Characterizing Plasma Diffusion in Transitional Hypersonic Flows

by

Marisa Elena Petrusky

B.S., Stony Brook University, 2021

M.S., University of Colorado, 2023

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Aerospace Engineering Sciences

2026

Committee Members:

Iain Boyd, Chair

Hisham Ali

Brian Argrow

Robyn Macdonald

Greg Werner

Petrusky, Marisa Elena (Ph.D., Aerospace Engineering Sciences)

Characterizing Plasma Diffusion in Transitional Hypersonic Flows

Thesis directed by Professor Iain Boyd

Under hypersonic entry conditions into an atmosphere, gas may be ionized via shock waves to create a partially ionized plasma. Accurate modeling of these flows is essential for deriving macroscopic properties relevant to vehicle design. When simulating hypersonic plasmas, modelers conventionally enforce ions and electrons to diffuse at the same rate– known as the ambipolar diffusion approximation. The approximation assumes the charged species undergo ambipolar diffusion: an idealized plasma diffusion process of the continuum regime. This approach circumvents resolution of fast electron motion but subsequently neglects the complex plasma dynamics of ions and electrons and their interactions.

This dissertation work assesses the efficacy of the ambipolar diffusion approximation through the lens of charged species diffusion physics in argon flows. The transitional flow portion of a reentry trajectory (ranging approximately between 60 and 100 km for Earth) is chosen for study since it requires consideration of rarefaction in the freestream in analyzing plasma diffusion. Charged species diffusion is first investigated through a weakly ionized expanding gas flow. New criteria are proposed for identification of plasma diffusion regimes in arbitrary flows and use of the ambipolar diffusion approximation which account for the theory of ambipolar diffusion and the numerical results observed in the study. These criteria are then successfully applied in the context of transitional regime hypersonic flows. First-order plasma effects are singled out which provide a baseline to analyze more complex, chemically reacting hypersonic flows with electrostatic modeling. The efficacy of the ambipolar diffusion approximation in simulating transitional regime hypersonic flows is quantified in terms of predicting plasma density distributions, electron temperature, and stagnation point heat flux.

In memory of Eileen Josenhans and Else Grotefendt.

Acknowledgements

The stereotypical scientist walks alone, but nothing could be further from the truth. I am blessed to have so many wonderful people that supported me throughout my life and this Ph.D. journey.

First, I would like to thank my advisor, Professor Iain Boyd, for taking me on as a student and helping me become the researcher I am today. I admire the trust you place in all your students, and am truly grateful for the opportunities you provided for me to try novel ideas, make mistakes and learn from them, and share my work with people around the world. I also appreciate that you always encouraged me to pursue my passions in professional service and future career aspirations.

Second, I would like to thank the members of my committee for the time they have invested in reviewing my work and thinking about its implications for the rest of the research community. All of the conversations we have had over the years, no matter how brief, have influenced the direction of this work and my own growth. I am lucky to have had so many kind mentors who are also at the forefront of advancing research in the field.

I am proud to be a part of the Nonequilibrium Gas and Plasma Dynamics Laboratory family, and would like to thank everyone for their help, advice, and friendship over the years. In no particular order, I would specifically like to thank Wai Hong Ronald Chan, Kaelan Hansson, Kal Monroe, Mitch Wall, Tim Aiken, Amin Taziny, Sarah Kinney, and Daniel Jimenez. I also wish to mention Lisa Ventura, Ryan Draves, Andrew Morell, and Conor Rowan– each of our weekly discussions with Sarah gave me hope for a better world. I will also give a shout out to the members of the Next Generation of the International Symposium on Rarefied Gas Dynamics. It is an honor

to be part of such a close-knit research community, and I appreciate all of the detailed discussions we have had on our work. You all represent academia at its best, and I am excited to see what the future has in store for us all.

Especially while living in distant Colorado, I am grateful to have kept in touch with my incredible friends: Neomi, Ignacia, Andrew, Puneet, Daniel, Mike, Mary Kate, Hamdan, Pritha, Mark, Anastasiya, and Deanndra. I also cannot thank my family enough for their love and support. To my parents especially, thank you for raising me to have a strong work ethic, committment to faith, and love of community– all of which made completing this degree possible.

