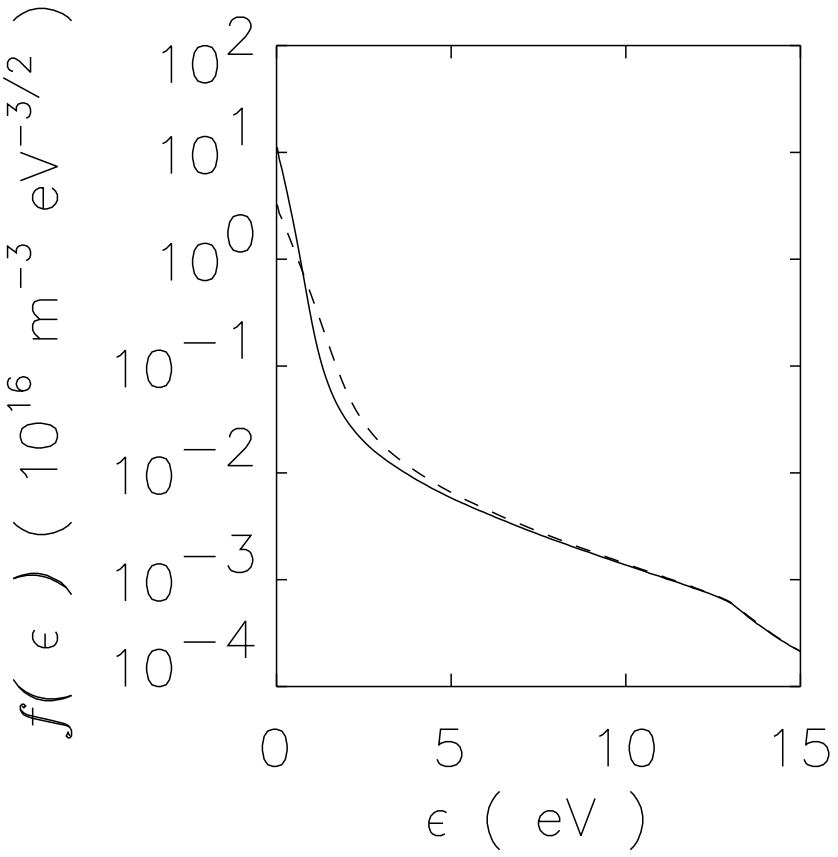

FIG. 9: Electron energy distributions at the discharge midplane, for the discharge conditions of fig. 6. Both curves are for $N_D \simeq 6000$. The solid curve is a calculation excluding electron-electron collisions, while the dashed curve shows the case where they are included. The temperatures of the low energy groups of electrons are approximately 0.25 eV for the case without electron-electron collisions and 0.49 eV with electron-electron collisions.

