

An adaptive Cartesian embedded boundary approach for fluid simulations of two- and three-dimensional low temperature plasma filaments in complex geometries

Robert Marskar^a

^a*SINTEF Energy Research*

Abstract

We review a scalable two- and three-dimensional computer code for low-temperature plasma simulations in multi-material complex geometries. Our approach is based on embedded boundary (EB) finite volume discretizations of the minimal fluid-plasma model on adaptive Cartesian grids, extended to also account for charging of insulating surfaces. We discuss the spatial and temporal discretization methods, and show that the resulting overall method is second order convergent, monotone, and conservative (for smooth solutions). Weak scalability with parallel efficiencies over 70% are demonstrated up to 8192 cores and more than one billion cells. We then demonstrate the use of adaptive mesh refinement in multiple two- and three-dimensional simulation examples at modest core counts. The examples include two-dimensional simulations of surface streamers along insulators with surface roughness; fully three-dimensional simulations of filaments in experimentally realizable pin-plane geometries, and three-dimensional simulations of positive plasma discharges in multi-material complex geometries. The largest computational example uses up to 800 million mesh cells with billions of unknowns on 4096 computing cores. Our use of computer-aided design (CAD) and constructive solid geometry (CSG) combined with capabilities for parallel computing offers possibilities for performing three-dimensional transient plasma-fluid simulations, also in multi-material complex geometries at moderate pressures and comparatively large scale.

Keywords: HPC, Plasma, Streamer, Cartesian AMR, Complex geometry

1. Introduction

Electrical discharges in gaseous media are of substantial significance in a wide range of applications. In the high-voltage industry, for example, electrical discharges in the form of coronas, streamers, and leaders, can be particularly detrimental and lead to permanent damage on electrical equipment. The damage is usually induced by a hot electrical arc which (i) permanently decomposes the molecular structure of the gas, or (ii) leads to irreparable thermal damage on solid surfaces. For example, arcing may lead to electrode welding, or decomposition of conductive particles on the surfaces of insulators that degrade their performance over time. In other applications, electrical discharges are used as catalysts. Streamers - which are fast transients that evolve due to self-enhanced electric fields at their tips - have found their use in sterilization of polluted gases and liquids by initiating chemical reactions or breaking up molecules [1, 2, 3, 4, 5, 6]. In the aerodynamics community, dielectric barrier discharges are used to control the gas flow in the boundary layer on airplane wings [7, 8, 9]. Excellent reviews on electrical discharges and their applications can be found in [10, 11] and references therein.

Computer simulations of electrical discharges are challenging for several reasons. Firstly, electrical discharges are inherently nonlinear phenomena which evolve transiently with exponential increases in mass densities, and the discharge chemistry is often complicated, requiring inputs of rate constants that may

URL: robert.marskar@sintef.no (Robert Marskar)

be more or less uncertain. Secondly, electrical discharges can evolve as filamentary plasmas with dynamic ranges that span several orders of magnitude. In air at standard atmospheric conditions, for example, a streamer channel is shielded by a space charge layer that requires a numerical resolution on the micrometer scale, while the streamer head itself moves as an ionization wave which can stretch over several centimeters. Numerical solutions can usually not be obtained on coarse grids due to the strong nonlinearities of the problem, and mesh resolution requirements are often quite extreme. Integration of the equations of motion require concurrent updates of electrodynamic, fluid (or kinetic), and radiative transport equations, possibly even coupled with nonlinear boundary conditions on electrodes and insulators. For this reason, software development of low-temperature plasma codes for large scale simulations is a challenge for researchers and industry alike, and production runs of three-dimensional discharges in engineering-relevant geometries may even be a challenge for today's largest parallel computers. Few results on three-dimensional propagation of streamers have therefore been presented to date. No reported results have so far combined 3D simulations of filamentary plasmas with the geometrical complexities encountered in e.g. applied high-voltage engineering, which can feature multiple types of materials, for example electrodes and insulators, which require separate treatment with respect to boundary conditions.

In three dimensions, all results reported in the scientific literature to date rely on geometrical simplicity, such as plane-parallel electrodes [12, 13, 14, 15] or point-plane gaps [16, 17, 18]. Some computer codes also rely on adaptive mesh refinement (AMR) for reducing the computational effort, a technique which becomes mandatory for large scale 3D simulations. Three-dimensional simulations have been performed by using a code based on Gerris in point-plane gaps [19]. Similar simulations, although on a very small length scale, have also been performed by using COMSOL Multiphysics [20], and an in-house code that was recently reported [18]. Interactions between multiple streamers have also recently been investigated [21]. Concept simulations for streamer discharges in transformer oil have also been presented [22, 23, 24, 25], but the simulations disregard the liquid-gas phase transition (which is observed in virtually all experiments), and must therefore be met with some skepticism. Teunissen and Ebert [15] also report on three-dimensional streamer simulations through atmospheric-pressure nitrogen in plane-parallel 4 cm gaps. The same authors have also presented fully 3D simulations in $(24\text{ cm})^3$ plane-parallel domains for oxygen-nitrogen mixtures at 133 mbar [14]. Note that for streamer discharges there is an inherent scaling with pressure; at decreasing pressures the physical length and time scales of the discharge increases. Decreased pressure therefore acts as a magnifying glass that allow simulations of larger domains and longer time scales. In our opinion, the simulations in [14, 15] ran quite fast, within 24 hours on two Xeon E5-2680v4 processors. The results clearly demonstrate the benefits of using AMR.

In the two-dimensional domain (either Cartesian or cylindrically symmetric coordinates), streamer simulations in complex geometries have progressed further. Many of these results arise from in-house codes from the aerodynamics community [7, 8, 26, 9, 27], but commercial and open source codes are also gaining traction. Other researchers have shown interest in engineering and atmospheric applications. For example, by using a static mesh, Dubinova et al. [28] has investigated streamer propagation along a cylindrical rod by using the Plasimo software [29], and the results qualitatively agree with experiments. Inception criteria for streamers near large ice particles have also been simulated [30], as well as various investigations of dielectric barrier discharges [31, 32, 33, 34]. Temporally adaptive methods have also been reported [35], and show promise.

Even though simulation capabilities of low-temperature plasmas have increased steadily over the past few decades, quantitative prediction using low-temperature plasma simulations remains challenging, particularly at the fidelity that is required for engineering applications. Some of these challenges are natural and arise due to a lack of physics in the mathematical models, uncertainty in rate constants for chemical reactions, or the inability of the code to scale to the desired length and time scales. Another uncertainty is the fact that computer codes used by various academic groups do not give the same answer to the same problem [36], demonstrating a broader need for code verification and validation efforts.

In this paper, we review a parallel two- and three-dimensional computer code for simulations of low-temperature plasmas in complex geometries. Our work uses adaptively refined structured Cartesian grids, and includes support for internal boundary conditions for both insulators and electrodes. We embed our

algorithms into Chombo [37, 38], which is a high-performance computing (HPC) framework that runs on systems ranging from desktops to supercomputers. Our current solvers include a multifluid Poisson solver; convection-diffusion-reaction solvers, and diffusion-approximation models for the radiative transport equation (RTE).

Our computational approach is based on structured adaptive mesh refinement (SAMR) with an embedded boundary (EB) formalism. Structured grids normally lead to good conditioning of the discretized equations and therefore resolve some of the disadvantages that are commonly encountered with 3D unstructured adaptive grids. These include a lack of efficient mesh adaptation and convergence loss with skewed elements, which is especially problematic for adaptive 3D unstructured grids near curved boundaries. In addition, unstructured grids use indirect memory addressing for data accessing, often leading to bandwidth limitations of the code. SAMR resolve these issues even for geometrically complex cases, at the cost of a significant increase in algorithm complexity, difficulties which are primarily related to spatial discretization and time stepping near the boundaries. We discuss how we resolve some of these difficulties.

The structure of this paper is as follows: In Sec. 2 we review the minimal plasma model and present the equations of motion and their coupling, and we outline our numerical methods in Sec. 3. Most of the methods discussed in Sec. 3 are known in the literature, and the main algorithmic work involved in this paper is therefore a demonstration of their successful linkage. Code verification studies are presented in Sec. 4 and weak scalability studies are presented in Sec. 5 for up to 8192 cores. Demonstrations of our computer code are given in Sec. 6, and some final remarks and a brief summary are given in Sec. 7 and Sec. 8.

2. Theory

Typically, streamers develop in the non-uniform electric field near curved electrodes. Their spatial and temporal development occur due to field driven acceleration of electrons to kinetic energies above the required ionization energy, leading to an avalanche of electrons that leave behind positive ions. For the most part, ions can be modeled as fluids whereas the most accurate electron descriptions are provided by kinetic models (e.g. particle models) that sample the electron phase space distribution. Kinetic models capture many of the physical mechanisms of streamers (e.g. stochastic branching and run-away electrons), but the spatial and temporal resolutions need be extremely fine which limit their applicability to comparatively small spatial and temporal scales. On the other hand, electron fluid models evolve the average electron density (or momentum) without regard to the velocity distribution function (which is generally not Maxwellian) and do not share this practical restriction. In fluid models, the electron evolution is described by its drift and diffusion velocities, and the production rate of electron ion-pairs is described by rate coefficients rather than stochastic evaluation of ionization probabilities. In the simplest case, which is the local field approximation (LFA) where the electrons are in equilibrium with the electric field, the ionization coefficient is simply a function of the electric field alone. More sophisticated physics is possible by using the local energy approximation (LEA) where this coefficient is a function of the average electron energy rather than the field. In either case, when compared to kinetic models the available level of physics in electron fluid models is quite limited. However, electron fluid models are applicable to much larger spatial scales and are therefore useful in the range of applications where kinetic models are unrealistic to apply. In this paper we will consider the minimal fluid model where the electrons and ions are described only by their average densities, and the momentum equations are closed by imposing drift and diffusion velocities to be functions of the electric field. Certainly, improvements to this model are possible by using Euler equations for the electrons, as is done in [39, 40, 41].

2.1. Electrodynamic equations

For a weakly ionized plasma such as a streamer, the electric current is sufficiently small so that one may disregard magnetic field effects. The remaining Maxwell equation for the evolution of the electromagnetic field is Poisson's equation:

$$\nabla \cdot (\epsilon \nabla \Phi) = -\frac{\rho}{\epsilon_0}, \quad (1)$$

where Φ is the electric potential, ϵ the relative permittivity, and ρ the charge density. The other Maxwell equation for $\mathbf{E} = -\nabla\Phi$ is $\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} = 0$, but we only use this for estimating a time step size which ensures numerical stability during our simulations. Eq. (1) must be supported by boundary conditions on the domain edges (faces in 3D) and internal surfaces.

2.2. Plasma-fluid equations

The spatiotemporal evolution of electrons, ions, and neutrals is solved in the continuum approximation. For each species μ we have

$$\frac{\partial n_\mu}{\partial t} + \nabla \cdot (\mathbf{v}_\mu n_\mu - D_\mu \nabla n_\mu) = S_\mu, \quad (2)$$

where n_μ is the volumetric density of species μ ; \mathbf{v}_μ its drift velocity, D_μ its diffusion coefficient, and S_μ a source term (describing the interplay between attachment, impact ionization, photoionization etc.). Equation (2) is a convection-diffusion-reaction (CDR) equation which describes the evolution of individual species densities.

Equation (2) is supplemented by boundary conditions describing the mass flux into or out of the computational domain. Our approach is to use outflow conditions on domain walls, and to expose the boundary conditions on embedded electrodes and dielectrics in an abstract plasma-kinetic framework which allows us to modify the surface kinetics of the plasma without affecting the underlying solver code.

2.3. Radiative transfer

Radiative transfer is handled by solving the radiative transfer equation (RTE) in the diffusion approximation. The RTE is

$$\frac{\partial f_\nu(\mathbf{x}, \boldsymbol{\Omega}, t)}{\partial t} + c\boldsymbol{\Omega} \cdot \nabla f_\nu(\mathbf{x}, \boldsymbol{\Omega}, t) = -c\kappa_\nu(\mathbf{x})f_\nu(\mathbf{x}, \boldsymbol{\Omega}, t) + \frac{1}{4\pi}\eta_\nu(\mathbf{x}), \quad (3)$$

where f_ν is the photon distribution function (i.e. the number of photon with frequency ν at (\mathbf{x}, t) traveling in direction $\boldsymbol{\Omega}$), κ_ν is the Beer's length for photons ν , and η_ν is an isotropic source term (hence the factor of 4π). Thus, η_ν is the total number of photons produced at (\mathbf{x}, t) .

2.3.1. The multigroup approximation

In the RTE, the frequency ν is a continuous variable. For most applications it becomes necessary to reduce the computational load by invoking the monochromatic multigroup approximation. That is, one assumes that f_ν consists of a number of frequency bands γ where each frequency band is sufficiently sharp-line in order to individually invoke a monochromatic approximation. We take $f_\nu = \sum_\gamma f_\gamma \delta(\nu - \gamma)$ which essentially replaces Eq. (3) with a finite set of equations for each frequency band γ .

The multigroup RTE is solved in the diffusive SP₁ approximation (i.e. the *Eddington* approximation) by closing the first order moment equation. That is, taking the first and second moments of Eq. (3) with respect to $\boldsymbol{\Omega}$ yields

$$\frac{\partial \Psi_\gamma}{\partial t} + \nabla \cdot \mathbf{F}_\gamma = -c\kappa_\gamma \Psi_\gamma + \eta_\gamma, \quad (4a)$$

$$\frac{\partial \mathbf{F}_\gamma}{\partial t} + \nabla \cdot \boldsymbol{\Pi}_\gamma = -\kappa_\gamma \mathbf{F}_\gamma, \quad (4b)$$

where $\Psi_\gamma = \int_{4\pi} f_\gamma d\boldsymbol{\Omega}$ is the radiative density, $\mathbf{F}_\gamma = c \int_{4\pi} f_\gamma \boldsymbol{\Omega} d\boldsymbol{\Omega}$ is the radiative flux, and $\boldsymbol{\Pi}_\gamma^{ij} = c \int_{4\pi} f_\gamma \Omega_i \Omega_j d\boldsymbol{\Omega}$ is the Eddington tensor. This system is closed by assuming

$$\boldsymbol{\Pi}_\gamma^{ij} = \frac{c}{3} \delta_{ij} \Psi_\gamma, \quad (5)$$

which is equivalent to the Eddington approximation. Note that the underlying assumption here is that the medium is primarily a scattering medium so that the radiation moves in a diffusion-like manner. In such cases the time required for substantial increase in the radiative flux is much longer than the time required

for traversing a mean free path. We can therefore impose $\partial_t \mathbf{F}_\gamma = 0$, and insertion of Eq. (5) into Eq. (4b) yields

$$\mathbf{F}_\gamma = -\frac{c}{3\kappa_\gamma} \nabla \Psi_\gamma, \quad (6)$$

which is expected from a diffusive approximation. Insertion of Eq. (6) into Eq. (4a) yields

$$\frac{1}{c} \frac{\partial \Psi_\gamma}{\partial t} + \kappa_\gamma \Psi_\gamma - \nabla \cdot \left(\frac{1}{3\kappa_\gamma} \nabla \Psi_\gamma \right) = \frac{\eta_\gamma}{c}. \quad (7)$$

This equation must be supplemented by appropriate boundary conditions for radiative transport. In the Eddington approximation, the appropriate outflow boundary is of the Robin type [42]:

$$\partial_n \Psi_\gamma + \frac{3\kappa_\gamma}{2} \frac{1+3r_2}{1-2r_1} \Psi_\gamma = \frac{g_\gamma}{\kappa_\gamma}, \quad (8)$$

where r_1 , and r_2 are reflection coefficients and g_γ is a surface source. In this paper, we do not consider photon reflection from boundaries nor injection of photons into the domain and the outflow boundary condition on photons on a boundary with outward (i.e. out of the gas-phase) normal vector \mathbf{n} simplifies to

$$\frac{\partial \Psi_\gamma}{\partial n} + \frac{3\kappa_\gamma}{2} \Psi_\gamma = 0, \quad (9)$$

where $\partial_n \Psi_\gamma = \mathbf{n} \cdot \nabla \Psi_\gamma$. Thus, for free outflow of photons the boundary flux is

$$\mathbf{F}_\gamma \cdot \mathbf{n} = \frac{c}{2} \Psi_\gamma. \quad (10)$$

Note that the diffusion assumption does not hold in a general manner. However, a successful strategy has been to combine solutions of (7) so that it represents the photon distribution of the exact integral solution [43]. We remark that diffusive approaches disregard even qualitative features (such as shadows) of the full RTE solutions. This is true even for improved diffusion models, such as the SP₃ approximation [42, 44, 45]. While we currently adopt the diffusion approximation, we can certainly foresee a necessity to account for higher-order moments of the RTE, or the use of Monte Carlo methods, sometime in the future. In fact, this will be necessary for studies of porous media, such as packed-bed DBD reactors. Currently, direct solutions of the RTE equation [46] in three dimensions appears infeasible, even on today's largest supercomputers.

2.4. Surface charge conservation

Our final equation of motion is local conservation of charge on *dielectric* surfaces. We do not solve an equivalent problem on electrodes because we assume that complete neutralization occurs on both anodes and cathodes, and that the voltage on the live and grounded electrodes are controlled by an external circuit. Charge conservation on dielectric surfaces is given by

$$\frac{\partial \sigma}{\partial t} = F_\sigma, \quad (11)$$

where F_σ is the charge flux. F_σ is always coupled to the boundary conditions for the species densities n_μ since charge must be conserved at the surface. Thus, F_σ is simply the sum of all species fluxes at the boundary. Currently, we do not include solid state kinetics and the surface state is therefore characterized only by the net charge density, without detailed knowledge of the distribution of the ionic states that are involved on the surface.

2.5. Minimal plasma model

To summarize, our plasma model consists of the following equations:

$$\nabla \cdot (\epsilon_r \nabla \Phi) = -\frac{\rho}{\epsilon_0}, \quad (12a)$$

$$\frac{\partial \sigma}{\partial t} = F_\sigma, \quad (12b)$$

$$\frac{1}{c} \frac{\partial \Psi_\gamma}{\partial t} + \kappa_\gamma \Psi_\gamma - \nabla \cdot \left(\frac{1}{3\kappa_\gamma} \nabla \Psi_\gamma \right) = \frac{\eta_\gamma}{c}, \quad (12c)$$

$$\frac{\partial n_\mu}{\partial t} + \nabla \cdot (\mathbf{v}_\mu n_\mu - D_\mu \nabla n_\mu) = S_\mu, \quad (12d)$$

where the final two equations denote *sets* of equations for μ and γ .

For streamer discharges Eq. (12) is nonlinearly coupled through source terms and species velocities. For the purposes of this paper, this coupling is restricted to the following:

$$\epsilon_r = \epsilon_r(\mathbf{x}), \quad (13a)$$

$$\kappa_\gamma = \kappa_\gamma(\mathbf{x}), \quad (13b)$$

$$\eta_\gamma = \eta_\gamma(\mathbf{E}, n_\mu), \quad (13c)$$

$$\mathbf{v}_\mu = \mathbf{v}_\mu(\mathbf{E}, n_\mu), \quad (13d)$$

$$D_\mu = D_\mu(\mathbf{E}), \quad (13e)$$

$$S_\mu = S_\mu(\mathbf{E}, n_\mu, \Psi_\gamma), \quad (13f)$$

where ϵ_r may additionally be discontinuous.

3. Numerical methods

SAMR exists in two separate categories, patch-based and tree-based AMR. Patch-based AMR is the more general type and contain tree-based grids as a subset; they can use refinement factors other than 2, as well as accommodate anisotropic resolutions and non-cubic patches. In patch-based AMR the domain is subdivided into a collection of hierarchically nested overlapping patches (or boxes). Each patch is a rectangular block of cells which, in space, exists on a subdomain of the union of patches with a coarser resolution. Patch-based grids generally do not have unique parent-children relations: A fine-level patch may have multiple coarse-level parents. An obvious advantage of a patch-based approach is that entire Cartesian blocks are sent into solvers, and that the patches are not restricted to squares or cubes. A notable disadvantage is that the overlapping grids inflate memory, and that additional logic is required when updating a coarse grid level from the overlapping region of a finer level. Tree-based AMR use quadtree or octree data structures that describe a hierarchy of unique parent-children relations throughout the AMR levels: Each child has exactly one parent, whereas each parent has multiple children (4 in 2D, 8 in 3D). For CPU cache performance reasons, the leaves of an octree are often cubic patches (e.g. 4^3 or 8^3 boxes), but the mesh can also be refined on a cell-by-cell basis. However, the use of single cell leaves becomes prohibitive at large scale for two reasons. The first is that special case must be taken in order to avoid memory inflation due to a growing tree structure. The second is that such trees, while still being SAMR, use indirect memory referencing, thus adding latency in data accessing and processing. This typically leads to poorer CPU performance since the data defined in neighboring cells may be stored on different cache lines.