Last but not least, I must thank Max. A more appropriate occasion to convey the depths of my love and appreciation is coming soon, so for now I will say: thanks for introducing me to Stargate.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2040434. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

This work utilized the Alpine and Blanca high performance computing resource at the University of Colorado Boulder. Alpine is jointly funded by the University of Colorado Boulder, the University of Colorado Anschutz, and Colorado State University. Blanca is jointly funded by computing users and the University of Colorado Boulder.

Contents

Chapter

Tables

Table

Figures

Figure

"I, therefore, did learn a lesson: The female mind is capable of understanding analytic geometry. Those people who have for years been insisting (in the face of all obvious evidence to the contrary) that the male and female are equally capable of rational thought may have something. The difficulty may just be that we have never yet discovered a way to communicate with the female mind. If it is done in the right way, you may be able to get something out of it."

– Richard Feynman, *What is Science?* (1966)

Chapter 1

Introduction

The objective of this dissertation work is to expand on the existing body of knowledge regarding simulation of rarefied and partially ionized hypersonic flows. Hypersonic flows involve many coupled physical processes, including, but not limited to: non-equilibrium chemistry, radiation, and gas-surface interactions. Many advancements have been made in simulating these processes with increasingly high-fidelity numerical models. However, comparatively less research has been done on the fundamentals of hypersonic plasma physics, especially in terms of numerical modeling. In this chapter, Section 1.1 describes the broader motivations for this work, Section 1.2 introduces the formal definition of a plasma, and Section 1.3 reviews current literature on plasma physics in hypersonic shock layers. Specific research objectives and the outline of this thesis are detailed in Section 1.4.

1.1 Motivations

The Artemis II mission in April 2026 marked many important milestones in humanity's quest for sustained space exploration. In addition to being the first crewed flight beyond low Earth orbit since Apollo 17 in 1972 and breaking the record for the furthest any human has been from Earth, the mission presented an invaluable opportunity to validate the Orion spacecraft's systems in advance of future lunar landings. Since 1972, countless advancements have been made in spacecraft design and optimization through both experimental and computational research.

One system inseparable from vehicle design is the flight environment of a spacecraft during

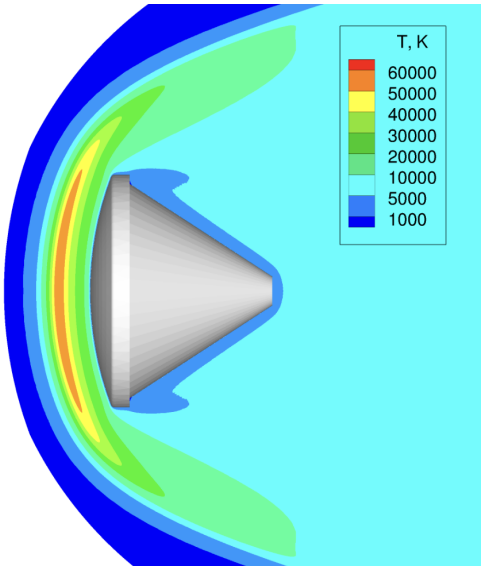

Figure 1.1: Schematic of the flow over a re-entry spacecraft capsule [1].

atmospheric entry. For reference, the Artemis II vehicle re-entered Earth's atmosphere at 25,000 miles per hour, or 11.2 kilometers per second. Accurate characterization of the flow conditions and heat loads a spacecraft will experience as it decelerates through the atmosphere is essential for vehicle design optimization, scientific instrumentation operation, and astronaut safety. Numerical simulation, enabled by decades of evolution in high-performance computing power, is an invaluable tool in this endeavor, as it provides a comparatively inexpensive way to study the many coupled physical phenomena around a hypersonic vehicle.

The flight environment of a spacecraft during atmospheric entry is classified as hypersonic. In aerodynamics, hypersonics refers to the regime of flight through an atmosphere with a Mach number greater than 5, where the Mach number is defined as

$$M = \frac{u}{a} \tag{1.1}$$

where u is the average velocity of flow and a is the local speed of sound. Around a hypersonic vehicle the flow is characterized by a shock wave (a sharp change in flow properties from the freestream) and the shock layer (the flowfield between the shock wave and vehicle surface). As the Mach number

increases, so does the translational kinetic energy of the flow, which is then converted into thermal and internal energy via shock wave compression and heating. Figure 1.1 displays a schematic of a hypersonic flow over a spacecraft capsule, with a translational temperature distribution representative of air above Mach 30 during re-entry flight conditions. What differentiates transonic (i.e., $0.8 < M \lesssim 1.2$) and supersonic (i.e. $1.2 < M < 5$) aerodynamics from hypersonic aerodynamics is that the significantly higher temperatures result in the presence of "real-gas effects" (as opposed to ideal gas behavior), such as vibrational excitation, dissociation, ionization, and radiation.