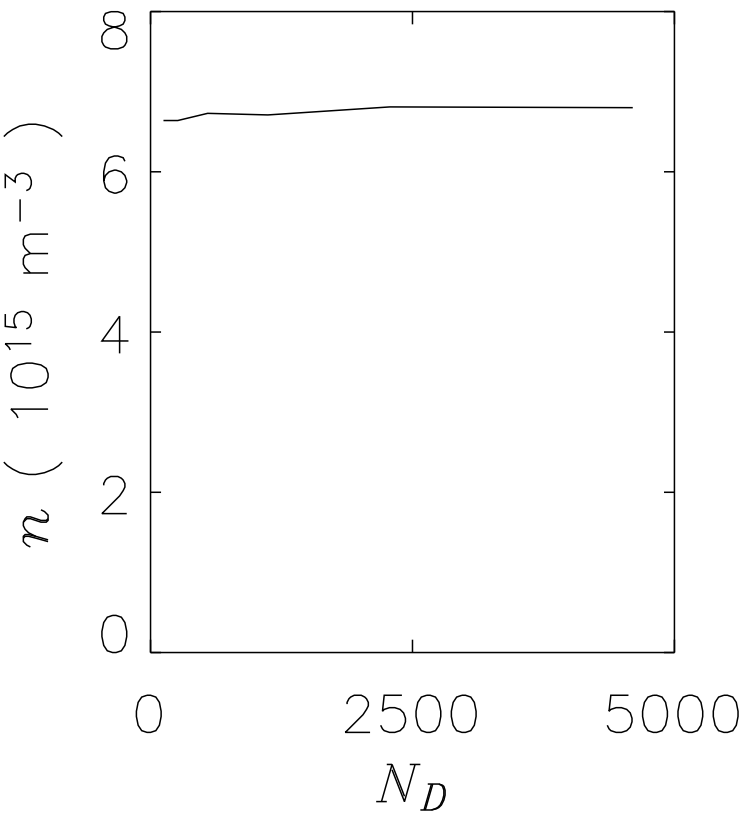

FIG. 10: Plasma density at the discharge mid-plane, as a function of the number of simulation particles per Debye length, for a radio frequency discharge in argon,with plane electrodes separated by 2 cm, an excitation frequency of 13.56 MHz, a pressure of 500 mTorr and a current density of 30 A m^{-2}

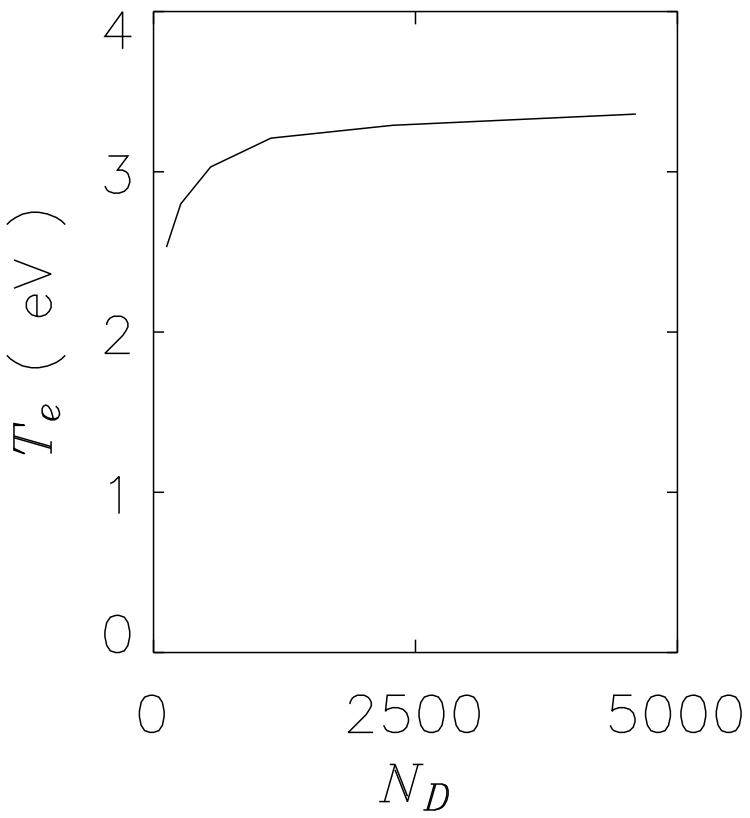

FIG. 11: Effective electron temperature at the discharge mid-plane, for a radio frequency discharge at the conditions of Fig. 10

