Hedging Direct Simulation Monte Carlo Bets via Event Splitting

G. Oblapenko[a,*], D. Goldstein[b], P. Varghese[b,c], C. Moore[d]

[a]*German Aerospace Center (DLR), Bunsenstraße 10, 37073 Göttingen, Germany*
[b]*ASE-EM Department, The University of Texas at Austin, 2617 Wichita St., Stop C0600, Austin, TX 78712*
[c] *Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, 2201 E 24th St, Stop C0200, Austin, TX 78712*
[d]*Sandia National Laboratories, Albuquerque, NM, USA*

Abstract

We propose a new scheme for simulation of collisions with multiple possible outcomes in variable-weight DSMC computations. The scheme is applied to a 0-D ionization rate coefficient computation, and 1-D electrical breakdown simulation. We show that the scheme offers a significant (up to an order of magnitude) improvement in the level of stochastic noise over the usual acceptance-rejection algorithm, even when controlling for the slight additional computational costs. The benefits and performance of the scheme are analyzed in detail, and possible extensions are proposed.

Keywords: Rarefied gas dynamics, Direct Simulation Monte Carlo, Boltzmann equation, ionization
2010 MSC: 76P05, 82D10, 82D05, 76X05

1. Introduction

Direct Simulation Monte Carlo [1] is one of the most ubiquitous tools for simulation of rarefied gas flows. The conceptually simple algorithm, which works by simulating the stochastic motion of a large number of computational particles (each representing a huge number of real-life atoms/molecules/electrons), has shown the ability to capture a very wide range of phenomena. Some examples of the wide range DSMC applications include modeling of microcombustion [2], hypersonic re-entry with ionization and radiation [3], planetary-scale phenomena [4–6], electrical breakdown [7, 8], and processes in electric propulsion systems [9]. Various extensions to DSMC include include hybridization with CFD solvers [10, 11], DSMC-based methods to model continuum flows via solution of the Fokker–Planck equation [12, 13], velocity-space hybridization with kinetic solvers [14, 15], and variable weight DSMC [16–19].

In some applications trace species (e.g. chemical radicals, high-velocity tails of the distribution function, or particles in high internal energy states) significantly affect the flow. "Vanilla" DSMC cannot model trace populations efficiently because it requires a single ratio of real-life particles to computational particles. The variable weight DSMC approach obviates the problem by allowing one to use a variable ratio of real-life particles to computational particles. Thus, by assigning a different computational weight to trace species, one can avoid an exorbitant number of total particles in the simulation in order to obtain adequate statistics on low-population species or

*Corresponding author
Email addresses: georgii.oblapenko@dlr.de (G. Oblapenko), david@oden.utexas.edu (D. Goldstein), varghese@mail.utexas.edu (P. Varghese), chmoore@sandia.gov (C. Moore)

states. Such an approach has found wide application in collisional plasma modeling (in Particle-in-Cell Direct Simulation Monte Carlo codes, also known as PIC-DSMC codes), where the molar fractions of electrons and ions can be several orders of magnitude lower than those of the neutral background species [7, 18, 20]. The drawback of variable weight DSMC is the particle splitting required to obtain the correct collision rate; this leads to a uncontrolled growth in the number of simulation particles and necessitates the occasional or periodic merging of particles [21]. Particle merging adds an additional step to the DSMC algorithm and reduces the accuracy of the simulation because some information about the system is irretrievably lost when one reduces the number of particles. A variety of algorithms have been developed to efficiently merge particles in a manner that reduces the errors introduced by merging [15, 21–29].

Variable weight DSMC tackles the problem of resolving low populations but it does not address another issue of DSMC – that of modeling low-probability collision events. Although one can achieve a desired level of accuracy by performing time-averaging, this is not feasible for unsteady simulations. Poor modeling of low-probability collision events affects both the overall noise in the simulation and can also lead to incorrect results if an unsteady simulation is run at conditions close to a bifurcation point. For example overestimation of a chemical reaction rate over the course of several timesteps might lead to transition of a deflagration into a detonation, or underestimation of an ionization rate leads to a drop in current instead of electrical breakdown.

In [30], a fractional collision scheme was proposed for variable weight DSMC. In the proposed scheme (later expanded to multi-species flows [20]), the number of collisions to be performed is increased by a integer factor of c_m (in [20] values of c_m range from 2 to 32), and only a fraction (equal to $1/c_m$) of the chosen particles actually undergoes the collision. This has been shown to greatly improve simulation accuracy, and it is obvious that increasing the number of collisions sampled improves the statistics when modeling low-probability collision processes. Instead of increasing the number of particles throughout the whole simulation, the fractional collision scheme only increases the number of post-collision particles. These can be merged down to a lower number of particles before performing collisions in another physical cell, and before convection and acceleration steps in the conventional time-splitting approach. A non-conservative fractional collision algorithm was also investigated in [20] to model charge-exchange collisions whilst minimizing the creation of additional particles. However, this fractional collision scheme also increases the number of simulated high-probability collisions that may already be adequately captured with a standard collision scheme.

The current work proposes an alternative collision algorithm: an event splitting algorithm which can be considered as a form of bet hedging — instead of simulating only a single outcome of a specific collision event, we simulate all possible processes that may occur, and weight their outcomes (i.e. the newly created particles) according to their probabilities. Such an idea was expressed in [31, 32] in the context of improvement of modeling of spontaneous photon emission, but has not received further attention to our knowledge.

The paper is structured as follows. First, we describe the proposed algorithm, and touch upon its potential benefits and drawbacks. Next, we verify this approach for a 0-dimensional case of an $Ar/Ar^+/e^-$ mixture accelerated by a constant electric field with elastic collisions between electrons and argon atoms and electron-impact ionization reactions. We investigate the gains in accuracy and reduction in variance provided by the event splitting algorithm, and propose several possible future modifications. Finally, we apply the collision-splitting algorithm to secondary emission boundary conditions as well, and perform a 1-dimensional breakdown simulation in an argon-filled gap. We assess the noise in the anode current density and investigate the computational cost and level-of-stochastic-noise benefits of treating collisions and/or boundary conditions using the event

splitting scheme.

2. Collision algorithm description

As mentioned in the introduction, we model a rarefied gas flow by simulating the motion of a large number of particles, each representing a much larger number of real-life atoms, molecules, or electrons. As in "vanilla" DSMC, we separate the collision, convection, and acceleration steps [1].

We use the No-Time-Counter collision scheme [1] adapted to variable-weight DSMC [33]. The fractional number of collisions in a physical cell is computed as follows:

$$N_{coll,frac} = \frac{1}{V} N_{p,c} N_{p,d} (n\sigma_{tot} g)_{\max} \Delta t. \tag{1}$$

Here V is the volume of the physical cell, $N_{p,c}$ and $N_{p,d}$ are the number of DSMC particles of species c and d respectively, $(n\sigma_{tot}g)_{\max}$ is the maximum value of $n \cdot \sigma_{tot} \cdot g$ in the physical cell, g is the relative velocity of the colliding particles, $\sigma_{tot}(g)$ is the total collision cross-section, n is the computational weight of the particles, and Δt is the timestep. For the case of same-species collisions, the quantity $N_{p,c}N_{p,d}$ is replaced by $\frac{1}{2}N_{p,c}(N_{p,c}-1)$. The integer-valued number of collisions is then computed via stochastic rounding as

$$N_{coll} = \lfloor N_{coll,frac} + \mathcal{R} \rfloor, \tag{2}$$

where $\mathcal{R} \sim \mathcal{U}(0,1)$ is a random number sampled from the uniform distribution over the interval $[0,1]$.

When a particle pair (with weights n_1 and n_2 and a relative velocity g) is tentatively chosen for a collision, it is accepted with a probability of

$$P = \frac{\mathrm{MAX}(n_1, n_2)\sigma_{tot} g}{(n\sigma_{tot} g)_{\max}}. \tag{3}$$

In DSMC simulations, it is often assumed that $\sigma_{tot} = \sigma_{el}$, even when inelastic collisions can occur in the flow. This is done to minimize time spent on computing inelastic collision cross-sections, even though the actual process may not be computed during the collision loop (due to acceptance-rejection sampling of the different processes). However, this simplifying approximation does not always hold and can lead to significant error both in the elastic and inelastic collision rates [34, 35]. Therefore in our work all possible collision cross-sections are summed in order to compute σ_{tot}. For the test problem considered in the present work, we only consider elastic and electron-impact ionizing collisions in an ionized argon flow. The elastic electron-argon collision cross-section and the electron-impact ionization cross-section used in the present work (based on BSR data [36, 37]) are shown on Fig. 1, and it can be seen that at high energies (and equivalent relative velocities), the ionization cross-section is comparable in magnitude to the elastic cross-section. Therefore, we compute the total cross-section accounting for all the possible processes occurring in the flow.

We assume without loss of generality that the particle with computational weight n_1 has the greater weight. When a pair is accepted for a collision, the heavier particle is split into two particles: one with weight n_2 (the weight of the lighter collision partner) and one with weight $n_1 - n_2$. The particles with weights n_2 collide and their post-collision velocities are computed based on the cross-section model used in the simulation, while the particle with weight $n_1 - n_2$ does not have its velocity affected.

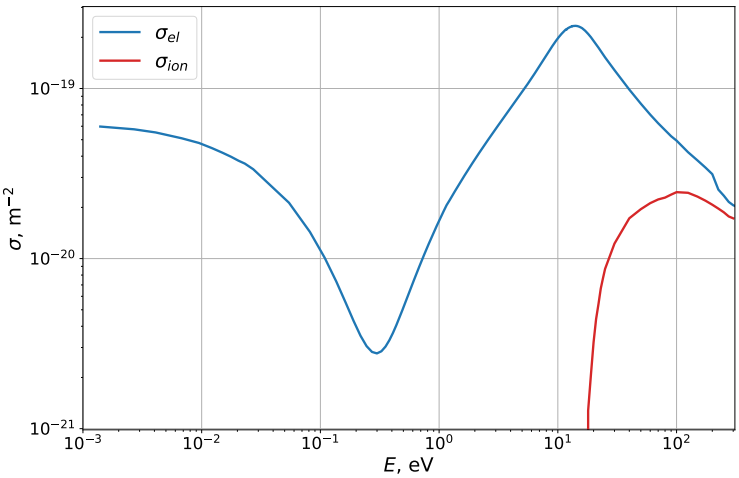

Figure 1: Elastic and ionization cross-sections for electron-argon collisions as a function of energy, BSR data [36, 37].