We have implemented Eq. (12) into Chombo [37, 38], which is a mixed-language C++ & Fortran library for performing patch-based AMR that also includes support for embedded boundaries. In EB applications, the mesh is additionally described by a graph near the cut-cells that describes the connectivity and geometric moments of the cell fragments. The cells in the AMR hierarchy are related to one another by refinement and coarsening operations on that graph [38]. Figure 1 shows an example of patch-based refinement near a complex boundary.

3.1. Spatial discretization

We now give a brief summary of the Chombo spatial discretization [38]. We discretize the physical domain by a set of Cartesian cells, where some cells are cut by a level-set function $s(\mathbf{x})$ where $s(\mathbf{x}) = 0$ describes the boundary interface. The geometric information is generated by a graph on the finest level grid first, which is then coarsened throughout the hierarchy. This information is stored in a sparse data format for all levels, with a memory overhead that depends on the complexity of the geometry. Inside each cut-cell, the boundary is linearized so that it intersects through the cell as lines in 2D, and planes in 3D. Control volume faces that are cut by this surface are referred to as cut-faces, and the boundary itself is referred to as an embedded boundary (EB).

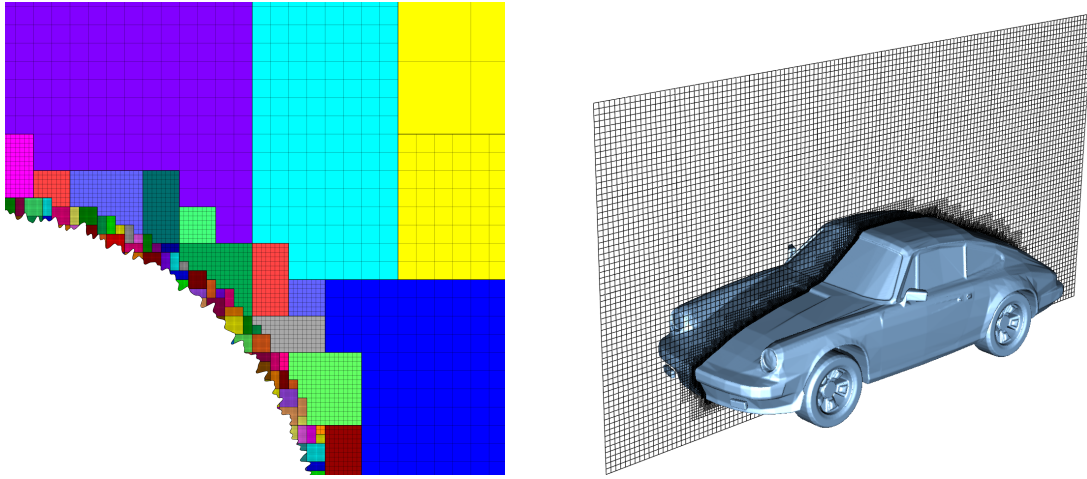


Figure 1: Left: Patch-based refinement near a complex boundary. The colors indicate individual patches and the refinement factor between each level is 4. Right: EB data generation from a surface tessellation. The surface mesh of the car contains 5168 vertices and 10316 facets. The complete EB mesh contains three levels of refinement (factor 2 between levels) of a base mesh of $128 \times 64 \times 64$.

Figure 2 shows the cut cells in greater detail. Here, a cell is identified by its index \mathbf{i} . We will take $\mathbf{x}_{\mathbf{i}}$ to be the *cell center*, and $\bar{\mathbf{x}}_{\mathbf{i}}$ to be the *cell centroid*. The volume fraction is $0 \leq \kappa_{\mathbf{i}} \leq 1$: Cells with $\kappa_{\mathbf{i}} = 0$ are denoted as *covered* cells and lie completely outside the computational domain; cells with volume fractions $0 < \kappa_{\mathbf{i}} < 1$ are called *irregular* cells and lie partially inside the computational domain. If an irregular cell contains more than one cell fragment, it is termed a multi-valued cell. By construction, multi-valued cells occur due to graph coarsening and do not exist on the finest graph level. Cell faces are denoted by $f_d^{\pm}(\mathbf{i})$ where \pm indicates the high (+) or low (-) face of the cell \mathbf{i} in the coordinate direction d . A face center for face $f_d^{\pm}(\mathbf{i})$ is therefore located at $\mathbf{x}_{\mathbf{i} \pm \frac{1}{2} \mathbf{e}_d}$ where \mathbf{e}_d is a unit vector in the d -direction. Face centroids are likewise denoted $\bar{\mathbf{x}}_{\mathbf{i} \pm \frac{1}{2} \mathbf{e}_d}$ and the corresponding area fractions are denoted by $\alpha_{\mathbf{i} \pm \frac{1}{2} \mathbf{e}_d}$. Finally, the EB centroid is denoted by $\bar{\mathbf{x}}_{\mathbf{i}}^{\text{EB}}$ with a corresponding area fraction $\alpha_{\mathbf{i}}^{\text{EB}}$ and an outward normal vector $\hat{\mathbf{n}}_{\mathbf{i}}^{\text{EB}}$. Explicitly defined expressions for these quantities can be found in [38].

3.2. Geometry generation

In addition to using the native Chombo tools for geometry generation [38], we have developed tools for generating level-set functions from surface tessellations. Our surface mesh implementation uses a doubly-connected edge list (DCEL) that supports mixtures of arbitrarily-sided planar polygons, but requires that the tessellation is a two-manifold non-intersecting, watertight surface. There are large performance penalties in using surface tessellations for generating the signed distance field that Chombo requires for EB generation since we must provide the signed distance to the closest polygon, vertex, or edge. Brute force searching through all polygons is unacceptably slow, so we accelerate the polygon search by embedding the polygonal tessellation in a boundary volume hierarchy that allows orders of magnitude faster traversal through the

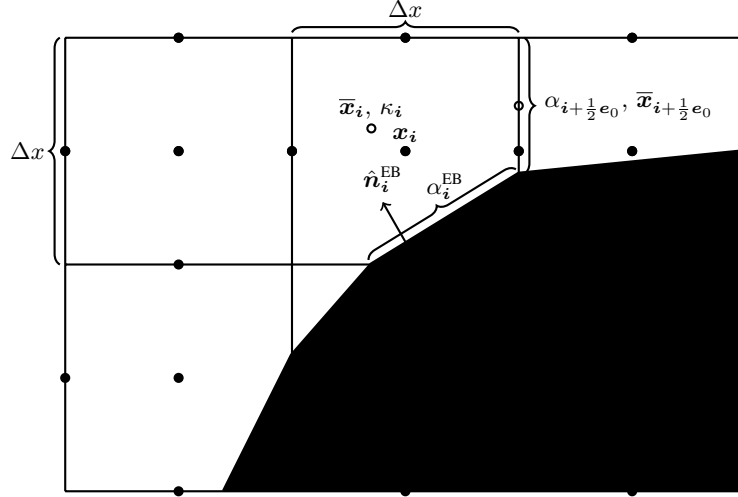


Figure 2: Cut-cell discretization.

tessellation. We use a k -d tree with axis-aligned boundary boxes for aggregated polygons, and the closest-feature search over the tessellation is performed by building a priority list as the tree is traversed from the root to the leaves. Figure 1 shows an example of an EB mesh that was generated from polygonal data.

3.3. Helmholtz equations

Next, we discuss the discretization of the Helmholtz equation

$$\alpha a(\mathbf{x})\phi + \beta \nabla \cdot (b(\mathbf{x})\phi) = \rho. \quad (14)$$

Helmholtz equation solvers are an important part of this paper. For example, the Poisson equation is represented by $\alpha = 0$, $\beta = -\epsilon_0$, $b(\mathbf{x}) = \epsilon_r(\mathbf{x})$. Furthermore, temporal discretizations of parabolic equations are also underpinned by a Helmholtz solver. For example, discretizing Eq. (7) by using the backward Euler method yields

$$(1 + c\Delta t\kappa_\gamma)\Psi_\gamma^{n+1} - \nabla \cdot \left(\frac{c\Delta t}{3\kappa_\gamma} \nabla \Psi_\gamma^{n+1} \right) = \Psi_\gamma^n - \frac{\Delta t\eta_\gamma}{c}, \quad (15)$$

which defines a Helmholtz problem for Ψ_γ^{n+1} .

We follow the work in [47, 48, 49, 50] and use the finite volume method for the Helmholtz equation. For ease of notation, we restrict the discussion below to the case $a = 0$ which yields the Poisson equation. Extensions to the full Helmholtz problem is straightforward by adding in another diagonal term. Our implementation of the Helmholtz equation also supports multi-fluids, i.e. cases in which $b(\mathbf{x})$ is additionally discontinuous across a level-set surface. The multifluid problem needs additional encapsulation of a quasi-boundary condition on the interface between two materials p and p' , given by

$$b_p \frac{\partial \phi}{\partial n_p} + b_{p'} \frac{\partial \phi}{\partial n_{p'}} = \sigma, \quad (16)$$

where \mathbf{n}_p and $\mathbf{n}_{p'}$ are unit normals that point into each fluid, with $\mathbf{n}_{p'} = -\mathbf{n}_p$, and σ is a surface source term. Recasting Eq. (14) in integral form yields

$$\oint_A b(\mathbf{x}) \nabla \phi \cdot d\mathbf{A} = \frac{1}{\beta} \int_V \rho dV. \quad (17)$$

We consider the cell shown in Fig. 3. Here, the volume V_i is a cut-cell at a domain boundary. Integration of Eq. (17) over the irregular cell yields

$$\oint_A b(\mathbf{x}) \nabla \phi \cdot d\mathbf{A} = (\alpha_1 F_1 + \alpha_2 F_2 + \alpha_3 F_3 + \alpha_D F_D + \alpha_{EB} F_{EB}) \Delta x, \quad (18)$$

where the fluxes are centroid-centered on their respective faces and α_i are face area fractions. The centroid fluxes are evaluated by constructing second order accurate face-centered fluxes, which are then interpolated to the respective centroids. For example, for the flux through the top face in Fig. 3 we find a standard expression for second order accurate approximations of the first derivative:

$$F_3 = F_{i,j+1/2} = b_{i,j+1/2} \frac{\phi_{i,j+1} - \phi_{i,j}}{\Delta x}, \quad (19)$$

For fluxes through face centroids we interpolate the face-centered fluxes. For example, the flux F_2 in Fig. 3 is given by

$$F_2 = [F_{i+1/2,j}(1-s) + sF_{i+1/2,j+1}], \quad (20)$$

where s is the normalized distance from the face center to the face centroid, and $F_{i+1/2,j}$ and $F_{i+1/2,j+1}$ are face-centered fluxes.

Flux evaluation on coarse-fine boundaries is slightly more involved. The AMR way of handling this is to reflux the coarse side by setting the flux into the coarse cell to be the sum of fluxes from the abutting finer cells. In Chombo, this is done by precomputing a set of flux registers that hold the face centered fluxes on both sides of the coarse-fine interface (see [38] for additional details). Refluxing is then a matter of subtracting the coarse flux from the divergence computation, and adding in the sum of the fine face fluxes. I.e. let $\{f_f(f_c)\}$ be the set of fine faces that are obtained when coarsening of a coarse face f_c . In the reflux step, the divergence operator in the coarse cell is modified as

$$\nabla \cdot \mathbf{F} \rightarrow \nabla \cdot \mathbf{F} + \frac{1}{\Delta x} \left(\sum_f F_f - F_c \right), \quad (21)$$

where F_c and F_f are the coarse and fine-face fluxes, and the sum runs over all the fine faces that abut the coarse face. For further details regarding the discretization and conditioning of the Poisson equation, see [47, 48, 49, 50].

3.3.1. Boundary conditions

Next, we discuss four types of boundary conditions for the Helmholtz equation: Neumann, Dirichlet, Robin, and multifluid type boundary conditions. For Neumann boundary conditions the domain and embedded boundary fluxes are specified directly. For Dirichlet boundary conditions the process is more involved. For Dirichlet conditions on domain faces we apply finite differences in order to evaluate the flux through the face. For example, for a constant Dirichlet boundary condition $\phi = \phi_0$ the face-centered flux at the bottom face in Fig. 3 is, to second order,

$$F_{i,j-1/2} = -\frac{b_{i,j-1/2}}{\Delta x} \left(3\phi_{i,j+1} - \frac{1}{3}\phi_{i,j} - \frac{8}{3}\phi_0 \right) \quad (22)$$

As with the flux F_2 on the interior face, fluxes on domain faces are also interpolated to face centroids. Thus, F_D becomes

$$F_D = [F_{i,j-1/2}(1-t) + tF_{i-1,j-1/2}], \quad (23)$$

where t is the distance from the face center to the face centroid.

The evaluation of Dirichlet boundary conditions on the EB is more complicated because the EB normal does not align with any of the coordinate directions. To evaluate the flux on the boundary we construct ray

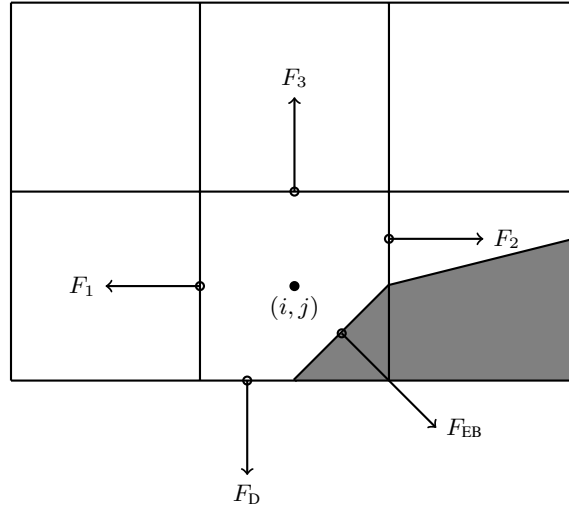


Figure 3: Cut cell at domain face.

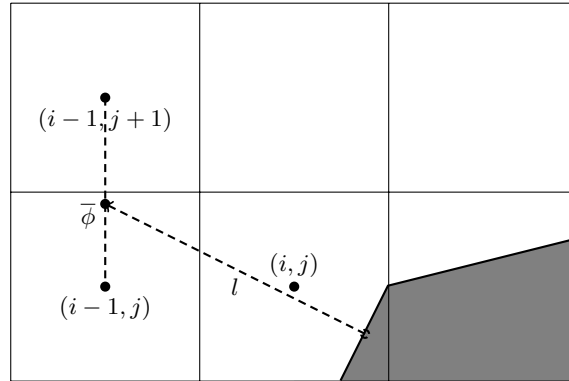


Figure 4: Ray casting at the EB for obtaining the normal gradient.

based or least squares based stencils for evaluating $\partial_n \phi$ (see [47] or [38] for details). Regardless of which approach is used, we have

$$\frac{\partial \phi}{\partial n} = w_0 \phi_0 + \sum_{i \in \psi} w_i \phi_i, \quad (24)$$

where ϕ_0 is the Dirichlet value on the boundary, w_0 is a boundary weight and ψ is a stencil that contains only interior points. The weights w_i are weights for these points. As an example, consider the flux depicted in Fig. 4. The first order accurate partial derivative on the boundary is given by

$$\frac{\partial \phi}{\partial n} = \frac{\phi_0 - \bar{\phi}}{l}, \quad (25)$$

where $\bar{\phi}$ is the interpolated value at the intersection of the ray and the line that connects $\mathbf{x}_{i-1,j}$ and $\mathbf{x}_{i-1,j+1}$. Since $\bar{\phi}$ can be linearly interpolated by using these two interior points only, this is clearly in the form of Eq. (24). The boundary derivative stencils are well separated from the boundary (i.e. they do not use the values of the irregular cell itself). For the Poisson equation this is a requirement in order to achieve good conditioning of the discretized system as the volume fraction approaches zero [47].

Higher-order approximations to the flux are built in a similar way by including more interior cells. In our experience, the best convergence results come from using second order accurate ray-based boundary stencils, which requires 3 ghost cells in the general case. If we cannot find a stencil for computing the normal derivative by ray-casting, which can occur if there aren't enough cells available, we use quadrant-based least squares for computing the normal derivative (again, see [47] or [38]).

We have also implemented Robin boundary conditions of the type

$$a_1\phi + a_2\frac{\partial\phi}{\partial n} = a_3, \quad (26)$$

which is an appropriate type of boundary condition for the radiative transfer equation. The normal derivative is given by $\partial_n\phi = (a_3 - a_1\phi)/a_2$ so that extrapolation of ϕ to the boundary is sufficient for imposing the boundary flux. Our way of doing this is simply to extrapolate ϕ to the boundary by using either least squares or Taylor-based stencils.

On multifluid boundaries the boundary condition is neither Dirichlet, Neumann, or Robin. Multifluid boundaries are more complex since the state at the boundary is not known, but rather depends on the solution inside both fluids. Our approach follows that of [50] where we first compute stencils for the normal derivative on each side of the boundary,

$$\frac{\partial\phi}{\partial n_q} = w_0^q\phi_B + \sum_{i \in \psi_q} w_i^q\phi_i, \quad (27)$$

where $q = p$ or $q = p'$ and ϕ_B is the solution on the surface centroid, and the stencil ψ_q only reaches into fluid q . The linear nature of this equation allows one to obtain the surface state ϕ_B from the matching condition Eq. (16), which can then be eliminated in order to evaluate $\partial\phi/\partial n_q$ on both sides of the boundary.

3.3.2. Multigrid solver

To solve the discretized Helmholtz equation we use the geometric multigrid (GMG) solver template that ships with Chombo [38]. GMG involves smoothing of the solutions on progressively coarsened grids and is compatible with AMR. Smoothing on each level involves relaxation (e.g. Jacobi or Gauss-Seidel), which primarily reduces the magnitude of high frequency errors. Removal of low-frequency errors from the solution is much slower. Because of this, multigrid accelerates convergence by projecting the error onto a coarser grid where the error has, from the viewpoint of the grid, a shorter wavelength, making relaxation more efficient. Once a bottom grid level has been reached and an approximate bottom-level solution has been found, the error is prolonged onto a finer grid and relaxation is then re-applied. Geometric multigrid works best when the long wavelength modes of the fine grid operator are well represented as short wavelength modes on the coarse grid operator. For EB applications however, coarsening can result in the removal of finer geometric features so that the relaxation step cannot sufficiently dampen the error modes at which GMG is aimed at. Because of this, geometric multigrid for EB applications usually involve lower convergence rates between each multigrid cycle than it does for geometry-less domains and, moreover, typically involves dropping to the bottom solver sooner. Currently, we only support relaxation solvers as the bottom solver for multi-phase problems, whereas we use the built-in BiCGStab and GMRES solvers in Chombo [38] for single-phase elliptic problems. In the future, we would like to use algebraic multigrid from e.g. PETSc as a bottom solver in the V-cycle in order to enhance solver efficiency for very complex geometries. In fact, this will be necessary for studies of highly porous media, such as packed bed DBD reactors [51].

Our elliptic operators are embedded into the multigrid template that comes with the Chombo library, which relies on a series of C++ abstractions that implement the residual-correction form of multigrid [37, 38]. Our multifluid solver for the Poisson equation is an extension of the functionality discussed in [50] where we additionally incorporate (i) simultaneous handling of types of boundary conditions, and (ii) continuous and discontinuous permittivities. In the multigrid implementation, we use either red-black or multi-colored Gauss-Seidel relaxation. For the multi-fluid case, the boundary state is updated before each relaxation step. Our best results for the multifluid operator comes from using multi-colored (4/8 colors in 2D/3D) Gauss-Seidel relaxation. As with red-black ordering, odd colors only reach into even colors, which allows us to

reduce the number of exchange operations to 2 in 2D and 4 in 3D. We remark that for patches that intersect multi-fluid boundaries, it is necessary to perform relaxation on the patch twice (once for each phase), since both patches represent individual degrees of freedom. Currently, our load-balancing algorithms do not take this into account, and so we expect that a speedup of the Poisson equation is possible by more careful estimation of computational loads.

Finally, our convergence criteria for GMG rely on sufficient reduction of an initial residual. Let the discretized Helmholtz equation be given by $L(\phi) = \rho$ where $L(\phi)$ is the left hand side of Eq. (14) which also accounts for boundary conditions. The residual is

$$r = L\phi - \rho. \quad (28)$$

We exit multigrid iteration if the residual is sufficiently small, i.e. if

$$\|r\| < \lambda \|r_0\| \quad (29)$$

where r_0 is the computed residual with $\phi = 0$ and λ is some predefined tolerance.