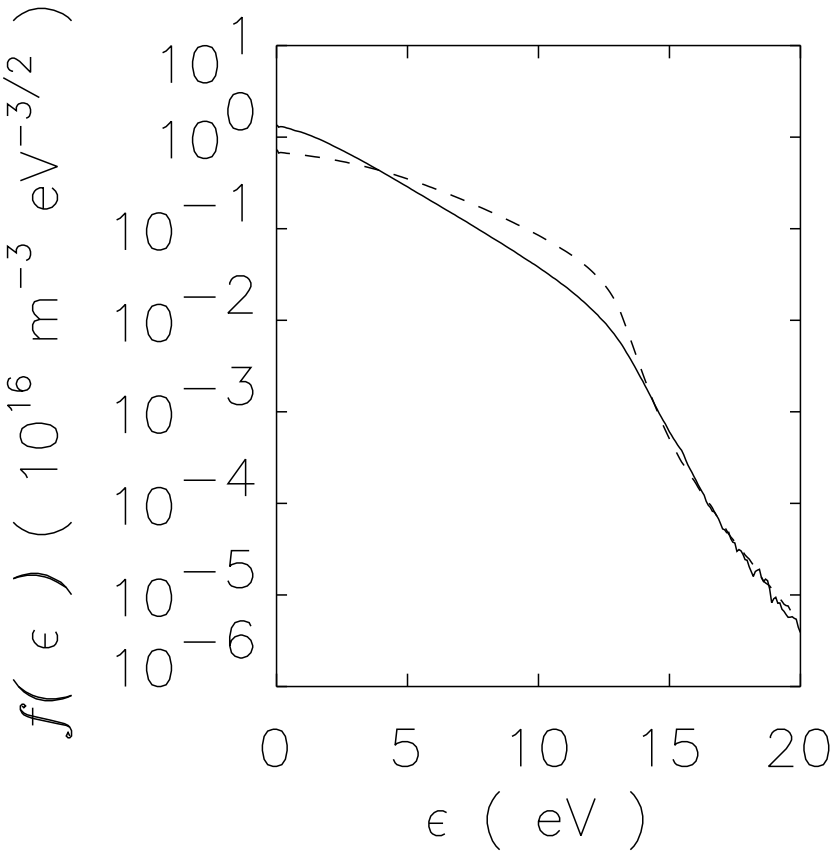

FIG. 12: Electron energy distributions at the discharge midplane, for the discharge conditions of fig. 10. The solid curve is a calculation for $N_D \simeq 100$, while the dashed curve is for $N_D \simeq 4600$.

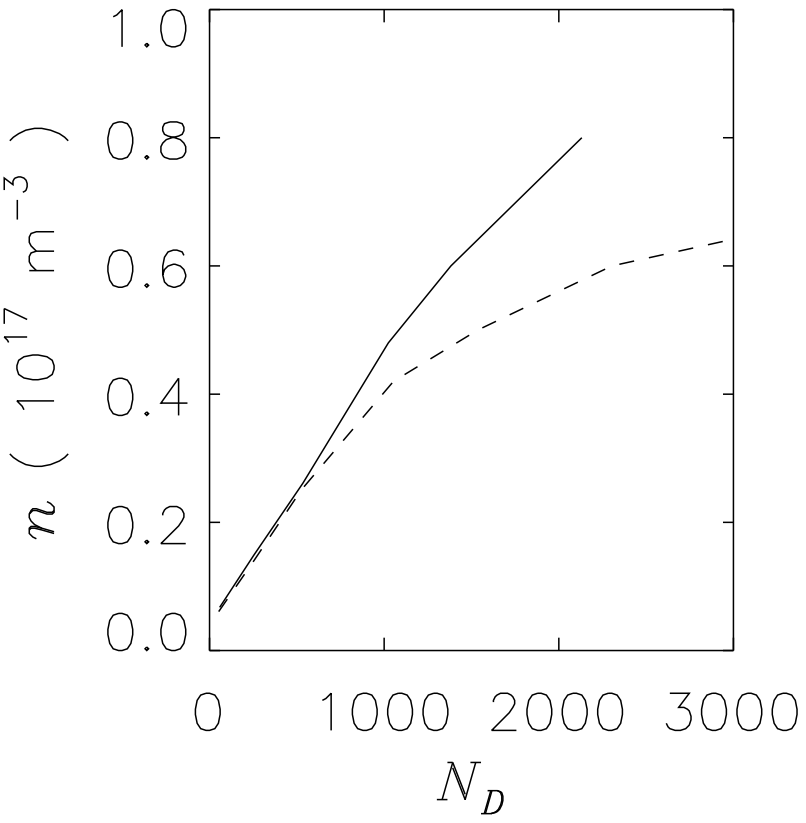

FIG. 13: Peak plasma density in the negative glow of a direct current discharge in helium as a function of the number of superparticles per Debye length, N_D. The discharge is formed between plane electrodes separated by 1 cm with a gas density of $3.2 \times 10^{22}\mathrm{m}^{-3}$ and with 300 V applied. The solid line denotes the result without electron-electron collisions; the dashed line is the corresponding result with electron-electron collisions.

