



Stochastic and self-consistent 3D modeling of streamer discharge trees with Kinetic Monte Carlo

Robert Marskar

SINTEF Energy Research, Sem Sælands vei 11, 7034 Trondheim, Norway

ARTICLE INFO

Keywords:

Streamer
Particle-In-Cell
Cartesian AMR
Parallel computing

ABSTRACT

This paper contains the foundation for a new Particle-In-Cell model for gas discharges, based on Îto diffusion and Kinetic Monte Carlo (KMC). In the new model the electrons are described with a microscopic drift-diffusion model rather than a macroscopic one. We discuss the connection of the Îto-KMC model to the equations of fluctuating hydrodynamics and the advection-diffusion-reaction equation which is conventionally used for simulating streamer discharges. The new model is coupled to a particle description of photoionization, providing a non-kinetic all-particle method with several attractive properties, such as: 1) Taking the same input as a fluid model, e.g. mobility coefficients, diffusion coefficients, and reaction rates. 2) Guaranteed non-negative densities. 3) Intrinsic support for reactive and diffusive fluctuations. 4) Exceptional stability properties. The model is implemented as a particle-mesh model on cut-cell grids with Cartesian adaptive mesh refinement. Positive streamer discharges in atmospheric air are considered as the primary application example, and we demonstrate that we can self-consistently simulate large discharge trees.

1. Introduction

Substantial efforts have been made in order to understand the nature of streamer discharges [1], which is a specific type of transient and filamentary plasma. Streamers occur naturally as precursors to electric sparks, and also appear as sprite discharges in the upper atmosphere [2–5]. They are also highly useful for CO₂ conversion [6–9], in plasma assisted combustion [10–15], plasma catalysis [16–21], and plasma medicine [22]. Streamers are inherently three-dimensional structures that usually appear in bundles or in the shape of discharge trees. These develop through repetitive branching, which is a fundamental property of streamers [23,24]. For positive streamer discharges in air, the amount of photoionization in front of the streamer strongly affects the degree of branching [25,26]. When more photoelectrons are generated in front of positive streamers the amount of streamer branching is reduced.

Although single streamers are now relatively well understood, discharge trees are the more relevant structures since they appear in virtually all applications involving streamer discharges. Fig. 1 shows an example of such a structure for a positive streamer discharge in air at atmospheric pressure and temperature. Clearly, this structure requires full 3D modeling over many orders of magnitude in both space and time, and is therefore quite difficult to describe quantitatively. Most contemporary computational models can only solve for at most a few filaments. See e.g. [25–30] for recent results with fluid models, or [31–35] for kinetic particle models.

Several researchers have questioned the feasibility of using fluid and particle models for obtaining numerical solutions that describe entire discharge trees [1,36,37], such as those in Fig. 1. Given the difficulties in simulating even just an isolated streamer

E-mail address: robert.marskar@sintef.no.

<https://doi.org/10.1016/j.jcp.2024.112858>

Received 16 January 2023; Received in revised form 10 September 2023; Accepted 10 February 2024

Available online 20 February 2024

0021-9991/© 2024 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. Example images of positive streamer discharge trees in technical air at 1 bar pressure. The images show the accumulated UV and optical light emission from the discharge (each frame shows a different discharge).

filament [38], it is easy to understand why such claims are made. In a recent review Nijdam et al. [1] remarked that although fluid and particle simulations of streamers with tens or hundreds of branches are computationally unfeasible, reduced-order models of single filaments [36,39] are candidates for improved tree and fractal based models [40–43]. Although these models have been heralded for quite some time and could be used for simulating discharge trees, they are still in their infancy and they are unfortunately also excessively simplified. Many natural phenomena like branching and charge transport do not self-consistently evolve from the model itself, and the lack of a density (or particle) description also complicates quantitative descriptions of the chemistry in the streamer channels.

Currently, drift-diffusion fluid models in the local field approximation (LFA) are most frequently used for studying streamer discharges. With fluid models the advection-diffusion-reaction equation for the electron density is discretized on an Eulerian grid, and the plasma density is updated in time using either explicit or implicit time integration [38]. There are several well-known numerical restrictions for fluid models, such as the existence of a Courant-Friedrichs-Lewy (CFL) condition on the time step Δt , or restriction by the dielectric relaxation time ϵ_0/σ where σ is the conductivity of the plasma. The latter can be avoided by using semi-implicit formulations [44]. Infrequently mentioned is the fact that there is a rather fundamental requirement on the spatial resolution Δx as well [45], which applies to both explicit and implicit temporal discretizations. One issue that is often faced in simulation codes is that explicit codes at best have time steps $\Delta t \propto \Delta x$, which leads to an undesired scaling of computational resources. For example, refining the grid $\Delta x \rightarrow \Delta x/2$ doubles the amount of grid cells per coordinate direction, and requires twice as many time steps. Implicit codes can decouple Δt from Δx , but it is not clear how to obtain a scalable implicit discretization in the context of the frequent regridding that is a de-facto requirement for large scale 3D streamer discharge simulations [26–29].

Cognizant of the above issues, we have developed a new model based on a *microscopic* drift-diffusion model rather than a macroscopic one. This is combined with mesoscopic reaction algorithms for describing the stochastic plasma chemistry, i.e. we replace the conventionally used deterministic chemistry by a Kinetic Monte Carlo (KMC) algorithm. The new model takes particle discreteness, random collisions, and stochastic reactions into account. Fundamentally, the model is a non-kinetic Particle-In-Cell (PIC) model, and it is indeed ironic that this is actually a helpful model since the switch from a fluid to a particle description is usually associated with an increased computational cost. But the new model has no fundamental restriction on Δx or Δt , and numerical tests show that it is exceptionally stable in both space and time. Importantly, we are also achieving a decoupling of Δt from Δx without requiring an implicit discretization. This renders the new model capable of obtaining numerical solutions not only for single filaments, but also for comparatively large discharge trees. In this paper we use the model to investigate laboratory discharges, but the model itself is applicable to many other types of streamers (e.g. sprites).

This paper has two main goals: 1) A thorough presentation of the model, with details as to how it can be implemented with robust and scalable computer algorithms. 2) A capability demonstration for self-consistent simulation of discharge trees at the laboratory scale, similar to the ones shown in Fig. 1. The organization of this paper is as follows. Section 2 presents a computational prelude that focuses on finite-volume discretization issues for fluid models, for both explicit and implicit time discretizations. In section 3 we formulate the new model and discuss its connection to the conventional fluid model. Section 4 contains the numerical discretization of the model. In section 5 we provide some numerical tests of the model, and some concluding remarks are provided in section 6.

2. Prelude

We first consider an underlying issue facing the stability properties of discretized fluid models in the LFA. Our line of reasoning follows Villa et al. [45] who proved that the spatial resolution for fluid models must essentially resolve the avalanche length for the solution to remain bounded in time. Consider a one-dimensional advection-reaction model for the electron density n , for the moment ignoring electron diffusion:

$$\partial_t n = -v \partial_x n + \alpha v n, \quad (1)$$

where $v > 0$ is a constant electron velocity and $\alpha > 0$ is a constant ionization coefficient. The exact solution is

$$n(x, t) = \tilde{n}(x - vt) \exp(\alpha x), \quad (2)$$

where $\tilde{n}(x - vt)$ is some initial function. If \tilde{n} is bounded in time, so is $n(x, t)$.

We now consider the numerical discretization of equation (1) on a one-dimensional Cartesian grid with grid points $x_i = i\Delta x$ where Δx is the grid point spacing and i is the grid index. Each grid cell spans the volume $[x_i - \Delta x/2, x_i + \Delta x/2]$. A first order finite-volume upwind discretization in space with an implicit Euler discretization in time for equation (1) yields

$$\begin{aligned} n_i^{k+1} &= \frac{n_i^k}{1 + \frac{v\Delta t}{\Delta x} - \alpha v\Delta t} + \frac{v\Delta t}{\Delta x} n_{i-1}^{k+1} \\ &\geq \frac{n_i^k}{1 + \xi(1 - \alpha\Delta x)}, \end{aligned} \quad (3)$$

where $\xi \equiv v\Delta t/\Delta x \geq 0$ is the Courant number. The solution n_i^{k+1} is bounded in time only if

$$\alpha\Delta x \leq 1. \quad (4)$$

Here, the discretization is fully implicit but it is only conditionally stable and non-negative.

More generally, the underlying stability issues are related to the advective-reactive coupling. Using a Godunov splitting for equation (1) with explicit fractional Euler steps yields

$$\begin{aligned} n_i^\dagger &= n_i^k(1 - \xi) + \xi n_{i-1}^k, \\ n_i^{k+1} &= n_i^\dagger + \alpha v\Delta t n_i^\dagger \\ &= [n_i^k(1 - \xi) + \xi n_{i-1}^k](1 + \xi\alpha\Delta x) \\ &\geq n_i^k(1 - \xi)(1 + \xi\alpha\Delta x). \end{aligned} \quad (5)$$

The stability region is now $\alpha\Delta x \leq 1/(1 - \xi)$, i.e. the discretization is more stable for larger time steps. Splitting methods expand the stability region because they advect electrons out of the grid cell before they react.

Several codes [27,38] use second order slope-limited discretizations, but these discretizations are not fundamentally more stable. An analysis is more difficult in this case, but one only needs to observe that slope-limited schemes default to piecewise constant reconstruction if there is a local maximum or a large gradient in the solution, and in this case one again obtains the stability limit $\alpha\Delta x \leq 1$. The problem dimensionality and presence of diffusion also affects the stability region. However, since equation (1) is a subset of multi-dimensional simulations where $v_y = v_z = 0$ and $v_x \neq 0$, the stability limit applies to multi-dimensional simulations as well.

The analysis above is quite simplified and ignores the fact that the streamer is a moving structure and thus that any potential instability regions $\alpha\Delta x > 1$ move with the solution. In practice, the situation is far less dire and one may still observe that n_i^k remains bounded even for quite significant violations of $\alpha\Delta x \leq 1$. It is nonetheless clear that caution is needed for fluid simulations since numerical underresolution is fundamentally capable of enabling unbounded growth in the plasma density and corresponding non-physical diverging growth of the electric field E , i.e. $E \rightarrow \infty$ as $n \rightarrow \infty$. Such instabilities have been scrutinized in recent years, and they are particularly relevant in the cathode sheath [46] and in the context of so-called stagnant positive streamers [47]. Although the LFA is often identified as the culprit [46], the underlying issue is also present in the transport equation itself. Furthermore, the ionization coefficient α increases with E , and in practice the requirement on Δx introspectively depends on the numerical solution itself, complicating the selection of a spatial step size.

Various resolutions to the spatial stability restriction have been proposed in the literature. Villa et al. [45] showed that the mesh-dependent stability criterion $\alpha\Delta x \leq 1$ can be removed by treating reactions with an upwind method. Marskar [26] softened it by using a Godunov splitting like equation (5), and used a Corner Transport Upwind (CTU) scheme [48] for maintaining an overall larger time step for multi-dimensional simulations. Model corrections to the ionization term have also been considered. Niknezhad et al. [46] change the characteristic length scale of the ionization term by applying a smoothing operator that changes α without net loss in the number of reactions. Marskar [26], Soloviev and Krivtsov [49], Teunissen [50], and Li et al. [51] have considered alterations to α based on electron energy considerations.

For explicit codes, the requirement on Δx can be quite penalizing for the time steps that can be used. On structured Cartesian 3D grids $\Delta x = \Delta y = \Delta z$, a fully explicit discretization using a first order upwind method and centered finite differencing of the diffusion operator yields the time step restriction

$$\Delta t \leq \left(\frac{|v_x| + |v_y| + |v_z|}{\Delta x} + \frac{6D}{\Delta x^2} \right)^{-1}, \quad (7)$$

where we now also include the electron diffusion coefficient D . Fig. 2 shows how this time step varies with Δx for different selections of the field strength E . Marskar [26] found that $E \sim 25$ kV/mm for streamer discharges in air, and with a spatial resolution of $\Delta x \lesssim 2$ μm the time step is approximately 1 ps. A reasonable simulation time for streamer discharges in atmospheric air is around 100 ns, requiring roughly 100 000 time steps, at which point even fluid simulations become numerically expensive.