3.4. Parabolic equations

In the presence of embedded boundaries we use implicit methods for parabolic equations (e.g. for diffusion advances). Embedded boundaries complicate the construction of explicit schemes for parabolic equations due to severe restrictions on the time step. For example, for the diffusion equation $\partial_t \phi = \nu \nabla \cdot (\nabla \phi)$, the stability constraint on the time step is

$$\Delta t \leq \frac{\Delta x^2 (\kappa_{\min})^{2/D}}{2D\nu}, \quad (30)$$

where D is the spatial dimension. For cut-cells, κ_{\min} can come arbitrarily close to 0, which leads to an unacceptable constraint. Although stabilization is possible by redistribution of mass in the vicinity of the embedded boundaries (see Sec. 3.5.2), this is not encouraged due to the $(\Delta x)^2$ scaling.

Our implicit method is as follows. For a parabolic equation

$$\frac{\partial \phi}{\partial t} + L(\phi) = \rho \quad (31)$$

where L is an elliptic operator, we use the implicit Runge-Kutta scheme by Twizell, Gumel, and Arigu (TGA) [52]. The TGA scheme discretizes Eq. (31) as

$$(I - \mu_1 L)(I - \mu_2 L)\phi^{n+1} = (I + \mu_3 L)\phi^n + (I + \mu_4 L)\rho^{n+1/2}, \quad (32)$$

whose solution requires the solution of two Helmholtz problems, one for each operator on the left hand side. The constants μ_i are

$$\mu_1 = \frac{a - \sqrt{a^2 - 4a + 2}}{2} \Delta t, \quad (33a)$$

$$\mu_2 = \frac{a + \sqrt{a^2 - 4a + 2}}{2} \Delta t, \quad (33b)$$

$$\mu_3 = (1 - a)\Delta t, \quad (33c)$$

$$\mu_4 = \left(\frac{1}{2} - a\right) \Delta t, \quad (33d)$$

$$a = 2 - \sqrt{2} - \lambda, \quad (33e)$$

where λ is the machine precision. Note that the centering of the source terms in Eq. (32) is $n + 1/2$ and $n + 1$, respectively. The details of the implementation of this scheme can be found in [48, 37, 38].

3.5. Convection-Diffusion-Reaction equations

Next, we discuss the spatial discretization of the CDR equation (2). Since the discretization of the diffusion operator was given in the preceding sections, it is sufficient to discuss the discretization of the advective term $\nabla \cdot (\mathbf{v}\phi)$.

3.5.1. Advective discretization

The advective term is solved by using finite volumes, such that

$$\int_V \nabla \cdot (\mathbf{v}\phi) dV = \sum_{f \in f(V)} (\mathbf{v}_f \cdot \mathbf{n}_f) \phi_f \alpha_f \Delta x^{D-1}, \quad (34)$$

where $f(V)$ denotes the set of faces that enclose the control volume V , and α_f is the face aperture fraction. For our purposes, the velocity \mathbf{v} is known so that it is sufficient to find a way to obtain the face-centered states ϕ_f . Once these are obtained, they are interpolated to the face centroids and the discrete divergence is then computed by summing the fluxes over all cell faces, as was done in Sec. 3.3

Currently, we use an unsplit Godunov method with van Leer limiting in order to obtain ϕ_f . The state at the boundary is then given by the solution to a Riemann problem with initial data given by the left and right slope-limited states; the solution to this particular problem (linear advection) is simply the upwind state.

3.5.2. Hybrid divergence

A problem with cut-cell grids is that naive discretization of conservation laws $\partial_t n + \nabla \cdot (\mathbf{v}\phi) = 0$ yields

$$\phi_i^{k+1} = \phi_i^k - \frac{\Delta t}{\kappa_i \Delta x^D} \oint (\mathbf{v}\phi) \cdot d\mathbf{A}_i, \quad (35)$$

which gives a CFL condition on the time step

$$\Delta t = \mathcal{O} \left(\frac{\Delta x}{\mathbf{v}_i^{\max} \kappa_i^{\frac{1}{D}}} \right). \quad (36)$$

This restriction is unacceptable since κ_i can be arbitrarily close to zero. Various remedies to this problem exist, such as merging of small cell volumes into larger control volumes, or the use of specialized stencils. The Chombo approach [38] is a simple one; we stabilize the advective discretization by computing a hybrid discretization which expands the range of influence of small control volumes to include a larger neighborhood around them. Let D_i^c be the conservative divergence of $\nabla \cdot \mathbf{F}$ for a cell i , where e.g. $\mathbf{F} = \mathbf{v}\phi$. For an irregular cell i we then also compute a non-conservative divergence D_i^{nc} as the average of the conservative divergence in a neighborhood \mathcal{N}_i of i :

$$D_i^{\text{nc}} = \frac{\sum_{j \in \mathcal{N}_i} \kappa_j D_j^c}{\sum_{j \in \mathcal{N}_i} \kappa_j}. \quad (37)$$

Note that there are other ways to compute D_i^{nc} , such as monotone extrapolation also to *covered faces*, in which D_i^{nc} is computed as if it was a regular cell. For streamer simulations, either approach is feasible.

We then use a hybridization of these two estimates,

$$D_i^{\text{H}} = \kappa_i D_i^c + (1 - \kappa_i) D_i^{\text{nc}}, \quad (38)$$

which cancels the small denominator in $\nabla \cdot \mathbf{F}$ and we obtain a stable method with a standard CFL condition. However, this method fails to conserve by the mass

$$\delta M_i = \kappa_i (1 - \kappa_i) (D_i^{\text{nc}} - D_i^c). \quad (39)$$

To enforce strict conservation, the excess mass is redistributed into neighboring cells $j \in \mathcal{N}_i$

$$\delta M_i = \sum_{j \in \mathcal{N}_i} \delta M_{j,i}, \quad (40)$$

where $\delta M_{i,j}$ is the distributed mass from i to j . For weighted redistribution we use

$$\delta M_{i,j} = \frac{\delta M_i \kappa_j W_j}{\sum_{k \in \mathcal{N}_i} \kappa_k W_k}, \quad (41)$$

where W_i are weights (for example mass). Then, for cells $j \in \mathcal{N}_i$ the final divergence is

$$D_j = D_j^H + \delta M_{i,j}. \quad (42)$$

For additional information on the redistribution process, see [38].

3.6. Time stepping

3.6.1. A second order fractional step scheme with implicit diffusion

Below, we provide a formulation of a second order accurate fractional step method with implicit diffusion for advancing $k \rightarrow (k+1)$ as follows:

1. Advance $n^* = n^k + \Delta t [S^k - \nabla \cdot (\mathbf{v}^k n^k)]$ and $\sigma^* = \sigma^k + \Delta t F_\sigma^k$
2. Compute the new electric field \mathbf{E}^* by solving the Poisson equation with the new space and surface charge densities ρ^* , σ^* .
3. Compute radiative transfer source terms $\eta^* = \eta(n^*, \mathbf{E}^*)$.
4. Obtain Ψ_γ^* by solving the RTE equations
5. Compute $S^* = S(\mathbf{E}^*, n^*, \Psi_\gamma^*, \nabla n^*)$ and $\mathbf{v}^* = \mathbf{v}(\mathbf{E}^*, n^*)$. Also recompute boundary fluxes on domains and internal boundaries.
6. Advance $n^\dagger = \frac{1}{2} (n^k + n^* + \Delta t [S^* - \nabla \cdot (\mathbf{v}^* n^*)])$ and $\sigma^{k+1} = \frac{1}{2} (\sigma^k + \sigma^* + \Delta t F_\sigma^*)$.
7. Obtain the new electric field \mathbf{E}^\dagger by solving the Poisson equation with ρ^\dagger and σ^{k+1} .
8. For diffusive species, obtain n^{k+1} with a implicit diffusion advance (see Eq. (32)). Otherwise, $n^{k+1} = n^\dagger$.
9. Obtain the final electric field \mathbf{E}^{k+1} by solving the Poisson equation with n^{k+1} and σ^{k+1} .
10. Compute radiative transfer source terms $\eta^* = \eta(n^{k+1}, \mathbf{E}^{k+1})$.
11. Obtain Ψ_γ^{k+1} by solving the RTE equations

In the above, steps 1 through 7 describe a consistent SSP Runge-Kutta method of order two (Heun's method). Step 8 and 9 describe an implicit diffusion advance and step 10 and 11 are the final updates for the radiative transfer equations. Note that charge injection into the domain is a part of the advective discretization, and the injected charge is therefore also redistributed in the vicinity of cells. The rationale for this design is that one avoids a possible division by κ when normalizing the injected charge by the volume fraction of a cut cell.

We have not yet explored the use of subcycling in time with AMR. Our experience with production runs of our solvers is that the finest AMR level typically contains a factor of ten more patches on the finest level than any other AMR level, and the benefits of subcycling in time are therefore unlikely to lead to drastic improvements in performance. Furthermore, streamer phenomena are physically unstable, and with the current operator splitting it is not clear if the coarse-grid equations can be advanced on their own with larger time steps without impacting numerical stability. In our approach, the coarse grid solutions are always synchronized with the finer levels. In other words, if a coarse grid is covered by a finer grid, the solution data on the coarse grid is the conservatively averaged version of the fine grid solution.

Restrictions on the time step are supplied through known criteria. For advection we restrict the time step on the CFL condition on the advection part by

$$\Delta t \leq \frac{\Delta x}{|\mathbf{v}|}. \quad (43)$$

In addition, we follow the physical time scale associated with the evolution of the electric field by

$$\epsilon_0 \frac{\partial \mathbf{E}}{\partial t} = \mathbf{J}, \quad (44)$$

where discretization by using the forward Euler rule yields the estimate

$$\Delta t < \text{Min} \left(\frac{\epsilon_0 |\mathbf{E}|}{|\mathbf{J}|} \right). \quad (45)$$

4. Verification tests

To verify that our code works as expected we now present standard test problems. All of our test problems use the full code framework for the verification studies. Internally, there are no differences between the numerical cases discussed here and the more complex cases discussed in Sec. 6.

4.1. The square wave advective problem

Our first problem, which verifies that our advective schemes are conservative and monotone and shows the amount of numerical diffusion that is involved, considers the pure advection of a single scalar with initial conditions taken as a square wave

$$n_0(x) = \phi_0 \left(\frac{x - x_0}{L} \right), \quad (46)$$

where L is the pulse width of the pulse

$$\phi_0(\xi) = \begin{cases} 1, & \text{if } |\xi| \leq 1, \\ 0, & \text{otherwise.} \end{cases} \quad (47)$$

Numerically, we consider a domain $(-1, -1) \times (1, 1)$ with $x_0 = 0.5$ and $L = 0.2$. The velocity of the wave is taken as $v = (v_x, v_y) = (1, 0)$ and we integrate for one time period, i.e. $T = L/v_x$. For this problem, where v_x is constant, the problem defined by Eq. (2) has solution $n(x, y, t) = n_0(x - v_x t)$, representing pure translation of the initial wave.

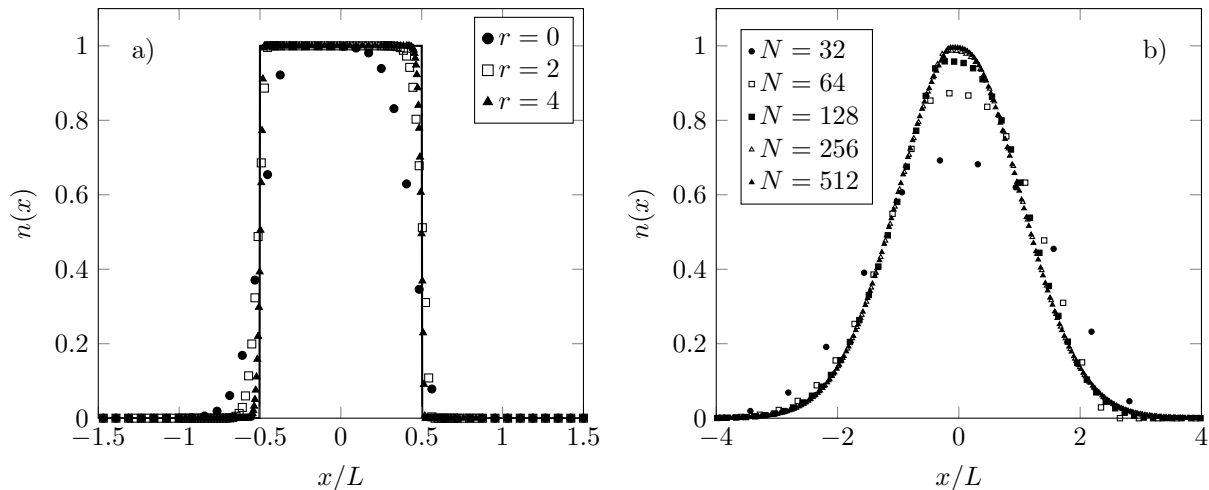


Figure 5: a) Comparison between numerically integrated solutions of the square wave problem with an analytic solution. The labels indicate the number of refinement levels used. b) Gaussian pulse advection problem. Labels represent the number of grid cells in each coordinate direction.

Figure 5a) shows the comparison between the analytic solution and the numerically computed solution for three different data sets. The data sets use a base mesh of $(128)^2$ and 0, 2, and 4 levels of refinement with a factor of two between levels. In the figure, the data has been translated to the center of the analytic solution, and are plotted versus the pulse width. All three cases work as expected, and the results also show the improvement that is expected with AMR: The numerical diffusion length of a discontinuity decreases with increasing resolution. For our scheme, the diffusion length is approximately 7 cells. All three simulation cases are also conservative: Integration of the data shown in Fig. 5 shows that the maximum mass fluctuations are within a factor of 10^{-8} of the initial mass. In addition, all three cases are monotone. That is, there are

no spurious oscillations or new extrema in the form of negative values or positive values exceeding the initial maxima.

Next, we perform a convergence test for our convection solver. We advect a Gaussian pulse

$$n_1(x) = \exp\left(-\frac{x^2}{2L^2}\right) \quad (48)$$

using five different resolutions. Since we want to test the convergence of our discrete scheme, we do not use AMR for these tests. The pulses have width $L = 0.1$, and the cases vary in resolution from $(32)^2$ to $(512)^2$ cells. Figure 5b) shows the advected pulses after one time period. We find that the numerical solutions approximate the analytic solution better when we increase the resolution. To investigate the rate of convergence of the advective solver, we compute the L_2 norm

$$L_2(n) = \sqrt{\frac{1}{\sum_i \kappa_i} \sum_i \kappa_i [n_i - n_e(\mathbf{x}_i)]^2}, \quad (49)$$

where n_i is the numerical solution and n_e is the exact solution $n(x, t) = n_1(x - v_x T)$. Figure 6 shows the resulting L_2 norm for the five cases, and we have also plotted a line $L_2 \propto \Delta x^2$. The figure shows second order convergence in the L_2 norm.

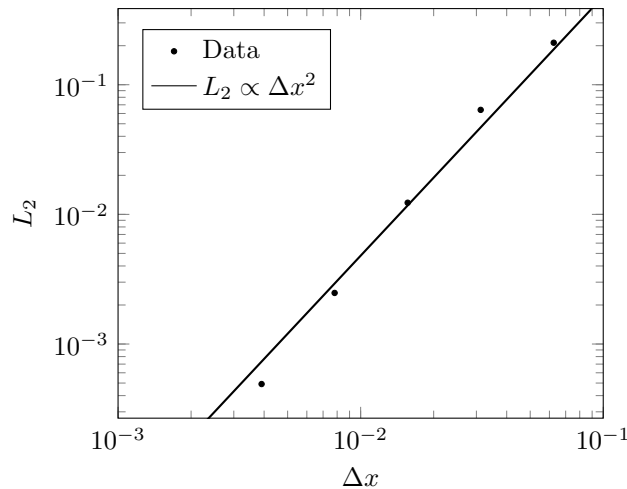


Figure 6: Convergence plot for the Gaussian wave problem.

4.2. The multi-fluid Poisson problem

Our next problem considers a multi-fluid Poisson problem on a two-dimensional coaxial cable geometry. The setup has translation symmetry along the z -coordinate, and rotational symmetry around the z -axis. This problem verifies that our multi-fluid operators produce reliable solutions. We do not explicitly verify our single-phase elliptic operators, which are used for the e.g. the diffusion advancements with Euler or TGA schemes, or for the Eddington equations. In fact, we consider verification of the multi-fluid operator to be sufficient for verification purposes. The reason for this is that, internally in our code, the multi-fluid solvers are high-level C++ abstractions that couple two single-phase elliptic solvers by using the matching condition (16). Thus, if the multi-fluid operator is correct, we can assert with some confidence that the single-fluid operators are also correct.

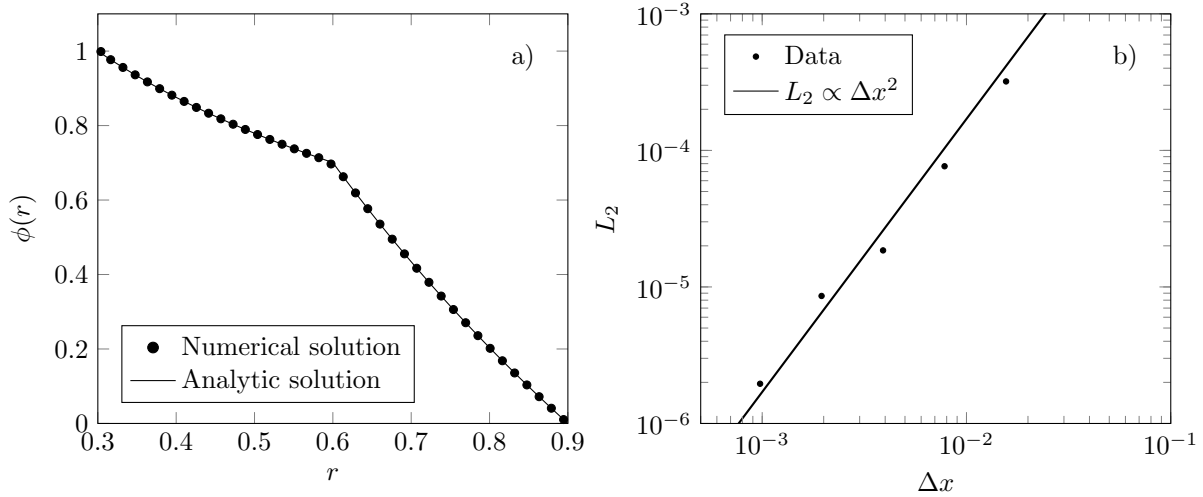


Figure 7: a) Comparison between numerically computed and analytic solutions to the multi-fluid coaxial Poisson problem. b) Convergence diagram.

In the following we do not consider space charge, in which case the Poisson problem for the case mentioned above reduces to solving

$$\frac{\partial}{\partial r} \left(r \epsilon \frac{\partial \phi}{\partial r} \right) = 0. \quad (50)$$

We consider a geometry consisting of three concentric shells; an outer electrode, an inner electrode, and a dielectric in between the two. On the outer electrode, we take the boundary condition $\phi = 0$ and on the inner electrode we take $\phi = 1$. On the dielectric between the two, the boundary condition is given by Eq. (16). For radii, we take the outer electrode radius $R_2 = 0.9$, the inner electrode radius is $R_0 = 0.3$, and the dielectric radius is $R_1 = 0.6$. The relative permittivity of the dielectric is $\epsilon = 4$ and the permittivity of the gas phase is taken to be ϵ_0 . We do not consider surface charge on the dielectric, in which case the exact solution to Eq. (50) is

$$\phi(r) = \begin{cases} \phi_0 + \frac{a}{\epsilon} \ln \left(\frac{r}{R_0} \right), & R_0 \leq r \leq R_1, \\ a \ln \left(\frac{r}{R_2} \right), & R_1 \leq r \leq R_2, \end{cases} \quad (51)$$

where

$$a = \phi_0 \left[\ln \left(\frac{R_1}{R_2} \right) - \frac{1}{\epsilon} \ln \left(\frac{R_1}{R_0} \right) \right]^{-1}. \quad (52)$$

We first compute the numerical solutions on a $(-1, -1) \times (1, 1)$ domain using $(256)^2$ cells. Note that while the solution here is effectively one-dimensional, our numerically computed solution is 2D. This also means that the data that we extract is off-axis by half a grid cell. Figure 7a) shows a comparison of Eq. (51) with the numerically computed solution. We find that our Poisson solver works as expected, essentially reproducing the analytic solution.