Real-gas effects in hypersonic flows are highly coupled to each other and all aspects of the vehicle. That is, when building a hypersonic vehicle, components such as wings and nacelles cannot be treated as separate aerodynamic bodies: they are all closely integrated, and must be carefully designed according to the expected real-gas effects in the flow regimes of interest [2]. Hypersonic flows can be further categorized according to their degree of continuum behavior. A formal definition of continuum is provided in Section 2.1.1– to summarize, 'continuum' refers to a flow where the characteristic length scale of collisions between particles is much smaller than the length scale being observed, such as in the nose radius of the vehicle or the width of a shock layer. At low altitudes on Earth, the air is denser, therefore collisions occur more frequently and at much shorter length scales. As the altitude increases the air becomes more rarefied, and the average distance a particle travels between collisions increases. If the distance a particle travels in between collisions is larger than the length scale of interest, the flow can be treated as 'collisionless' and is classified into the free molecular regime. The regime between free molecular and continuum is known as the 'transitional flow regime'. Figures 1.2 and 1.3 mark the changes in hypersonic flow regime along the re-entry trajectory of the Stardust Sample Return Capsule, which returned to Earth from a cometary dust collection mission in 2006 with an atmospheric entry speed of 12.6 km/s [3]. The transitional flow regime is increasingly relevant to the aerospace community due to interest in deploying air-breathing vehicles to very low Earth orbit (VLEO), utilizing aerobraking maneuvers for future missions to outer planets, and the optimization of vehicle design during the fastest portion of a re-entry trajectory.

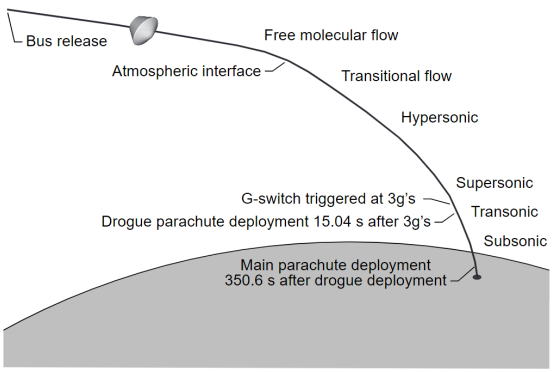

Figure 1.2: Stardust sample return capsule entry, descent, and landing sequence [4].

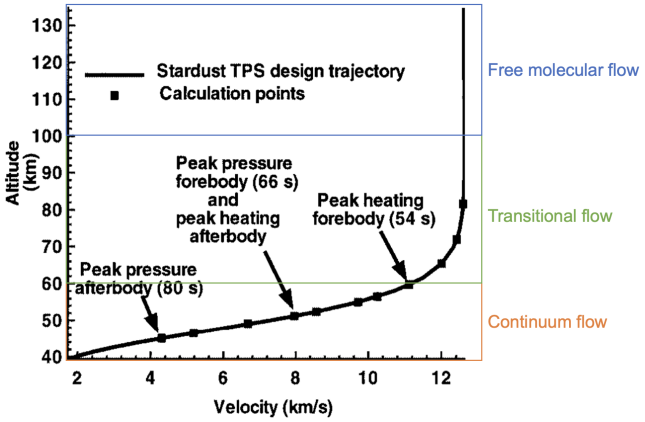

Figure 1.3: Stardust heat shield design trajectory [3] with gas dynamic flow regimes labeled.