Implicit methods are subject to the same requirement on Δx as explicit methods, but they remain attractive since they do not impose fundamental limitations on Δt . Unfortunately, 3D simulations often use hundreds of millions of grid points [26–28] and billions of degrees of freedom. Full Newton methods [54] are not very practical at this scale since the full Jacobian must be factored




Fig. 2. Computed time step for 3D fluid models using equation (7) for discharges in atmospheric air. The combined line-mark style indicates the region $\alpha\Delta x \leq 1$ while the marks-only style indicates the region $\alpha\Delta x > 1$. Velocities are given by $v_x = v_y = v_z = \mu E / \sqrt{3}$ where E is the electric field magnitude, and the electron mobility μ and diffusion coefficient D are obtained using BOLSIG+ [52] and the SIGLO database [53].

at every time step. Jacobian-Free Newton-Krylov (JFNK) is a more attractive computational strategy, but it is not clear if JFNK methods remain computationally feasible at this scale, particularly when adaptive mesh refinement (AMR) is required.

Next, consider the evolution of the microscopic version of equation (1):

$$d_t X = v, \quad (8a)$$

$$d_t W = \alpha v W, \quad (8b)$$

where X is a one-dimensional electron position and W is the (average) number of electrons sharing this position. For demonstration purposes, we are using a deterministic reaction rate equation for W . This is less meaningful when dealing with a particle method but we improve on this aspect later in the paper. Consider a single starting electron, $X(0) = 0$, $W(0) = 1$, in which case the solutions to equation (8) using the explicit Euler rule until time $t = k\Delta t$ are

$$X^k = vt, \quad (9)$$

$$W^k = (1 + \alpha v \Delta t)^k. \quad (10)$$

The number of particles per unit length for a grid cell i is then

$$n_i^k = \begin{cases} \frac{1}{\Delta x} (1 + \alpha v \Delta t)^k & \text{if } |vt - x_i| \leq \Delta x/2, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

This is to be contrasted with the exact solution to equation (1) with the initial condition $\tilde{n}(x - vt) = \delta(x - vt)$,

$$n_i(t) = \frac{1}{\Delta x_i} \int_{x_i - \Delta x/2}^{x_i + \Delta x/2} \delta(x - vt) \exp(\alpha x) dx \quad (12)$$

$$= \begin{cases} \frac{1}{\Delta x} \exp(\alpha vt) & |vt - x_i| \leq \Delta x/2, \\ 0 & \text{otherwise.} \end{cases}$$

The two solutions (equations (11) and (12)) differ only due to the way we approach the numerical integration of equation (8). Notably, numerical discretizations that start from equation (8) do not require $\alpha\Delta x \leq 1$, and we can identify why: Equation (3) is numerically diffusive and the electron density only asymptotically tends to zero as $t \rightarrow \infty$, and thus there is always some fraction of n that will react in the grid cell [45]. On the other hand, there is no numerical diffusion involved in equation (11) and the discretization is also stable for any time step, i.e. it does not have a CFL condition. These are the two basic properties that we exploit in the new PIC model.

3. The new model

3.1. Particle transport

Rather than using a macroscopic drift-diffusion model for the electrons, which is subject to fairly strict requirements on Δx and Δt , we consider a *microscopic* model based on Itô diffusion

$$d\mathbf{X}_p = \mathbf{V}_p dt + \sqrt{2D_p} d\mathbf{W}_t^p, \quad (13)$$

where \mathbf{X}_p is the position of a particle p , \mathbf{V}_p is the drift velocity of the particle and $\sqrt{2D_p}$ is the diffusion coefficient of the particle. Here, $d\mathbf{W}_t^p$ is a Wiener process over a time dt . It can be represented as $d\mathbf{W}_t^p = \sqrt{dt} \mathcal{N}_p$ where \mathcal{N}_p is a normal distribution with standard deviation of 0 and variance of 1 in d -dimensional physical space. The noise is uncorrelated in time and space, and independent of noise acting on other particles. The representation of the particle diffusion coefficient as $\sqrt{2D_p}$ is due to a convenient normalization when coarse-graining the model onto a continuum representation where the macroscopic diffusion coefficient D appears instead.

Averaging equation (13) over many identical particles, i.e. $\mathbf{V}_p = \mathbf{v}$, $D_p = D$, yields

$$\langle \mathbf{X}_p(t + \Delta t) - \mathbf{X}_p(t) \rangle = \mathbf{v} \Delta t, \quad (14a)$$

$$\langle [\mathbf{X}_p(t + \Delta t) - \mathbf{X}_p(t) - \mathbf{v} \Delta t]^2 \rangle = 2Dd \Delta t, \quad (14b)$$

where $\langle \dots \rangle$ indicates the expectation value. Thus, by taking \mathbf{V} and D to be the macroscopic electron drift velocity and diffusion coefficients, the Itô model recovers macroscopic drift-diffusion statistics. In this paper we adopt the LFA and take \mathbf{v} and D to be functions of E . The velocity and diffusion coefficients are found by interpolation of the macroscopic quantities \mathbf{v} and D to the particle positions, i.e.

$$\mathbf{V}_p = \mathbf{v}(\mathbf{X}_p), \quad (15)$$

$$D_p = D(\mathbf{X}_p). \quad (16)$$

Extensions to the local mean energy approximation where the coefficients are given as functions of the average electron energy are not examined in this paper.

In a formal derivation Dean [55] showed that the evolution of the global density

$$n(\mathbf{x}, t) = \sum_p \delta[\mathbf{x} - \mathbf{X}_p(t)] \quad (17)$$

yields the advection-diffusion equation of fluctuating hydrodynamics:

$$\frac{\partial n}{\partial t} = \nabla \cdot (-\mathbf{v}n + D\nabla n + \sqrt{2Dn}\mathbf{Z}), \quad (18)$$

where $\mathbf{Z}(\mathbf{x}, t)$ is a Gaussian random field without space-time correlations,

$$\langle \mathbf{Z}(\mathbf{x}, t) \mathbf{Z}(\mathbf{x}', t') \rangle = \delta(\mathbf{x} - \mathbf{x}') \delta(t - t'). \quad (19)$$

The term $\sqrt{2Dn}\mathbf{Z}$ is a stochastic flux that accounts for fluctuations from Brownian motion. This term is usually ignored in studies of non-equilibrium gas discharges. In the macroscopic limit of vanishing fluctuations, equation (18) yields the deterministic advection-diffusion equation which is the conventional starting point for fluid models. The recovery of equation (18) from equation (13) in the macroscopic limit is hardly surprising since equation (13) is a *microscopic* drift-diffusion model.

3.2. Kinetic Monte Carlo

For the plasma chemistry we compute reactions locally within each grid cell, using KMC. Suppose that we are provided with a set of reactions that evolve a system of M different chemical species S_i , $i \in [1, 2, \dots, M]$. The number of particles for each species S_i is given by a state vector

$$\vec{Q}(t) = \begin{pmatrix} Q_1(t) \\ Q_2(t) \\ \vdots \\ Q_M(t) \end{pmatrix}, \quad (20)$$

where $Q_i(t)$ is the number of particles of type S_i in some computational volume ΔV at time t . Reactions are represented stoichiometrically, e.g.



where k is the reaction rate. The set of such reactions is called the *reaction network* \vec{R} . Let \vec{v}_r denote the state change in \vec{Q} caused by a single firing of a reaction of type r , i.e.

$$\vec{Q} \Rightarrow \vec{Q} + \vec{v}_r. \quad (22)$$

For example, if $\vec{Q} = (Q_1, Q_2)^T$ and the reaction network consists of a single reaction $S_1 \rightarrow S_2$ then $\vec{v}_1 = (-1, 1)^T$.

Propensity functions $a_r(\vec{Q}(t), t) dt$ are defined as the probability that exactly one reaction of type r occurs in the infinitesimal interval $[t, t + dt]$. In other words, a_r can loosely be interpreted as the number of reactions of type r per unit time in a computational volume. The rates k that occur in reactions like equation (21) are not equivalent to the conventional reaction rates that are used in the deterministic reaction rate equation (RRE). For a unipolar reaction of the type $S_1 \xrightarrow{k} \emptyset$ the propensity function is $a_r = kQ_1$ and in this case k is numerically equal to the rate that occurs in the RRE (see equation (27)). However, for bipolar reactions of the type $S_1 + S_1 \xrightarrow{k} \emptyset$ the propensity is $k\frac{1}{2}Q_1(Q_1 - 1)$ since there are $\frac{1}{2}Q_1(Q_1 - 1)$ unique pairs of particles of type S_1 .

We use the KMC algorithm as proposed by Cao et al. [56,57]. This algorithm advances \vec{Q} over a time step Δt using a sequence of adaptive smaller steps $\Delta\tau$ where the reaction network is advanced using either the SSA [58], tau-leaping, or a combination of these. This algorithm is discussed further in section 3.2.3, but we first provide some context to the SSA/KMC and tau-leaping algorithms.

3.2.1. Stochastic simulation algorithm (SSA) and tau-leaping

The SSA (or Gillespie algorithm [58]), is a next-reaction model which advances $\vec{Q}(t)$ one reaction at a time. Given a total propensity $A(t) = \sum_r a_r(\vec{Q}(t), t)$, the time until the next reaction is randomly determined from

$$\Delta T_{\text{next}} = \frac{1}{A(t)} \ln\left(\frac{1}{u_1}\right), \quad (23)$$

where $u_1 \in [0, 1]$ is a random number sampled from a uniform distribution. The reaction type is further determined with

$$j = \text{smallest integer satisfying } \sum_{r=0}^{j-1} a_r(\vec{Q}(t), t) > u_2 A(t), \quad (24)$$

where $u_2 \in [0, 1]$ is another uniformly sampled random number. The system is then advanced as

$$\vec{Q}(t + \Delta T_{\text{next}}) = \vec{Q}(t) + \nu_j. \quad (25)$$

The SSA resolves one reaction at a time, and the algorithm becomes increasingly inefficient as the number of reactions per unit time grows. In its isolated form, the algorithm is not very useful for discharge simulations.

The tau-leaping method advances the entire reaction network in a single step over time Δt using Poisson sampling:

$$\vec{Q}(t + \Delta t) = \vec{Q}(t) + \sum_r \vec{\nu}_r \mathcal{P}\left[a_r(\vec{Q}(t), t) \Delta t\right], \quad (26)$$

where $\mathcal{P}(\mu)$ is a random number sampled from a Poisson distribution with mean μ . If the propensity functions a_r do not change significantly on the time interval $[t, t + \Delta t]$ then reaction events are statistically independent, which is the condition for the validity of the tau-leaping scheme. A tau-leaping method has been considered by Luque and Ebert [59] in the context of streamer discharges (although the authors do not use the tau-leaping terminology). Unlike the SSA, equation (26) does not guarantee a physically valid state since Poisson sampling of reactions that consume reactants can yield negative population numbers, and thus needs to be combined with time step selection and rejection sampling [56].