Usually, if inelastic collisions can occur, the exact type of collision to be performed is determined by an acceptance-rejection scheme. For our considered test-case (of an $\mathrm{Ar}/\mathrm{Ar}^{+}/\mathrm{e}^{-}$ mixture), for electron-neutral collisions (with only elastic and ionizing collisions), the type of collision to be performed is determined by the ratio $P_{el} = \frac{\sigma_{eL}}{\sigma_{tot}}$ (where σ_{el} is the elastic cross-section). A random number $\mathcal{R}$ uniformly distributed over the interval $[0, 1]$ is sampled, and checked against P_{el}. If $\mathcal{R} < P_{el}$, then an elastic collision is performed, otherwise, an ionizing collision is performed. From Fig. 2 we can see at energies below approximately 30 eV, the ratio of the ionization cross-section to the total collision cross-section (computed as $\sigma_{tot} = \sigma_{el} + \sigma_{ion}$) reaches a value of 0.1 at most. Thus, if the electron temperature is sufficiently low (i.e. very few electrons have an energy higher than 30 eV), for an accepted electron-neutral collision pair an acceptance-rejection-based scheme will perform an ionization reaction for less than 10% of the collisions. For even colder electron distribution functions, the fraction of ionizing collisions will be even lower. If one takes processes with even smaller cross-sections (such as electron-impact excitation reactions) into consideration, then collisions for those processes will be performed even less frequently than ionizing collisions. This is the motivation for our collision scheme modification, which improves the fidelity of modeling low-probability processes.

Our proposed event splitting algorithm is the following: we perform an additional particle split based on the probabilities of the possible collision events that can occur for the considered collision pair. After the initial split (of the heavier particle with weight n_1 into particles with weights n_2 and $n_1 - n_2$), the colliding particles of weight n_2 are split into two particles, of weights $P_{el}n_2$ and $(1 - P_{el})n_2$. The first of these two particle pairs undergoes an elastic collision, while the second pair undergoes an ionizing collision. In a general case, where for a given collision pair (of species c and d) $N_{proc,cd}$ various processes can take place, after the initial split of the heavier particle (into particles with weight n_2 and $n_1 - n_2$), the event splitting algorithm would further split the particles with weights n_2 into $N_{proc,cd}$ smaller particles, with weights $n_2 P_i$, $i = 1, \ldots, N_{proc,cd}$, where $P_i = \sigma_i / \sum_i \sigma_i$. Here P_i is the probability of a process i, and σ_i is the cross-section of the corresponding process.

Figures 3, 4 show a schematic of how the variable-weight DSMC collision routines are performed (for an argon-electron collision that can be either elastic or ionizing) when one uses acceptance-rejection to select the type of collision to perform (Fig. 3), and when one uses fractional collisions (Fig. 4). For the purposes of the illustration, it is assumed that the statistical weight of the argon

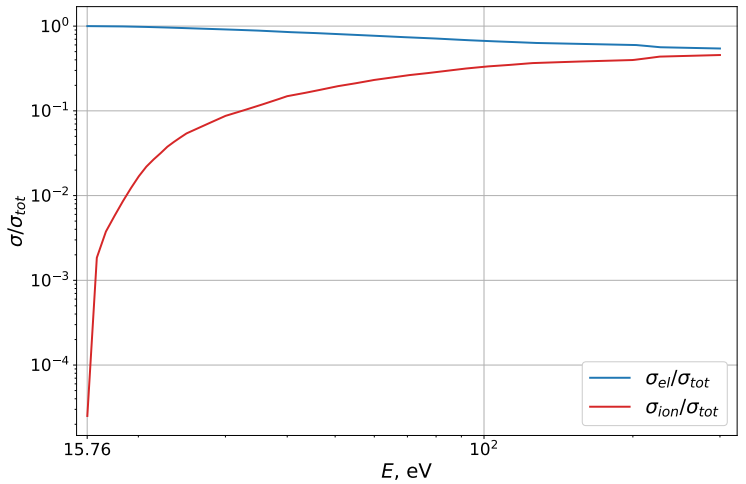

Figure 2: Ratios of the elastic and ionization cross-sections to the total electron-argon collision cross-section as a function of energy, BSR data.

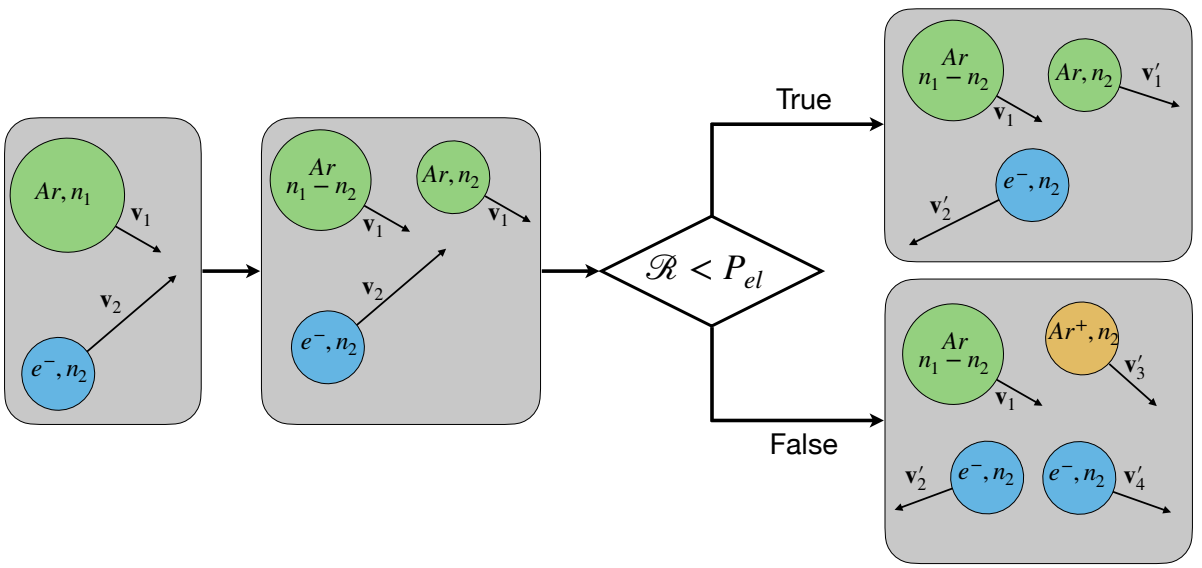

Figure 3: Schematic showing how argon-electron collisions are performed in variable weight DSMC when acceptance-rejection sampling is used.

particle (n_1) is larger than that of the electron particle (n_2). As mentioned previously, the first step is to split the heavier particle (in this case, the argon particle) into two particles, one of which has the same weight as the electron particle and will be the one undergoing the collision (the velocity, computational weight, and all other properties of the second argon particle with weight $n_1 - n_2$ remain unchanged). Then, the collision is performed. In the case of the acceptance-rejection scheme, either three or four computational particles are produced from the initial two particles, depending on which collision type is chosen. For the case of event splitting both the elastic and inelastic collisions take place and six particles are produced unless one of the probabilities (P_{ion} or P_{el}) is zero.

A pseudocode description of the standard variable-weight DSMC collision procedure (using

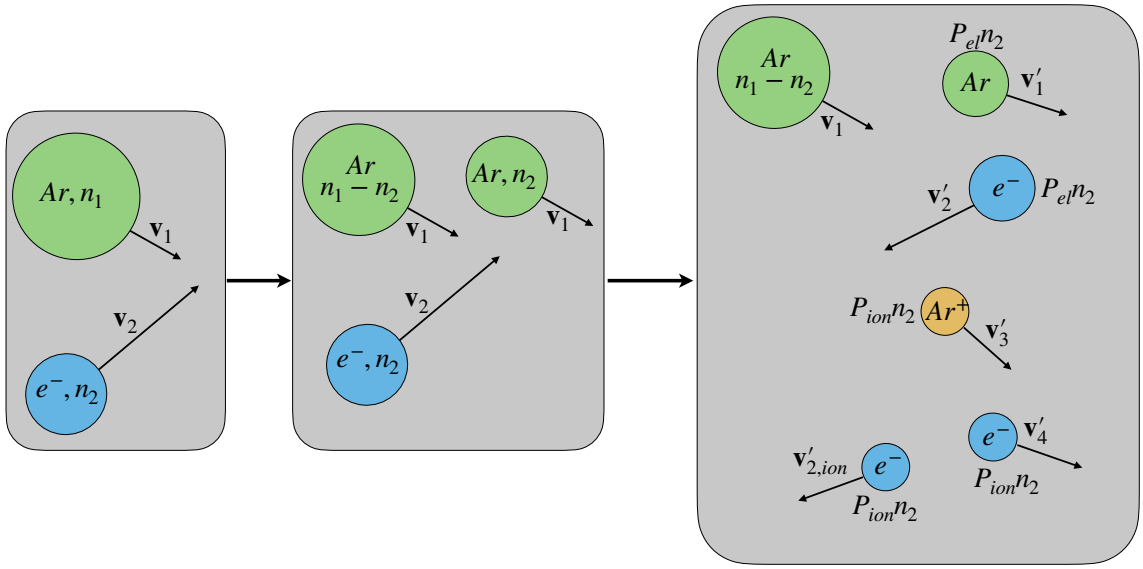

Figure 4: Schematic showing how argon-electron collisions are performed in variable weight DSMC when the fractional collision method is used (here, $P_{ion} = 1 - P_{el}$.)

the variable-weight NTC formulation [33]) and the event splitting collision procedure is given by Algorithms 1, 2, respectively.

This proposed event splitting approach provides a significant improvement in the level of stochastic noise in computed inelastic rate coefficients, as will be shown in the numerical results section. First, however, we consider a theoretical analysis of the benefits of the proposed scheme in a simplified setting. Let us assume that

1. We have a single inelastic process that has a threshold (activation) energy E_{min}, the probability of sampling a particle of species c with sufficient energy is p_E
2. The probability of a particle undergoing the inelastic process is $p_i = \text{const}$ if the relative kinetic energy of the colliding particles is larger than E_{min}; otherwise, the probability of a particle undergoing the inelastic process is 0
3. The velocity distribution of only one of the two colliding species affects the rate. For example, in electron-impact processes the velocities of the incident electrons are significantly higher than the velocities of the neutral species, and it is mostly the electron velocity distribution that affects the rate. We shall denote this species c
4. We sample $N_{p,c}$ particles (with uniform weight n_c)
5. Our total cross-sections are Maxwell-molecule-like, so $\sigma_{tot} g = \text{const}$
6. We test each particle for a collision exactly once (so $N_{coll} = N$ and each particle is selected exactly once)
7. Each particle is accepted for a collision (since $(n\sigma g)_{\max} = n\sigma g = \text{const}$)

If one uses acceptance-rejection, the rate coefficient (denoted k_i) of the process is proportional to $n_c N_*$, where N_* is the number of particles of species c whose energy exceeds the threshold energy E_{min} and that are accepted for an inelastic process (probability of acceptance is p_i). So each particle undergoes two independent Bernoulli trials: first with probability of success p_E, second with probability of success p_i. Thus, N_* follows a binomial distribution $B(N_{p,c}, p_E p_i)$. Therefore,