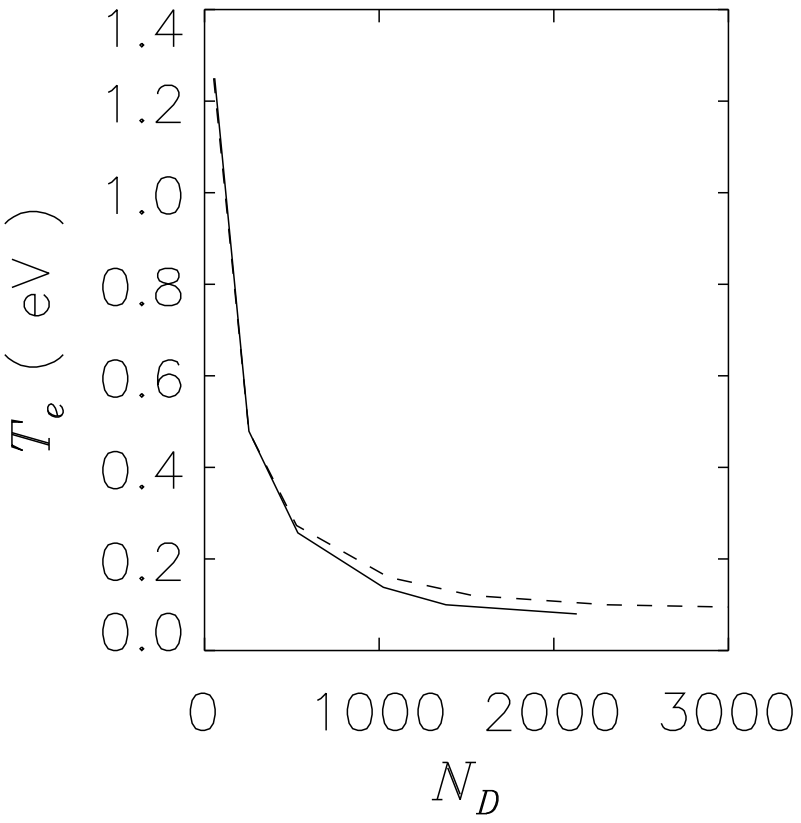

FIG. 14: Effective electron temperature at the peak of the plasma density in the negative glow of a direct current discharge, as a function of the number of superparticles per Debye length, N_D, for the conditions of fig. 13. The solid line denotes the result without electron-electron collisions; the dashed line is the corresponding result with electron-electron collisions.