At sufficiently high speeds ionization rates can exceed recombination rates resulting in a sustained population of charged species. These ions and electrons carry electric charge, leading to electrostatic interactions among them which strongly differ from neutral particle interactions. Depending on the flight conditions and location around the vehicle the ionization fraction for a hypersonic flowfield can range between 10^{-6} to 10^{-1}. Though these fractions may seem small, given the high densities characteristic of hypersonic flows, they can indicate non-negligible populations of charged particles from a plasma physics perspective (see Figure 1.4). In the presence of a density gradient the charged species will form an electric field in an attempt to preserve local charge neutrality– otherwise the lightweight electrons quickly diffuse away from the heavier ions. The magnitude of this generated electric field is directly proportional to the magnitude of the density gradient [5]. Classical motion of charged species through an electric field, known as electrostatics, leads to the acceleration or deceleration of particles, thereby altering the distribution of kinetic energy among the chemical species in the shock layer. If the partially ionized gas meets certain requirements (explained in Section 1.2), it is defined as a plasma.

Study of plasmas in hypersonic flows is vital to the development of advanced aerospace technologies. A challenge that has persisted since the earliest days of space travel is communications blackout: the period of time during atmospheric entry when communication between the vehicle and ground stations is lost. This typically corresponds with the portion of an entry trajectory when the plasma densities are highest, as radio waves will be attenuated and/or reflected if the frequency of a signal is lower than the local plasma frequency. Similarly, hypersonic vehicle remote observation depends on the identification of flight states via gas radiation spectra, which itself relies heavily upon the characteristics of the plasma species itself [6, 7, 8].

Other technologies seek to take advantage of the generated plasma. Since 1958 [9], engineers have explored the use of magnetohydrodynamic (MHD) body forces for high-speed acceleration, power generation, and flow control (particularly for reducing surface heating). MHD-augmented systems would use onboard electromagnets or electrodes rather than chemical fuel, thereby reducing the vehicle's propellant mass fraction. Since MHD forces depend on characteristics of the plasma

rather than atmospheric density [10], they are potentially employable in highly rarefied atmospheres such as that of Mars [11, 12], or in Earth's upper atmosphere. Significant advancements have been made in superconductivity and magnetic-field generation technology since the 1950s which make such technologies feasible to develop today. Electron transpiration cooling (ETC) is a proposed thermal management technique where the vehicle surface is cooled via thermionic emission of electrons [13, 14]. The efficiency of ETC depends heavily on the kinetics of the plasma immediately adjacent to the surface of the vehicle, known as a plasma sheath [15], which is naturally influenced by the plasma's properties throughout the shock layer [16].

At a more fundamental level, there is growing interest in high-fidelity characterization of non-equilibrium plasma chemistry and electron kinetics due to the role plasma species play in re-distributing energy of the gas via excitation and ionization processes [17, 18]. At super-orbital Earth entry flow conditions, electron impact ionization of nitrogen is the primary source of plasma species [19, 20]. Atomic oxygen and nitrogen dominate in thermal heat flux contributions to the surface of the vehicle for flow conditions with post-shock temperatures above 5,000 K [21]. Additionally, the energy distribution function of these electrons is strongly non-Maxwellian [22], that is, the energy distribution of electrons does not follow a normal Gaussian distribution with respect to their temperature. A major challenge in characterizing non-equilibrium plasma chemistry is propagating the large uncertainties associated with hypersonic chemistry rate coefficients [23]. This has led to an increase in uncertainty quantification and sensitivity analysis work in the literature [24, 25, 26].

Ultimately, as space exploration objectives become more complex and the need for highly optimized hypersonic vehicles grows, so does the need for accurate modeling of plasma dynamics. In particular, there is a gap in the literature with respect to the nature of electrostatic fields in hypersonic plasmas, how those electric fields impact plasma diffusion processes, and the impacts resolution of those physics have on quantities of interest relevant to hypersonic flowfield modeling.

1.2 Definition of a Plasma

The mere presence of charged particles does not always result in complex electromagnetic interactions associated with plasmas. One key characteristic of a plasma is its ability to 'shield' out charge disturbances, such that sufficiently far from the disturbance the plasma is approximately charge neutral. The Debye length is best understood as a characteristic measure of this distance. If the background medium can be treated as a vacuum (appropriate for hypersonic flowfields), the full definition of the Debye length is