Algorithm 1: Standard variable-weight DSMC collision algorithm

for $l \leftarrow 1$ **to** N_{coll} **do**
  sample particle of species c in current cell (with index j_1 and weight n_1);
  sample particle of species d in current cell (with index k_1 and weight n_2);
  $g \leftarrow$ compute relative velocity of particles j, k;
  $\{\sigma_i(g)\}_{i=1}^{N_{proc,cd}} \leftarrow$ compute cross-sections of all possible processes between species c, d;
  $\sigma_{tot}(g) \leftarrow \sum_i \sigma_i(g)$;
  $\mathcal{R} \leftarrow \mathcal{U}(0,1)$;
  if $\mathrm{MAX}(n_1, n_2)\sigma_{tot} > (n\sigma_{tot}g)_{\max}$ **then**
    $(n\sigma_{tot}g)_{\max} \leftarrow \mathrm{MAX}(n_1, n_2)\sigma_{tot}$;
  end
  $P \leftarrow \mathrm{MAX}(n_1, n_2)\sigma_{tot}g/(n\sigma_{tot}g)_{\max}$;
  if $\mathcal{R} < P$ **then**
    $\{P_i\}_{i=1}^{N_{proc,cd}} \leftarrow \{\sigma_i/\,\sigma_{tot}\}_{i=1}^{N_{proc,cd}}$;
    if $n_1 > n_2$ **then**
      create particle of species c with weight $(n_1 - n_2)$ and velocity of particle j_1;
      set weight of particle j_1 to n_2;
    else if $n_2 > n_1$ **then**
      create particle of species d with weight $(n_2 - n_1)$ and velocity of particle k_1;
      set weight of particle k_1 to n_1;
    end
    sample process i to model based on probabilities $\{P_i\}_{i=1}^{N_{proc,cd}}$;
    simulate process i for collision of particles j_1, k_1;
  end
end

for the expectation of the rate coefficient computed using acceptance-rejection (denoted k_i^{AR}) we have

$$\mathbb{E}[k_i^{AR}] \propto n_c N_{p,c} p_E p_i, \tag{4}$$

and for the variance of the rate coefficient we have

$$\mathrm{Var}(k_i^{AR}) \propto n_c^2 N_{p,c} p_E p_i (1 - p_E p_i). \tag{5}$$

If one uses event splitting, all particles with sufficiently high energy contribute towards the rate. So only one Bernoulli trial is performed, and the rate is proportional to $n_c p_i M_E$, where M_E is the number of particles whose energy exceeds the threshold energy E_{min}. Thus, M_E follows a binomial distribution $B(N_{p,c}, p_E)$. Therefore, for the expectation of the rate coefficient computed using event splitting (denoted k_i^{CS}) we again have

$$\mathbb{E}[k_i^{CS}] \propto n_c N_{p,c} p_E p_i, \tag{6}$$

and for the variance of the rate coefficient we now have

$$\mathrm{Var}(k_i^{CS}) \propto n_c^2 N_{p,c} p_E p_i (1 - p_E). \tag{7}$$

Algorithm 2: Event splitting variable-weight DSMC collision algorithm

```
for l ← 1 to N_coll do
    sample particle of species c in current cell (with index j_1 and weight n_1);
    sample particle of species d in current cell (with index k_1 and weight n_2);
    g ← compute relative velocity of particles j, k;
    {σ_i(g)}_{i=1}^{N_proc,cd} ← compute cross-sections of all possible processes between species c, d;
    σ_tot(g) ← Σ_i σ_i(g);
    R ← U(0,1);
    if MAX(n_1, n_2)σ_tot > (nσ_tot g)_max then
        (nσ_tot g)_max ← MAX(n_1, n_2)σ_tot;
    end
    P ← MAX(n_1, n_2)σ_tot g/(nσ_tot g)_max;
    if R < P then
        n_min ← MAX(n_1, n_2);
        {P_i}_{i=1}^{N_proc,cd} ← {σ_i/ σ_tot}_{i=1}^{N_proc,cd};
        if n_1 > n_2 then
            create particle of species c with weight (n_1 − n_2) and velocity of particle j_1;
            set weight of particle j_1 to n_2;
        else if n_2 > n_1 then
            create particle of species d with weight (n_2 − n_1) and velocity of particle k_1;
            set weight of particle k_1 to n_1;
        end
        for i ← 1 to N_proc,cd do
            reduce weight of particle j_1 by n_min P_i;
            reduce weight of particle k_1 by n_min P_i;
            create particle of species c with weight n_min P_i and velocity of particle j_1;
            create particle of species d with weight n_min P_i and velocity of particle k_1;
            collide the newly created particles, simulating collision process i;
        end
    end
end
```

Comparing the variances of the rate coefficients computed using the two different algorithms, we obtain

$$\frac{\mathrm{Var}(k_i^{AR})}{\mathrm{Var}(k_i^{CS})} = \frac{1 - p_E p_i}{1 - p_E}. \tag{8}$$

Since $p_i < 1$, we see that $\mathrm{Var}(k_i^{AR}) > \mathrm{Var}(k_i^{CS})$. Thus, in this simplified case the collision-splitting algorithm leads to a lower variance in the computed inelastic rate coefficient. The above analysis should not qualitatively change when relaxing the simplifications for more realistic simulations, and it is expected that the variance reduction is larger the more rare the event is that is being split.

We note that the proposed event splitting scheme has a computational cost. First, in the acceptance-rejection scheme, only a single computation of post-collision velocities is performed for each collision, either for the elastic or the inelastic process chosen to be simulated. In the splitting scheme, $N_{proc,cd}$ post-collision velocities are computed for each collision. Of course, the value of

$N_{proc,cd}$ may vary from timestep to timestep and from cell to cell, depending on the velocities of the particles in a given cell and the processes themselves – especially if the processes have an activation energy. Depending on the complexity of the scattering laws for the inelastic processes, the computation of post-collision velocities for all possible processes might incur a noticeable computational cost. Second, as seen from Fig. 4, the event splitting scheme leads to creation of a larger number of particles when compared to the acceptance-rejection scheme. The number of post-collision computational particles in the acceptance-rejection scheme leads is at most four (in case of an ionizing collision), whereas use of the splitting scheme leads to six particles instead of two (assuming the electron has sufficient energy to ionize the atom). For example, if one were to take N_{exc} electron-impact excitation reactions into consideration, then the maximum number of post-collision particles in the acceptance-rejection scheme will not change (an excitation reaction will produce three particles), but the splitting scheme will produce $6 + 2\,N_{exc}$ particles for each collision (again, assuming an incident electron with sufficiently high energy). This in turn will lead to the need for more frequent merging to control the particle counts. Merging not only requires additional computational effort, but also introduces additional error into the simulation (such as error in the heat flux and other higher-order moments of the distribution function [15, 38]).

One possible modification of the event splitting scheme that would reduce the number of post-collision particles and scattering computations involves process binning. The outline of such a scheme is as follows: different inelastic processes are grouped based on some criteria, and either an acceptance-rejection scheme is used to choose a collision group, and all the possible collision processes within that group are performed via event splitting, or the colliding particles are split according to the probabilities of each group, but only one collision from each collision group is chosen (via acceptance-rejection). The investigation of such a modification is left for future work. On the other hand, a possible computational benefit of the proposed algorithm is improved parallelization efficiency. After we split particles based on elastic and inelastic process probabilities, these split collision pairs are collided independently of each other, and this can be parallelized. On hybrid computational architectures, which aim to make the most efficient use of the large number of threads available on GPUs, this can provide an additional parallelization option, especially when the total number of physical grid elements is comparable to the number of available threads, and physical domain decomposition-based parallelization reaches its limit. In the present work, we focus on assessing the performance and properties of the splitting scheme for the simple test-case described above: an ionized argon flow with electron-impact ionization reactions (so an argon-electron collision can have two possible outcomes: an elastic scattering, or an ionization reaction). The questions of how the performance of the algorithm scales with the number of considered processes will be considered in future work.

3. Numerical results

All simulations were performed using an in-house variable-weight DSMC code developed at UT Austin. The simulations were performed serially on a workstation with an i7-4790 CPU and 32 GB of RAM.

3.1. 0-dimensional ionization rate computation

We assess the validity and performance of the event splitting scheme by using it to simulate a 0-dimensional $Ar/Ar^+/e^-$ mixture accelerated by a constant electric field. We consider elastic argon-argon collisions (modelled with the VHS model [1]), elastic electron-argon scattering, and electron-impact ionization reactions. The cross-sections for the electron-argon collisions are based

on BSR data [36] (see Fig. 1), and the Okhrimovskyy scattering model is used for argon-electron collisions [39]. The two electrons after an ionization event are assumed to have equal energy. The initial number density of the neutral background gas was set to $10^{23}\,\mathrm{m}^{-3}$, and the initial number density of the ions and electrons was $10^{16}\,\mathrm{m}^{-3}$. The simulation timestep was $10^{-14}\,\mathrm{s}$. The initial background gas temperature was set at 300 K, and the initial temperature of the electrons was 4.5 eV. We considered two values of the applied reduced electric field E/n: 50 Td and 500 Td to assess our collision scheme in two regimes: with a low ionization rate and a high ionization rate – the ionization rate differs by two orders of magnitude between these cases. We also varied the number of argon and electron particles used in the simulation.

To simplify the analysis, we used the same number of DSMC particles for the argon neutrals as for the electrons. The simulations were run with the following nominal numbers of particles for each species: 100, 200, 500, 1000, 2000, 5000, 10000, 30000. (The overall number of particles is twice each value). Since we use a variable-weight DSMC scheme, the number of particles varies from timestep to timestep. The numbers given above are the target numbers of particles passed onto the octree-based merging algorithm used for the present simulation [24]. The adaptive merge algorithm stops refinement of velocity space once the post-merge number of particles is less than the target number of particles. Merging was performed once the number of particles for a given species exceeded 110% of the target particle number. Thus, the number of particles in the simulation never exceeds 110% of the target particle number, but can occasionally drop lower than the target particle number, depending on where the particles are located in velocity space.

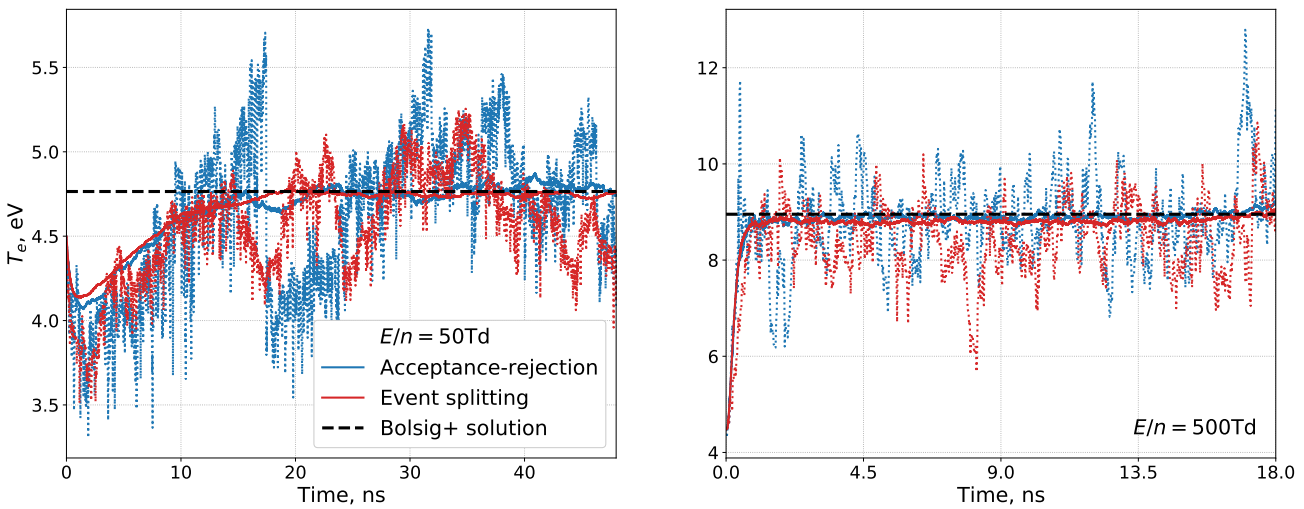

Figure 5: Temporal evolution of the electron temperature for an applied reduced field of 50 Td (left) and 500 Td (right). Dotted lines denote solutions obtained with target number of particles set to 100 for each species; solid lines denote solutions with a target number of particles set to 30000 for each species. The black dashed line denotes the solution given by the Bolsig+ solver.