Next, we perform a convergence study on the same problem by computing numerical solutions on mesh sizes $(128)^2$, $(256)^2$, $(512)^2$, $(1024)^2$, and $(2048)^2$. Figure 7b) shows the computed norms (points) plotted together with a curve $L_2 \propto \Delta x^2$. The data points show second order convergence in the L_2 norm. Deviations from the curve are most likely due to the data-centering that we use in our software: In our output methods, we use the cell-centered potential from the gas side in the multifluid cells, whose center may fall on either side of the embedded boundary for different resolutions, and these errors can dominate.

4.3. The surface charge conservation problem

Our final test problems demonstrate charge conservation at insulating surfaces. This problem verifies that our implementation of charge transport on the gas-matter interface is physically correct.

Our first conservation problem considers the one-dimensional case as in 4.1 but with a dielectric slab filling the half space $x > 0.0$. Extrapolated outflow conditions are applied to the advected scalar. We use a base mesh of $(128)^2$ cells and use a single level of refinement with refinement factor 2.

Figure 8a) shows the evolution of the normalized volume and surface charge before and after the square wave pulse is absorbed by the insulator. The figure also shows the sum of charges. For this problem, where the pulse has a constant amplitude and velocity, one expects that the charging rate of the surface is close to linear. In the simulations, this is only achieved away from the rising and falling flanks, which is due to numerical diffusion which tends to smooth out the initial discontinuities (see Fig 5). Furthermore, we find excellent charge conservation on both the surface and in the volume (to within a factor 10^{-8}), and therefore conclude that our coupled gas-surface kinetics work as expected.

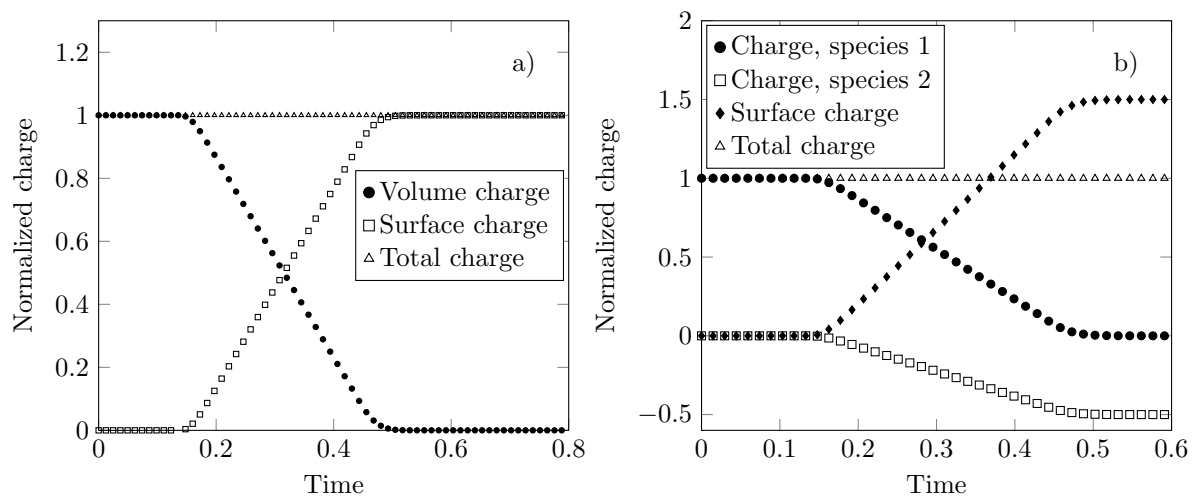


Figure 8: Charge conservation test on embedded boundaries. The figure shows the total volume charge, surface charge, and their sum. The y -axis has been normalized against the total initial charge.

Next, we perform a second surface conservation verification problem by coupling the boundary conditions of two advected species on the dielectric slab. We consider opposite signs of the two species' charge, i.e. the two species propagate in opposite directions. For the first species, its initial density is set to the square wave form above (recall Eq. (46)), and the other species' initial density is set to zero. On the dielectric boundary, we couple the boundary fluxes F_1 and F_2 of the two species by imposing

$$F_2 = -\beta \text{Max}(0, F_1) \quad (53)$$

where $\beta = 1/2$ is the emission coefficient of the second species. In our implementation, a positive flux is directed out of the cell so that F_2 provides an influx only if F_1 provides an outflux. This problem verifies that ion bombardment coupling works as expected. In this particular problem, we expect that the dielectric boundary reflects the incoming square pulse into an outgoing square wave with half the amplitude.

Figure 8b) shows the integrated charges for the double advection test. The figure shows that the total charge in the domain is conserved to a very good degree (open triangles). Before the initial square wave strikes the dielectric, the total charge is contained only in the incoming species (filled circles). After the incoming square wave is absorbed by the dielectric, the ejected charge back into the gas phase is a factor of $-1/2$ times the initial charge (open squares), which is what we expect with an emission coefficient of $1/2$

and opposite signs of the species' velocities. Charge conservation then implies that the surface must have absorbed a charge of 3/2 times the initial charge, which is what we observe (filled diamonds).

4.4. Model verification

We now verify that the combined algorithm is second order convergent in time by means of two-dimensional experiments. The convergence tests use Richardson extrapolation; i.e. we use a fine resolved solution as a replacement for the exact solution. We then compute the norm between the reference solution and a solution with coarser resolution. We consider a $(2\text{mm})^2$ domain with a blade electrode protruding 1 mm from the top domain edge. The electrode has a $100\ \mu\text{m}$ radius and is live with a voltage of 5 kV. We impose homogeneous Neumann boundary conditions for the Poisson equation on the left and right side edges, whereas the bottom edge is grounded. The plasma kinetics that we use for this test is summarized in Appendix A and solves for three species: electrons n_e , positive ions n_+ , and negative ions n_- , with only electrons being diffusive.

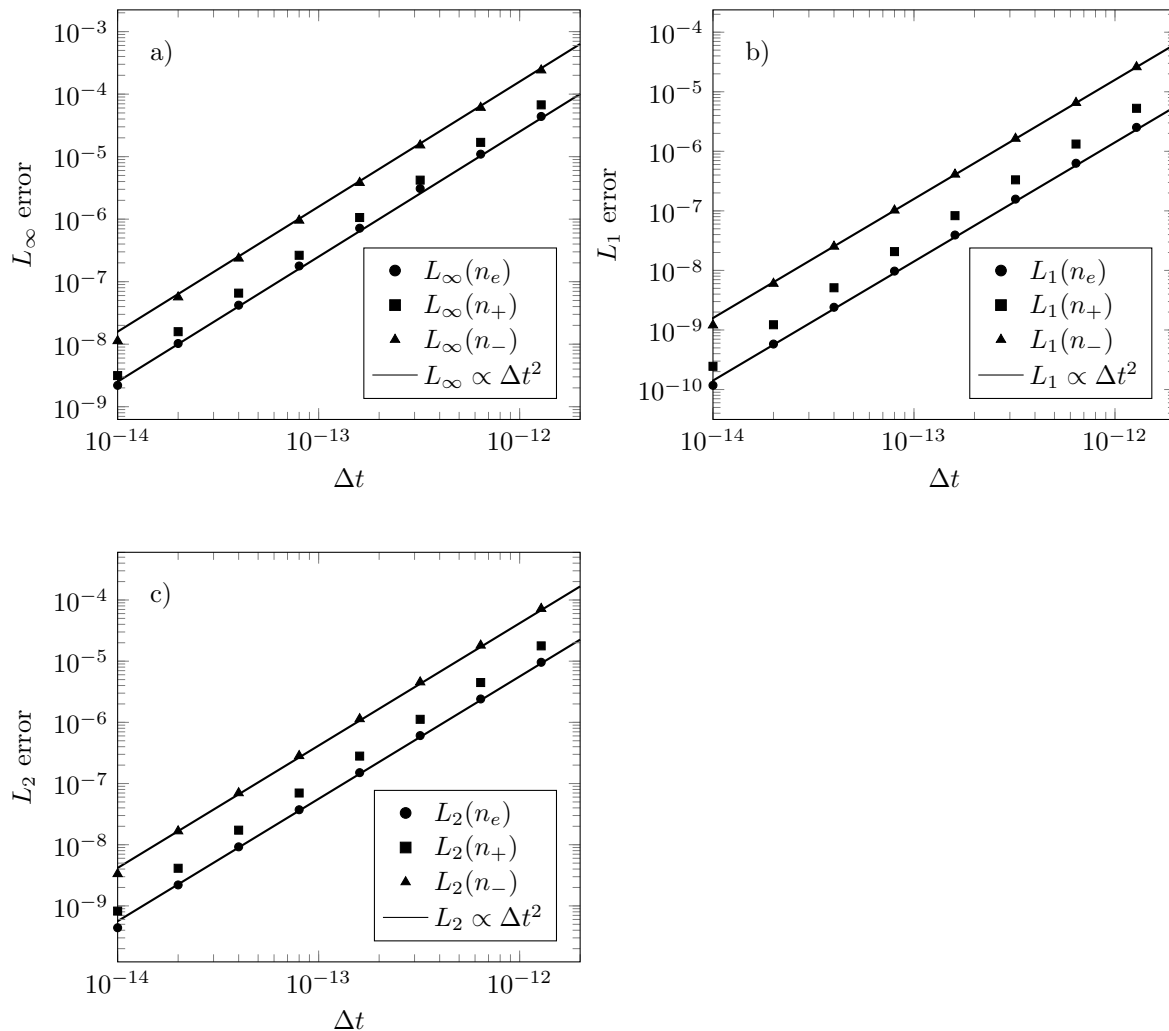


Figure 9: Temporal 2D convergence tests. The data points show the normalized error norms between the reference solution and additional solutions with coarser time steps.

For initial conditions, we consider an electron-ion seed

$$n_e = n_0 \exp\left(-\frac{(\mathbf{x} - \mathbf{x}_0)^2}{R^2}\right), \quad (54)$$

$$n_+ = n_e, \quad (55)$$

$$n_- = 0, \quad (56)$$

where $n_0 = 10^{18} \text{ m}^{-3}$, $R = 100 \mu\text{m}$, and \mathbf{x}_0 is the position of the electrode tip. We compute the solution error by comparison with a reference solution, which is taken to be the solution with a step size of $\Delta t_{\text{ref}} = 5 \times 10^{-15} \text{ s}$ which is advanced for 1024 time steps on a $(2048)^2$ cell grid. This is done for time steps $\Delta t_n = 2^n \Delta t_{\text{ref}}$, $n = 1, 2, \dots, 8$ which ranges up to $256 \Delta t_{\text{ref}}$. The relative L_∞ , L_1 , and L_2 error norms are then computed for all three species. The results of these computations are presented in Fig. 9 and are plotted together with the line $L_i \propto \Delta t^2$. We find second order convergence in the L_∞ , L_1 , and L_2 norms.

Next, we present spatial convergence tests in 2D using the same geometry as above, and in 3D using a domain of $(1\text{mm})^3$ with a rod electrode protruding from the top domain edge. In 2D and 3D, the solution is advanced 16 time steps using a step size of $\Delta t_{\text{ref}} = 10^{-14} \text{ s}$. In 2D, the reference solution is taken to be the solution on a $(2048)^2$, and in 3D the reference solution uses a $(512)^3$ grid. The reference solution is denoted $\phi^{\Delta x}$ in both cases, and we compare this solution with coarser solutions $\phi^{2\Delta x}$ and $\phi^{4\Delta x}$. On the coarser grids, the errors are computed by comparison with a coarsening of $\phi^{\Delta x}$, i.e. the error is computed as

$$E_{2\Delta x} = \phi^{2\Delta x} - A(\phi^{\Delta x}), \quad (57)$$

where $A(\phi^{\Delta x})$ indicates a coarsening of $\phi^{\Delta x}$ with a coarsening factor of 2. Likewise, $E_{4\Delta x} = \phi^{4\Delta x} - A(A(\phi^{\Delta x}))$ indicates the error on the grid with resolution $4\Delta x$. The convergence rates are computed as

$$p = \frac{\log\left(\frac{E_{4\Delta x}}{E_{2\Delta x}}\right)}{\log(2)}, \quad (58)$$

and results of these computations are given in Table 1 and Table 2. We find second order convergence in all three norms in both 2D and 3D.

Norm	Variable	$E_{2\Delta x}$	$E_{4\Delta x}$	p
L_∞	n_e	2.3399×10^{-5}	1.1063×10^{-4}	2.24
L_∞	n_+	2.1426×10^{-5}	1.0646×10^{-4}	2.31
L_∞	n_-	5.7657×10^{-10}	2.8769×10^{-9}	2.32
L_1	n_e	2.7712×10^{-7}	1.3826×10^{-6}	2.31
L_1	n_+	2.7711×10^{-7}	1.3827×10^{-6}	2.31
L_1	n_-	7.3172×10^{-12}	3.6902×10^{-11}	2.33
L_2	n_e	1.4245×10^{-6}	7.0809×10^{-6}	2.31
L_2	n_+	1.4240×10^{-6}	7.0804×10^{-6}	2.31
L_2	n_-	3.6675×10^{-11}	1.8744×10^{-10}	2.35

Table 1: 2D spatial convergence tests based on a $(2048)^2$ grid with $\Delta x = 0.976 \mu\text{m}$.

5. Weak scalability

We now present results on weak scalability of the combined algorithm code on uniform grids. Weak scaling involves increasing the number of processor cores by the same factor as the geometry refinement. In 3D, increasing the resolution by a factor of 2 requires 8 times as many processor cores. These tests are performed by evolving the discharge model defined in Appendix A for 10 time steps and then taking the

Norm	Variable	$E_{2\Delta x}$	$E_{4\Delta x}$	p
L_∞	n_e	9.7570×10^{-4}	4.7717×10^{-3}	2.29
L_∞	n_+	9.7156×10^{-4}	4.7153×10^{-3}	2.28
L_∞	n_-	1.3212×10^{-8}	6.7158×10^{-8}	2.34
L_1	n_e	7.5879×10^{-7}	3.7863×10^{-6}	2.32
L_1	n_+	7.5695×10^{-7}	3.7847×10^{-6}	2.32
L_1	n_-	1.9289×10^{-11}	9.7338×10^{-11}	2.33
L_2	n_e	9.8833×10^{-6}	4.8467×10^{-5}	2.29
L_2	n_+	9.8760×10^{-6}	4.8471×10^{-5}	2.30
L_2	n_-	2.4219×10^{-10}	1.2474×10^{-9}	2.36

Table 2: 3D spatial convergence tests based on a $(512)^2$ grid with $\Delta x = 3.906\mu\text{m}$.

average of the execution time. We consider the three geometries in Fig. 10 with initial conditions on the charged species taken as $n_e = n_+ = 10^{10}\text{m}^3$, $n_- = 0$, whereas the boundary conditions on the Poisson equation are: Live Dirichlet on the electrodes and top domain faces; homogeneous Dirichlet on the bottom domain face, and homogeneous Neumann boundary conditions elsewhere. The resolution on each geometry is then successively doubled and the core count is increased by a factor of 8. The computational domains range in size from $128 \times 128 \times 256$ to $1024 \times 1024 \times 1024$ and are decomposed with 32^3 grid patches. We distribute the patches uniformly to all cores (i.e. the patch volumes are used as proxies for load balancing). These tests are run at core counts up to 8192 on Fram, which is a Lenovo NeXtScale nx360M5 with dual 16-core Intel Xeon E5-2683v4 (2.1GHz) processors per node and up to 896 cores on SuperMUC Phase 2, which is a Lenovo NeXtScale nx360M5 WCT with dual 14-core Intel Xeon E5-2697v3 (2.6GHz) processors per node. On Fram, we have 4 patches per CPU core whereas on SuperMUC Phase 2 we have at maximum 3 patches per CPU core.

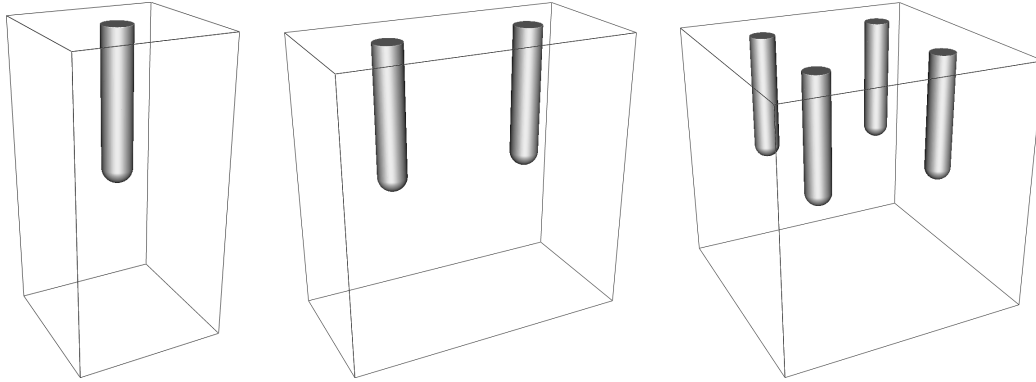


Figure 10: Geometries used for weak scalability study for concurrencies C . Left: $128 \times 128 \times 256$ grid. Used for $C = 256, 2048$. Middle: $256 \times 128 \times 256$ grid. Used for $C = 512, 4096$. Right: $256 \times 256 \times 256$ grid. Used for $C = 1024, 8192$. All geometries are decomposed with 32^3 patches.

The results of the scalability studies on Fram and SuperMUC are shown in Fig. 11 and show reasonable scalability. On Fram, we observe a parallel efficiency above 70% from 256 to 8192 cores and we observe about the same efficiency on SuperMUC Phase 2 from 112 cores to 896 cores. The scaling tests were not run further due to availability of each machine, but further scaling tests are likely to be performed in the future.

6. Numerical examples

In this section we provide basic numerical examples. Note that the examples are intended as demonstrations of capabilities, and that some of these examples may be recalculated or analyzed more rigorously in the

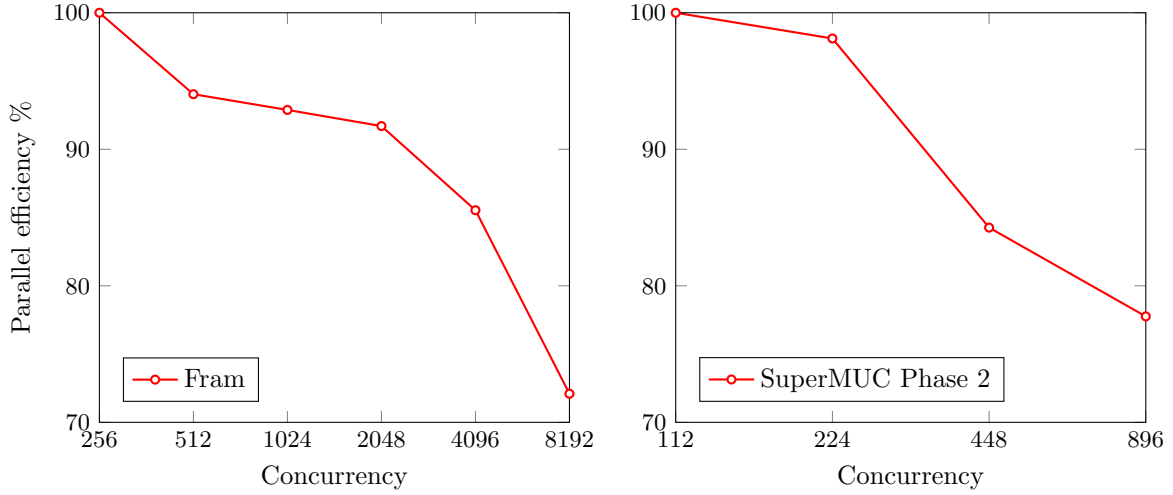


Figure 11: Weak scaling on Lenovo NeXtScale nx360M5 (Fram) and SuperMUC Phase 2. The vertical axis is the parallel efficiency and the horizontal axis is the number of processor cores. For reference, the average time per step was around 14 seconds at 256 cores on Fram and about 10 seconds at 112 cores on SuperMUC Phase 2.

future. The examples below use refinement criteria based on resolution of space charge layers and impact ionization regions. At each regrid stage, we compute a tracer field

$$T = \frac{|\mathbf{E}|}{\max(|\mathbf{E}|)}. \quad (59)$$

We refine if

$$\frac{1}{T} |\nabla T| \Delta x > \epsilon_1 \vee T > \epsilon_2. \quad (60)$$

where ϵ_i are tagging thresholds. The first of these criteria is responsible for resolving gradients in the electric field, which typically arise due to space charge formation. The second test resolves the streamer head.