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B / e^2}{n_e/T_e + \sum_j z_j^2 n_j/T_j}} \tag{1.2}$$

where ϵ_0 is the permittivity of free space, k_B is the Boltzmann constant, e is the elementary charge, n_e is the number density of electrons, T_e is the temperature of electrons, and n_j is the number density of ionic species j with charge number z_j and translational temperature T_j. Often, the mobility of ions is negligible compared to the timescale of the disturbance and the mobility of electrons, so the ion contribution term is dropped. Equation 1.2 then reduces to the formula for the electron Debye length, defined as

$$\lambda_{D,e} = \sqrt{\frac{\epsilon_0 k_B T_e}{n_e e^2}} \tag{1.3}$$

The term 'Debye length' can be used to refer to either Equation 1.2 or Equation 1.3. From the Debye length stems the three criteria required for a partially ionized gas to constitute a plasma [27]:

(1) $\lambda_D \ll L$, where L is a characteristic length scale for the system,

(2) $N_D \ggg 1$, where N_D is the number of plasma particles in a sphere of radius λ_D, known as the plasma parameter,

(3) $\omega\tau > 1$, where ω is the frequency of typical plasma oscillations and τ is the mean time between collisions of plasma particles and neutrals

The first requirement, known as quasineutrality, comes from the nature of charge shielding. When the dimensions of the plasma are much greater than the Debye length, as local charge disturbances or external potentials are introduced, they are shielded out such that the bulk of the plasma remains largely charge neutral. It is important to distinguish that "charge neutral" refers to the net charge of an ionized gas being zero, while "quasineutral" refers to the entire plasma at length scales of L being sufficiently close to charge neutral such that $n_i \simeq n_e$. Therefore, it is possible that within a quasineutral plasma at length scales $L \lesssim \lambda_D$, local charge neutrality may not be conserved. In fact, it is these local charge imbalances that lead to many interesting plasma physics phenomena.

The second requirement states that there must be a sufficient number of charged particles available in order to shield out a charge disturbance. Having many particles within a Debye sphere means that the collective effect of Coulomb forces from nearby particles is greater than Coulomb forces from binary charged particle interactions, which can potentially lead to rare-large angle scattering collisions. However, there is not a standard minimum value of N_D applicable to any plasma physics problem.

The third requirement indicates if the frequency of neutral collisions is larger than the frequency of plasma oscillations then hydrodynamic forces dominate over electrodynamic ones and the charged species will behave like a neutral gas rather than a plasma. Generally, the typical plasma frequency equals the characteristic electron oscillatory frequency, given by

$$\omega_{p,e} = \sqrt{\frac{n_e e^2}{\epsilon_0 m_e}} \tag{1.4}$$

where m_e is the mass of the electron. The third requirement typically holds true for rarefied gases, which are the subject of interest of this work.

Thus, if a partially ionized gas meets these three criteria, it can be considered a proper plasma. In a spatially or temporarily varying flow, it is possible for a partially ionized gas to only be a plasma at certain points in space or time. Generally speaking, a plasma is considered

'weakly ionized' if the ionization fraction is 10^{-4} or lower, and 'strongly ionized' if the ionization fraction is greater than 10^{-4}. However, the effectiveness of this language in describing partially ionized hypersonic flows warrants further investigation. Additionally, even if a partially ionized gas does not constitute a formal plasma, plasma physics theory may still be effective in describing the charged species dynamics.

Figure 1.4 shows the general range of number densities and temperatures found in plasmas formed by conventional hypersonic flight conditions (i.e. flight conditions found along realistic vehicle trajectories for an air atmosphere) compared to plasmas in other areas of research or in nature. In hypersonic flows, plasma density and temperature tend to increase with vehicle speed and altitude. A high speed gas flow has more translational kinetic energy available for energetic ionization reactions, and at higher altitudes, the probability of recombination is lessened due to the rarefied atmosphere. Of note, some hypersonic flows may have plasma densities comparable to those present in magnetic fusion reactors or the solar corona. The main difference is that hypersonic plasmas are considered 'low temperature' with respect to other areas of plasma physics.