After an initial transient period, the system reaches a quasi-steady state, characterized by a constant electron temperature and a constant rate of growth of the electron number density, as shown on Figs. 5–6. From this rate of growth, the value of the ionization rate coefficient at each timestep can be computed, and that is the main quantity of interest we will be considering in the present analysis. We can also compare both the quasi-steady state ionization rate coefficient and electron temperature to the values obtained via the Bolsig+ solver [40], which we treat as a reference solution. After the system had reached the quasi-steady state we gathered statistics over 600,000 timesteps (for an applied reduced field of 50 Td), and 300,000 timesteps (for an applied

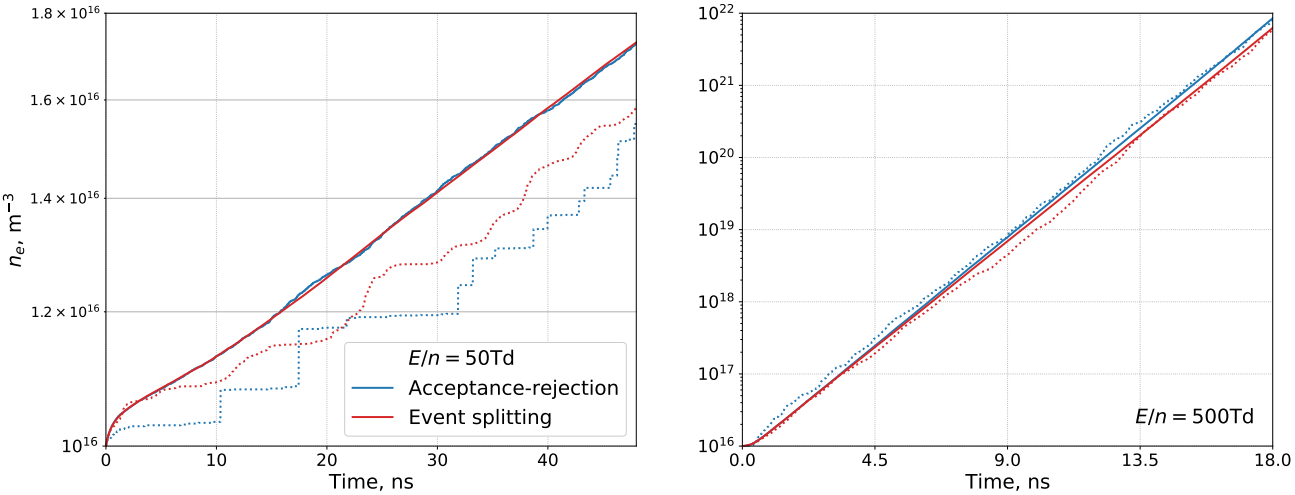

Figure 6: Temporal evolution of the electron number density for an applied reduced field of 50 Td (left) and 500 Td (right). Dotted lines denote solutions obtained with target number of particles set to 100 for each species; solid lines denote solutions with a target number of particles set to 30,000 for each species.

reduced field of 500 Td) to quantify the computational effort and the error in the ionization rate coefficient.

Assuming a Maxwellian velocity distribution of the electrons, and taking into account that 1) neutral species are much heavier than the electrons, and 2) neutral species are much slower than the electrons, the fraction x_{ion} of the electron distribution function that can cause an ionization reaction can be computed as

$$x_{ion} \approx \left(\frac{m_e}{2\pi k T_e}\right)^{3/2} \int_{v_{min}}^{\infty} 4\pi v_e^2 \exp\left(-\frac{m_e v_e^2}{2kT_e}\right) dv_e = \frac{4}{\sqrt{\pi}} \int_{\hat{v}_{min}}^{\infty} \hat{v}_e^2 \exp\left(-\hat{v}_e\right) d\hat{v}_e. \tag{9}$$

Here m_e is the electron mass, T_e is the electron temperature, v_{min} is the minimum velocity an electron should possess in order to ionize a neutral: $v_{min} = \sqrt{2E_{ion}/m_e}$, where E_{ion} is the ionization energy of the neutral species (15.76 eV for argon), and $\hat{v}_{min} = v_{min}/\sqrt{2kT_e/m_e}$. Thus, for the 500 Td applied reduced field roughly 30% of the electron distribution function has a sufficiently high enough energy to cause an ionization reaction after quasi-steady state has been reached, whereas for the 50 Td applied reduced field, only 8-9% of the electrons can cause an ionization reaction. From Figs. 5–6 it is evident that for the weak applied reduced field of 50 Td simulations with a low number of particles exhibit a higher level of noise. Both schemes are similar in that respect, which is to be expected. However, on Fig. 6 we can clearly see the all-or-nothing nature of the acceptance rejection scheme when the external field is weak and the number of particles in the simulation is low. The electron number density stays constant over noticeable periods of time, and exhibits rapid jumps during timesteps at which ionization events occur. Use of the event splitting scheme leads to a smoother behavior of the density — while the level of variation is still high, no rapid jumps can be seen. The variance in the electron temperature profiles (as seen on Fig. 5) are more similar between the event splitting and acceptance-rejection collision schemes, since the elastic electron-neutral collisions (which are not a low-probability event) cause changes in the electron temperature as well. Thus, the variance reduction achieved with the event splitting scheme depends on how dependent the quantity of interest is on the split event.

We consider the low-particle case (100 particles for each species) of the weak applied field (reduced field value of 50 Td) in more detail. Figure 7 shows the mean electron number density

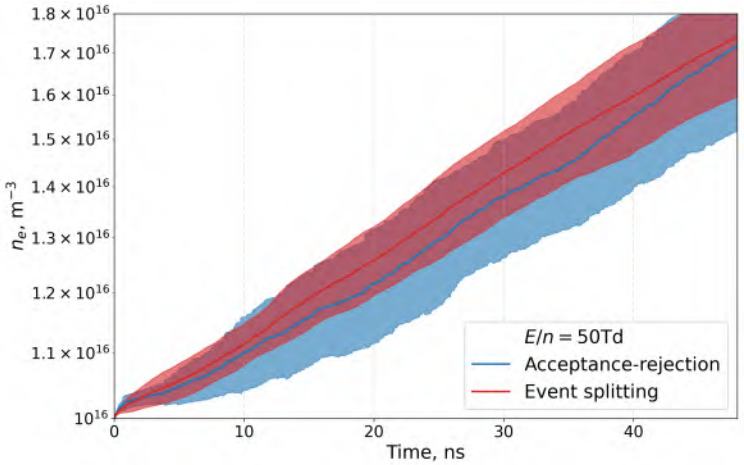

Figure 7: Temporal evolution of the ensemble-averaged electron number density for an applied reduced field of 50 Td, with the shaded regions showing a range between the 5-th and 95-th percentiles.

computed from 50 ensemble simulations (that used different random seeds), as well as its 5th-95th percentile range. It is evident that the range of values of the electron number density at a given time is narrower when one uses event splitting. Conversely, one can consider another quantity of interest: the time at which a certain electron number density is exceeded. Again, the uncertainty in the time required to achieve a specific electron number density value is lower when event splitting is used.

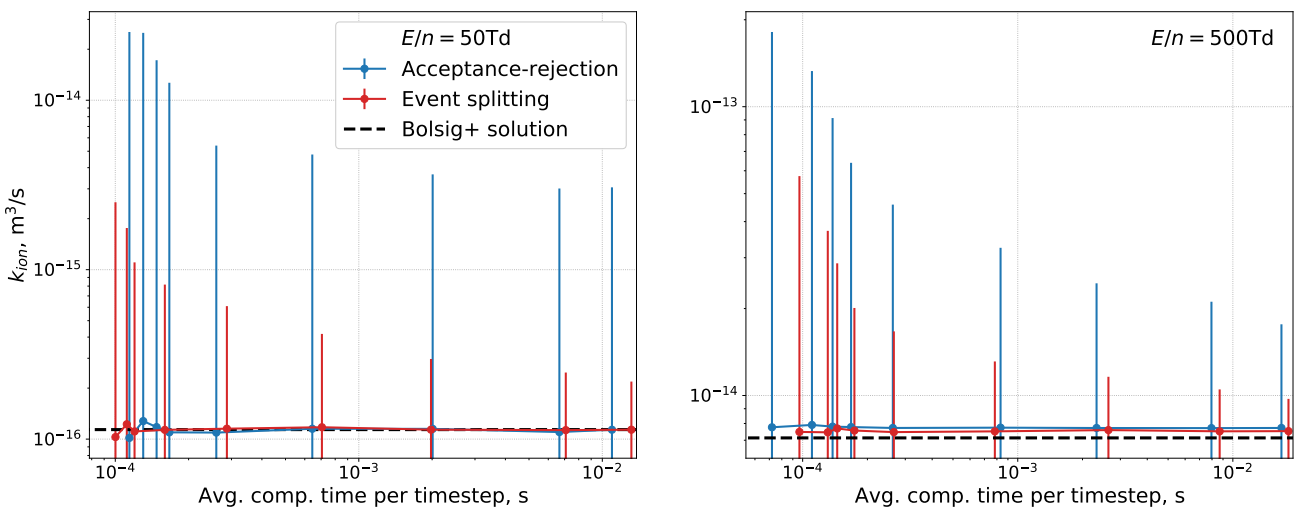

Figure 8: The ionization rate coefficient plotted against the computational time spent on a single timestep for an applied reduced field of 50 Td (left) and 500 Td (right) for different numbers of particles used in the simulation. Vertical bars denote a range of plus-minus standard deviation. The black dashed line is the solution given by the Bolsig+ solver.

We now focus on analyzing how the choice of a collision scheme and the number of particles used in a simulation affects the computed instantaneous rate coefficient. Figure 8 shows the average computed rate coefficient and the range (± 1 standard deviation) as a function of the average computational time spent on a single timestep for different numbers of particles used in the simulation. This computational time involves simulation of the collisions, the merging of particles (if required), and their acceleration. It should be noted that the value of the computed

standard deviation is up to 1-2 orders of magnitude larger than the value of the rate coefficient itself. Therefore, due to the log scale of the Y-axis, the range of ± 1 standard deviation is cut off at the lower end, as negative values cannot be represented on the plot. Such a plot allows us to not only verify the convergence of our simulations to a single value (which is close to that given by the Bolsig+ solver), but also assess the level of stochastic noise in the rate coefficient and compare the computational costs of the two collision algorithms. We can see that use of the event splitting scheme leads to a noticeable reduction in the noise in the ionization rate coefficient. As the number of particles is increased, the computational cost increases and the noise is reduced for both schemes. At higher numbers of particles, the event splitting scheme has slightly higher computational cost due to the larger number of scatterings computed at each timestep. However, achieving a comparable level of accuracy with the acceptance-rejection scheme would require many more particles and incur an even higher computational cost.