Likewise, we coarsen a cell if

$$\frac{1}{T} |\nabla T| \Delta x < \epsilon_3 \wedge T < \epsilon_4. \quad (61)$$

For the simulations below, we take $\epsilon_1 = \epsilon_3 = 0.1$, $\epsilon_2 = 0.8$, and $\epsilon_4 = 0.2$.

6.1. Creeping streamers along rough surfaces

We first consider the propagation of positive streamers along a rough surface. The computational domain is $2\text{cm} \times 2\text{cm}$, with a blade electrode protruding 1cm from the top plane; the thickness and curvature of the electrode is $250\ \mu\text{m}$. The dielectric is represented by a planar level-set surface that sits about 1mm away from the electrode, and whose surface has been displaced with Perlin noise [53] with an amplitude of roughly $100\ \mu\text{m}$. The relative permittivity of the dielectric is 4. For initial conditions we take $n_e = n_+ = 10^{10}\ \text{m}^3$.

In this example we use a base mesh of 128^2 with 4 AMR levels with a refinement factor of 4 between levels, yielding a base resolution of $312.5\ \mu\text{m}$ and a finest level resolution of $1.22\ \mu\text{m}$. Computationally, the sweet spot for this simulation is found by using a minimum box size of 16 and a maximum box size of 32. The integration is done with a CFL number of 0.8. For boundary conditions on the Poisson equation, we use homogeneous Neumann conditions on the left and right side edges. We impose a voltage of 30 kV on the top domain edge and on the electrode, whereas the bottom domain edge is grounded.

Figure 12 shows the initial field for this configuration. Arrangements such as these (often called "triple points") are particularly problematic in high-voltage engineering where streamer flashover is unwanted.

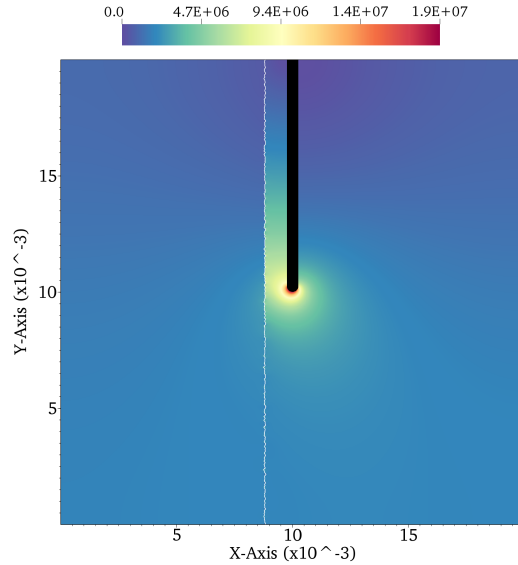


Figure 12: Initial field distribution for the DBD arrangement. The color coded data shows the electric field magnitude.

The degraded dielectric performance occurs because the field lines that cross into the dielectric experience enhanced voltage drops in the gas-phase due polarization of the dielectric. In turn, this can lead to rapid streamer inception in the dielectric-gas gap.

Figure 13 shows snapshots of the temporal evolution of the electric field, plasma and charge densities during the first few nanoseconds. In general, the streamer follows the topology of the surface and there is no plasma detachment into the gas phase (at least for the surface roughness considered here). We observe a height-of-flight of the positive streamer above the surface and find that the sheath region $l < 30 \mu\text{m}$ from the surface contains fewer electrons and is incompletely shielded by space and surface charge effects. The topological features that we observe is also observed in planar simulations by others [31, 54, 9, 27, 55, 56], also with PIC codes [57], and the plasma sheath thickness that we observe is consistent with analytical estimations [55].

Through the simulation, we always observe electron impact ionization in the plasma sheath closest to the dielectric, although the ionization rate is modest when compared to the streamer head. Inside the sheath the plasma dynamics occurs in the following way. Emission of secondary electrons from the dielectric appears through ion bombardment or photoemission. In turn, these electrons ionize the gas in the sheath; propagate through the sheath boundary and into the channel where they essentially become immobile due to field screening of a positively charged space charge layer that floats some tens of microns above the surface. For longer integration times, the residual positive ions in the sheath drift toward the dielectric, which increases the surface charge density and decreases space charge effects closest to the dielectric. Likewise, negative ions drift away from the surface. As a result, the sheath remains active until sufficient charge deposition onto the surface has occurred. An order of magnitude estimation of the sheath relaxation time can be provided by considering the drift time of ions over the thickness of the sheath. I.e. $t = L/(\mu_{\text{ion}}E) \approx 15 \text{ ns}$, where $\mu_{\text{ion}} \approx 2 \times 10^{-4} \text{ m}^2/(\text{Vs})$ is the ion mobility. In practice, this time is a maximum estimate since the choice secondary emission coefficients strongly influence the number of electron-ion pairs that are generated in the sheath, and therefore also the sheath relaxation time.

Furthermore, we find that the plasma essentially follows the topology of the surface throughout the simulation. Figure 14 shows the streamer after a time 5.25 ns; the discharge has propagated about 7 mm in this time, which yields an average velocity of roughly 1.3 mm/ns. In particular, we observe that the plasma density has local maxima at the protrusions on the dielectric, reaching densities over 10^{21} m^{-3} . Correspondingly, the space charge density is lower in these regions (see Fig. 14).

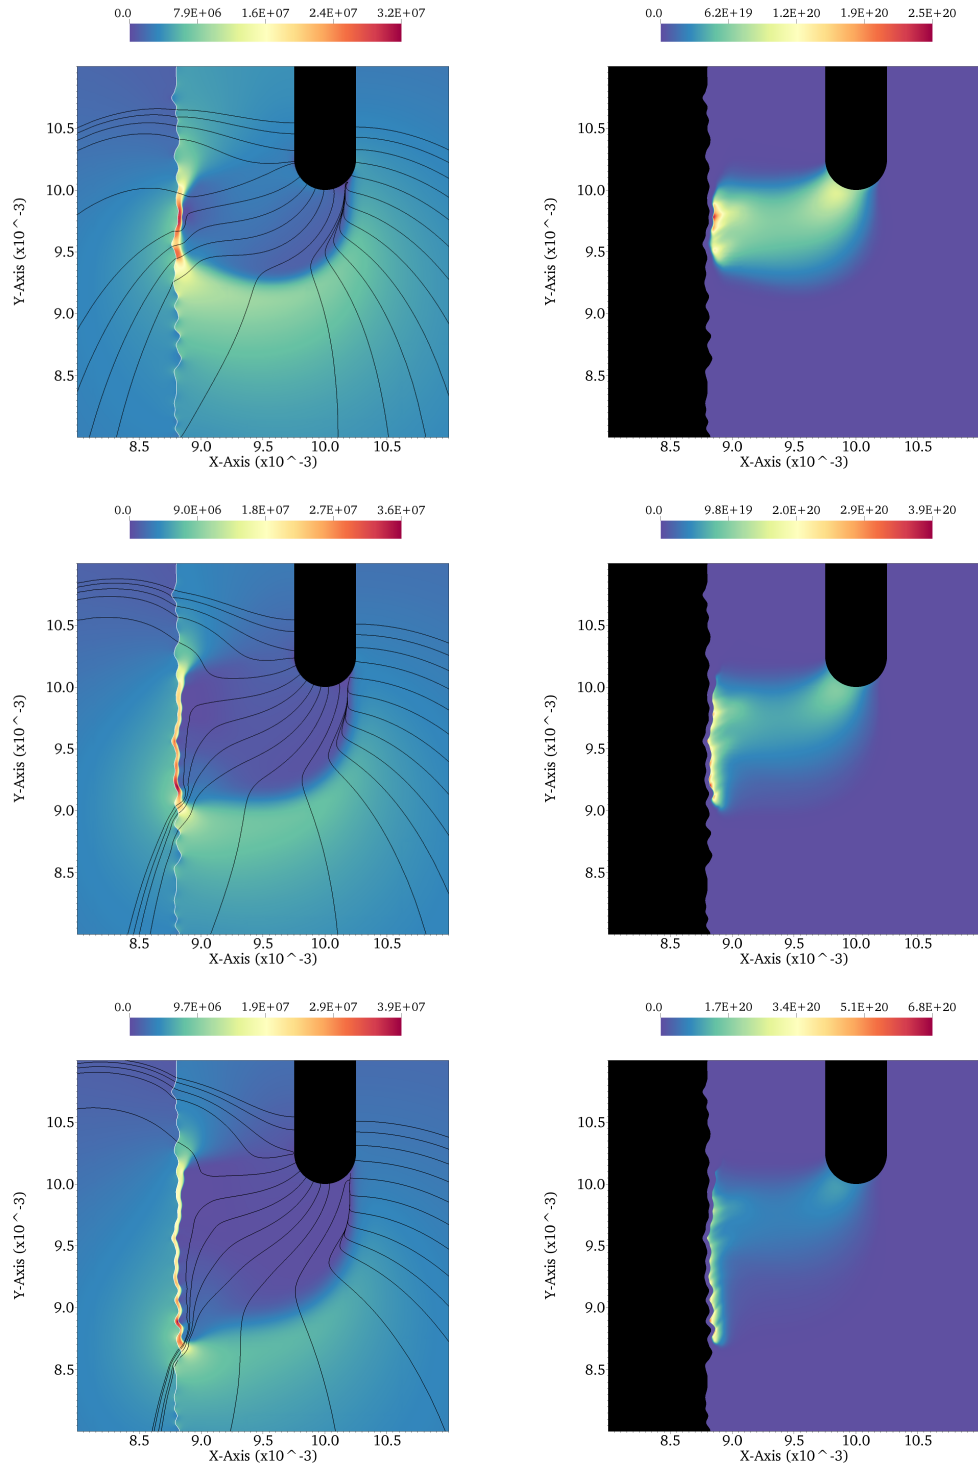


Figure 13: Snapshots of the evolution after inception. The figure shows the field magnitude (left column) and plasma density (right column). Top row: Snapshot at 1 ns. Middle row: Snapshot at 1.25 ns. Bottom row: Snapshot at 1.5 ns.

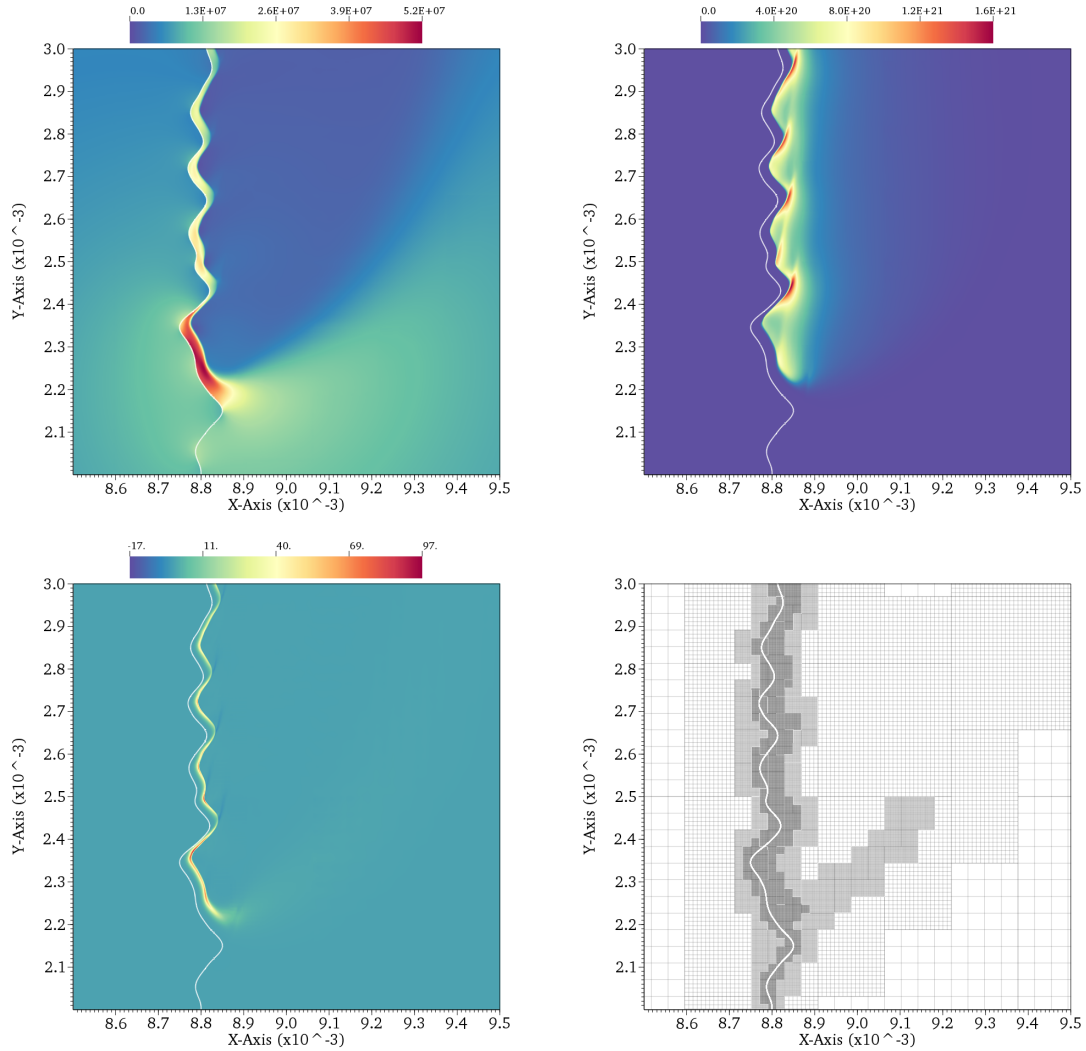


Figure 14: Snapshots of the streamer after 5.25 ns. Top left: Electric field magnitude. Top right: Electron density: Bottom left: Space charge density: Bottom right: Mesh

6.2. 3D simulations in pin-plane gaps

In our next example we consider inception and propagation of a positive streamer in a pin-to-plane geometry. The computational domain is a $(2\text{ cm})^3$ domain with a cylindrical needle electrode with a $500\ \mu\text{m}$ radius protruding 1 cm down from the upper z -plane. We impose homogeneous Neumann conditions on the potential on the domain faces in the xz and yz planes. The lower xy plane is grounded, and we apply a live potential of 15 kV to the needle and upper xy plane. For initial conditions we take $n_e = n_+ = 10^{10}\ \text{m}^3$. Both the geometry and initial conditions are rotationally symmetric around the needle axis for this case, and this simulation could be run in a 2D cylindrically symmetric configuration, as is done by many other groups. In three dimensions, no currently published results are available for comparison. However, inception and propagation of streamers from a $(1\text{ mm})^3$ cubic volume was considered recently [18], and an equivalent setup has also been considered in the past [20]. The difference between the simulations in [18, 20] and ours lies primarily in the volume that is simulated; we consider a volume that is several thousand times bigger. In two-dimensional cylindrically symmetric coordinates, equivalent simulations are too numerous to explicitly mention here. However, the intention of this example is not a comparison with the two-dimensional world,

but rather a demonstration of the use of AMR with internal boundaries for moderately large configurations that are of experimental interest, using only a moderate amount of CPU cores. Certainly, deviations in the initial conditions can lead to non-symmetric solutions, as has been for example demonstrated in [18], where plasma spots were used to provoke branching of the streamer. For geometric configurations that lack symmetry, a full 3D description is obviously required.

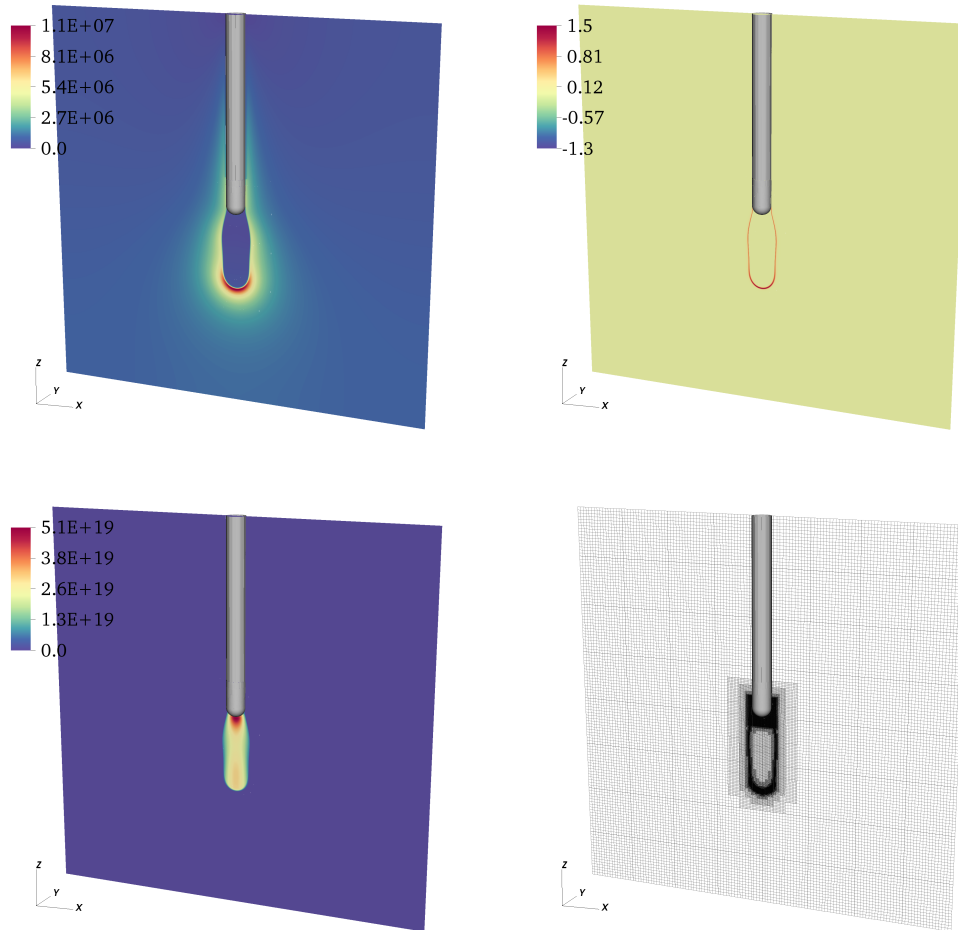


Figure 15: Simulation state after 10.8 ns. Top left: Field magnitude. Top right: Space charge. Bottom left: Electron density. Bottom right: Mesh.

For this application we use a base mesh of 128^3 cells and use 5 levels of refinement with a factor 2 between levels. The effective domain is $(4096)^3$, and the finest level resolution is roughly $5 \mu\text{m}$. For this application we refine the region around the needle tip, which is where we expect streamer inception to occur. In the bulk gas, we use the refinement and coarsening criteria discussed above. The simulation was run with a CFL number of 0.5 throughout the simulation, yielding a time step of roughly 5 ps. The strong scaling limit for this simulation yields an execution time per time step of approximately 6 seconds, depending on the simulation state. The initial avalanche phase is faster to compute since there are no space charge effects, leaving the potential largely unaffected. Once a streamer starts propagating, the execution time per time step increases. We remark that our temporal scheme requires 11 elliptic solves per time step. In addition, unlike [14, 15, 18, 20], every species is convected so that the slope-limited upwind states that are required for the advection advance must be determined at every Runge-Kutta stage for both electrons, positive ions, and negative ions. The reason for this is that we require ion mobility for modeling of appropriate physical

effects near electrodes and dielectrics, examples being surface charging of insulating surfaces or dynamics of cathode sheaths. However, the numerical cost of advancing convection equations is generally less than for elliptic equations. Towards the end of the simulation at $t = 20$ ns, the mesh is composed of roughly 80 million cells.