3.2.2. Connection to the reaction rate equation

Tau-leaping is related to the RRE as follows: If a sufficiently large number of reactions occur within Δt , i.e. $a_r(\vec{Q}(t), t) \Delta t \gg 1$, then we can approximate the Poisson process by a Gaussian process. Furthermore, if reactive fluctuations are negligible, i.e. $\sqrt{a_r(\vec{Q}(t), t) \Delta t} \ll a_r(\vec{Q}(t), t) \Delta t$ then we can replace the Gaussian process by its mean value. It can then be shown [60] that equation (26) yields

$$\frac{d\vec{Q}}{dt} = \sum_r \vec{\nu}_r a_r(\vec{Q}(t), t), \quad (27)$$

which we recognize as the deterministic reaction rate equation for the particle density $\vec{n}(t) = \vec{Q}(t)/\Delta V$. Equation (27) now allows us to identify the usual rate constants from the propensities. For example, for the bimolecular reaction $S_i + S_j \xrightarrow{k} \emptyset$ the reaction rate constant is $2k/\Delta V$ for $i = j$ and $k/\Delta V$ for $i \neq j$.

3.2.3. Reaction algorithm outline

Complete details regarding the reactive algorithm that we use are found in [56–58]. We are interested in advancing $\vec{Q}(t)$ from time t to time $t + \Delta t$ for a set of reactions \vec{R} . Letting τ be the simulated time within Δt , this proceeds as follows:

1. Partition \vec{R} into critical and non-critical reaction sets \vec{R}_c and \vec{R}_{nc} . The critical reactions are defined as the set of reactions that are within N_c firings of consuming its reactants. We take $N_c = 5$ in this paper.

2. Compute all propensities, the total propensity A and the critical propensity A_c :

$$A = \sum_{r \in \bar{R}} a_{r,\tau}, \quad (28a)$$

$$A_c = \sum_{r \in \bar{R}_c} a_{r,\tau}, \quad (28b)$$

where $a_{r,\tau} \equiv a_r(\bar{Q}(t + \tau), t + \tau)$.

3. Compute the time $\Delta\tau_c$ until the next critical reaction:

$$\Delta\tau_c = \frac{1}{A_c} \ln\left(\frac{1}{u_1}\right), \quad (29)$$

where $u_1 \in [0, 1]$ is a uniformly distributed random number.

4. Compute a permitted time step $\Delta\tau_{nc}$ such that non-critical reaction propensities do not change by a relative factor greater than ϵ :

$$\Delta\tau_{nc} = \min_{i \in I_{rs}} \left(\frac{\max\left(\frac{\epsilon Q_i}{g_i}, 1\right)}{|\xi_i|}, \frac{\max\left(\frac{\epsilon Q_i}{g_i}, 1\right)^2}{\sigma_i^2} \right), \quad (30)$$

where I_{rs} is the set of reactant species in \bar{R}_{nc} and

$$\xi_i = \sum_{r \in \bar{R}_{nc}} v_{ri} a_{r,\tau}, \quad (31)$$

$$\sigma_i^2 = \sum_{r \in \bar{R}_{nc}} v_{ri}^2 a_{r,\tau}. \quad (32)$$

Here, v_{ri} is the state change of Q_i due to one firing of reaction r . The factor g_i depends on the highest order of reaction where the reactant i appears [56]. In this paper we consider only first order reactions and then $g_i = 1$.

5. To halt integration at time Δt , select a reactive substep $\Delta\tau$ within Δt from

$$\Delta\tau = \min[\Delta t - \tau, \min(\Delta\tau_c, \Delta\tau_{nc})]. \quad (33)$$

6. Resolve reactions as follows:

- (a) If $\Delta\tau_c < \Delta\tau_{nc}$ and $\Delta\tau_c < \Delta t - \tau$: One critical reaction fires. Determine the critical reaction r_c from

$$r_c = \text{smallest integer satisfying } \sum_{r=0}^{r_c-1} a_r > u_2 A_c, \quad (34)$$

where $u_2 \in [0, 1]$ is sampled from a uniform distribution. The sum only runs over the critical reactions. Advance the state \bar{Q} with the results from the SSA and tau-leaping reaction firings:

$$\bar{Q} \rightarrow \bar{Q} + \vec{v}_{r_c} + \sum_{r \in \bar{R}_{nc}} \mathcal{P}(a_{r,\tau} \Delta\tau). \quad (35)$$

- (b) Otherwise: No critical reactions fire. Advance \bar{Q} over $\Delta\tau$ with the non-critical reactions only:

$$\bar{Q} \rightarrow \bar{Q} + \sum_{r \in \bar{R}_{nc}} \mathcal{P}(a_{r,\tau} \Delta\tau). \quad (36)$$

An exception is made if $A\Delta\tau$ is smaller than some factor (we take $A\Delta\tau \leq 1$) since tau-leaping is inefficient in this limit. In this case we switch to SSA stepping using the whole reaction network, taking up to $N_{SSA} = 10$ steps for a total integration time $\Delta\tau'$. Obviously, we restrict this integration to $\Delta\tau' \leq \Delta\tau$.

7. Check if \bar{Q} is a valid state:

- (a) If any particle numbers in \bar{Q} are negative, reject the update. Let $\Delta\tau_{nc} \rightarrow \Delta\tau_{nc}/2$ and return to step 5.

- (b) Otherwise, increment τ by $\Delta\tau$, or by $\Delta\tau'$ if triggering use of the SSA in step 6(b).

8. If $\tau < \Delta t$, return to step 1.

The above algorithm is a well-tested procedure which uses the SSA and tau-leaping algorithms in their respective limits. The factor ϵ determines the maximum permitted relative change in the propensities during one tau-leaping step, and therefore adjusts the accuracy and number of reactive substeps that the algorithm will take.




Fig. 3. Comparison between the KMC algorithm and the RRE. a) Breakdown conditions with $E = 10$ kV/mm and a single starting electron. b) Sub-breakdown conditions with $E = 1.7$ kV/mm and five initial electrons.

3.2.4. Comparing reaction algorithms

To highlight the potential importance of reactive fluctuations, we compare the stochastic reaction algorithm with the RRE for a zero-dimensional test case that illustrates basic electron-neutral interactions in atmospheric pressure air. For simplicity we only consider electron impact ionization and attachment, i.e.



where the ionization and attachment rates are computed using BOLSIG+ (see section 5 for further details). The breakdown field is $E_b \approx 3$ kV/mm, and the exact solution to the RRE (equation (27)) is

$$Q_e(t) = Q_e(0) \exp[(k_\alpha - k_\eta)t], \quad (38)$$

where $Q_e(0)$ is the number of starting electrons.

Fig. 3a) shows the results under breakdown conditions with an electric field $E = 10$ kV/mm and a single starting electron. We have advanced for 50 ps which from the RRE yields 4-5 electrons. This value is compared with the predictions of eight independent runs using the stochastic algorithm. The stochastic algorithm shows considerable variation in the final number of electrons, including one case where the initial electron attached.

Fig. 3b) shows a similar case for sub-breakdown conditions with an electric field $E = 1.7$ kV/mm and five initial electrons. In this case attachment processes dominate the evolution. We find that the hybrid algorithm eventually leads to attachment of all five initial electrons while the RRE yields a solution which only asymptotically decays to zero, i.e. it contains fractional electrons. This latter point is particularly pertinent to positive streamers in highly attaching gases (e.g., sulphur-hexafluoride). In a fluid model the seed electrons ahead of the streamer never completely attach and form negative ions. Rather, the asymptotic decay of the electron density means that a computational fraction of the electron is always available for further seeding the streamer, artificially reducing fluctuations at the streamer tip.

3.3. Model remarks

The Îto-KMC model presented above is a microscopic drift-diffusion model with stochastic chemistry, and in the above we have shown that it recovers the standard drift-diffusion-reaction fluid model in the coarse-grained deterministic limit. In principle, one can think of the Îto-KMC model as a non-kinetic PIC method that samples the macroscopic evolution using computational particles that represent average electrons. Our model rectifies some shortcomings of macroscopic drift-diffusion models, in particular those that pertain to numerical stability and efficiency, but also by maintaining a particle description in regions where the plasma is rarefied.

The Îto-KMC model is qualitatively similar to the model by Luque and Ebert [59], which is essentially a reaction-diffusion master equation (RDME) model supplemented with electron drift. The RDME model evolves the total number of particles in a cell using stochastic sampling of transfer rates. However, the Îto-KMC model has a few important distinctions. Firstly, we expand beyond pure tau-leaping for the plasma chemistry, with the primary benefits being guaranteed non-negativeness and adjustable accuracy. The same algorithm could be used in the Luque and Ebert [59] model. Secondly, the RDME model [59] does not generalize very well to the strong drift regime of streamers where negative transfer probabilities between grid cells can appear [59,61]. The Îto-KMC model resolves these issues, but the cost is the adoption of a microscopic model rather than a mesoscopic one.

We have not included any energy description for the electrons or ions. In fact, it is generally not clear if the model can be extended to include energy transport in such a way that one also recovers the fluid electron energy transport equation when coarse-graining the model. However, an excellent alternative is to combine Îto-KMC with a kinetic electron description. The Îto-KMC method is




Fig. 4. Classic cartoon of a cut-cell AMR grids with a solid boundary (shaded region). The coarsest grid covers a domain of 16×16 cells with two-levels of properly nested refined grids. Bold lines represent grid boundaries.

already a particle model and so the inclusion of kinetic electrons is possible, and is most likely algorithmically simpler than existing hybrid models based on fluid-particle couplings [62,63].

4. Computer implementation

In this section we present our implementation of the new PIC model on cut-cell Cartesian AMR grids. The equations of motion are the Îto-KMC-Poisson system

$$d\mathbf{X}_p = \mathbf{V}_p dt + \sqrt{2D_p} d\mathcal{N}_p, \quad (39a)$$

$$\vec{Q}(t) \xrightarrow{\bar{R}} \vec{Q}(t + dt), \quad (39b)$$

$$\nabla^2 \Phi = -\frac{\rho}{\epsilon_0}, \quad (39c)$$

where equation (39c) is the Poisson equation for the electric field $\mathbf{E} = -\nabla\Phi$ where Φ is the electrostatic potential and ρ is the space charge density. This model has been implemented into the chombo-discharge¹ code, which we have used for streamer simulations in the past [26–28,64–66].

4.1. Spatial discretization

We discretize the equations over a Cartesian grid with patch-based AMR, see Fig. 4. With AMR, the equations of motion are solved over a hierarchy of grid levels $l \in 0, 1, \dots, l_{\max}$. When refining a grid level, the resolution increases by a factor two, i.e. $\Delta x_{l+1} = \Delta x_l/2$. Each grid level consists of a union of properly nested rectangular grid boxes. That is, the valid region of levels $l-1$ and $l+1$ are separated by at least one grid cell on level l , and boxes are disjoint (non-overlapping) on each level. As in previous publications relating to non-equilibrium gas discharges [26–28,64–66] we use the Chombo [67] library for handling the AMR infrastructure. Algorithmic details that are not specific to the Îto-KMC model are found in references [26–28,64,67].

4.2. Charge deposition & interpolation

Deposition of particles and interpolation to the particle positions is done using a cloud-in-cell (CIC) scheme with Cartesian AMR. The mesh densities n are given on cell centers $n(\mathbf{x}_i) = n_i$ where i is a multi-dimensional index and $\mathbf{x}_i = i\Delta\mathbf{x}$. The mesh density is given by

$$n_i = \sum_{p \in R(i)} \left(\frac{w_p}{\Delta V_i} \right) \mathcal{W}_{\text{cic}} \left(\frac{\mathbf{x}_i - \mathbf{X}_p}{\Delta \mathbf{x}} \right), \quad (40)$$

where $p \in R(i)$ indicates particles whose clouds extend into cell i , w_p is the particle weight, ΔV_i is the volume of the grid cell and

$$\mathcal{W}_{\text{cic}}(\mathbf{x}) = \prod_{s=1}^d \mathcal{W}_{\text{cic}}(x_s), \quad (41)$$

¹ <https://github.com/chombo-discharge/chombo-discharge>.