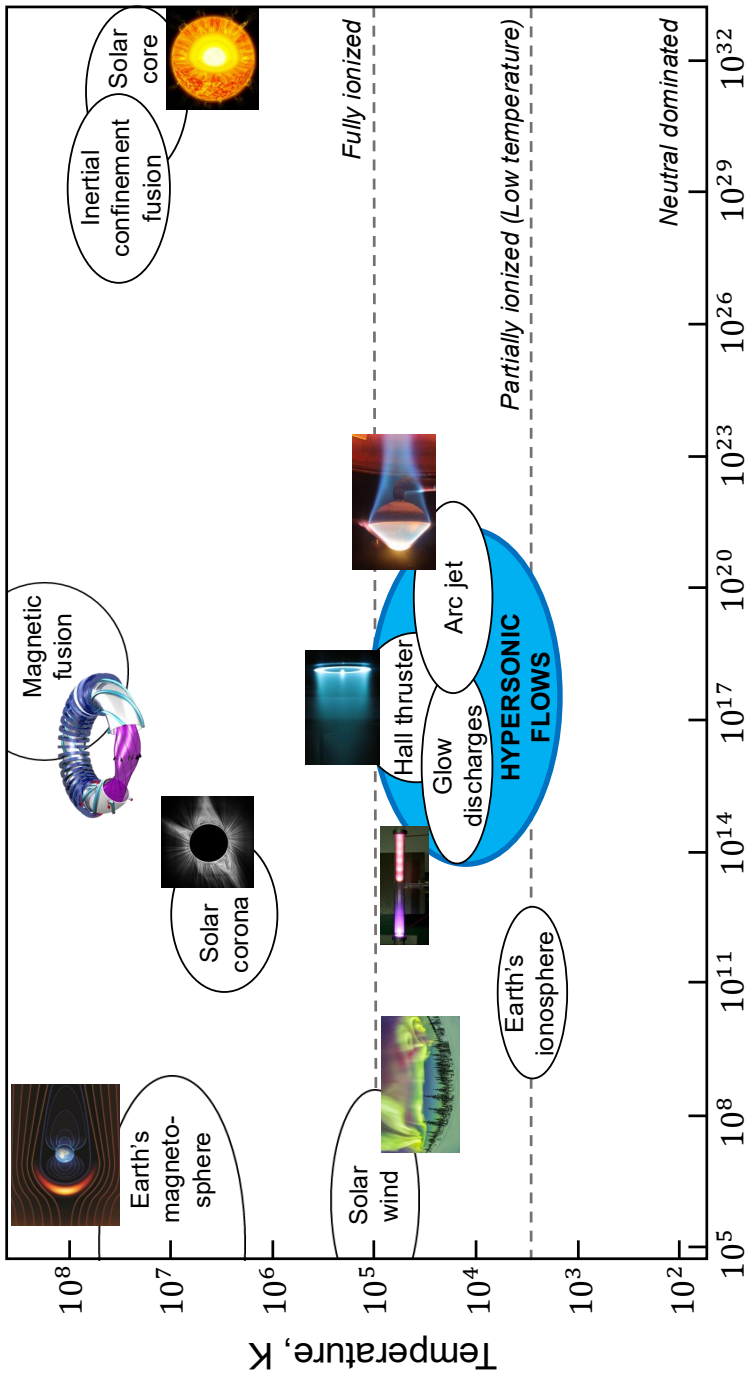

Figure 1.4: Map of typical number densities and temperatures of various plasmas.