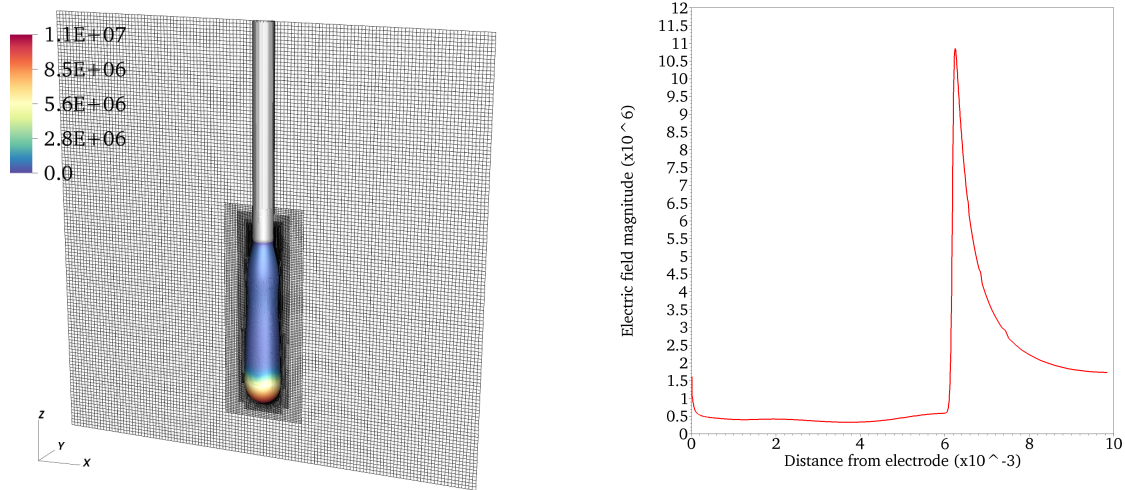


Figure 16: Left: Three-dimensional evolution of a cylindrically symmetric streamer in a pin-plane gap after a time 17.7 ns. The surface shows the isosurface of the plasma density at $n_e = 10^{18} \text{ m}^{-3}$ and the color coding shows the electric field magnitude. Right: Field magnitude on the needle axis after a time of 15.8 ns.

Figure 15 shows the simulation state after roughly 10.8 ns. The various panels show the electric field magnitude, space charge layer, electron density, and mesh; all quantities have been sliced through the xz plane that passes through the needle tip. We observe generic features that are associated with propagation of positive streamers, such as a positive space charge layer and screening of the channel interior. More careful evaluation of the data sets shows that the streamer diameter, taken as the width of the space charge layer (i.e. the electrodynamic diameter) at its widest point, measures roughly 1.5 mm. The radiative diameter, however, is smaller and measures less than 1 mm. Furthermore, the streamer has propagated roughly 5 mm within the first 10.8 ns, yielding an average velocity of 0.5 mm/ns. Finally, note that our simplified refinement and coarsening criteria resolve the space charge layer without over-resolving the channel volume. In fact, our experience is that appropriate refinement and coarsening criteria are essential for optimum use of computer resources, and can lead to an order of magnitude reductions in computation time or resources. However, we emphasize the need for caution in the grid coarsening stage. Under-resolution of the streamer head, or the region around it, can lead to solution errors that amplify during the propagation process, leading to unphysical solutions (such as numerical branching, for example).

Finally, we remark that the simulation data here is consistent with experiments and two-dimensional rotationally symmetric simulations performed by others. We observe a mean velocity of roughly 0.5 mm/ns and an optical diameter less than 1 mm (the diameter defined by the space charge layer is larger, up to 1.5 mm). Pancheshnyi et al. [58] has examined the development of positive streamers in slightly larger gaps, and the authors observed a mean velocity of roughly 0.5 mm/ns and optical diameters of roughly 0.5 mm, which aligns fairly well with our computational results. However, we remark that this comparison is only indicative. Indeed, streamer velocities and diameters change as a function of the applied voltage, and a broader sampling of simulations and experiments is generally necessary for validating the computational results. Finally, the field magnitude in the channel is not constant, it has an average value of roughly 0.5 kV/mm (see Fig. 16), which corresponds to the streamer stability field computed by others [59, 60, 61, 62], and is also the value that is used by engineering models [63].

6.3. 3D simulations in complex geometries

Our final simulation example considers a mock up model of a mechanical switch where a rotating mechanical shaft protrudes through a mechanical bushing. The domain is a $(4\text{ cm})^3$ domain with a single electrode and an insulating dielectric shaft, as shown in Fig. 18. The electrode is represented by a hollow cylinder that protrudes 1 cm from the live wall and has inner and outer radii of 6 mm and 1 cm, respectively. The electrode corners have been rounded with a rounding radius of $500\ \mu\text{m}$. The dielectric is a 4 cm long rod with a hexagonal cross section and relative permittivity of 4. The cross-sectional diameter of the insulator is 1 cm, and the corners have been rounded with a $500\ \mu\text{m}$ rounding radius. For fast parametrization, the geometry was generated by means of constructive solid geometry.

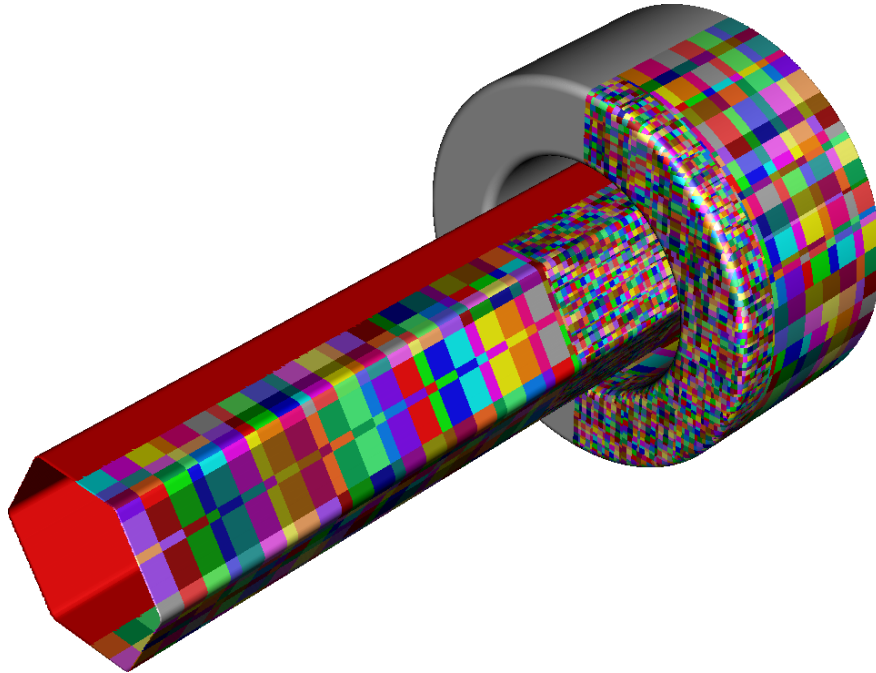


Figure 17: Initial computational patches for the insulating shaft geometry. Each patch represents a unit of computational work and has minimum size $(16)^3$ and maximum size $(32)^3$.

For this simulation we have decided to reduce the gas pressure. Current hardware limitations (4096 cores) prevent us from simulating this geometry at atmospheric pressure where the simulation would probably require several billion cells. Because of those limitations, the gas pressure used for this simulation is $p = 0.5\text{ atm}$, and we only simulate the inception and initial propagation phases.

We use a base mesh of $(128)^3$ cells and include five levels of refinement. The refinement factor between each level is 2, which yields a base resolution of $312.5\ \mu\text{m}$ and a finest level resolution of $10\ \mu\text{m}$. We use a reduced EB mesh for this simulation, initially only refining the EB region $z \in [2.5\text{ cm}, 3.2\text{ cm}]$ down to the finest AMR level; the remaining part of the EB is refined three times down to a resolution of $39\ \mu\text{m}$. The simulation is run with a maximum patch size of $(32)^3$ and a minimum patch size of $(16)^3$, and the initial mesh contains around 300 million cells for a total of roughly 2.1 billion unknowns (potential, three ion densities, and three photon densities). A uniform mesh would require almost 70 billion cells. For the types of applications that we attempt to simulate here, grids based on unstructured meshes or SAMR are, from both run-time and memory consumption point of views, the only realistic alternatives today. Figure 17 shows the distribution of the initial mesh on the boundary. In this figure, each colored square represents a patch of minimum size $(16)^3$ and maximum size $(32)^3$.

For boundary conditions on the Poisson equation, we apply a positive step potential of 45 kV on the electrode and on the wall it protrudes from, the opposite domain face is grounded, and homogeneous Neumann

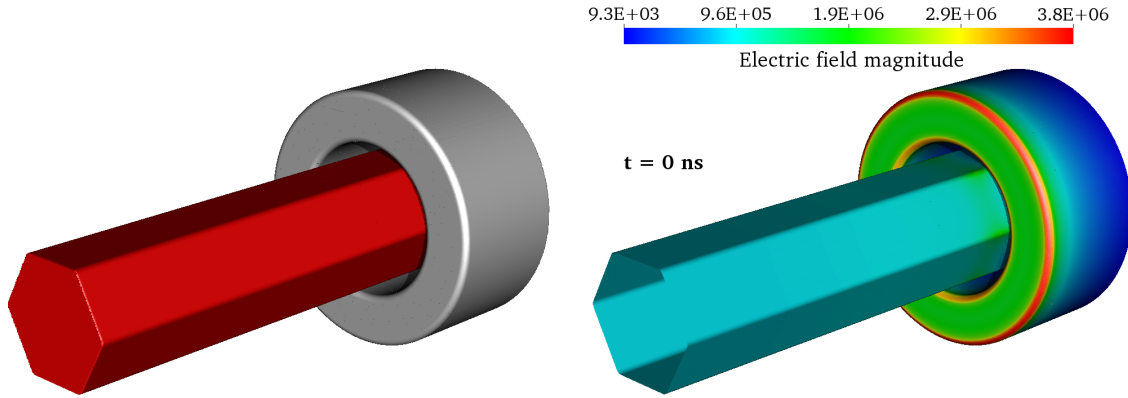


Figure 18: Left: Mechanical shaft simulation geometry. Right: Initial field stress $|\mathbf{E}|$ at the boundaries at $t = 0$.

conditions are used on the side walls. The initial charge on the rod is set to zero whereas the ion densities are $n_e = n_+ = 10^{10} \text{ m}^{-3}$, $n_- = 0$. Figure 18 shows the initial field stress with these boundary conditions. The maximum field stress is found on the outer edge of the electrode and equals roughly 5.6 kV/mm , whereas the field stress on the inner edge of the electrode equals roughly 3.5 kV/mm . For distances greater than 1 cm away from the electrode and towards the ground plane, the field is otherwise comparatively homogeneous. On the dielectric boundary, we find slightly enhanced field stresses on the corners of the dielectric shaft. Enhancement in these regions are due to a reduced distance between the electrode and insulator. For field lines that penetrate into these regions, polarization of the dielectric enhances the voltage drop in the gas phase, leading to an increased electric field in these positions.

Figure 19 shows the spatial evolution of the electric field sliced through two symmetry planes in the domain. The left and right columns in this figure show the electric field magnitude sliced through planes that are rotated 60 degrees with respect to one another. One of these planes intersect the "flat" side of the insulator and the other intersects one of the corners. Simulation times are indicated in each frame. We find that after streamer inception at $t \sim 3.5 \text{ ns}$ the maximum electric field magnitude is roughly 4.9 kV/mm for the streamer that initiates from the outer electrode edge. At this point, there are no visible space charge effects on the inner edge, nor do we find significant field screening on that edge (see Fig. 20). After roughly $t \sim 7.27 \text{ ns}$, a toroidal streamer has disconnected from the outer electrode edge. Streamers have also initiated from the inner electrode edge, but only on the six locations on the electrode that are closest to the insulating surface (Fig. 19). The maximum field strength are found in the head of the toroidal streamer, with a field magnitude of $\sim 5.0 \text{ kV/mm}$. After $t \sim 11.37$, the maximum field stress is found in the head of the streamers that propagate towards the insulator, with field strengths reaching 4.6 kV/mm , see bottom row in Fig. 19. These streamers have not yet reached the surface of the insulator.

Figure 20 shows the field magnitude and surface charge density on the boundaries for the same times as in Fig. 19. At the time of inception (3.5 ns), surface charging of the insulator is insignificant, indicating that non-local photoemission of seed electrons from the insulator leads to negligible charging in the inception stage. For example, Fig. 20 shows that the surface charge density reaches a value of $1.5 \times 10^{-14} \text{ C/m}^2$ after 3.5 ns . For comparison, the initial surface charge density on the electrode is $\sigma = \epsilon_0 E \sim 3.4 \times 10^{-5} \text{ C/m}^2$. Note that the negative surface charge densities that are observed for $t = 3.5 \text{ ns}$ (top right panel in Fig. 20) are due to electron impact of the initial preionization that was assigned in order to start the simulation. The inflection point where the surface charge goes from positive to negative (approximately 1 cm from the electrode) is due to electrostatic field effects. The normal field switches sign at this position such that $\mathbf{E} \cdot \mathbf{n}$ is negative for distances $\lesssim 1 \text{ cm}$ from the electrode, and positive otherwise. After roughly 7.27 ns , inception of positive streamer has also occurred on the inner electrode edge. In this region, the streamer is not rotationally symmetric but instead initiates on the positions on the electrode that are closest to the six outer corners on the insulator. This can be seen in e.g. Fig. 20 by appearance of a field-screened spot on the inner electrode edge after 7.27 ns . After 11.37 ns , the maximum field stress on the boundary has been transferred to the

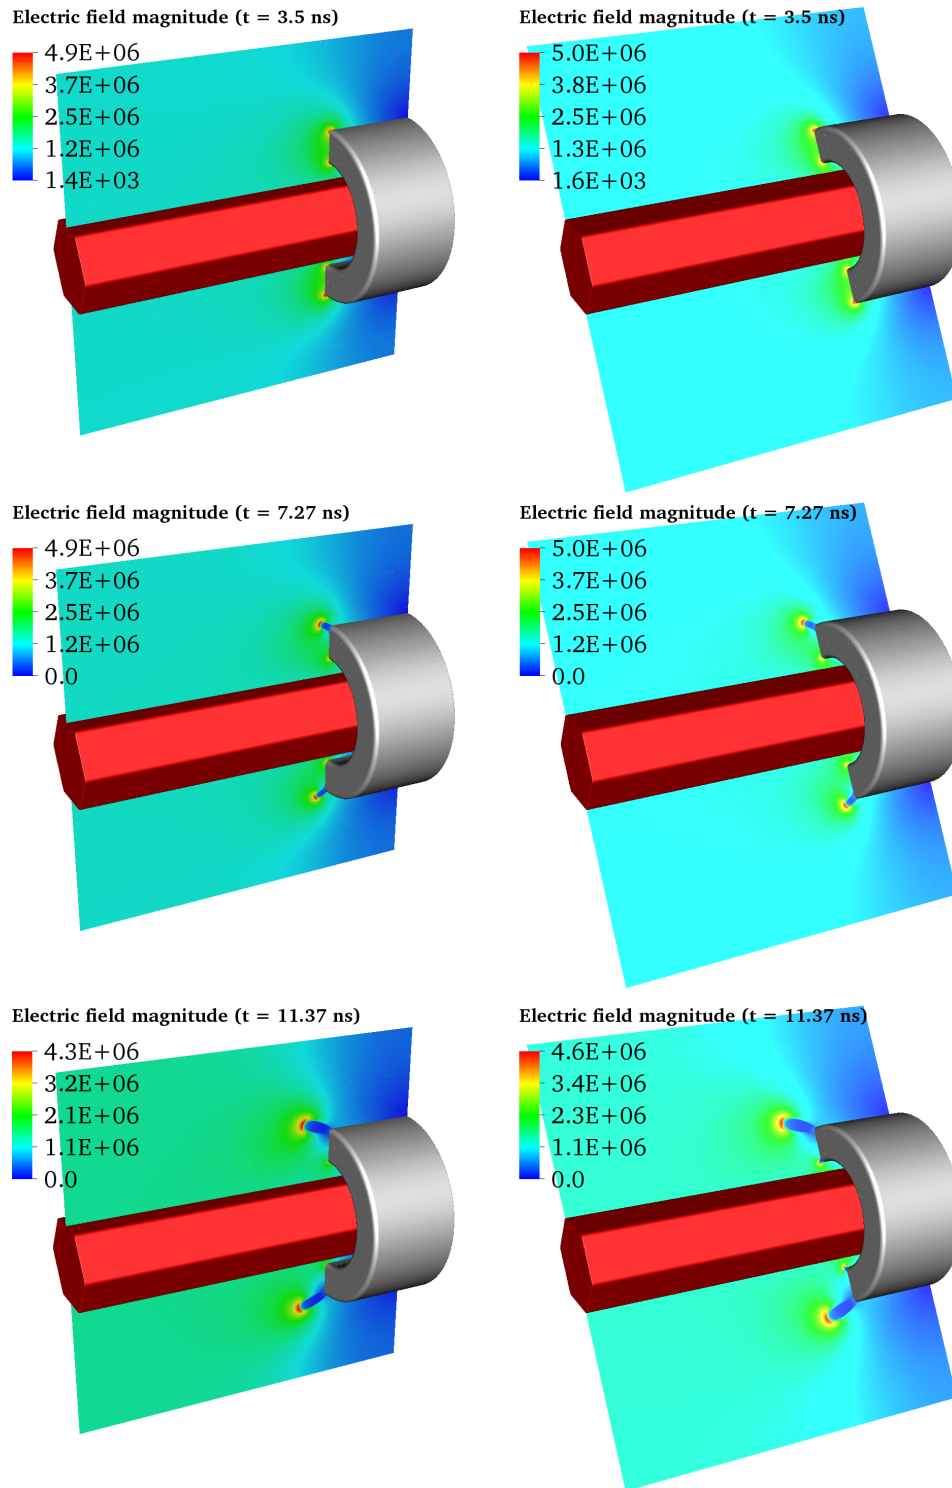


Figure 19: Snapshots of the electric field magnitude in the gas volume. The left column shows the electric field magnitude sliced through the yz -plane, and the right column shows the electric field magnitude sliced through one of the corners of the insulator shaft. Times are indicated in each plot.

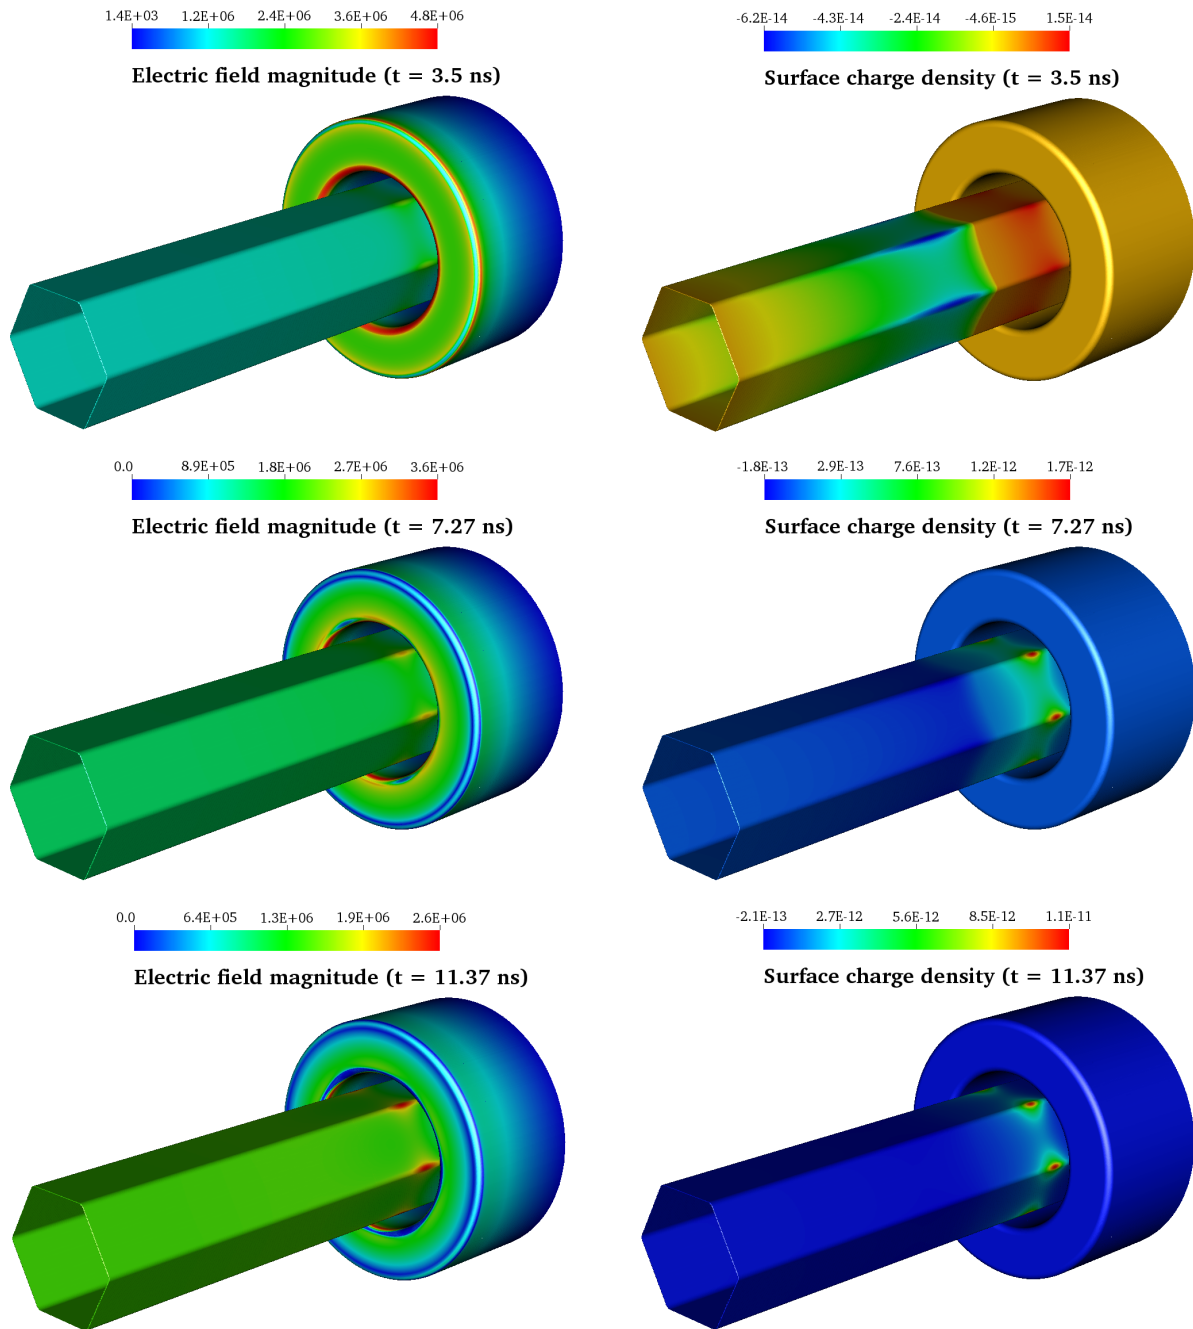


Figure 20: Simulation state after 1.1 ns. The color coded data is in SI units. Left: Field stress at the boundary. Inception has occurred on the outer circular edge of the electrode, and is also occurring on the inner edges where the electrode-insulator distance is smallest. Right: Surface charge density on the boundaries.

insulator, predominantly due to space charge effects in the electrode-insulator gap. At this time, both the outer and inner electrode edges are screened by space charges.