Fig. 5. Refinement boundary deposition procedure at the coarse-fine interface between levels l and $l + 1$. Particles on level l can deposit into level $l + 1$; they retain their original particle shape. Particles on level $l + 1$ can also deposit into level l ; the deposition weight across the refinement boundary is added to the corresponding coarse-side grid cells.

$$W_{\text{cic}}(x) = \begin{cases} 1 - |x|, & |x| < 1, \\ 0, & \text{otherwise.} \end{cases} \quad (42)$$

Deposition of particles near refinement boundaries are sketched in Fig. 5 and handled as follows: If the particle cloud for a particle at level $l + 1$ hangs over the refinement boundary into the coarse level l , the deposition weight is added to the corresponding level l cells and then normalized by the appropriate volume fraction. Particles that live on a coarse grid level l but whose clouds hang into level $l + 1$ have particle widths Δx_l on level l , and for factor two refinements they may extend into the first strip of cells on level $l + 1$ (see Fig. 5). In order to ensure that this weight ends up in the correct cells on level $l + 1$, these particles are also deposited on level $l + 1$, but using the original particle width Δx_l . These particles thus have a width of $2\Delta x_{l+1}$ and in 3D they can deposit into at most 9 grid cells on level $l + 1$. Their deposition function on level $l + 1$ is

$$W'_{\text{cic}}(x) = \begin{cases} \frac{1}{2}, & |x| \leq \frac{1}{2} \\ \frac{1}{2} \left(\frac{3}{2} - |x| \right), & \frac{1}{2} < |x| \leq \frac{3}{2} \\ 0, & \text{otherwise.} \end{cases} \quad (43)$$

4.3. Semi-implicit Euler-Maruyama method

The standard Euler-Maruyama method for equation (13) is

$$\mathbf{X}_p^{k+1} = \mathbf{X}_p^k + \Delta t \mathbf{V}_p^k + \sqrt{2D_p^k \Delta t} \mathcal{N}_p. \quad (44)$$

Coupled to equation (39c), the discretization must be restricted by the dielectric relaxation time in order to be stable, i.e. $\Delta t \leq \sigma/\epsilon_0$. We remove this limitation by using a semi-implicit formulation as follows:

$$\mathbf{X}_p^{k+1} = \mathbf{X}_p^k + \Delta t \mathbf{V}_p^{k+1} + \sqrt{2D_p^k \Delta t} \mathcal{N}_p, \quad (45)$$

where

$$\mathbf{V}_p^{k+1} = \text{sgn}(Z_s) \mu_p^k \mathbf{E}^{k+1}(\mathbf{X}_p^k) \quad (46)$$

and Z_s is the charge number for species s , and sgn is the sign operator. I.e., we have $\text{sgn}(Z_s) = -1$ for electrons and $\text{sgn}(Z_s) = 1$ for positive ions. We achieve this coupling by first solving the Poisson equation

$$\nabla \cdot \mathbf{E}^{k+1} = \frac{\rho^{k+1}}{\epsilon_0}, \quad (47)$$

which to first order in Δt can be written

$$\nabla \cdot \mathbf{E}^{k+1} = \frac{1}{\epsilon_0} \rho^\dagger - \frac{\Delta t}{\epsilon_0} \nabla \cdot \mathbf{J}_{\text{adv}}^k, \quad (48)$$

where $\mathbf{J}_{\text{adv}}^k$ is the advective current density and ρ^\dagger is the space charge density computed from the update

$$\mathbf{X}_p^\dagger = \mathbf{X}_p^k + \sqrt{2D_p^k \Delta t} \mathcal{N}_p. \quad (49)$$

For a species s the advective current density at a grid point \mathbf{x}_i is

$$\mathbf{J}_{i,\text{adv.}}^{s,k} = q_e |Z_s| \sum_p \mu_s^k \mathbf{E}^{k+1}(\mathbf{X}_p^k) \mathcal{W}(\mathbf{x}_i - \mathbf{X}_p^k). \quad (50)$$

Expanding $\mathbf{E}^{k+1}(\mathbf{X}_p^k)$ as a polynomial around the grid point \mathbf{x}_i yields

$$\mathbf{E}^{k+1}(\mathbf{X}_p^k) \approx \mathbf{E}_i^{k+1} - (\mathbf{x}_i - \mathbf{X}_p^k) \cdot \nabla \mathbf{E}_i^{k+1} + \mathcal{O}\left((\mathbf{x}_i - \mathbf{X}_p^k)^2\right), \quad (51)$$

where $\mathbf{E}_i = \mathbf{E}(\mathbf{x}_i)$. Equation (50) yields

$$\mathbf{J}_{i,\text{adv.}}^{s,k} \approx \left[q_e |Z_s| \sum_p \mu_p^k \mathcal{W}(\mathbf{x}_i - \mathbf{X}_p^k) \right] \mathbf{E}_i^{k+1} - \left[q_e |Z_s| \sum_p \mu_p^k \mathcal{W}(\mathbf{x}_i - \mathbf{X}_p^k) (\mathbf{x}_i - \mathbf{X}_p^k) \right] \cdot \nabla \mathbf{E}_i^{k+1}, \quad (52)$$

where the first term is the conventional Ohmic contribution that we recognize from semi-implicit formulations for fluid models [44]. The support of \mathcal{W} is Δx , and so the second term scales as Δx as well. When used together with equation (48) this term scales as $\mathcal{O}(\Delta x \Delta t)$ in the semi-implicit Poisson equation, and it also has a small error constant: The moments $\mathcal{W}(\mathbf{x}_i - \mathbf{X}_p^k) (\mathbf{x}_i - \mathbf{X}_p^k)$ are anti-symmetric in \mathbf{X}_p^k with respect to the grid cell center \mathbf{x}_i , so when particles distribute uniformly over a grid cell the summation yields the zero vector. This is, for example, the case in the discharge channels where there are many electrons per grid cell, whereas outside of the channel the current is negligibly small. Summing over all species s to leading order yields

$$\mathbf{J}_{\text{adv.}}^k = \left(q_e \sum_s |Z_s| \mu_s^k n_s^k \right) \mathbf{E}^{k+1}, \quad (53)$$

where μ_s^k and n_s^k are mesh variables for species s (we have suppressed the index i). Thus, ignoring higher-order moments, equation (48) can be written in the familiar form [44]

$$\nabla \cdot \left[\left(1 + \frac{\sigma^k \Delta t}{\epsilon_0} \right) \mathbf{E}^{k+1} \right] = \frac{1}{\epsilon_0} \rho^\dagger, \quad (54)$$

where

$$\sigma^k = q_e \sum_s |Z_s| \mu_s^k n_s^k \quad (55)$$

is the conductivity of the plasma.

Equation (54) is discretized with finite volumes, using a standard 9-point stencil in the interior points (and flux matching at the coarse-fine interface). The embedded boundary fluxes are also constructed to second order, using additional interior points when evaluating the normal derivative on the cut-cell boundary centroids. The corresponding linear system is solved using geometric multigrid with V-cycling, using red-black Gauss-Seidel relaxation as the smoother on each grid level and a biconjugate gradient stabilized method (BiCGSTAB) as a bottom solver. A relative exit tolerance of 10^{-10} is used as a convergence criterion for multigrid. Further details regarding the finite volume discretization of the variable-coefficient Poisson equation and its embedding into geometric multigrid in the presence of embedded boundaries and Cartesian AMR are given in e.g. [67–70]. After obtaining the electric field \mathbf{E}^{k+1} we compute \mathbf{V}_p^{k+1} from equation (46) and complete the particle update (equation (45)).

4.4. KMC-particle coupling

The reaction algorithm solves for the total number of particles in a grid cell and leaves substantial freedom in how one assigns the chemistry products into new computational particles. Since we are concerned with methods that potentially use very large time steps, the creation of computational particles with physical weights $w = 1$ is not possible due to the large number of physical particles generated in a time step. In this paper, if the reaction step led to net creation of particles we instead create at most $N_{\text{ppc}}^{\text{new}} = 64$ new computational particles in the cell. If the KMC solver gave $N_{\text{phys}} > N_{\text{ppc}}^{\text{new}}$ physical particles in the cell, we construct $N_{\text{ppc}}^{\text{new}}$ computational particles with weights

$$w = N_{\text{phys}} \div N_{\text{ppc}}^{\text{new}}, \quad (56)$$

where \div denotes integer division. The remainder $\text{mod}(N_{\text{phys}}, N_{\text{ppc}}^{\text{new}})$ is assigned to one of the new particles. These particles are later merged with the computational particles that already exist in the cell. If the KMC algorithm led to net loss of particles, we remove the weight directly from the existing computational particles.

Production of particles in cut-cells only takes place in the valid region of the cell. Particle positions \mathbf{X}_p are drawn from a uniform distribution in each coordinate direction with the requirement

$$(\mathbf{X}_p - \mathbf{x}_c) \cdot \hat{\mathbf{n}}_c \geq 0, \quad (57)$$

where \mathbf{x}_c is the cut-cell boundary centroid and $\hat{\mathbf{n}}_c$ is the cut-cell boundary normal. Since cut-cell volume fractions can be arbitrarily small, we optimize this step by only drawing the position \mathbf{X}_p inside the minimum bounding box that encloses the valid region of the cut-cell.




Fig. 6. Concept sketch of Cartesian 2D bounding volume hierarchy generation. Solid lines inside the cell indicate splitting lines (planes in 3D). From left to right: Original particles, two bounding volumes, four bounding volumes, and 16 bounding volumes.

4.5. Photon generation and transport

Photons are also treated with a particle method, and for simplicity we consider instantaneous transport where we don't have to track the photons in time. As with the particles, the KMC algorithm provides the number of physical photons that is generated in the reaction step, which we limit to $N_{\text{ppc}}^{\text{new}} = 64$ computational photons. Photon weights and emission positions are assigned in the same way as we do for the particles.

Photon absorption positions are determined individually for each (super-)photon. For example, assume that a photon has some frequency f and is emitted from an initial position \mathbf{Y}_f^0 . This photon is absorbed at position

$$\mathbf{Y}_f = \mathbf{Y}_f^0 + r_f \hat{c}, \quad (58)$$

where r_f is a random number drawn from an exponential distribution with parameter $\kappa(f)$, and \hat{c} is a uniformly distributed random point on the unit sphere. Here, $\kappa(f)$ is the mean absorption coefficient in the gas for a photon with frequency f , i.e. $1/\kappa(f)$ is the mean absorption length. Since spectral absorption lines also have a spectral width, the mean absorption coefficient $\kappa(f)$ is frequency dependent. When we sample the photon generation, we begin by sampling the spectral line by stochastically determining f according to some distribution. We then use a known expression for $\kappa(f)$ for stochastically determining r_f for each photon. Depending on the photoionization model that is used, one can sample multiple spectral lines [34] in combination [71] or individually [26]. Photoemission is disregarded in this paper, so if a photon trajectory intersects an internal boundary (e.g. an electrode) or a domain boundary, it is removed from the simulation.

4.6. Superparticle management

In order to maintain a manageable number of particles, only computational particles that represent many physical particles are tracked. Our particle merging and splitting strategy uses a bounding volume hierarchy with k - d trees for locating spatial clusters of particles, and particles are merged/split within each cluster. We use a standard tree structure which uses top-down construction, i.e. it is hierarchically built from the root node and downwards. However, the algorithm that we use for splitting a leaf node is new and it is therefore discussed in detail.

Initially, a leaf node L contains a list \vec{P}_L of M particles that are each identified by a tuple $\langle \mathbf{X}_p, w_p \rangle$, where w_p is the particle weight. Then, \vec{P}_L is sorted based on one of the axis coordinates and split into two bounding volumes such that the total weight of the two halves differ by at most one physical particle. This process is shown in Fig. 6 and proceeds as follows:

1. Pick a splitting direction. We choose the coordinate direction where the minimum bounding box enclosing the particles has the largest extent.
2. Sort the particle list \vec{P}_L from smallest to largest coordinate in splitting direction.
3. Locate the median particle with index p' in the list, where $p' > 1$ is the smallest index satisfying

$$\sum_{p=1}^{p'-1} w_p + w_{p'} > \sum_{p=p'+1}^M w_p. \quad (59)$$

The index p' indicates the position of the particle on the splitting plane, i.e. all particles $p < p'$ are found on the left hand side of the splitting plane and all particles $p > p'$ are found on the right-hand side.