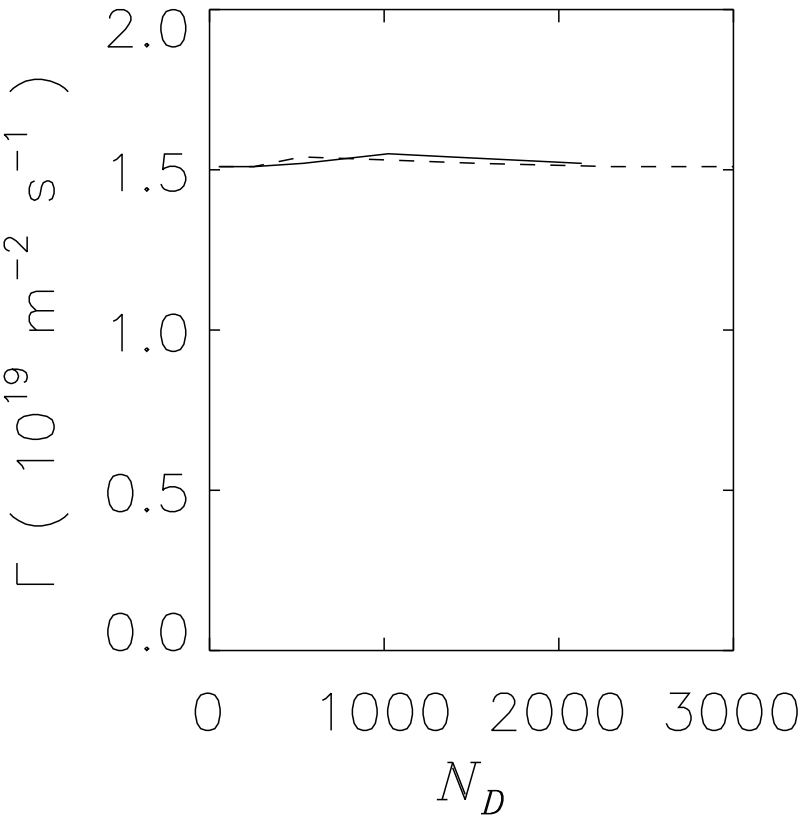

FIG. 15: Ion flux collected at the cathode of a direct current discharge, as a function of the number of superparticles per Debye length, N_D, for the conditions of fig. 13. The solid line denotes the result without electron-electron collisions; the dashed line is the corresponding result with electron-electron collisions.

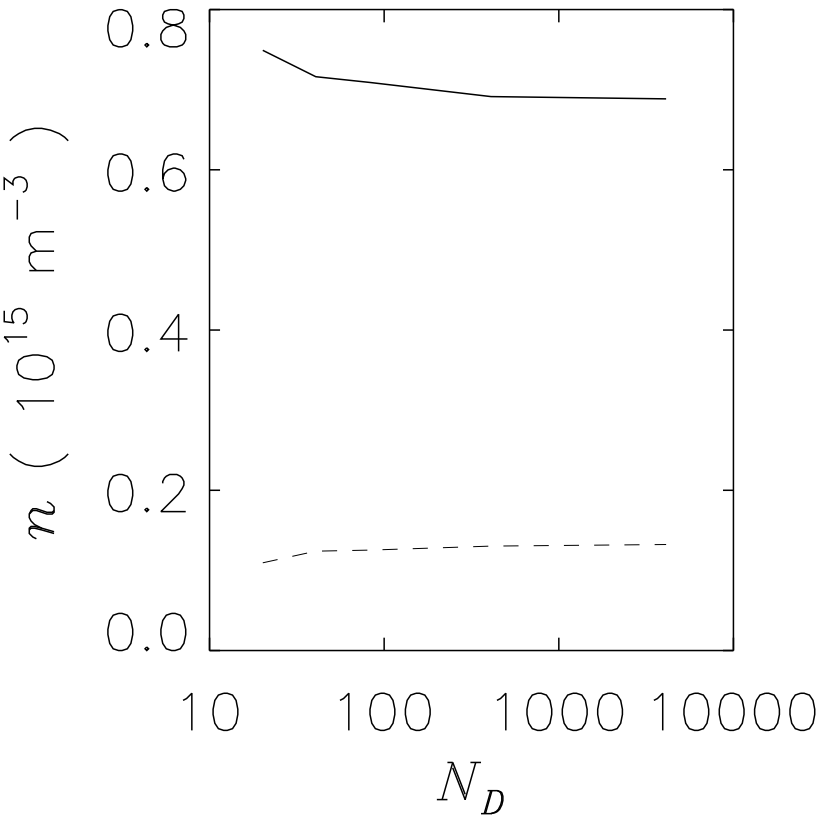

FIG. 16: Electron density (solid line) and negative ion density (dashed line), as a function of the number of superparticles per Debye length, N_D, in an oxygen positive column formed between plane electrodes separated by 5 cm with a neutral gas density of $3.2 \times 10^{21}\mathrm{m}^{-3}$ and ax average axial current density of 20A m^{-2}.