To demonstrate the streamer in greater detail, Fig. 21 shows a volume rendering of the space charge density at $t = 11.37$ ns. We have cut out a part of the simulation domain and adjusted the data range (only positive space charge densities are plotted) for enhanced visibility. We clearly observe a toroidal space charge

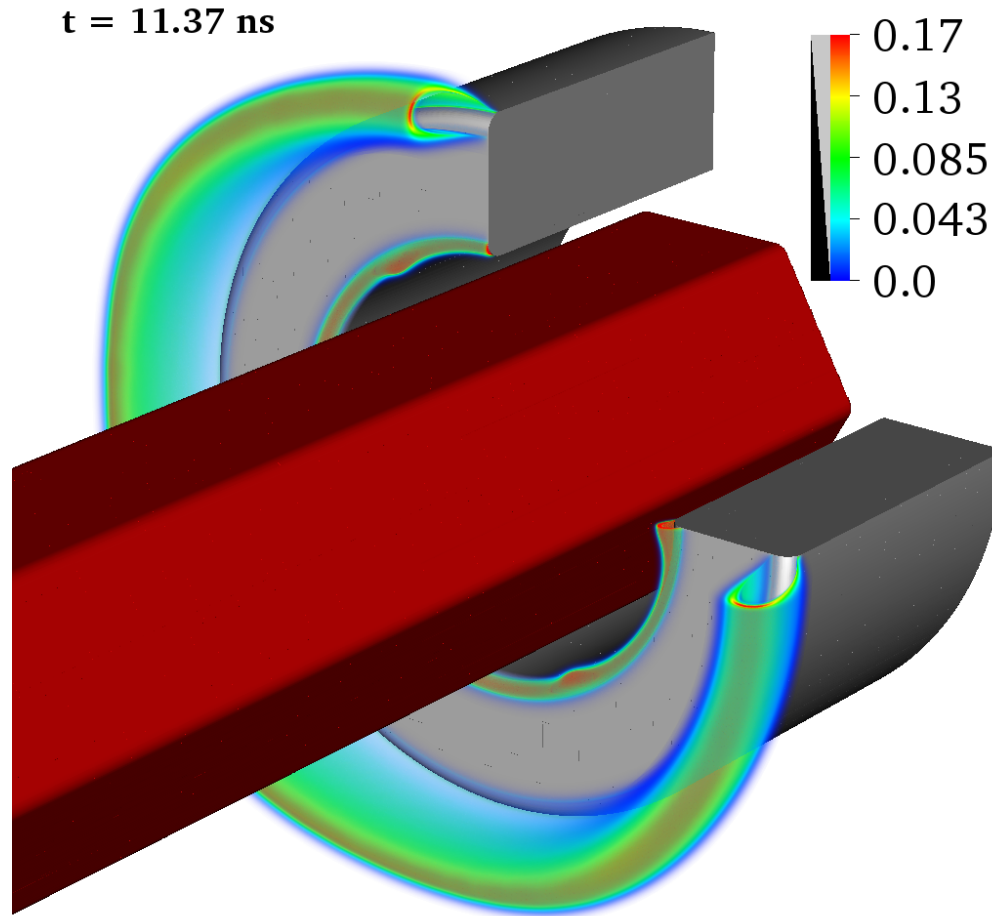


Figure 21: Volume rendered space charge density for the simulation state at 11.37 ns. The data range has been adjusted for improved visibility.

density layer that has initiated from the outer edge of the electrode. We remark that this space charge does not have perfect rotational symmetry; the streamer is thinner where domain boundaries are closer (e.g. at the top and right-hand sides in Fig. 21).

The final mesh for this simulation contains around 800 million cells, or roughly 1.5 billion when one accounts for ghost cells. Since there are three species densities and three photon densities that need to be solved in addition to the electric potential, the final mesh contains 5.6 billion unknowns. The mesh inflation from the initial 300 million cells is primarily due to the toroidal streamer on the outer electrode edge, which required refinement in an annular ring with a comparatively large radius. Plot files for this simulation ranged in sizes up to 250 GB, which had to be visualized in parallel.

7. Final remarks

We have presented an implementation of a low-order plasma fluid model in 2D and 3D, for which several improvements can be made. Our implementation is parallelized with flat MPI and is therefore not expected to scale very well on heterogeneous HPC architectures. We currently do not support fine-grained parallelism (for example through cache blocking inside each patch), nor do we support the use of GPU accelerators. Ideally, we'd like to reduce memory redundancies across MPI ranks by using inter-node MPI parallelism and on-node fine-grained parallelism by means of loop tiling or GPU accelerators. For the models considered in

this paper, the computational bottleneck is due to multiple elliptic solves at every time step, and we expect that we can improve the strong scaling limit quite substantially by use of GPU accelerators.

Finally, this paper does not provide explicit validation studies and the numerical plasma examples that we compute in this paper only indicate physical correctness. For example, one of our cases (streamer along rough surface) provide velocities and sheath thicknesses that are consistent with other computations, while one of our other examples (3D rod simulation) provide velocities and diameters that are consistent with experimental observations at larger gaps. For our largest simulation, no comparison with experiments is currently available.

8. Conclusions

We have presented a parallel code for large-scale simulations of low-temperature filamentary plasmas with inclusion of internal insulators and electrodes in the computational domain. Our code is based on structured adaptive mesh refinement with embedded boundaries for representation of solids, and currently features a multi-fluid Poisson solver; convection-diffusion-reaction solvers, and diffusive radiative transport solvers. We have implemented support for CAD-generated geometries and thus provide a means to perform plasma-fluid simulations in large industrial-grade complex geometries with realistic boundary conditions on dielectric and electrode surfaces. Two- and three-dimensional verification tests show that the overall model is second order accurate for smooth solutions. Weak scalability with parallel efficiencies above 70% were demonstrated at two different machines (Fram and SuperMUC Phase 2), using up to 8192 cores at one of them. Finally, we presented several two- and three-dimensional simulation examples at up to 4096 cores. These simulation examples include creeping streamers over insulators with surface roughness, pin-plane simulations of plasma filaments, and fully transient 3D simulations in industrial grade complex problems. Our code is flexible with regards to the plasma kinetics, making it applicable to a broad range of research fields, ranging from applied high-voltage engineering, aerodynamics, sterilization of polluted gases, and upper atmosphere lightning (sprites). Future efforts will focus on large scale applications on a subset of these, further scaling, as well as improving the overall efficiency of the code.

Acknowledgements

This work was financially supported by the Research Council of Norway through project 245422 and industrial partner ABB AS, Norway. The computations were performed on resources provided by UNINETT Sigma2 - the National Infrastructure for High Performance Computing and Data Storage in Norway. We acknowledge PRACE for awarding us access to SuperMUC at GCS@LRZ, Germany for the scalability studies. The author expresses his gratitude to D. T. Graves at Lawrence Berkeley National Laboratory (LBNL) for help with the Chombo infrastructure.

Appendix A. Discharge model

Appendix A.1. Plasma kinetics

We consider a simplified kinetic scheme for air at standard atmospheric conditions. We solve for three charged species: electrons n_e , positive ions n_+ , and negative ions n_- , with only electrons being diffusive:

$$\frac{\partial n_e}{\partial t} = -\nabla \cdot (\mathbf{v}_e n_e) + \nabla \cdot (D_e \nabla n_e) + S_e, \quad (\text{A.1a})$$

$$\frac{\partial n_+}{\partial t} = -\nabla \cdot (\mathbf{v}_+ n_+) + S_+, \quad (\text{A.1b})$$

$$\frac{\partial n_-}{\partial t} = -\nabla \cdot (\mathbf{v}_- n_-) + S_-, \quad (\text{A.1c})$$

where the source terms are

$$S_e = (\alpha - \eta) |\mathbf{v}_e| n_e - \beta n_e n_+ + S_{\text{ph}}, \quad (\text{A.2a})$$

$$S_+ = \alpha |\mathbf{v}_e| n_e - \beta n_e n_+ - \beta n_- n_+ + S_{\text{ph}}, \quad (\text{A.2b})$$

$$S_- = \eta |\mathbf{v}_e| n_e - \beta n_+ n_-, \quad (\text{A.2c})$$

Here, S_{ph} is a photonization source term that is discussed below, $\alpha = \alpha(|\mathbf{E}|)$ and $\eta = \eta(|\mathbf{E}|)$ are impact ionization and attachment coefficients, and β is a recombination coefficient.

The kinetic coefficients are taken from Morrow and Lowke [64] and are as follows:

$$\mathbf{v}_e = -\frac{\mathbf{E}}{|\mathbf{E}|} \times \begin{cases} (6.87 \times 10^{24} E + 3.38 \times 10^2), & E \leq 2.6 \times 10^{-21}, \\ (7.29 \times 10^{23} E + 1.63 \times 10^4), & 2.6 \times 10^{-13} \leq E \leq 10^{-20}, \\ (1.03 \times 10^{24} E + 1.30 \times 10^4), & 10^{-12} \leq E \leq 2 \times 10^{-19}, \\ (7.40 \times 10^{23} E + 7.10 \times 10^4), & E \geq 2 \times 10^{-19} \end{cases} \quad (\text{A.3a})$$

$$\mathbf{v}_+ = 2.34 \times 10^{-4} \mathbf{E}, \quad (\text{A.3b})$$

$$\mathbf{v}_- = -\frac{\mathbf{E}}{|\mathbf{E}|} \times \begin{cases} 2.7 \times 10^{-4} & E \leq 5 \times 10^{-20}, \\ 1.86 \times 10^{-4} & \text{otherwise,} \end{cases} \quad (\text{A.3c})$$

where $N = 2.45 \times 10^{25} \text{ m}^{-3}$ is the neutral number density and $E = |\mathbf{E}|/N$. The electron diffusion coefficient is

$$D_e = 4.86 \frac{|\mathbf{v}_e|}{|\mathbf{E}|} \times 10^6 \quad (\text{A.4})$$

The explicit expressions for the ionization, attachment, and recombination coefficients are

$$\alpha = N \begin{cases} 6.619 \times 10^{-17} \exp\left(-\frac{7.25 \times 10^{-11}}{E}\right), & E < 1.5 \times 10^{-19}, \\ 2 \times 10^{-16} \exp\left(-\frac{5.59 \times 10^{-11}}{E}\right), & \text{otherwise,} \end{cases} \quad (\text{A.5a})$$

$$\eta = \eta_1 + \eta_2, \quad (\text{A.5b})$$

$$\eta_1 = N \begin{cases} (6.09 \times 10^{-12} E + 2.983 \times 10^{-23}), & E \leq 1.05 \times 10^{-19}, \\ (8.89 \times 10^{-13} E + 2.567 \times 10^{-23}), & \text{otherwise,} \end{cases} \quad (\text{A.5c})$$

$$\eta_2 = 3.79 N^2 E^{-1.275} \times 10^{-74}, \quad (\text{A.5d})$$

$$\beta = 2 \times 10^{-13}. \quad (\text{A.5e})$$

In the above, the entries on the right hand side of Eqs. (A.3) through (A.5) are dimensionless, whereas the left hand sides are in standard (SI) units.

Appendix A.2. Photoionization

The radiative transport module is implemented as a 3-group diffusive model [45]. That is, we consider Eq. (7) in the stationary approximation with $\kappa_\gamma = \lambda_\gamma p_{\text{O}_2} / \sqrt{3}$ for $\gamma = 1, 2, 3$, where λ_γ is a constant and $p_{\text{O}_2} = 0.2 \text{ bar}$ is the oxygen partial pressure. The equations of motion are

$$\kappa_\gamma \Psi_\gamma - \nabla \cdot \left(\frac{1}{3\kappa_\gamma} \nabla \Psi_\gamma \right) = \frac{\eta_\gamma}{c}, \quad \gamma = 1, 2, 3. \quad (\text{A.6})$$

The source terms are the same for all RTE equations and are taken as

$$\eta_\gamma = \nu_{\text{exc}} \frac{p_q}{p + p_q} \alpha n_e |\mathbf{v}_e|, \quad (\text{A.7})$$

where the excitation efficient is taken as $\nu_{\text{exc}} = 0.6$. The term $p_q/(p + p_q)$ is a collisional quenching term that takes into account collisional quenching of radiative transitions. We take $p_q/p = 0.03947$, whereas α and \mathbf{v}_e are as given above. The photoionization rate of O_2 is then

$$S_{\text{ph}} = \nu_{\text{eff}} p_{\text{O}_2} \sum_{\gamma=1,2,3} A_{\gamma} \Psi_{\gamma}, \quad (\text{A.8})$$

where $\nu_{\text{eff}} = 0.1$ is the photoionization efficiency, and A_{γ} are parameters tabulated in Tab. A.3

γ	λ_{γ}	A_{γ}
1	$4.15 \times 10^{-2} \text{ m}^{-1} \text{ Pa}^{-1}$	$1.12 \times 10^{-4} \text{ m}^{-1} \text{ Pa}^{-1}$
2	$1.09 \times 10^{-1} \text{ m}^{-1} \text{ Pa}^{-1}$	$2.88 \times 10^{-3} \text{ m}^{-1} \text{ Pa}^{-1}$
3	$6.69 \times 10^{-1} \text{ m}^{-1} \text{ Pa}^{-1}$	$2.76 \times 10^{-1} \text{ m}^{-1} \text{ Pa}^{-1}$

Table A.3: Parameters use for the photoionization model of oxygen.

Appendix A.3. Charged species boundary conditions

We consider comparatively simple boundary conditions on both dielectrics and electrodes [65]. The outflow of a species i is

$$\mathbf{F}_i \cdot \mathbf{n} = a_i (\mathbf{v}_i n_i) \cdot \mathbf{n} \quad (\text{A.9})$$

where

$$a_i = \begin{cases} 1 & \mathbf{v}_i \cdot \mathbf{n} > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.10})$$

In addition to charged species outflow, we consider injection of electrons from dielectrics due to secondary emission from ion bombardment and photoemission. The electron flux is then

$$\mathbf{F}_e \cdot \mathbf{n} \rightarrow \mathbf{F}_e \cdot \mathbf{n} - \sum_{\text{sign}(Z_i) > 0} \alpha_i \mathbf{F}_i \cdot \mathbf{n} - \sum_{\gamma} \beta_{\gamma} \mathbf{F}_{\gamma} \cdot \mathbf{n}, \quad (\text{A.11})$$

where $\alpha_i = 10^{-6}$ is the second Townsend coefficient for species i , and $\beta_{\gamma} = 10^{-6}$ is the quantum efficiency for photon transition γ . The condition $\text{sign}(Z_i) > 0$ indicates that only positively charged ions lead to secondary electron emission.

References

- [1] A Abou-Ghazala, S Katsuki, K H Schoenbach, F C Dobbs, and K R Moreira. Bacterial decontamination of water by means of pulsed-corona discharges. *IEEE Transactions on Plasma Science*, 30(4):1449–1453, aug 2002. ISSN 0093-3813. doi: 10.1109/TPS.2002.804193.
- [2] J S Clements, A Mizuno, W C Finney, and R H Davis. Combined removal of SO_2 , NO_x , and fly ash from simulated flue gas using pulsed streamer corona. *IEEE Transactions on Industry Applications*, 25(1):62–69, jan 1989. ISSN 0093-9994. doi: 10.1109/28.18870.
- [3] J S Clements, M Sato, and R H Davis. Preliminary Investigation of Prebreakdown Phenomena and Chemical Reactions Using a Pulsed High-Voltage Discharge in Water. *IEEE Transactions on Industry Applications*, IA-23(2):224–235, mar 1987. ISSN 0093-9994. doi: 10.1109/TIA.1987.4504897.
- [4] G Dinelli, L Civitano, and M Rea. Industrial experiments on pulse corona simultaneous removal of NO_x and SO_2 from flue gas. *IEEE Transactions on Industry Applications*, 26(3):535–541, may 1990. ISSN 0093-9994. doi: 10.1109/28.55956.