4. Transfer particles $p \in [1, p' - 1]$ to a new list \vec{P}_l in the left leaf node, and particles $p \in [p' + 1, M]$ to another list \vec{P}_r in the right leaf node.
5. Assign the median particle p' :
 - (a) If the median particle is a physical particle the particle is assigned to whichever child list (\vec{P}_l or \vec{P}_r) has the lowest total weight.




Fig. 7. Particle merging example. a) Initial particles. b) Merged particles with $N_{\text{ppc}} = 64$. The labels on the colorbar indicate particle weights w_p . (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

- (b) If the median particle is a superparticle, i.e. $w_{p'} \geq 2$, it is split into two new particles with the same position $\mathbf{X}_{p'}$ but with new weights. Due to the median selection in equation (59), these weights can be constructed such that the weight of \vec{P}_l and \vec{P}_r differ by at most one physical particle.

Because we merge particles by groups rather than in pairs [72], the process above only proceeds until we have N_{ppc} leaves in the tree. Choosing the final number of computational particles to be a factor of two gives a balanced tree where all the leaves exist on the same tree level, and in this case the number of physical particles between any two arbitrary leaves in the tree differs by at most one. At the end of the tree-building algorithm each leaf node represents a bounding volume with a list \vec{P} of computational particles. The particles in this list become a new superparticle with weight and position

$$w = \sum_{p \in \vec{P}} w_p \quad (60a)$$

$$\mathbf{X} = \frac{1}{w} \sum_{p \in \vec{P}} w_p \mathbf{X}_p. \quad (60b)$$

Particle merging is done on a cell-by-cell basis in order to prevent creation of particles that lie inside the embedded boundary. The algorithm also handles splitting of superparticles. If a cell contains a single particle with a large weight then step (v) in the above algorithm ensures that this particle is hierarchically split until we have created N_{ppc} new particles.

Fig. 7 shows an example of merging 1000 initial particles whose positions are uniformly distributed within a cut-cell and whose weights are uniformly distributed on the interval $w_p \in [1, 100]$. These particles are merged into $N_{\text{ppc}} = 64$ particles, and as expected the final particle weights differ by at most one.

4.7. Final algorithm

The final algorithm for integration over a time step Δt is as follows:

1. Compute the conductivity σ^k .
2. Perform the diffusive advance:

$$\mathbf{X}_p^\dagger = \mathbf{X}_p^k + \sqrt{2D_p^k \Delta t} \mathcal{N}_p.$$

3. Compute the space charge density $\rho^\dagger = \rho(\mathbf{X}_p^\dagger)$ and solve for \mathbf{E}^{k+1} using equation (54).
4. Interpolate particle velocities $\mathbf{V}_p^{k+1} = \mathbf{v}[\mathbf{E}^{k+1}(\mathbf{X}_p^k)]$.
5. Advect particles $\mathbf{X}_p^{k+1} = \mathbf{X}_p^\dagger + \mathbf{V}_p^{k+1} \Delta t$.
6. Move photons $\mathbf{Y} = \mathbf{Y}_0 + r_\kappa \hat{\mathbf{c}}$ using equation (58).
7. Advance the reaction network over Δt , see section 3.2 and section 4.4.
8. Manage superparticles, section 4.6.

Conceptually, the above algorithm uses a Godunov splitting between particle transport (steps 1 through 5) and plasma chemistry (step 7). The particle transport step is a first-order accurate semi-implicit discretization, and the plasma chemistry is solved with a stochastic reaction algorithm with adjustable accuracy through the factor ϵ (section 3.2). Setting $\epsilon = \infty$ will accept any tau-leaping




Fig. 8. Example of dual mesh load balancing for different kernels. Each colored square represents a grid patch (of e.g. 16^3 cells). The MPI rank ownership is indicated by numbers inscribed in each square. a) Example MPI rank assignment and load distribution for kernels whose load scale with the number of grid points. b) Example MPI rank assignment and load distribution for kernels that scale with the number of particles.

step, i.e. the chemistry is resolved with a time step Δt . On the other hand, setting $0 < \epsilon \ll 1$ yields a highly accurate chemistry algorithm which may potentially take many substeps within Δt . But if high order chemistry is used together with a large splitting step Δt , the overall stability of the algorithm can deteriorate. The reason for this is that when transport and field updates are performed rarely but the chemistry integration is highly accurate, the number of free electrons that are generated in a grid cell is overestimated. This issue is not unique for Îto-KMC but also occurs for deterministic fluid models when using operator splitting methods with large splitting steps. Since we use fixed time steps in this paper, we therefore resolve the transport and chemistry with the same time step, i.e. we use $\epsilon = \infty$, and rely on the SSA steps primarily to avoid negative particle numbers. In the future, we will be extending our methodology to dynamic time stepping (either CFL or physics based), at which point we will be able to leverage the adjustable accuracy features in the KMC integrator.

4.8. Parallelization

Our computer implementation is parallelized with flat MPI, using the natural domain decomposition offered by the AMR grids where each MPI rank solves for a subset of the grids on each level. The simulations are performed away from the strong scaling limit, which left room for load balancing our application.

The field and particle updates have different computational metrics. A reasonable proxy for the computational load of the discretized Poisson equation is the number of grid cells in a grid patch, while for the particles the load is better estimated by the number of particles that are assigned to the patch. We have load balanced our simulations with dual grids. In this approach we use two sets of AMR grids where the grid levels consist of the same grid patches, but where the assignment of grid subsets among the MPI ranks differ, see Fig. 8. One AMR grid set is load balanced with the grid patch volume as a proxy for the computational load, and is used for grid kernels that scale with the number of grid points, e.g. the discretized Poisson equation or advancing the reaction network. On the other grid, we advance kernels that scale with the number of grid particles, i.e. transport kernels, mesh deposition and interpolation, and superparticle handling. The dual grid approach adds some computational complexity, but these drawbacks were offset by reductions in simulation times which were up to 40%.

5. Numerical tests

5.1. Simulation conditions

We consider a 10 cm^3 computational domain with a vertical needle-plane gap. A cross section of the computational domain and the boundary conditions is shown schematically in Fig. 9. A 5 cm long cylindrical electrode with a spherical cap at the end sticks out of the live electrode plane, and the opposite plane is grounded. The electrode diameter is 1 mm and the vertical distance between the live electrode and the ground plane is 5 cm. Homogeneous Neumann boundary conditions are used for the Poisson equation on the side faces, and all simulations start from a step voltage of 20 kV. The peak initial electric field magnitude is roughly 11 kV/mm on the anode tip. All simulations use a coarsest AMR level of 128^3 cells, but use up to another eight levels of refinement, i.e. up to effective domains of 32768^3 cells.

We use a three-species model for discharges in air, consisting of electrons, positive ions, and negative ions. The plasma kinetics that we use is summarized in Table 1. We focus on using a simple and well-known reaction set for our example simulations. Using more elaborate plasma chemistry is possible, but not required for our simulation examples. The electron diffusion coefficient D_e , mobility μ_e , temperature T_e , ionization frequency k_a , and attachment frequency k_η are field-dependent and are computed by using BOLSIG+ [52] and the SIGLO database [53]. The electron-ion and ion-ion recombination rates are

$$k_{ep} = \beta_{ep}/\Delta V, \quad (61)$$

$$k_{np} = \beta_{np}/\Delta V, \quad (62)$$

where ΔV is the grid cell volume and




Fig. 9. Cross-section of the simulation domain, also showing the electrostatic boundary conditions. The dimensions are not to scale.

Table 1

Simplified air plasma chemistry used for the example simulations. The notation \emptyset indicates an untracked species (e.g., N_2 or O_2) incorporated directly into the rate constant for the reaction.

Reaction	Rate	Propensity	Ref.
$e + \emptyset \rightarrow e + e + M^+$	$k_\alpha(E)$	$k_\alpha Q_e$	[52]
$e + \emptyset \rightarrow M^-$	$k_\eta(E)$	$k_\eta Q_e(E)$	[52]
$e + M^+ \rightarrow \emptyset$	$k_{ep}(E)$	$k_{ep} Q_e Q_{M^+}$	[73]
$M^- + M^+ \rightarrow \emptyset$	k_{np}	$k_{np} Q_{M^+} Q_{M^-}$	[74]
$e + \emptyset \rightarrow e + \gamma + \emptyset$	$k_\gamma(E)$	$k_\gamma Q_e$	[75,76]

$$\beta_{ep} = 1.138 \times 10^{-11} T_e^{-0.7} \text{ m}^3/\text{s}, \quad (63)$$

$$\beta_{np} = 2 \times 10^{-13} (300/T)^{0.5} \text{ m}^3/\text{s}, \quad (64)$$

where $T = 300$ K is the gas temperature and $T_e = T_e(E)$ is the electron temperature.

We use the Zheleznyak photoionization model [75] including the corrections by Pancheshnyi [76] for modeling photon transport for the reaction $e + \emptyset \xrightarrow{k_\gamma} e + \gamma + \emptyset$. The rate constant is

$$k_\gamma = \frac{p_q}{p + p_q} v_Z(E) k_\alpha, \quad (65)$$

where $v_Z(E)$ is a lumped function that accounts for excitation efficiencies and photoionization probabilities [76]. The quenching pressure is $p_q = 40$ mbar and the gas pressure is $p = 1$ bar. When a photon is generated within the reaction step we draw a random absorption coefficient as

$$\kappa_f = K_1 \left(\frac{K_2}{K_1} \right)^{\frac{f-f_1}{f_2-f_1}}, \quad (66)$$

where $K_1 = 530 \text{ m}^{-1}$, $K_2 = 3 \times 10^4 \text{ m}^{-1}$, $f_1 = 2.925 \text{ PHz}$, $f_2 = 3.059 \text{ PHz}$, and f is a random number sampled from a uniform distribution on the interval $[f_1, f_2]$. The propagation distance of each photon is then determined by drawing a random number from an exponential distribution with parameter κ_f .

All simulations start by drawing 10^4 initial electron-ion pairs uniformly distributed in a sphere with a $500 \mu\text{m}$ radius centered at the needle tip. Electron-ion pairs whose positions end up inside the electrode are removed before the simulation starts.

In the computer simulations we refine cells on level l if

$$\alpha \Delta x_l \geq 1, \quad (67)$$




Fig. 10. Comparison of the field distribution and electron density in the hydrodynamic and Îto-KMC descriptions. Left: Field magnitude. Right: Electron density.

and coarsen if

$$\alpha \Delta x_l \leq 0.2, \quad (68)$$

where Δx_l is the grid spacing on level l and α is the Townsend ionization coefficient.

5.2. Comparison with hydrodynamics

In this section we present a comparison between the Îto-KMC model and an equivalent drift-diffusion-reaction model based on deterministic hydrodynamics (equation (18) without the stochastic term), using the same photoionization model and transport data. Since we use cut-cell Cartesian AMR grids, the discretization of the fluid model is bit involved, and is therefore not discussed in detail here. For the fluid transport equation we follow the discretization that we used in [26]. There, we used a Godunov splitting between plasma transport and reactions, and employed a CTU scheme [48] that also include transverse slopes in the advective term, permitting a softer CFL constraint as discussed in Sec. 2. Here, the only major difference between that discretization and the current one is that we here use explicit diffusion and a semi-implicit coupling to the electric field.