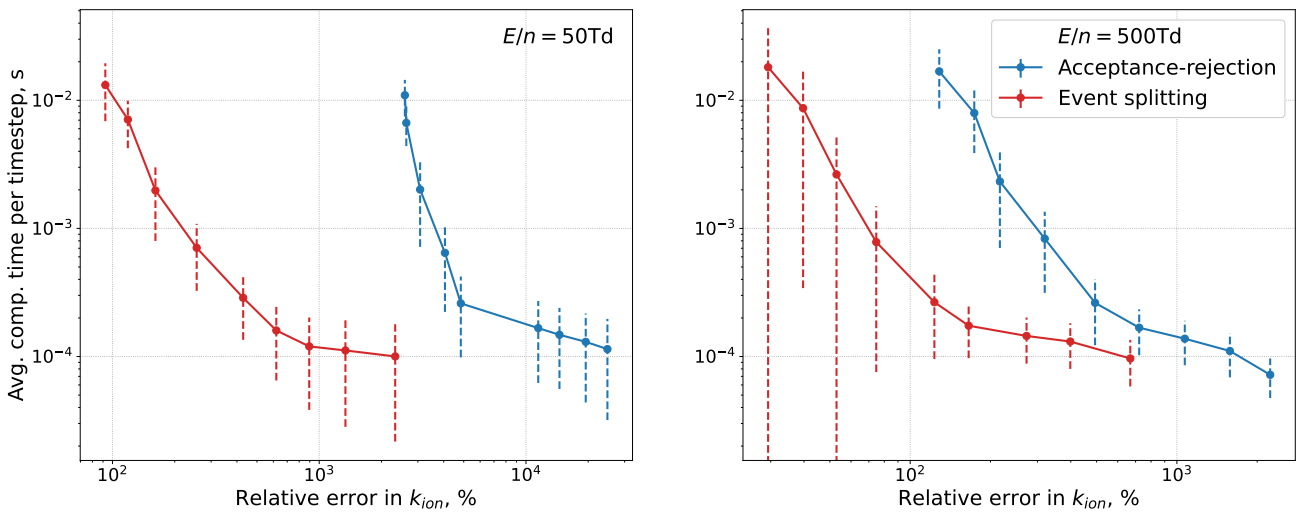

Figure 9: The average computational time spent on a single timestep plotted against the relative error in the ionization rate coefficient for an applied reduced field of 50 Td (left) and 500 Td (right). Vertical bars denote a range of plus-minus standard deviation in the computational time

To better understand the relation between computational cost and noise in the computed rate coefficient, we plot the average computational time per timestep against the relative error in the ionization rate coefficient. We define the relative error as

$$\text{Error}(k_{ion}) = 100\% \times \frac{1}{\overline{k}_{ion}} \sqrt{\frac{\sum_t \left(k_{ion}^{(t)} - \overline{k}_{ion}\right)^2}{N_t}}, \tag{10}$$

where $k_{ion}^{(t)}$ is the instantaneous ionization rate coefficient at timestep t, N_t is the number of timesteps the statistics are computed over, and $\overline{k}_{ion} = \sum_t k_{ion}^{(t)}/N_t$ is the mean of the computed ionization rate coefficient. Figure 9 shows the computational cost associated with a specific desired/acceptable level of noise in the ionization rate coefficient. From the figure we see that the event splitting scheme allows for a 10 to 25-fold reduction in the stochastic noise compared to the acceptance rejection scheme for the 50 Td case, and a 4 to 5-fold reduction for the 500 Td case, without a significant increase in computational cost. Conversely, achieving a specific desired/acceptable level of noise would require up to 100 times as many particles when using the acceptance-rejection scheme as compared to the event splitting scheme. We note that the variations

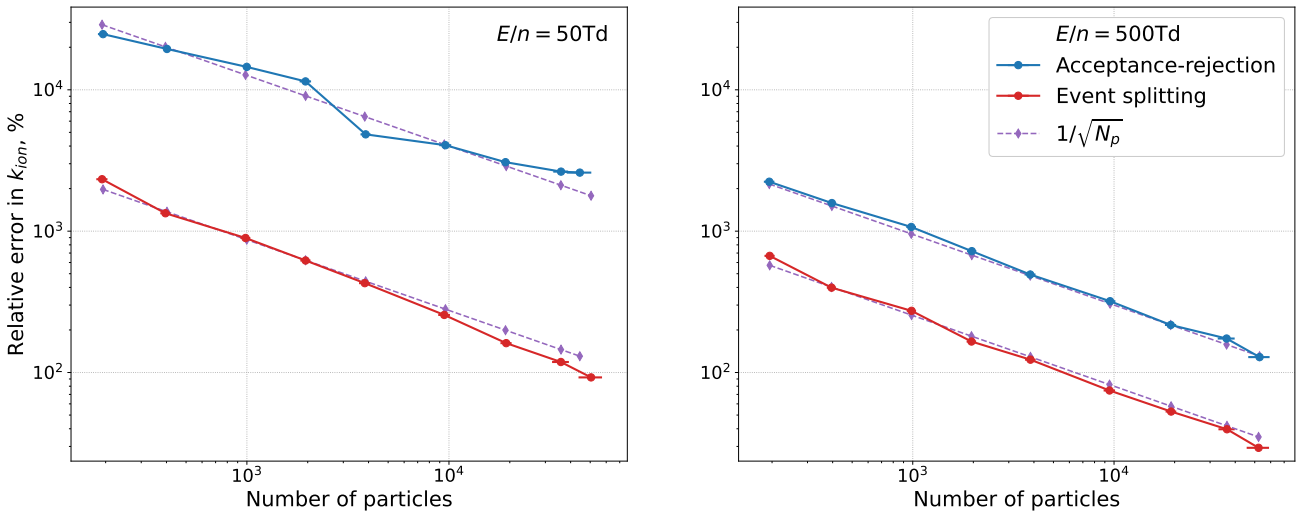

Figure 10: The relative error in the ionization rate coefficient plotted against the total number of particles in the simulation for an applied reduced field of 50 Td (left) and 500 Td (right). Horizontal bars denote a range of plus-minus standard deviation in the number of particles. Dashed lines show an inverse-square-root scaling trend.

in computational cost seen on Fig. 10 are due to the variations in the number of particles in the simulation, as well as the occasional particle merging (which has a cost of order $N_p \log(N_p)$ [24]).

Modern hybrid computer architectures require both computationally-efficient and also memory-efficient algorithms in order to reduce data transfer between nodes and CPU/GPU units. Hence, it is useful to see how the error scales with the (average) number of particles used. From Fig. 10 we see that the error of both collision schemes scales as $1/\sqrt{N_p}$ as expected, where N_p is the total number of relevant collisional particles in the simulation (the sum of argon neutral particles and electron particles). Thus, achieving the same of level of noise with the acceptance-rejection scheme as with the event splitting scheme would require the use of 100-200 times more particles for the 50 Td field case and 16-25 times more particles for the 500 Td field case.

Next, we examine the distribution of the instantaneous rate coefficients. As was shown on Fig. 6, depending on the flow conditions and the number of particles in the simulation, the acceptance-rejection scheme may have very few ionizing collision occurring over the course of the whole simulation. Given its all-or-nothing nature, when these collisions do occur, they correspond to a very large instantaneous ionization rate. Figure 11 shows the distributions (probability density functions obtained via binning) of the instantaneous ionization rate coefficients for different numbers of particles used in the simulation for the case of a 50 Td reduced field. We see that the acceptance-rejection collision scheme indeed produces a very wide distribution of rates (reaching up to approximately 5×10^{-13} m^3/s, even though the nominal rate is approximately 1.14×10^{-16} m^3, i.e. 3 orders of magnitude lower). The distribution of the ionization rate obtained via event splitting is narrower, and while its range is of the same order of magnitude when 200 particles are used in the simulations, it is rapidly reduced as the simulation particle count is increased. In fact, the distribution obtained with the event splitting algorithm is so narrow compared to that obtained with the acceptance-rejection algorithm that it can hardly be seen on the two lower subplots on Fig. 11. We therefore add insets to these subplots with a greatly expanded scale for the x-axis. We observe that the range of the distribution of instantaneous rate coefficients is comparable to the actual value of the ionization rate coefficient (the value obtained with the Bolsig+ solver is plotted as a black line). We also observe that the distribution still has a significant number of

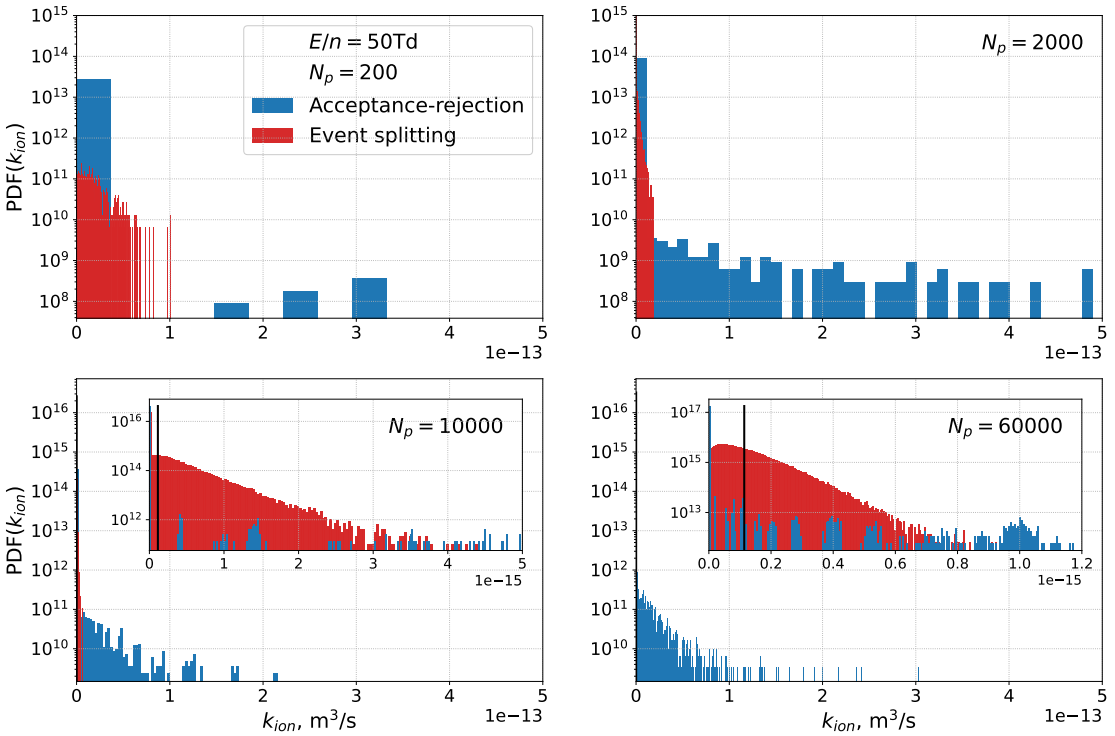

Figure 11: Probability density functions of the ionization rate coefficient for different total numbers of particles used in the simulation. The value of the applied reduced field is 50 Td. Insets have a decreased range on the x-axis in order to improve legibility. The solid black line denotes the value of the ionization rate coefficient given by the Bolsig+ solver.

zero-rate timesteps, as evidenced by the height of the leftmost red bar (which is proportional to the percentage of timesteps with a zero instantaneous rate coefficient).