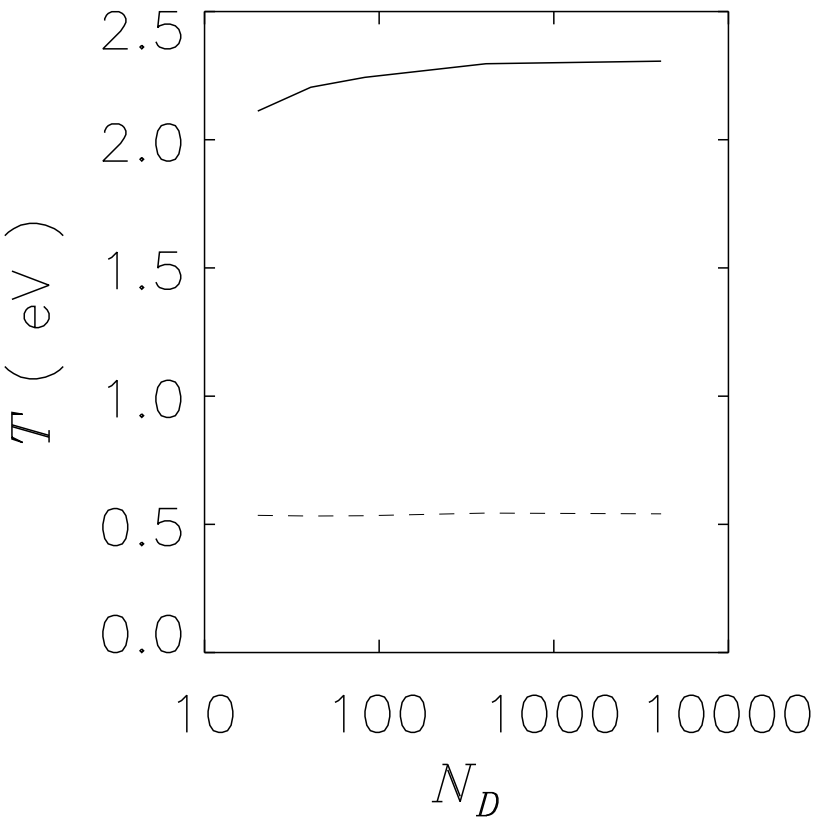

FIG. 17: Effective electron temperature (solid line) and effective negative ion temperature (dashed line), as a function of the number of superparticles per Debye length, N_D, in an oxygen discharge under the conditions of fig 16.

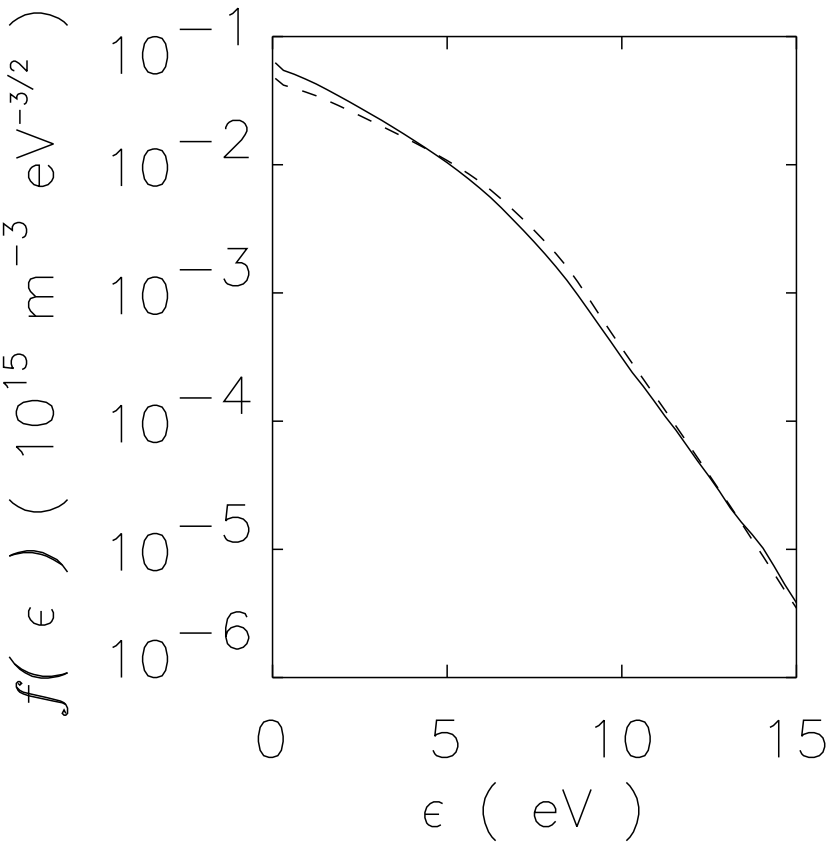

FIG. 18: Electron energy distribution functions corresponding to the extreme cases shown in figs. 16 and 17. The solid line is the case with the smallest value of N_D and the dashed line is the case with the largest value of N_D.