Because both the Îto-KMC model and the fluid model are both stochastic, the latter due to discrete photoionization, the two models are compared in terms of ensemble averages where we quantitatively compare the average of ten computer simulations for each model. To initialize the fluid model we include draw particles from same particle distribution as for the Îto-KMC model and deposit them as a density using a nearest-grid-point scheme when the simulation begins. In the Cartesian 2D comparison we have also raised the potential on the electrode to 80 kV in order to facilitate propagation of the streamer. The Cartesian 2D version is included because fluctuations are reduced which leads to a single streamer, and the models can then be both qualitatively and quantitatively compared.

We compare the two models in 2D planar coordinates, using fixed time steps of $\Delta t = 2.5$ ps for a total integration time of $t = 10$ ns. The field magnitude and electron density after 10 ns are shown in Fig. 10. We find that the two solutions are qualitatively very similar. Figs. 11a) and b) show the temporal evolution of the maximal electric field and the streamer propagation distance for the simulations. The streamer head position is here defined as the position where the electric field is at its maximum. Note that the data shows the average of many computer simulations, with error bars indicating standard deviations. Figs. 11c) and d) show the electron density and electric field measured along the streamer axis after $t = 10$ ns.

Quantitatively, only minor differences are found between the two models: For example, the largest difference between the maximum electric field in the two models is about 4%, while the average streamer velocities agree to within 0.01 mm/ns.

5.3. Discrete particle noise

To determine if particle noise impacts the simulations, we consider three-dimensional numerical solutions obtained using a varying number of N_{ppc} but fixed Δx and Δt . Note that if there are N_{phys} physical particles in a cell that are represented by N_{ppc} computational particles, then the computational particles have on average weights $w = N_{\text{phys}}/N_{\text{ppc}}$. Particle fluctuations (i.e.,




Fig. 11. Quantitative comparison between the fluid and Îto-KMC model for ensembles of 10 simulations (per model). Lines and points indicate average values and error bars indicate standard deviations. a) Maximum field in the streamer head. b) Streamer head position (distance from electrode). c) Ensemble-averaged electron density in the streamer channel. d) Ensemble-averaged electric field in the streamer channel.

discrete particle noise) therefore increase by a factor of \sqrt{w} compared to their physical level. We select $\Delta x \approx 12 \mu\text{m}$ and $\Delta t = 20 \text{ ps}$ and integrate for 5 ns. The same initial particles are used in these tests. Fig. 12 shows the resulting electron density in the neighborhood of the electrode using between 8 and 256 computational particles per cell (per species). There is no apparent disagreement or numerical artifacts for these simulations, which suggests that even $N_{\text{ppc}} = 16$ is sufficiently accurate for these particular simulations. Similar results were obtained by Teunissen and Ebert [31]. This finding can not be extrapolated to coarser grids because if we keep N_{ppc} fixed and reduce Δx by a factor of two, the average particle weights increase by a factor of 8. It also warrants mention that there is no particle noise in the KMC algorithm itself, as it operates with the number of physical particles. However, elevated particle noise still arises due to transport and splitting/merging of superparticles.

5.4. Grid sensitivity

In this section we perform a grid sensitivity study by varying the spatial and temporal resolutions. We set $N_{\text{ppc}} = 32$ and consider temporal resolutions ranging from 5 ps to 80 ps, and spatial resolutions ranging from 3 μm to 195 μm , and integrate for 50 ns. Fig. 13 shows the final state for the 35 different simulations, and we make several observations:

1. Simulations are stable for all time steps, also when the time step is orders of magnitude larger than the dielectric relaxation time. For the simulations with $\Delta t = 80 \text{ ps}$ and $\Delta x \approx 3 \mu\text{m}$ the time step is equivalent to using an advective CFL number of > 30 .
2. Bounded solutions are obtained for all spatial resolutions, i.e. we do not have $n_e \rightarrow \infty$ or $E \rightarrow \infty$ for any Δx we investigate. Here, the spatial resolutions range from moderately fine to extremely coarse. From experience with fluid simulations [26–28] we have found that instabilities can occur for branching streamers, particularly if some of the branches become stagnant.
3. Numerical branching occurs on too coarse grids, which can be seen on the column $\Delta t = 80 \text{ ps}$ where the streamer initially splits into four branches, but these four initial branches disappear on finer grids. Similar phenomena are seen on the row $\Delta x \approx 24 \mu\text{m}$ where we find small protrusion needles for $\Delta t \leq 20 \text{ ps}$, and similarly for $\Delta x \approx 12 \mu\text{m}$ for $\Delta t \leq 10 \text{ ps}$. We believe that these branches appear due to spatial underresolution since for $\Delta x \approx 12 \mu\text{m}$ we only have about 5-10 grid cells for resolving the cross section of the streamers with the smallest radii.
4. The degree of branching decreases when larger time steps are used, which is particularly evident for the row $\Delta x \approx 3 \mu\text{m}$. This can be understood in terms of the photoionization-induced noise ahead of the streamer. Bagheri and Teunissen [25] and Marskar [26] have shown that the branching behavior depends on the amount of photoionization ahead of the streamer. Increasing the amount of photoionization reduces noise in the plasma density ahead of the streamer, which also reduces the amount of




Fig. 12. Electron density in the neighborhood of the rod electrode computed using various thresholds for the maximum number of computational particles, indicated in each frame by N_{ppc} .

branching. As larger time steps are used there are more photoelectrons generated during time steps, which artificially suppresses fine-grained temporal variations in the plasma density ahead of the streamer.

5. The velocity of the streamers increase with increasing temporal resolution. There are at least two reasons for this:
 - (a) The number of electrons in the ionization zone in the streamer grows exponentially, and larger time steps lead to numerical underestimation of the electron impact ionization in the streamer tip.
 - (b) When large time steps are used the electrons in the reaction zones at the tip of the streamers can be moved completely out of it. For example, the simulation with $\Delta x \approx 3 \mu\text{m}$ and $\Delta t = 80 \text{ ps}$ used an effective CFL number of > 30 . The reaction zone is just a few grid cells thick, and moving the electrons too far out of the reaction zone reduces the amount of ionization in it, and thus also velocity of the streamer.
6. The velocities of the streamers increase slightly when the resolution increases, which we can see on the column $\Delta t = 10 \text{ ps}$. We believe this occurs because coarse grids lead to under-resolution of the electric field ahead of the streamer. For coarser Δx the electric field on the tips is therefore lower, and this reduces the amount of ionization.
7. Streamer radii agree with experimental observations only on the fine grids ($\Delta x \lesssim 6 \mu\text{m}$). Briels et al. [77] have measured streamer diameters in atmospheric air and found that diameters range from $100 \mu\text{m}$ to 3 mm , depending on experimental conditions like applied voltage, gap inhomogeneity, and various other factors. However, for $\Delta x = 97 \mu\text{m}$ to $195 \mu\text{m}$ we only find streamers with radii $R \sim 1 \text{ mm}$. On finer grids streamers with smaller radii also emerge. E.g. for $\Delta x \approx 3 \mu\text{m}$ we find streamer radii ranging from $100 \mu\text{m}$ to 1 mm , which agree with experimental observations.
8. The electric field at the streamer tips vary by streamer radius. On the finest grid $\Delta x \sim 3 \mu\text{m}$ we find $E \sim 25 \text{ kV/mm}$ to 30 kV/mm for filaments with radii on the order of $R \sim 50 \mu\text{m}$ to $150 \mu\text{m}$. This is in agreement with fluid simulations of positive streamers [26] as well as analytical estimates [78]. On grids $\Delta x \geq 24 \mu\text{m}$ we find that the electric field strength at the streamer tips is $E \lesssim 10 \text{ kV/mm}$, but these solutions are quite clearly underresolved and thus have no practical relevance.

In summary, we find that the grid resolution should be $\Delta x \leq 6 \mu\text{m}$, which is about the same requirement as in fluid models [26]. The time step should be $\Delta t \leq 10 \text{ ps}$, which at $\Delta x \approx 6 \mu\text{m}$ is about a factor of 5 larger than that permitted through a conventional CFL condition like equation (7), as shown in Fig. 2. We also point out that the $\hat{\text{I}}\text{to-KMC}$ method can maintain this time step even for finer grids, which is not possible for explicit fluid codes. Although $\hat{\text{I}}\text{to-KMC}$ is quite forgiving for larger time steps, it is clear from Fig. 13 that lack of temporal resolution leads to suppression of several morphological features in the discharge, while underresolved grids lead to numerical branching. Also note that the fastest time scales in the $\hat{\text{I}}\text{to-KMC}$ method are the same as in fluid methods, which for the reaction set in Table 1 is the electron ionization impact frequency $\alpha\mu_e E$. We have not been able to run numerical convergence tests due to the stochastic component that is involved, which would require ensemble studies at high spatial and temporal resolutions. We nonetheless observe that the solutions convergence to physically meaningful solutions with streamer diameters and velocities that quantitatively agree with experimental observations [77]. Section 5.3 showed that discrete particle noise had a comparatively low qualitative impact on the simulations, so the artificial branching for $\Delta x \geq 12 \mu\text{m}$ is probably mesh-based. This finding can not be automatically extrapolated to different pressure due to very different plasma densities. For example, in our simulations with




Fig. 13. Simulation outputs after integrating for 50 ns. The rod diameter is 1 mm and the spatial scale is otherwise indicate in the top left and bottom right frames. Along the first column we include the effective number of grid cells along the indicated spatial direction.

$n_e \approx 10^{18}/\text{m}^3$, $\Delta x \approx 10 \mu\text{m}$ and $N_{\text{ppc}} = 64$, particles have an average weight $w \approx 15.6$ in the streamer head. But in sprite discharges one may have $n_e \approx 10^{10}/\text{m}^3$ and $\Delta x \approx 1 \text{ m}$, so $N_{\text{ppc}} = 64$ yields particle weights $w \approx 1.56 \times 10^8$. Discrete particle noise is therefore much higher for sprites than for atmospheric pressure streamers.

5.5. Positive streamer evolution

We now run the simulation with $\Delta x \approx 6 \mu\text{m}$ and $\Delta t = 10 \text{ ps}$, using $N_{\text{ppc}} = 64$ particles per cell until $t = 150 \text{ ns}$. This simulation case is probably not completely grid converged, which we can see from the fact that the discharge tree in Fig. 13 for $\Delta x = 6 \mu\text{m}$ and




Fig. 14. Snapshots of the electron density at various time instants. The circle shown in panels with $t \geq 90$ ns indicates one of the stagnant streamer branches.

$\Delta t = 5$ ps is approximately 25% faster. Part of this difference can also be due to a natural variation in the front velocity, depending on how the discharge tree develops. Since we only ran a single simulation per spatial and temporal resolution, we do not know the size of these deviations. Comparing with the panel $\Delta x = 3 \mu\text{m}$ and $\Delta t = 5$ ps in the same figure, we find that we are probably also missing some fine-scale features in the discharge tree as well. Ideally, this is the simulation that we would have run further, but the case $\Delta x = 6 \mu\text{m}$, $\Delta t = 10$ ps was the largest case we could fit in our current compute quota.

The simulation case is nonetheless quite challenging. Positive streamers in air have smaller radii and branch more frequently than negative ones, and the branching behavior is also voltage-dependent as positive streamers in air branch more frequently at lower background fields. In our case the average background electric field measured along the symmetry axis is just 0.4 kV/mm, which is below the so-called positive streamer stability field of ~ 0.5 kV/mm. Under these conditions, experiments show that repeated branching leads to development of a discharge tree consisting of multiple small-diameter positive streamers [77]. Streamers may also stagnate, which has been identified as a challenge for fluid models [26,46,47,51].

Fig. 14 shows temporal snapshots of the positive streamer evolution every ten nanoseconds until $t = 150$ ns. We have not (yet) been able to skeletonize the discharge structure for quantitative analysis, but from the figure(s) we extract the following information:

- The streamer radii vary between ~ 1 mm and $\sim 100 \mu\text{m}$. The thicker streamers appear closest to the anode, and as they propagate they branch into thinner filaments.
- In addition to propagating towards to the ground plane, electrostatic repulsion between the filaments yields numerous sideways branches, which determines the radius of the tree.