Figure 12 shows the distributions of the instantaneous ionization rate coefficients for the case of a 500 Td external electric field. For low particle counts the overall qualitative behavior of the distributions is similar to the 50 Td case, except that the ranges of the two distributions produced by the different collision algorithms are more comparable to each other. For the highest number of particles used in the simulation (bottom-right), the distribution obtained with the event splitting algorithm has no zero-rate values and is centered nicely around the value given by the Bolsig+ solver. A Poisson distribution of ionization events is shown on the lower-right subplot to better highlight the range of values one can expect to obtain in the limiting case [41] and to show that at sufficiently high particle counts the event splitting scheme provides quite close results.

Finally, we look at the fraction of timesteps with a zero value of the instantaneous rate coefficient as a function of the number of particles used in the simulation. From Fig. 13 we see that this fraction strongly depends on number of particles for both collision algorithms, and that the slopes of these lines differ significantly between the two collision algorithms. For the 50 Td case, it is hard to discern the change in the fraction of timesteps with a zero-rate for the acceptance-rejection algorithm when plotted simultaneously with the curve obtained from the simulations that used the event splitting algorithm. When plotted separately (see inlay on Fig. 13), we can see the linear decrease, but the fraction of timesteps during which no ionization happens is still extremely high – at least 99%. For the 500 Td case, the behavior of the two algorithms is somewhat more comparable (which is to be expected, given the higher population of the high-velocity tails of the electron distribution function and a larger average ionization probability due to the higher temperature of

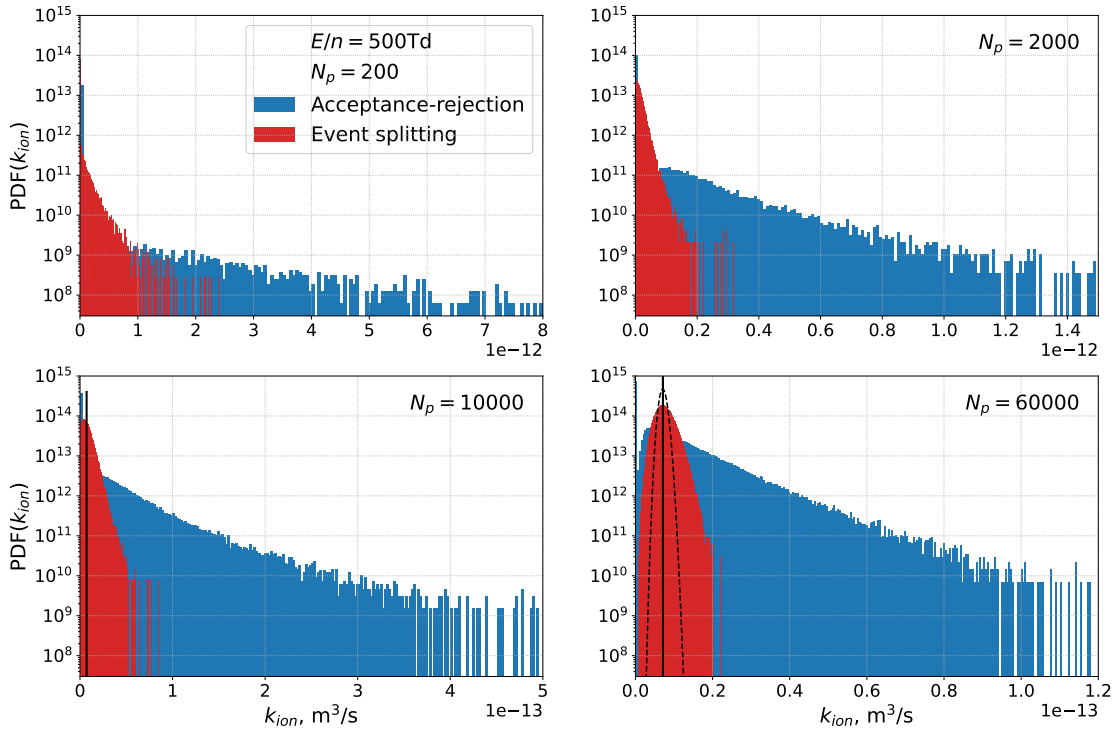

Figure 12: Probability density functions of the ionization rate coefficient for different total numbers of particles used in the simulation. The value of the applied reduced field is 500 Td. The solid black line denotes the value of the ionization rate coefficient given by the Bolsig+ solver, and the dashed black line in the lower right subplot corresponds to a Poisson distribution of ionization events.

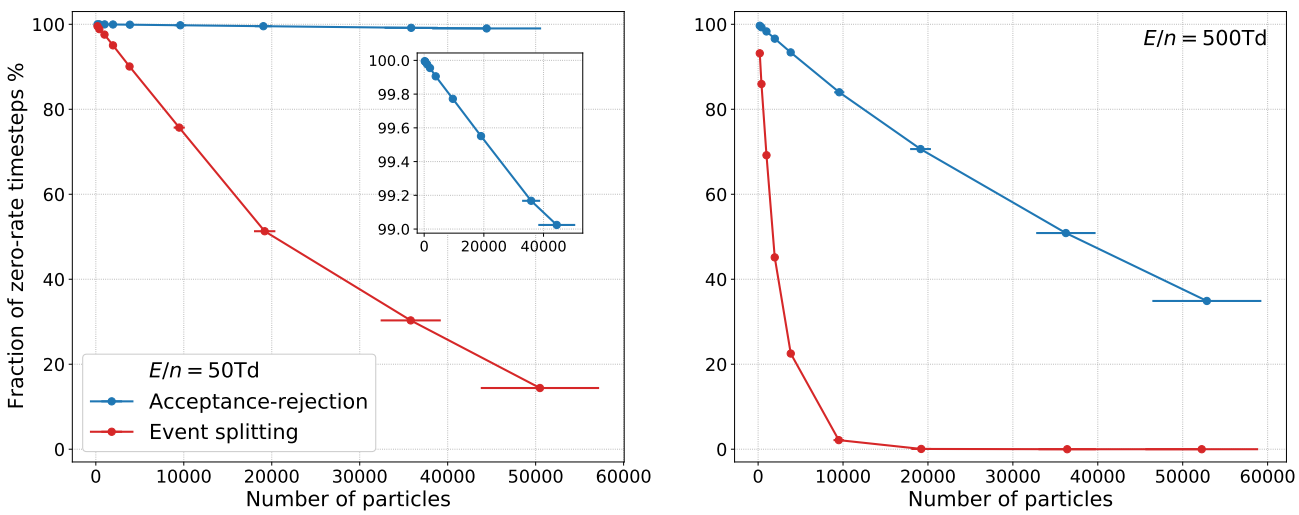

Figure 13: Fraction of timesteps during which no ionization events occur (instantaneous value of the ionization rate coefficient is zero) as a function of the average number of simulation particles for an applied reduced field of 50 Td (left) and 500 Td (right).

the electron distribution function). However, in this case the event splitting algorithm actually reaches a zero fraction of the zero-rate timesteps for the considered numbers of simulation particles, meaning that an ionization reaction occurs at each timestep in the simulation. Of course, it is not feasible to use the high number of simulation particles required to achieve no timesteps with zero ionization events except in model 0-dimensional simulations. However, a significant improvement

over the acceptance-rejection scheme can be observed in accurately modeling the ionization process even for low particle counts. We emphasize that it is not an increased particle count from event splitting that improves the accuracy of the simulations but the treatment of the collision process itself. Recall that we perform merging operations to limit the maximum possible number of particles in the simulation so the numbers of particles for the event splitting and acceptance-rejection cases are in the same range.

Lastly, returning to Figs. 6 and Figs. 7, we should also note not only the difference in the variance of the ionization rate coefficients, but also the differences in the electron number densities themselves computed with the different codes. As evidenced by the right-hand side of Fig. 6 (case of applied field of 500 Td), the differences in the number densities may reach up to 40%. In line with previous DSMC studies on the role of variance reduction in reducing the error in unsteady simulations [42], this points to another benefit of the event splitting scheme, namely, more accurate time-resolved values of quantities of interest.

Summarizing the analysis, we conclude that the proposed event splitting scheme leads to a significant improvement over the standard all-or-nothing acceptance-rejection scheme in terms of accurately representing the rate of a trace collision event. For the considered test case, the computational overhead of the scheme was evident only at high numbers of simulation particles, and was still insignificant. Overall, depending on the value of the external electric field considered, achieving the same level accuracy in the representation of the ionization rate coefficient using the acceptance-rejection scheme would require from 16 to 200 times as many simulation particles. We conclude that the computational benefit of the event splitting scheme increases as the process being split becomes more and more rare. For more complicated flows this gives a way to decide which processes can be treated using acceptance-rejection without a significant loss in accuracy.

3.2. 1-dimensional breakdown modeling

In this numerical test case, we simulate electric breakdown in an argon-filled gap between two plates held at a constant voltage difference. The model assumes no external circuit, i.e. an infinite power supply capable of holding the anode voltage at a fixed value, no matter the current density, is available. In the present study, we assume a constant electric field throughout the domain, and do not use a coupled particle-in-cell simulation. We seed the system with a small number density (3.1×10^{-6} that of the density of the neutral argon gas) of electrons and ions in the cell closest to the cathode in order to trigger breakdown when the voltage is sufficiently high. These seed electrons can result from many processes including cosmic radiation and other sources of photoemission from the electrode surface.

As electrons are accelerated across the gap, they ionize the argon atoms. The ions are accelerated towards the cathode and as they hit the cathode, they cause secondary emission. For the purposes of the present simulation, we assume a constant secondary emission probability of 10%, which is a reasonable approximation for low-energy Auger neutralization at a metal cathode surface [43]. Thus, we can perform event splitting not only for the electron-neutral collisions (as investigated in the 0-dimensional case), but also for the secondary-emission boundary conditions. Therefore, we perform 4 different simulations: one that uses acceptance-rejection for both collisions and the secondary emission boundary conditions, one that uses acceptance-rejection for the electron-argon collisions and event splitting for the boundary conditions, one that uses event splitting for the electron-argon collisions and acceptance-rejection for the boundary conditions, and one that uses event splitting for both the collisions and the secondary emission boundary conditions.

We take the background gas pressure to be 100 Torr and the gap width to be 30 μm. The applied voltage is taken to be 450 V. A timestep of 5.18×10^{-15} s was used for the simulations,

and 200 spatial cells were used to discretize the simulation domain. The target particle numbers (used in the merging algorithm) were taken to be 300 particles/cell for the argon neutrals, 450 particles/cell for the argon ions, and 900 particles/cell for the electrons.