- [5] S.A. Nair, K. Yan, A.J.M. Pemen, G.J.J. Winands, F.M. van Gompel, H.E.M. van Leuken, E.J.M. van Heesch, K.J. Ptasinski, and A.A.H. Drinkenburg. A high-temperature pulsed corona plasma system for fuel gas cleaning. *Journal of Electrostatics*, 61(2):117–127, jun 2004. ISSN 0304-3886. doi: 10.1016/J.ELSTAT.2004.02.002. URL <https://www.sciencedirect.com/science/article/pii/S0304388604000397?via=IIS>.
- [6] David R. Grymonpré, Amit K. Sharma, Wright C. Finney, and Bruce R. Locke. The role of Fenton’s reaction in aqueous phase pulsed streamer corona reactors. *Chemical Engineering Journal*, 82(1-3):189–207, mar 2001. ISSN 1385-8947. doi: 10.1016/S1385-8947(00)00345-4. URL <https://www.sciencedirect.com/science/article/pii/S1385894700003454?via=IIS>.
- [7] J. P. Boeuf and L. C. Pitchford. Electrohydrodynamic force and aerodynamic flow acceleration in surface dielectric barrier discharge. *Journal of Applied Physics*, 97(10):103307, may 2005. ISSN 0021-8979. doi: 10.1063/1.1901841. URL <http://aip.scitation.org/doi/10.1063/1.1901841>.
- [8] Eric Moreau. Airflow control by non-thermal plasma actuators. *Journal of Physics D: Applied Physics*, 40(3):605–636, feb 2007. ISSN 0022-3727. doi: 10.1088/0022-3727/40/3/S01. URL <http://stacks.iop.org/0022-3727/40/i=3/a=S01?key=crossref.91028d7f08306b56889e7cfc3abce903>.
- [9] V R Soloviev and V M Krivtsov. Surface barrier discharge modelling for aerodynamic applications. *Journal of Physics D: Applied Physics*, 42(12):125208, jun 2009. ISSN 0022-3727. doi: 10.1088/0022-3727/42/12/125208. URL <http://stacks.iop.org/0022-3727/42/i=12/a=125208?key=crossref.6a72c0ae60b829dd552a56b7f9dfa90c>.
- [10] U Ebert, C Montijn, T M P Briels, W Hundsdorfer, B Meulenbroek, A Rocco, and E M van Veldhuizen. The multiscale nature of streamers. *Plasma Sources Science and Technology*, 15(2):S118–S129, may 2006. ISSN 0963-0252. doi: 10.1088/0963-0252/15/2/S14. URL <http://stacks.iop.org/0963-0252/15/i=2/a=S14?key=crossref.7155d333648ba9011d162625b47fd9d4>.
- [11] Annemie Bogaerts, Erik Neyts, Renaat Gijbels, and Joost van der Mullen. Gas discharge plasmas and their applications. *Spectrochimica Acta Part B: Atomic Spectroscopy*, 57(4):609–658, apr 2002. ISSN 0584-8547. doi: 10.1016/S0584-8547(01)00406-2. URL <https://www.sciencedirect.com/science/article/pii/S0584854701004062?via=IIS>.
- [12] S. Pancheshnyi, P. Ségur, J. Capeillère, and A. Bourdon. Numerical simulation of filamentary discharges with parallel adaptive mesh refinement. *Journal of Computational Physics*, 227(13):6574–6590, jun 2008. ISSN 0021-9991. doi: 10.1016/J.JCP.2008.03.020. URL <https://www.sciencedirect.com/science/article/pii/S0021999108001691>.
- [13] A. Luque, U. Ebert, and W. Hundsdorfer. Interaction of Streamer Discharges in Air and Other Oxygen-Nitrogen Mixtures. *Physical Review Letters*, 101(7):075005, aug 2008. ISSN 0031-9007. doi: 10.1103/PhysRevLett.101.075005. URL <https://link.aps.org/doi/10.1103/PhysRevLett.101.075005>.
- [14] S. Nijdam, J. Teunissen, E. Takahashi, and U. Ebert. The role of free electrons in the guiding of positive streamers. *Plasma Sources Science and Technology*, 25(4), 2016. ISSN 13616595. doi: 10.1088/0963-0252/25/4/044001.
- [15] Jannis Teunissen and Ute Ebert. Simulating streamer discharges in 3D with the parallel adaptive Afivo framework. *Journal of Physics D: Applied Physics*, 50(47):474001, nov 2017. ISSN 0022-3727. doi: 10.1088/1361-6463/aa8faf. URL <http://stacks.iop.org/0022-3727/50/i=47/a=474001?key=crossref.3482bbfb872e5442dc361fe881ee81c>.
- [16] A Hallac, G E Georghiou, and A C Metaxas. Secondary emission effects on streamer branching in transient non-uniform short-gap discharges. *Journal of Physics D: Applied Physics*, 36(20):2498–2509, oct 2003. ISSN 0022-3727. doi: 10.1088/0022-3727/36/20/011. URL <http://stacks.iop.org/0022-3727/36/i=20/a=011?key=crossref.95b34ecfeb9638a2b12be116909f0fdf>.

- [17] L Papageorgiou, A C Metaxas, and G E Georghiou. Three-dimensional numerical modelling of gas discharges at atmospheric pressure incorporating photoionization phenomena. *Journal of Physics D: Applied Physics*, 44(4):045203, feb 2011. ISSN 0022-3727. doi: 10.1088/0022-3727/44/4/045203. URL <http://stacks.iop.org/0022-3727/44/i=4/a=045203?key=crossref.2e01f7ab6053426830fd6b81864bc84f>.
- [18] J-M Plewa, O Eichwald, O Ducasse, P Dessante, C Jacobs, N Renon, and M Yousfi. 3D streamers simulation in a pin to plane configuration using massively parallel computing. *Journal of Physics D: Applied Physics*, 51(9):095206, mar 2018. ISSN 0022-3727. doi: 10.1088/1361-6463/aaa91b. URL <http://stacks.iop.org/0022-3727/51/i=9/a=095206?key=crossref.2196e446ec5a2c2015c60946e522821a>.
- [19] V.I. Kolobov and R.R. Arslanbekov. Towards adaptive kinetic-fluid simulations of weakly ionized plasmas. *Journal of Computational Physics*, 231(3):839–869, feb 2012. ISSN 0021-9991. doi: 10.1016/J.JCP.2011.05.036. URL <https://www.sciencedirect.com/science/article/pii/S0021999111003482>.
- [20] S Singh, Y V Serdyuk, and R Summer. Streamer branching in air: Physical model and simulations in fully 3D spatial domain. In *2015 IEEE 11th International Conference on the Properties and Applications of Dielectric Materials (ICPADM)*, pages 220–223, jul 2015. doi: 10.1109/ICPADM.2015.7295248.
- [21] Feng Shi, Ningyu Liu, and Joseph R. Dwyer. Three-Dimensional Modeling of Two Interacting Streamers. *Journal of Geophysical Research: Atmospheres*, 2017. ISSN 21698996. doi: 10.1002/2017JD026935.
- [22] Jouya Jadidian, Markus Zahn, Nils Lavesson, Ola Widlund, and Karl Borg. Stochastic and deterministic causes of streamer branching in liquid dielectrics. *Journal of Applied Physics*, 114(6):063301, aug 2013. ISSN 0021-8979. doi: 10.1063/1.4816091. URL <http://aip.scitation.org/doi/10.1063/1.4816091>.
- [23] J Jadidian, M Zahn, N Lavesson, O Widlund, and K Borg. Effects of Impulse Voltage Polarity, Peak Amplitude, and Rise Time on Streamers Initiated From a Needle Electrode in Transformer Oil. *IEEE Transactions on Plasma Science*, 40(3):909–918, mar 2012. ISSN 0093-3813. doi: 10.1109/TPS.2011.2181961.
- [24] J. Jadidian, M. Zahn, N. Lavesson, O. Widlund, and K. Borg. Effects of electrode size and solid barrier orientation on streamer discharge in transformer oil. *Journal of Applied Physics*, 115(14):143304, apr 2014. ISSN 0021-8979. doi: 10.1063/1.4870477. URL <http://aip.scitation.org/doi/10.1063/1.4870477>.
- [25] J G Hwang, M Zahn, L A A Pettersson, O Hjortstam, and R Liu. Modeling streamers in transformer oil: The transitional fast 3rd mode streamer. In *2009 IEEE 9th International Conference on the Properties and Applications of Dielectric Materials*, pages 573–578, jul 2009. doi: 10.1109/ICPADM.2009.5252362.
- [26] Alexandre V. Likhanskii, Mikhail N. Shneider, Sergey O. Macheret, and Richard B. Miles. Modeling of dielectric barrier discharge plasma actuator in air. *Journal of Applied Physics*, 103(5):053305, mar 2008. ISSN 0021-8979. doi: 10.1063/1.2837890. URL <http://aip.scitation.org/doi/10.1063/1.2837890>.
- [27] V R Soloviev and V M Krivtsov. Mechanism of streamer stopping in a surface dielectric barrier discharge. *Plasma Physics Reports*, 40(1):65–77, jan 2014. ISSN 1562-6938. doi: 10.1134/S1063780X1312009X. URL <https://doi.org/10.1134/S1063780X1312009X>.
- [28] A Dubinova, D Trienekens, U Ebert, S Nijdam, and T Christen. Pulsed positive discharges in air at moderate pressures near a dielectric rod. *Plasma Sources Science and Technology*, 25(5):055021, sep 2016. ISSN 1361-6595. doi: 10.1088/0963-0252/25/5/055021. URL <http://stacks.iop.org/0963-0252/25/i=5/a=055021?key=crossref.512b925b2701d5829dfe2d2540d0675b>.
- [29] Jan Van Dijk, Kim Peerenboom, Manuel Jimenez, Diana Mihailova, and Joost Van Der Mullen. The plasma modelling toolkit Plasimo, 2009. ISSN 00223727.

- [30] Anna Dubinova, Casper Rutjes, Ute Ebert, Stijn Buitink, Olaf Scholten, and Gia Thi Ngoc Trinh. Prediction of Lightning Inception by Large Ice Particles and Extensive Air Showers. *Physical Review Letters*, 115(1):015002, jun 2015. ISSN 0031-9007. doi: 10.1103/PhysRevLett.115.015002. URL <https://link.aps.org/doi/10.1103/PhysRevLett.115.015002>.
- [31] S. Celestin, Z. Bonaventura, O. Guaitella, A. Rousseau, and A. Bourdon. Influence of surface charges on the structure of a dielectric barrier discharge in air at atmospheric pressure: experiment and modeling. *The European Physical Journal Applied Physics*, 47(2):22810, aug 2009. ISSN 1286-0042. doi: 10.1051/epjap/2009078. URL <http://www.epjap.org/10.1051/epjap/2009078>.
- [32] Jaroslav Jánký, Fabien Tholin, Zdeněk Bonaventura, and Anne Bourdon. Simulation of the discharge propagation in a capillary tube in air at atmospheric pressure. *Journal of Physics D: Applied Physics*, 43(39):395201, oct 2010. ISSN 0022-3727. doi: 10.1088/0022-3727/43/39/395201. URL <http://stacks.iop.org/0022-3727/43/i=39/a=395201?key=crossref.28031ba27f2de8e7d299ec3cdd25ef8c>.
- [33] François Pechereau, Jaroslav Jánký, and Anne Bourdon. Simulation of the reignition of a discharge behind a dielectric layer in air at atmospheric pressure. *Plasma Sources Science and Technology*, 21(5):055011, oct 2012. ISSN 0963-0252. doi: 10.1088/0963-0252/21/5/055011. URL <http://stacks.iop.org/0963-0252/21/i=5/a=055011?key=crossref.a1800ff7ed358b8df5b184c66dbe40e3>.
- [34] François Pechereau, Zdeněk Bonaventura, and Anne Bourdon. Influence of surface emission processes on a fast-pulsed dielectric barrier discharge in air at atmospheric pressure. *Plasma Sources Science and Technology*, 25(4):044004, aug 2016. ISSN 0963-0252. doi: 10.1088/0963-0252/25/4/044004. URL <http://stacks.iop.org/0963-0252/25/i=4/a=044004?key=crossref.061b2b8e05235f8f415a721a4e45a23d>.
- [35] Max Duarte, Zdeněk Bonaventura, Marc Massot, Anne Bourdon, Stéphane Descombes, and Thierry Dumont. A new numerical strategy with space-time adaptivity and error control for multi-scale streamer discharge simulations. *Journal of Computational Physics*, 231(3):1002–1019, feb 2012. ISSN 0021-9991. doi: 10.1016/J.JCP.2011.07.002. URL <https://www.sciencedirect.com/science/article/pii/S0021999111004062>.
- [36] Behnaz Bagheri, Jannis Teunissen, Ute Ebert, Markus M Becker, She Chen, Olivier Ducasse, Olivier Eichwald, Detlef Loffhagen, Alejandro Luque, Diana Mihailova, Joseph-Marie Plewa, Jan van Dijk, and Mohammed Yousfi. Comparison of six simulation codes for positive streamers in air. *Plasma Sources Science and Technology*, aug 2018. ISSN 0963-0252. doi: 10.1088/1361-6595/aad768. URL <http://iopscience.iop.org/article/10.1088/1361-6595/aad768>.
- [37] P Colella, D T Graves, T J Ligocki, G Miller, D Modiano, P O Schwartz, B Van Straalen, J Pilliod, D Trebotich, M Barad, B Keen, A Nonaka, and C Shen. EBChombo software package for cartesian grid, embedded boundary applications. Technical report, Lawrence Berkeley National Laboratory.
- [38] M Adams, P. Colella, D.T. Graves, J. N. Johnson, N. D. Keen, T. J. Ligocki, D. F. Martin, P. W. McCorquodale, D. Modiano, P. O. Schwartz, T. D. Sternberg, and B Van Straalen. Chombo Software Package for AMR Applications - Design Document. Technical report, Lawrence Berkeley National Laboratory.
- [39] O. Eichwald, O. Ducasse, N. Merbahi, M. Yousfi, and D. Dubois. Effect of order fluid models on flue gas streamer dynamics. *Journal of Physics D: Applied Physics*, 2006. ISSN 00223727. doi: 10.1088/0022-3727/39/1/015.
- [40] S. Dujko, A. H. Markosyan, R. D. White, and U. Ebert. High-order fluid model for streamer discharges: I. Derivation of model and transport data. *Journal of Physics D: Applied Physics*, 2013. ISSN 00223727. doi: 10.1088/0022-3727/46/47/475202.

- [41] A. H. Markosyan, S. Dujko, and U. Ebert. High-order fluid model for streamer discharges: II. Numerical solution and investigation of planar fronts. *Journal of Physics D: Applied Physics*, 2013. ISSN 00223727. doi: 10.1088/0022-3727/46/47/475203.
- [42] Edward W. Larsen, Guido Thömmes, Axel Klar, Mohammed Sead, and Thomas Götz. Simplified PN Approximations to the Equations of Radiative Heat Transfer and Applications. *Journal of Computational Physics*, 183(2):652–675, dec 2002. ISSN 0021-9991. doi: 10.1006/JCPH.2002.7210. URL <https://www.sciencedirect.com/science/article/pii/S0021999102972104>.
- [43] M. B. Zhelezniak, A. K. Mnatsakanian, and S. V. Sizykh. Photoionization of nitrogen and oxygen mixtures by radiation from a gas discharge. *Teplofizika Vysokikh Temperatur*, 20:423–428, nov 1982.
- [44] P. Ségur, A. Bourdon, E. Marode, D. Bessieres, and J. H. Paillol. The use of an improved Eddington approximation to facilitate the calculation of photoionization in streamer discharges. *Plasma Sources Science and Technology*, 15(4):648–660, nov 2006. ISSN 0963-0252. doi: 10.1088/0963-0252/15/4/009. URL <http://stacks.iop.org/0963-0252/15/i=4/a=009?key=crossref.e2f9e61e0c5a4082930c63e0777c096a>.
- [45] A. Bourdon, V. P. Pasko, N. Y. Liu, S. Célestin, P. Ségur, and E. Marode. Efficient models for photoionization produced by non-thermal gas discharges in air based on radiative transfer and the Helmholtz equations. *Plasma Sources Science and Technology*, 16(3):656–678, aug 2007. ISSN 0963-0252. doi: 10.1088/0963-0252/16/3/026. URL <http://stacks.iop.org/0963-0252/16/i=3/a=026?key=crossref.26470bbe9c1f765777a162c80aa57bb7>.
- [46] Julien Capeillère, Pierre Ségur, Anne Bourdon, Sébastien Célestin, and Sergey Pancheshnyi. The finite volume method solution of the radiative transfer equation for photon transport in non-thermal gas discharges: application to the calculation of photoionization in streamer discharges. *Journal of Physics D: Applied Physics*, 41(23):234018, dec 2008. ISSN 0022-3727. doi: 10.1088/0022-3727/41/23/234018. URL <http://stacks.iop.org/0022-3727/41/i=23/a=234018?key=crossref.d0be3c7c41b5fad3f08cbe991e123a84>.
- [47] Hans Johansen and Phillip Colella. A Cartesian Grid Embedded Boundary Method for Poisson’s Equation on Irregular Domains. *Journal of Computational Physics*, 147(1):60–85, nov 1998. ISSN 0021-9991. doi: 10.1006/JCPH.1998.5965. URL <https://www.sciencedirect.com/science/article/pii/S0021999198959654>.
- [48] Peter McCorquodale, Phillip Colella, and Hans Johansen. A Cartesian grid embedded boundary method for the heat equation on irregular domains. *Journal of Computational Physics*, 2001. ISSN 00219991. doi: 10.1006/jcph.2001.6900.
- [49] Peter Schwartz, Michael Barad, Phillip Colella, and Terry Ligocki. A Cartesian grid embedded boundary method for the heat equation and Poisson’s equation in three dimensions. *Journal of Computational Physics*, 2006. ISSN 00219991. doi: 10.1016/j.jcp.2005.06.010.
- [50] R.K. Crockett, P. Colella, and D.T. Graves. A Cartesian grid embedded boundary method for solving the Poisson and heat equations with discontinuous coefficients in three dimensions. *Journal of Computational Physics*, 230(7):2451–2469, apr 2011. ISSN 0021-9991. doi: 10.1016/J.JCP.2010.12.017. URL <https://www.sciencedirect.com/science/article/pii/S0021999110006819?via=I%3Dihub>.
- [51] Koen Van Laer and Annemie Bogaerts. Fluid modelling of a packed bed dielectric barrier discharge plasma reactor. *Plasma Sources Science and Technology*, 2015. ISSN 13616595. doi: 10.1088/0963-0252/25/1/015002.
- [52] E. H. Twizell, A. B. Gumel, and M. A. Arigu. Second-order, L₀-stable methods for the heat equation with time-dependent boundary conditions. *Advances in Computational Mathematics*, 6(1):333–352, 1996. ISSN 1019-7168. doi: 10.1007/BF02127712. URL <http://link.springer.com/10.1007/BF02127712>.

- [53] Ken Perlin. Improving noise. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques - SIGGRAPH '02*, page 681, 2002. ISBN 1581135211. doi: 10.1145/566570.566636. URL <http://portal.acm.org/citation.cfm?doid=566570.566636>.
- [54] Natalia Yu Babaeva and Mark J Kushner. Structure of positive streamers inside gaseous bubbles immersed in liquids. *Journal of Physics D: Applied Physics*, 42(13):132003, jul 2009. ISSN 0022-3727. doi: 10.1088/0022-3727/42/13/132003. URL <http://stacks.iop.org/0022-3727/42/i=13/a=132003?key=crossref.995a75b9fc435f5829b6c996f96a3fbd>.
- [55] V Soloviev and V Krivtsov. Positive Streamer in the Surface Dielectric Barrier Discharge in Air: Numerical Modelling and Analytical Estimations. *Journal of Physics: Conference Series*, 927:012059, nov 2017. ISSN 1742-6588. doi: 10.1088/1742-6596/927/1/012059. URL <http://stacks.iop.org/1742-6596/927/i=1/a=012059?key=crossref.d5589ea44c111555ce506e55eccdad68>.
- [56] Weizhuo Hua and Koji Fukagata. Influence of grid resolution in fluid-model simulation of nanosecond dielectric barrier discharge plasma actuator. *AIP Advances*, 8(4):045209, apr 2018. ISSN 2158-3226. doi: 10.1063/1.5013627. URL <http://aip.scitation.org/doi/10.1063/1.5013627>.
- [57] Ashish Kumar Jindal, Christopher Hudson Moore, Andrew S Fierro, and Roy E Jorgenson. Steamer Formation Near a Dielectric Surface with Variable Quantum Yield.
- [58] S. Pancheshnyi, M. Nudnova, and A. Starikovskii. Development of a cathode-directed streamer discharge in air at different pressures: Experiment and comparison with direct numerical simulation. *Physical Review E*, 71(1):016407, jan 2005. ISSN 1539-3755. doi: 10.1103/PhysRevE.71.016407. URL <https://link.aps.org/doi/10.1103/PhysRevE.71.016407>.
- [59] I. Gallimberti. A computer model for streamer propagation. *Journal of Physics D: Applied Physics*, 5(12):2179–2189, 1972. ISSN 00223727. doi: 10.1088/0022-3727/5/12/307.
- [60] N. L. Allen and A. Ghaffar. The conditions required for the propagation of a cathode-directed positive streamer in air. *Journal of Physics D: Applied Physics*, 28(2):331–337, 1995. ISSN 13616463. doi: 10.1088/0022-3727/28/2/016.
- [61] Natalya Y. Babaeva and George V. Naidis. Dynamics of positive and negative streamers in air in weak uniform electric fields. *IEEE Transactions on Plasma Science*, 25(2):375–379, 1997. ISSN 00933813. doi: 10.1109/27.602514.
- [62] Jianqi Qin and Victor P. Pasko. On the propagation of streamers in electrical discharges. *Journal of Physics D: Applied Physics*, 47(43), 2014. ISSN 13616463. doi: 10.1088/0022-3727/47/43/435202.
- [63] Thomas Christen. Nonstandard High-Voltage Electric Insulation Models. In *COMSOL Conference in Milan*, number October, pages 1–6, 2012.
- [64] R Morrow and J J Lowke. Streamer propagation in air. *Journal of Physics D: Applied Physics*, 30(4):614–627, feb 1997. ISSN 0022-3727. doi: 10.1088/0022-3727/30/4/017. URL <http://stacks.iop.org/0022-3727/30/i=4/a=017?key=crossref.7fd5909110367fd5ab39f05c2b4dfc92>.
- [65] G. J. M. Hagelaar, F. J. de Hoog, and G. M. W. Kroesen. Boundary conditions in fluid models of gas discharges. *Physical Review E*, 62(1):1452–1454, jul 2000. ISSN 1063-651X. doi: 10.1103/PhysRevE.62.1452. URL <https://link.aps.org/doi/10.1103/PhysRevE.62.1452>.