Fig. 15. Different views of the positive streamer after $t = 150$ ns. a) Side view. b) Looking into the streamer from below.

- The front velocity is on average 0.2 mm/ns, and is defined by the velocity of the front streamers.
- The total length of the discharge tree at $t = 150$ ns is approximately 3 cm and its radius is approximately 1.5 cm.
- Many streamer filaments stop propagating after some distance, but do not lead to unbounded growth in the plasma density. We have indicated one of these branches in Fig. 14, but many more can be identified.

Fig. 15 shows the discharge at $t = 150$ ns from additional perspectives. This data corresponds to the bottom-right panel in Fig. 14. As an amendment to Fig. 14 and Fig. 15, we have added an animation of the corresponding data to the supplemental material in this article.

In the simulations we do not observe streamer merging, which can occur if two streamer heads are sufficiently close in space and there is sufficient photoionization between them [79]. Although streamer merging is fundamentally possible under very specific conditions, the prerequisites for it to occur are obviously not present in our simulations. Streamer reconnection [80] is not observed in the simulations either. Reconnection occurs when a streamer connects into the wake of another streamer, and is thus a different phenomenon from streamer merging where two streamer heads merge directly. We have observed streamer reconnection in other 3D fluid simulations when one of the streamer filaments reach the ground plane and acts as a virtual ground for the other streamers. Further details regarding this finding will be reported elsewhere.

The simulation results presented here clearly can not be understood from axisymmetric simulations. Many axisymmetric studies of streamer discharges have been reported in the last decades, but these only show the emergence of a single filament and therefore have a limited range of applicability. The radius and velocity of the streamer then tend to increase with streamer length [38] (depending on the field conditions), but experiments generally show repeated branching into thinner filaments [77]. Because thinner streamers propagate slower than thicker streamers, it is reasonable to expect that repeated branching lowers the front velocity. It is obvious that the front velocity of the tree is determined by the velocity of the front streamers, but we point out that these streamers are also influenced by fields set up by neighboring branches. The dynamics of such streamers and corresponding single-filament streamers with the same radii are therefore not necessarily the same.

In our simulations the streamers branch into small-diameter streamers, and the most relevant parametric 2D studies are the ones pertaining to so-called minimal streamers. Such streamers have the smallest experimentally observed radius, do not branch, but still propagate over comparatively long distances. In order to compare our front velocity with that of single streamers, we consider the case in Li et al. [51] who presented a computational analysis of steady and stagnating positive streamers in air. The reported velocities in [51] varied from 0.25 mm/ns to 1.25 mm/ns when the streamer radii varied from 25 μm to 125 μm , and the corresponding electric field at the streamer tip varied from 220 kV/cm to 150 kV/cm. In our simulations we observe the same emerging radii and velocities, while the electric field is slightly higher at approximately ~ 250 kV/cm to ~ 280 kV/cm. We do not know the source of this discrepancy, but it could be due to the slightly larger numerical resolution used in our simulations. Another factor could be that Li et al. [51] use a correction to the α -coefficient while we do not. Both of these factors can facilitate small-scale features with higher fields. Regardless of these finer points, the streamer radii, velocities, and fields that emerge in our simulations are consistent with the parametric study in Li et al. [51].

5.6. Computational characteristics

We now present some of the computational characteristics for the simulation in section 5.5. The simulation was run on 32 nodes on the Norwegian supercomputer Betzy, and ran to completion in about 4 days. Each node on Betzy consists of two AMD Epyc 7742 CPUs for a total of 128 cores per node, so we used 4096 CPU cores in total. When we terminated the simulation it consisted




Fig. 16. Computational characteristics, showing the time (in seconds) spent in various computational kernels.

of approximately 500 million grid cells and 10^{10} computational particles, so there were approximately 15.6 million grid cells and 312.5 million particles per node. Although our simulations are firmly footed in the realm of high-performance computing, they do not represent a large burden for modern supercomputers (which currently can have more than one million CPU cores).

Fig. 16 shows a breakdown of the kernel costs for the time step at $t = 150$ ns, showing the wall clock time spent in various computational routines. This particular time step took around 26 s, but time steps varied down to around 1 s at the very beginning of the simulation. Radiation transport for our simulations has a negligible cost, but this would not be the case e.g. for sprite simulations where many more physical photons are generated per cell. The KMC algorithm has a cost of about 5% to 10%, whereas setting up and solving the Poisson equation had a relative cost of approximately 35%. Particle-related routines like superparticle handling, particle-mesh operations, transport, and spatial binning of particles had an accumulated cost of around 50%. There was also significant load imbalance for the super-particle handling since we load-balanced using the number of particles. This distributed the load along the streamer channels, but did not account for particle merging/splitting which mainly occurred in the streamer head.

Regridding the solution had about the same cost as 1-2 time steps, which was done every ten time steps. The majority of the regrid cost comes from 1) recomputing cut-cell stencils and re-solving the Poisson equation on the new grids, and 2) redistributing particles on the new grids. Particle redistribution is, unfortunately, quite expensive since virtually all computational particles change MPI rank ownership during regrids.

Two forms of I/O were used in the simulations: Plot files and checkpoint files. Plot files contained data for analysis and ranged up to 160 GB per file. Checkpoint files contained data for restarting simulations, e.g. in case of hardware failures or for allocating more nodes as the simulation mesh grows. These ranged up to 360 GB per file. Plot and checkpoint files took about two minutes to write, and were written every 100th time step.

6. Conclusions and outlook

6.1. Main findings

We have presented the foundation of a new type of model for streamer discharges based on a microscopic drift-diffusion model with a Kinetic Monte Carlo solver for the plasma chemistry. A thematic discussion on the role of $\hat{\text{Ito}}$ -KMC and its connection to conventional fluid models was presented. The model was coupled to photoionization with Monte-Carlo radiative transport, and a particle merging and splitting algorithm was presented. Suitable algorithms for integrating the equations of motion were then presented. These algorithms were implemented in 2D and 3D and adapted to cut-cell Cartesian AMR grids. We then implemented an example model for streamer discharges in air in needle-plane gaps, and showed that the $\hat{\text{Ito}}$ -KMC model agrees qualitatively and quantitatively with conventional fluid models. Example simulations that demonstrate the output and stability of the $\hat{\text{Ito}}$ -KMC model were then provided. The simulations presented in this paper demonstrate the feasibility of simulating discharge trees containing many streamer branches.

There are some advantages to using the new model, which are listed below:

- The model takes the same input as a fluid model, e.g. mobility and diffusion coefficients, and reaction rates.
- It is inherently a PIC model and maintains particle discreteness.
- $\hat{\text{Ito}}$ -KMC incorporates both reactive and diffusive fluctuations.
- The model is exceptionally stable in both space and time, even on very coarse grids and for large time steps. The absence of a CFL condition is particularly liberating.

We conjecture that the $\hat{\text{I}}$ to-KMC method will be suitable for hybrid modeling [62,63] where some of the electrons are treated kinetically. Unlike hybrid models based on a fluid description, transfer of electrons between $\hat{\text{I}}$ to-KMC and PIC-MCC descriptions can be done without disturbing the original charge distribution.

6.2. Future work

In this paper our focus has been on an all-discrete approach where also the ions are treated using $\hat{\text{I}}$ to diffusion. Future works will benefit from a mixed description where the electrons are treated using $\hat{\text{I}}$ to diffusion while (some of) the heavy species are treated using a continuum model. This can substantially reduce the computational load when more species of ions are tracked. However, the relative cost of the Poisson solver is already quite high, and even with this improvement the Poisson equation will remain a computational bottleneck.

As the $\hat{\text{I}}$ to-KMC is highly stable, higher order algorithms become particularly attractive. In the absence of solid boundaries, fourth order deposition and interpolation methods exist [81], but these techniques have not been extended to cases where embedded boundaries are involved. For the Poisson equation, fourth order discretizations that include embedded boundary formulations have been reported [82]. In time, only second-order convergence can be expected in the context of splitting methods and, furthermore, higher-order integration for particle transport is a significant challenge due to the presence of the Wiener process. Suppressing numerical streamer branches for coarse-grid simulations is also desirable. Moving forward, we will explore strategies for suppressing these by filtering the solutions [83].

The KMC-particle reactive coupling used in this paper adds some numerical diffusion into the system. Here, the algorithm is set up to uniformly distribute reactive products over a grid cell, and so it will end up placing secondary electrons in the wake of primary electrons. This can become a source of numerical instability quite similar to the one seen in fluid simulations, but fortunately we have not (yet) observed these types of instabilities in our simulations. Regardless, future efforts may benefit from reducing this source of numerical diffusion by introducing sub-grid models for the KMC-particle coupling.

We expect that 3D streamer simulations will become increasingly more sophisticated in the future. In parallel with this development, there is an emergent need for improving the tools that we use for analyzing such simulations. While two-dimensional simulations are relatively straightforward to analyze, 3D discharge trees are much more complex. Quantitative analysis requires us to skeletonize the discharge trees in full 3D for extraction of branching ratios and angles, filament lengths, velocities, and so on. In the future, we will also focus on establishing such analysis procedures.

CRedit authorship contribution statement

Robert Marskar: Conceptualization, Methodology, Formal analysis, Software, Investigation, Writing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Code availability statement

The computer code that was used to perform the calculations in this paper is publically available at <https://github.com/chombo-discharge/chombo-discharge>. Input scripts are available upon reasonable request.

Data availability

Data will be made available on request.

Acknowledgements

This study was partially supported by funding from the Research Council of Norway through grants 319930 and 321449. The computations were performed on resources provided by UNINETT Sigma2 - the National Infrastructure for High Performance Computing and Data Storage in Norway. The author expresses his gratitude to Fanny Skirbekk for providing the images used in Fig. 1.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jcp.2024.112858>.

References

- [1] S. Nijdam, J. Teunissen, U. Ebert, Plasma Sources Sci. Technol. (2020), <https://doi.org/10.1088/1361-6595/abaa05>.
- [2] V.P. Pasko, U.S. Inan, T.F. Bell, Geophys. Res. Lett. 25 (1998) 2123, <https://doi.org/10.1029/98GL01242>.