In order to compare the level of noise of the different collision and boundary condition algorithms, we consider the current density at the anode as the quantity of interest – breakdown is indicated by an exponential rise in the current [7].

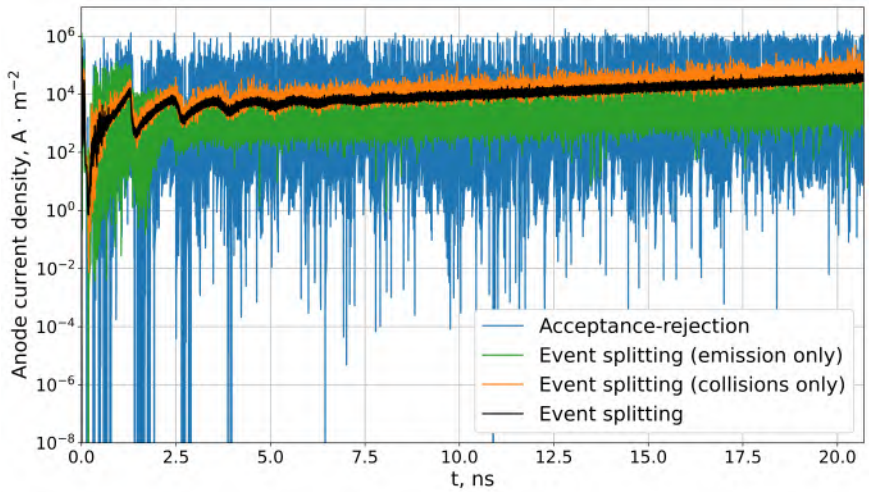

Figure 14: Anode current density as a function of time.

Figure 14 shows the anode current density j as a function of time. At the start of the simulation, the current exhibits oscillatory behavior with a period roughly equal to the ion transit time — the time it takes for an ion produced near the anode to accelerate and reach the cathode; for the given gap width and voltage across the gap, this time is 1.2 ns. We see that the level of noise is highest for the acceptance-rejection scheme, lower when boundary conditions are treated using the event splitting algorithm, even lower when ionization reactions are treated using the event splitting algorithm (and boundary conditions use acceptance-rejection), and is at its lowest when both the collisions and boundary conditions are treated using the event splitting scheme.

It should be noted that the results of our numerical simulations show significantly lower average numbers of electron and ion particles per cell in case acceptance-rejection is used for collisions ("Acceptance-rejection" and "Event splitting for emission only" curves on Fig. 14). This can be explained by the fact that due to the relatively low probability of ionization, few electron particles are produced. An insufficient number of electron particles leads to fewer collision events, which also contributes to a low particle count. In case acceptance-rejection is used to treat the emission boundary condition as well ("Acceptance-rejection" curve on Fig. 14), the situation is further exacerbated, since on average, only 1 in 10 ions hitting the cathode will produce an electron particle. Thus, the differences in the level of noise in the anode current seen on Fig. 14 are not only due to different levels of stochastic noise produced by the collision schemes, but also due to differing average particle counts. The paper on octree-binning-based merging [24] proposes a particle scheme to artificially increase the number of computational particles if too few are present in a grid cell. In order to provide a more consistent comparison between the different collision algorithms, we implemented a simple particle cloning scheme that increases the number of particles in a cell by a factor of M (where M is an integer) by creating $M-1$ copies of each particle at the same location in physical space and velocity space, with a weight of $1/M$ of the

original particle's weight. The original particle's weight is also set to $1/M$-th of its initial value. This cloning procedure is performed if the number of particles in a cell is less than 50.

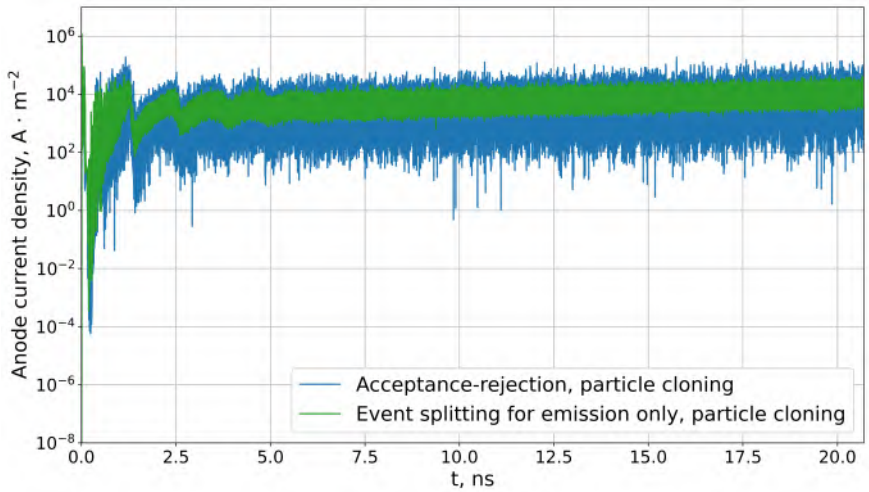

Figure 15: Anode current density as a function of time. Simulations performed using particle cloning of electrons and ions if particle counts in a cell became too low.

Figure 15 shows the anode current density as a function of time when particle cloning is used, and acceptance-rejection is used either for all processes (collisions and cathode emission) or for collisions only. We see that the level of noise is greatly reduced compared to Fig. 14 due to the imposed minimum threshold on the number of ions and electrons per cell.

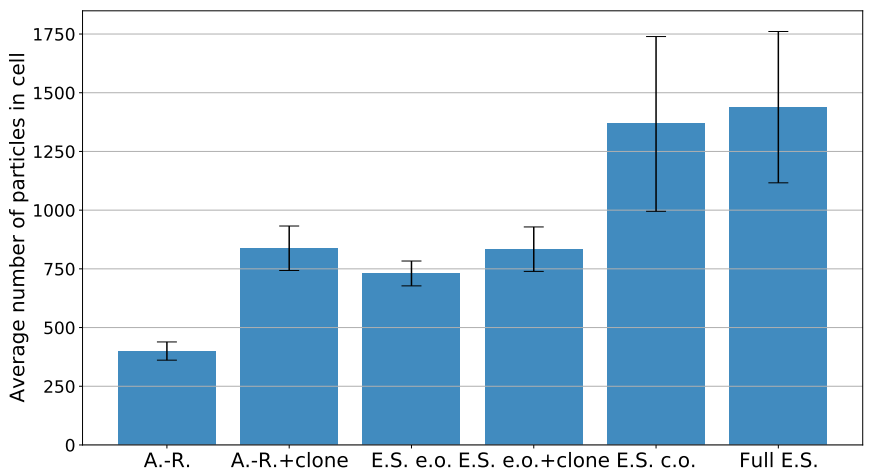

Figure 16: Average number of particles per cell computed using different collision and boundary condition algorithms. "A.R." denotes the acceptance-rejection algorithm, "E.S. e.o." denotes event splitting for emission only, "E.S. c.o." denotes event splitting for collisions only, "Full E.S." denotes event splitting for both collisions and emission. "+ clone" denotes simulations performed with particle cloning. Bars denote a range of one standard deviation.

Fig. 16 shows the average number of particles per cell computed using different collision and boundary condition algorithms. We see that when particle cloning is used in conjunction with acceptance-rejection treatment of collisions ("A.-R. + clone" and "E.-S. e.o. + clone" test cases), the average number of particles per cell is noticeably larger (and closer to the expected value) than when no cloning is performed. Use of particle cloning increases the average number of particles in a

cell, and as seen previously on Fig. 15, leads to a reduction in the stochastic noise. However, even in case more complicated particle cloning algorithms are used (such as the one given in [24]), they do not increase the accuracy of the simulation by themselves – it is only through the interaction of the newly created particles with other particles and with surfaces that the fidelity of the simulations is increased. Thus, another advantage of the event splitting algorithm is the reduced need for artificial particle cloning, since one can take advantage of the branching in the physical processes occurring in the flow (by splitting during collisions and gas-surface interactions) to maintain a sufficient particle count.

In order to quantify the level of noise in the anode current density, we perform a linear regression on the (log-scaled) computed anode current densities (after 16 ns have elapsed) in order to extract the trendline. Then we calculate the root-mean-square-log-error (RMSLE) of the computed anode current densities relative to the trend-line (we use RMSLE as a metric due to the order of magnitude variations in the current density over time):

$$\mathrm{RMSLE}(j) = \sqrt{\frac{\sum_t \left(\log_{10} j^{(t)} - \log_{10} \bar{j}(t)\right)^2}{N_t}}. \tag{11}$$

Here $j^{(t)}$ is the anode current density at time t, and the $\log_{10} \bar{j}(t)$ is the linear regression estimate of the logarithm of the anode current density at time t, given by:

$$\log_{10} \bar{j}(t) = A_0 + A_1 t. \tag{12}$$

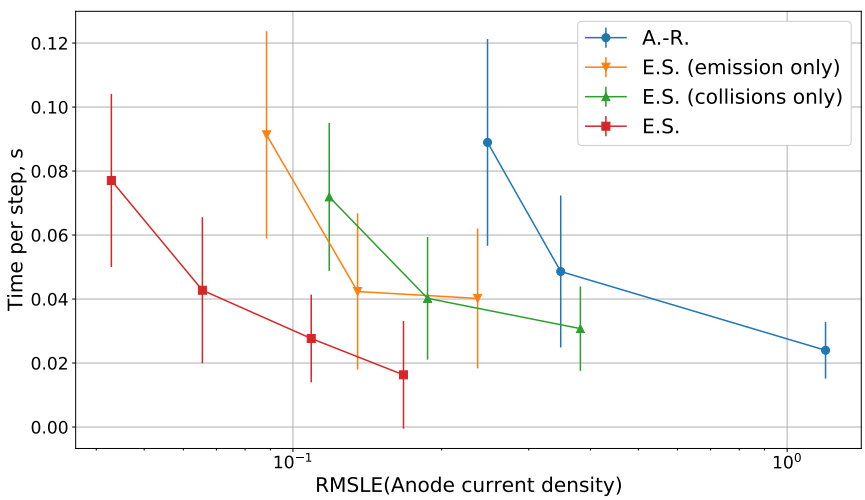

Figure 17: Computational time per step plotted against noise in the log-scaled anode current density. Bars denote a range of plus-minus one standard deviation.

Figure 17 shows the average computational time per step against the noise in the anode current density, as computed via Eqn. 11.

To provide a more detailed comparison of the cost-vs-accuracy, we ran additional simulations with different average numbers of particles per cell, ranging from approximately 110 particles/cell to 1450 particles/cell for the full event splitting simulations, from 330 particles/cell to 1360 particles/cell for the simulations that used event splitting for electron-neutral collisions, from 730 particles/cell to 1560 particles/cell for the simulations that used event splitting for secondary electron emission, and from 400 particles/cell to 1570 particles/cell for the simulations that used

acceptance-rejection for all collisional and boundary processes. We see that use of the event splitting scheme provides an reduction in the stochastic noise over the acceptance-rejection algorithm. The greatest benefit is observed when event splitting is used for both the electron-neutral collisions and boundary conditions. The gains in accuracy (over standard variable weight DSMC) obtained when event splitting is used for either the electron-neutral collisions or for the secondary emission process are similar. It should be also noted that for the given conditions (pressure of 100 Torr, gap width of 30 μm, and a constant applied voltage of 450 V), the reduced electric field is equal to 4660 Td. Due to the strong (and constant) field, the electrons are accelerated to a high speed, and have a high temperature (reaching up to 35-40 eV near the anode). Therefore, a large number of electron-neutral collisions will have a relatively high ionization probability, and the benefits of the event splitting scheme are somewhat diminished. However, for different simulation conditions such as a smaller applied voltage, or in gases with additional electron-collision processes which will reduce the average electron temperature, the gains in accuracy obtained with the event splitting scheme will be more significant.