- [3] H.C. Stenbae-Nielsen, D.R. Moudry, E.M. Wescott, D.D. Sentman, F.T.S. Sabbas, *Geophys. Res. Lett.* 27 (2000) 3829.
- [4] R.A. Marshall, U.S. Inan, *Radio Sci.* 41 (2006), <https://doi.org/10.1029/2005RS003353>.
- [5] U. Ebert, S. Nijdam, C. Li, A. Luque, T. Briels, E. van Veldhuizen, *J. Geophys. Res. Space Phys.* 115 (2010), <https://doi.org/10.1029/2009ja014867>.
- [6] K.V. Laer, A. Bogaerts, *Plasma Sources Sci. Technol.* 25 (2015), <https://doi.org/10.1088/0963-0252/25/1/015002>.
- [7] Q.Z. Zhang, A. Bogaerts, *Plasma Sources Sci. Technol.* 27 (2018) 035009.
- [8] W. Wang, H.H. Kim, K.V. Laer, A. Bogaerts, *Chem. Eng. J.* 334 (2018) 2467.
- [9] K.V. Laer, A. Bogaerts, *Energy Technol.* 3 (2015) 1038.
- [10] S.M. Starikovskaia, *Plasma Assisted Ignition and Combustion*, 2006.
- [11] I.V. Adamovich, I. Choi, N. Jiang, J.H. Kim, S. Keshav, W.R. Lempert, E. Mintusov, M. Nishihara, M. Samimy, M. Uddi, *Plasma Sources Sci. Technol.* (2009), <https://doi.org/10.1088/0963-0252/18/3/034018>.
- [12] N.L. Aleksandrov, S.V. Kindysheva, I.N. Kosarev, S.M. Starikovskaia, A.Y. Starikovskii, *Proc. Combust. Inst.* (2009), <https://doi.org/10.1016/j.proci.2008.06.124>.
- [13] D. Breden, L.L. Raja, C.A. Idicheria, P.M. Najt, S. Mahadevan, *J. Appl. Phys.* 114 (2013), <https://doi.org/10.1063/1.4818319>.
- [14] S.M. Starikovskaia, *Plasma-Assisted Ignition and Combustion: Nanosecond Discharges and Development of Kinetic Mechanisms*, 2014.
- [15] A. Starikovskiy, *Physics and Chemistry of Plasma-Assisted Combustion*, 2015.
- [16] A. Abou-Ghazala, S. Katsuki, K.H. Schoenbach, F.C. Dobbs, K.R. Moreira, *IEEE Trans. Plasma Sci.* 30 (2002) 1449.
- [17] J.S. Clements, A. Mizuno, W.C. Finney, R.H. Davis, *IEEE Trans. Ind. Appl.* 25 (1989) 62.
- [18] J.S. Clements, M. Sato, R.H. Davis, *IEEE Trans. Ind. Appl.* IA-23 (1987) 224.
- [19] G. Dinelli, L. Civitano, M. Rea, *IEEE Trans. Ind. Appl.* 26 (1990) 535.
- [20] S.A. Nair, K. Yan, A.J. Pemen, G.J. Winands, F.M. van Gompel, H.E. van Leuken, E.J. van Heesch, K.J. Ptasinski, A.A. Drinkenburg, *J. Electrostat.* 61 (2004) 117.
- [21] D.R. Grymonpré, A.K. Sharma, W.C. Finney, B.R. Locke, *Chem. Eng. J.* 82 (2001) 189.
- [22] M. Laroussi, *Plasma 1* (2018) 47.
- [23] U. Ebert, M. Arrayás, A. Rocco, W. Hundsdorfer, Spontaneous branching of negative streamers, in: *IEEE Conference Proceedings*, 2002, p. 148.
- [24] M. Arrayás, U. Ebert, W. Hundsdorfer, *Phys. Rev. Lett.* 88 (2002) 1745021.
- [25] B. Bagheri, J. Teunissen, *Plasma Sources Sci. Technol.* 28 (2019) 45013.
- [26] R. Marskar, *Plasma Sources Sci. Technol.* (2020), <https://doi.org/10.1088/1361-6595/ab87b6>.
- [27] R. Marskar, *Plasma Res. Express* 1 (2019) 015011.
- [28] R. Marskar, *J. Comput. Phys.* 388 (2019) 624.
- [29] J. Teunissen, U. Ebert, *J. Phys. D, Appl. Phys.* 50 (2017) 474001.
- [30] B. Lin, C. Zhuang, Z. Cai, R. Zeng, W. Bao, *J. Comput. Phys.* (2020), <https://doi.org/10.1016/j.jcp.2019.109026>.
- [31] J. Teunissen, U. Ebert, *Plasma Sources Sci. Technol.* 25 (2016), <https://doi.org/10.1088/0963-0252/25/4/044005>.
- [32] A. Fierro, J. Stephens, S. Beeson, J. Dickens, A. Neuber, *Phys. Plasmas* 23 (2016), <https://doi.org/10.1063/1.4939475>.
- [33] A. Fierro, C. Moore, B. Yee, M. Hopkins, *Plasma Sources Sci. Technol.* 27 (2018) 105008.
- [34] J. Stephens, M. Abide, A. Fierro, A. Neuber, *Plasma Sources Sci. Technol.* 27 (2018), <https://doi.org/10.1088/1361-6595/aacc91>.
- [35] C. Köhn, O. Chanrion, L.P. Babich, T. Neubert, *Plasma Sources Sci. Technol.* 27 (2018) 015017.
- [36] C. Pavan, M. Martínez-Sánchez, C. Guerra-García, *Plasma Sources Sci. Technol.* 29 (2020) 095004.
- [37] F.J. Gordillo-Vázquez, F.J. Pérez-Invernón, *Atmos. Res.* 252 (2021) 105432.
- [38] B. Bagheri, J. Teunissen, U. Ebert, M.M. Becker, S. Chen, O. Ducasse, O. Eichwald, D. Loffhagen, A. Luque, D. Mihailova, J.M. Plewa, J.V. Dijk, M. Yousfi, *Plasma Sources Sci. Technol.* 27 (2018), <https://doi.org/10.1088/1361-6595/aad768>.
- [39] A. Luque, M. González, F.J. Gordillo-Vázquez, *Plasma Sources Sci. Technol.* (2017), <https://doi.org/10.1088/1361-6595/aa987a>.
- [40] L. Niemyer, L. Pietronero, H.J. Wiesmann, *Phys. Rev. Lett.* (1984), <https://doi.org/10.1103/PhysRevLett.52.1033>.
- [41] M. Akyuz, A. Larsson, V. Cooray, G. Strandberg, *J. Electrostat.* 59 (2003) 115.
- [42] A. Luque, U. Ebert, *New J. Phys.* 16 (2014) 013039.
- [43] M. González, F.J. Gordillo-Vázquez, A. Luque, *Plasma Sources Sci. Technol.* (2019), <https://doi.org/10.1088/1361-6595/ab4e7a>.
- [44] P.L.G. Ventzek, R.J. Hoekstra, M. Kushner, *J. Vac. Sci. Technol., B Microelectron. Nanometer Struct.* 12 (1994) 461.
- [45] A. Villa, L. Barbieri, M. Gondola, A.R. Leon-Garzon, R. Malgesini, *Comput. Fluids* 105 (2014) 1.
- [46] M. Niknezhad, O. Chanrion, J. Holbl, T. Neubert, *Plasma Sources Sci. Technol.* 30 (2021), <https://doi.org/10.1088/1361-6595/ac3214>.
- [47] S.V. Pancheshnyi, A.Y. Starikovskii, *Plasma Sources Sci. Technol.* 13 (2004), <https://doi.org/10.1088/0963-0252/13/3/B01>.
- [48] P. Colella, *J. Comput. Phys.* 87 (1990), [https://doi.org/10.1016/0021-9991\(90\)90233-Q](https://doi.org/10.1016/0021-9991(90)90233-Q).
- [49] V.R. Soloviev, V.M. Krivtsov, *Plasma Phys. Rep.* 40 (2014) 65.
- [50] J. Teunissen, *Plasma Sources Sci. Technol.* 29 (2020), <https://doi.org/10.1088/1361-6595/ab6757>.
- [51] X. Li, B. Guo, A. Sun, U. Ebert, J. Teunissen, *Plasma Sources Sci. Technol.* 31 (2022) 65011.
- [52] G.J. Hagelaar, L.C. Pitchford, *Plasma Sources Sci. Technol.* 14 (2005) 722.
- [53] Siglo database, www.lxcat.net, retrieved on November 1, 2022.
- [54] R.R. Arslanbekov, V.I. Kolobov, *Plasma Sources Sci. Technol.* 30 (2021), <https://doi.org/10.1088/1361-6595/abeff4>.
- [55] D.S. Dean, *J. Phys. A, Math. Gen.* (1996), <https://doi.org/10.1088/0305-4470/29/24/001>.
- [56] Y. Cao, D.T. Gillespie, L.R. Petzold, *J. Chem. Phys.* (2005), <https://doi.org/10.1063/1.1992473>.
- [57] Y. Cao, D.T. Gillespie, L.R. Petzold, *J. Chem. Phys.* (2006), <https://doi.org/10.1063/1.2159468>.
- [58] D.T. Gillespie, Exact stochastic simulation of coupled chemical reactions, *J. Phys. Chem.* 81 (25) (1977) 2340–2361.
- [59] A. Luque, U. Ebert, *Phys. Rev. E, Stat. Nonlinear Soft Matter Phys.* 84 (2011) 046411.
- [60] D.T. Gillespie, *Annu. Rev. Phys. Chem.* 58 (2007), <https://doi.org/10.1146/annurev.physchem.58.032806.104637>.
- [61] A. Noel, D. Makrakis, *IEEE Trans. Nanobiosci.* (2018), <https://doi.org/10.1109/TNB.2018.2878065>.
- [62] C. Li, U. Ebert, W. Hundsdorfer, *J. Comput. Phys.* 229 (2010) 200.
- [63] C. Li, U. Ebert, W. Hundsdorfer, *J. Comput. Phys.* 231 (2012) 1020.
- [64] H.K. Meyer, F. Mauseth, R. Marskar, A. Pedersen, A. Blaszczak, *IEEE Trans. Dielectr. Electr. Insul.* 26 (2019) 1163.
- [65] H.K. Meyer, R. Marskar, H. Gjerdal, F. Mauseth, *Plasma Sources Sci. Technol.* (2020), <https://doi.org/10.1088/1361-6595/abbae2>.
- [66] H.K.H. Meyer, R. Marskar, F. Mauseth, *Plasma Sources Sci. Technol.* 31 (2022) 114006.
- [67] P. Colella, D.T. Graves, T.J. Ligocki, G. Miller, D. Modiano, P.O. Schwartz, B.V. Straalen, J. Pilliod, D. Trebotich, M. Barad, B. Keen, A. Nonaka, C. Shen, *Echombo Software Package for Cartesian Grid, Embedded Boundary Applications*, 2003.
- [68] H. Johansen, P. Colella, *J. Comput. Phys.* 147 (1998) 60.
- [69] P. McCorquodale, P. Colella, H. Johansen, *J. Comput. Phys.* 173 (2001) 620.
- [70] P. Schwartz, M. Barad, P. Colella, T. Ligocki, *J. Comput. Phys.* 211 (2006) 531.
- [71] O. Chanrion, T. Neubert, *J. Comput. Phys.* 227 (2008) 7222.
- [72] J. Teunissen, U. Ebert, *J. Comput. Phys.* (2014), <https://doi.org/10.1016/j.jcp.2013.12.005>.
- [73] X.M. Zhao, J.C. Diels, C.Y. Wang, J.M. Elizondo, *IEEE J. Quantum Electron.* (1995), <https://doi.org/10.1109/3.364418>.

- [74] I.A. Kossyi, A.Y. Kostinsky, A.A. Matveyev, V.P. Silakov, *Plasma Sources Sci. Technol.* 1 (1992) 207.
- [75] M.B. Zheleznyak, A.K. Mnatsakanyan, S.V. Sizykh, *High Temp.* 20 (1982) 357.
- [76] S. Pancheshnyi, *Plasma Sources Sci. Technol.* 24 (2015), <https://doi.org/10.1088/0963-0252/24/1/015023>.
- [77] T.M. Briels, J. Kos, G.J. Winands, E.M.V. Veldhuizen, U. Ebert, *J. Phys. D, Appl. Phys.* 41 (2008), <https://doi.org/10.1088/0022-3727/41/23/234004>.
- [78] S. Chen, R. Zeng, C. Zhuang, *J. Phys. D, Appl. Phys.* 46 (2013), <https://doi.org/10.1088/0022-3727/46/37/375203>.
- [79] A. Luque, U. Ebert, W. Hundsdorfer, *Phys. Rev. Lett.* 101 (2008) 075005.
- [80] S. Nijdam, C.G.C. Geurts, E.M. van Veldhuizen, U. Ebert, *J. Phys. D, Appl. Phys.* 42 (2009) 045201.
- [81] A. Myers, P. Colella, B.V. Straalen, *SIAM J. Sci. Comput.* (2017), <https://doi.org/10.1137/16M105962X>.
- [82] D. Devendran, D. Graves, H. Johansen, T. Ligoeki, *Commun. Appl. Math. Comput. Sci.* 12 (2017) 51.
- [83] C.K. Birdsall, A.B. Langdon, *Plasma Physics via Computer Simulation*, 2004.