In a general case, one can perform event splitting for specific processes in the flow, and perform a similar cost-vs-error analysis in order to obtain insight into the most significant processes driving the flow and which processes benefit most from application of the event splitting scheme.

4. Conclusions

We devised and tested a new DSMC collision scheme to improve the modeling of low-probability collision events and applied it to an ionized argon flow accelerated by a constant electric field. We investigated the behavior of the instantaneous ionization rate coefficient and the dependence of the magnitude of the stochastic noise on the number of particles used in the simulation and the type of collision scheme used. We show that the proposed event splitting scheme greatly reduces the noise in the modeling of the ionization reactions, especially when the external applied field is low and very few electrons have a sufficiently high energy to cause an ionization reaction. Achieving a similar level of noise with the acceptance-rejection scheme would require a significant increase in the number of simulation particles. Extending the scheme to modeling of secondary emission from a cathode, we model a simulate a breakdown in a narrow 1-dimensional gap. We show that use of event splitting for treatment of electron-neutral collisions and/or secondary emission boundary conditions greatly improves the level of noise in the anode current compared to acceptance-rejection sampling, and reduces the need for particle cloning.

Future work on the event splitting algorithm will consider its application to more complex flows, and investigate the scaling of its computational cost with the number of processes considered. Possible extension of the scheme include (i) allowing for partial collisions (currently, acceptance-rejection is still used to select or reject possible collision pairs), (ii) splitting based on certain collision process *groups* (and not individual collision processes), (iii) applying event splitting to scattering laws (e.g., diffuse boundary conditions, preferential scattering for electron-neutral collisions, etc.), and (iv) combining the scheme with the fractional collision method developed in [20].

5. Acknowledgments

This work was supported by Sandia National Laboratories. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

Georgii Oblapenko acknowledges the funding provided by the Alexander von Humboldt foundation for his stay as a guest postdoctoral researcher at the German Aerospace Agency (DLR).

[1] G. A. Bird, Molecular Gas Dynamics and the Direct Simulation of Gas Flows, Clarendon, Oxford, England, UK, 1994.

[2] H. Zhang, M. Xie, Dsmc simulation of non-premixed combustion of H2/O2 in a Y-shaped microchannel, Nanoscale Microscale Thermophys. Eng. 19 (2015) 31–62.

[3] I. Sohn, Z. Li, D. Levin, M. Modest, Coupled DSMC-PMC radiation simulations of a hypersonic reentry, J. Thermophys. Heat Transfer 26 (2012) 22–35.

[4] J. Zhang, D. Goldstein, P. Varghese, N. Gimelshein, S. Gimelshein, D. Levin, Simulation of gas dynamics and radiation in volcanic plumes on Io, Icarus 163 (2003) 182–197.

[5] D. Goldstein, P. Varghese, Rarefied gas dynamics on a planetary scale, Phys. Fluids 23 (2011) 030608.

[6] H. Deng, T. Ozawa, D. Levin, Analysis of chemistry models for DSMC simulations of the atmosphere of Io, J. Thermophys. Heat Transfer 26 (2012) 36–46.

[7] C. H. Moore, M. M. Hopkins, P. S. Crozier, J. J. Boerner, L. C. Musson, R. W. Hooper, M. T. Bettencourt, 1D PIC-DSMC simulations of breakdown in microscale gaps, in: AIP Conf. Proc, volume 1501, AIP, pp. 629–636.

[8] A. Fierro, C. Moore, B. Yee, M. Hopkins, Three-dimensional kinetic modeling of streamer propagation in a nitrogen/helium gas mixture, Plasma Sources Sci. Technol. 27 (2018) 105008.

[9] R. Jambunathan, D. A. Levin, Kinetic, 3-D, PIC-DSMC simulations of ion thruster plumes and the backflow region, IEEE Trans. Plasma Sci. 48 (2020) 2017–2034.

[10] F. Lumpkin, III, P. Stuart, G. Le Beau, Enhanced analyses of plume impingement during Shuttle-Mir docking using a combined CFD and DSMC methodology, in: 31st Thermophysics Conference, p. 1877.

[11] W.-L. Wang, I. Boyd, Hybrid DSMC-CFD simulations of hypersonic flow over sharp and blunted bodies, in: 36th AIAA Thermophysics Conference, p. 3644.

[12] M. H. Gorji, P. Jenny, Fokker–Planck–DSMC algorithm for simulations of rarefied gas flows, J. Comput. Phys. 287 (2015) 110–129.

[13] C. Hepp, M. Grabe, K. Hannemann, A kinetic Fokker–Planck approach to model hard-sphere gas mixtures, Phys. Fluids 32 (2020) 027103.

[14] T.-J. Pan, K. A. Stephani, Investigation of velocity-space coupling approach in DSMC for tail-driven processes, in: AIP Conf. Proc., volume 1786, AIP Publishing, p. 050017.

[15] G. Oblapenko, D. B. Goldstein, P. Varghese, C. Moore, A velocity space hybridization-based Boltzmann equation solver, J. Comput. Phys. 408 (2020) 109302.

[16] I. D. Boyd, Conservative species weighting scheme for the direct simulation Monte Carlo method, J. Thermophys. Heat Transfer 10 (1996) 579–585.

[17] S. Rjasanow, W. Wagner, A stochastic weighted particle method for the Boltzmann equation, J. Comput. Phys. 124 (1996) 243–253.

[18] V. V. Serikov, S. Kawamoto, K. Nanbu, Particle-in-cell plus direct simulation Monte Carlo (PIC-DSMC) approach for self-consistent plasma-gas simulations, IEEE Trans. Plasma Sci. 27 (1999) 1389–1398.

[19] R. S. Martin, J.-L. Cambier, Moment preserving adaptive particle weights using octree velocity distributions for PIC simulations, in: AIP Conf. Proc., volume 1501, AIP, pp. 872–879.

[20] S. Araki, R. Martin, Interspecies fractional collisions, Phys. Plasmas 27 (2020) 033504.

[21] S. Rjasanow, T. Schreiber, W. Wagner, Reduction of the number of particles in the stochastic weighted particle method for the Boltzmann equation, J. Comput. Phys. 145 (1998) 382–405.

[22] A. Vikhansky, M. Kraft, Conservative method for the reduction of the number of particles in the Monte Carlo simulation method for kinetic equations, J. Comput. Phys. 203 (2005) 371–378.

[23] D. R. Welch, T. C. Genoni, R. E. Clark, D. V. Rose, Adaptive particle management in a particle-in-cell code, J. Comput. Phys. 227 (2007) 143–155.

[24] R. S. Martin, J.-L. Cambier, Octree particle management for DSMC and PIC simulations, J. Comput. Phys. 327 (2016) 943–966.

[25] M. Vranic, T. Grismayer, J. L. Martins, R. A. Fonseca, L. O. Silva, Particle merging algorithm for PIC codes, Comput. Phys. Commun. 191 (2015) 65–73.

[26] M. Pfeiffer, A. Mirza, C.-D. Munz, S. Fasoulas, Two statistical particle split and merge methods for particle-in-cell codes, Comput. Phys. Commun. 191 (2015) 9–24.

[27] P. T. Luu, T. Tückmantel, A. Pukhov, Voronoi particle merging algorithm for PIC codes, Comput. Phys. Commun. 202 (2016) 165–174.

[28] D. Faghihi, V. Carey, C. Michoski, R. Hager, S. Janhunen, C.-S. Chang, R. Moser, Moment preserving constrained resampling with applications to particle-in-cell methods, J. Comput. Phys. (2020) 109317.

[29] H. Gorji, S. Küchlin, P. Jenny, Particle number control for direct simulation Monte-Carlo methodology using kernel estimates, Phys. Fluids 31 (2019) 062008.

[30] R. S. Martin, J.-L. Cambier, Low noise fractional NTC collisions for DSMC, in: AIP Conf. Proc, volume 1501, AIP, pp. 615–620.

[31] D. Petkow, M. Zaretskaya, A. Kling, G. Herdrich, Progress in probabilistic modelling of atomic spontaneous emission processes in DSMC, in: Proceedings of the 5th International Workshop on Radiation of High Temperature Gases in Atmospheric Entry, ESA Special Publication (2012), pp. 16–19.

[32] A. Fierro, J. Stephens, S. Beeson, J. Dickens, A. Neuber, Discrete photon implementation for plasma simulations, Phys. Plasmas 23 (2016) 013506.

[33] D. P. Schmidt, C. Rutland, A new droplet collision algorithm, J. Comput. Phys. 164 (2000) 62–80.

[34] J. Strand, D. Goldstein, Sensitivity analysis for DSMC simulations of high-temperature air chemistry, in: 49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, p. 535.

[35] J. Strand, D. Goldstein, Application of Bayesian statistical methods for the analysis of DSMC simulations of hypersonic shocks, in: 41st AIAA Fluid Dynamics Conference and Exhibit, p. 3705.

[36] L. Pitchford, L. Alves, K. Bartschat, S. Biagi, M. Bordage, A. Phelps, C. Ferreira, G. Hagelaar, W. Morgan, S. Pancheshnyi, et al., Comparisons of sets of electron–neutral scattering cross sections and swarm parameters in noble gases: I. Argon, J. Phys. D: Appl. Phys. 46 (2013) 334001.

[37] O. Zatsarinny, Y. Wang, K. Bartschat, Electron-impact excitation of argon at intermediate energies, Phys. Rev. A 89 (2014) 022706.

[38] S. Lama, J. Zweck, M. Goeckner, A higher order moment preserving reduction scheme for the Stochastic Weighted Particle Method, SIAM J. Sci. Comp. 42 (2020) A2889–A2909.

[39] A. Okhrimovskyy, A. Bogaerts, R. Gijbels, Electron anisotropic scattering in gases: A formula for Monte Carlo simulations, Phys. Rev. E 65 (2002) 037402.

[40] G. Hagelaar, L. Pitchford, Solving the Boltzmann equation to obtain electron transport coefficients and rate coefficients for fluid models, Plasma Sources Sci. Technol. 14 (2005) 722.

[41] L. C. Kröger, W. A. Kopp, M. Döntgen, K. Leonhard, Assessing statistical uncertainties of rare events in reactive molecular dynamics simulations, J. Chem. Theory Comput. 13 (2017) 3955–3960.

[42] W. McDoniel, On-Average Error in Zero-Dimensional DSMC., Technical Report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 2019.

[43] M. A. Lieberman, A. J. Lichtenberg, Principles of plasma discharges and materials processing, John Wiley & Sons, 2005.