1.3 Modeling Plasmas in Hypersonic Flows

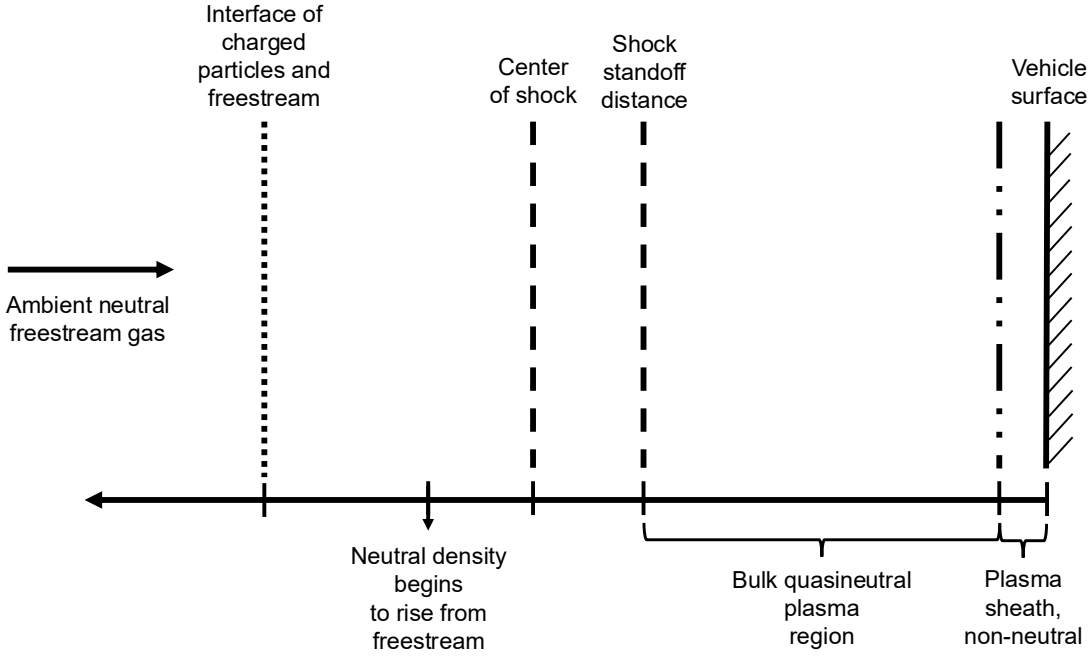

Figure 1.5: Notional schematic of plasma generated in a transitional regime hypersonic shock layer (not to scale).

Figure 1.5 provides a notional schematic of the structure of a partially ionized transitional regime shock layer with respect to a strong bow shock and the vehicle surface based on current literature. A key difference between hypersonic shock layers in the continuum versus the transitional flow regime is the width of the shock itself. The distance between where the density begins to rise from the freestream to the shock standoff distance* can be comparable to the width of the post-shock region itself due to the increased distance between particle collisions. The incoming freestream atmospheric gas comprises neutral species for an air atmosphere. Between the shock and the vehicle surface, the plasma is typically quasineutral and undergoes ambipolar diffusion– that is, ions and electrons diffuse at the same rate [28]. It is often assumed that because the plasma is quasineutral, the local charge separation is negligible, and thus the electrostatic fields of the

*The shock standoff distance is defined as the point where the flowfield density is four times that of the freestream for a monatomic gas, and six times that for a diatomic gas such as air.

shock layer are also negligible. Charged particles may diffuse upstream and out of the shock layer. The plasma within a transitional regime shock and upstream of the entire shock layer has not been extensively explored in the literature.

Finally, at the vehicle surface, lightweight electrons will reach the wall at a faster rate than ions. This creates a charge imbalance; subsequently, the bulk quasineutral plasma of the shock layer shields itself via the formation of a plasma sheath. Within just several Debye lengths, large electric fields form that ensure the ion and electron currents throughout the sheath and to the wall are equal. Charge neutrality is not conserved in this region. Note that in aerospace engineering, the term 'plasma sheath' is occasionally used to refer to the entire partially ionized shock layer. This work will exclusively refer to the plasma sheath as the thin, non-charge neutral region at the vehicle surface.

First principles modeling of a plasma is challenging, especially in a hypersonic flowfield where many other physical phenomena must be accounted for. The electric potential must be computed from some form of Maxwell's equations. The computational cost of this additional step scales with the number of grid points used to resolve the flowfield. Simulations must then incorporate the computed electromagnetic fields into flow advection routines. If electrons are resolved as a separate species (which is necessary to observe charge separation), then due to their lightweight mass and large thermal velocities relative to heavy particles, a small grid cell size and time step must be used to resolve the fast motion of the electrons. To avoid the computational expense associated with solving Maxwell's equations, convention in hypersonic numerical simulation is to utilize the ambipolar diffusion approximation, defined in the following section.

1.3.1 The Ambipolar Diffusion Approximation

The ambipolar diffusion approximation is a modeling approach that assumes charge separation in the flowfield is sufficiently close to zero so as to justify neglection of charge separation-induced electrostatic effects. There are several ways to enforce this in kinetic simulation. For the particle-based Direct Simulation Monte Carlo (DSMC) method [29], the first approach was developed by

Bird, where simulated electrons are restrained to always remain close to the ions they were generated with regardless of their high velocity components [30], ensuring exact charge neutrality. A second approach was developed by Boyd, which does not restrain the electrons' physical locations, but propagates them with the calculated ion bulk velocity in each cell [31]. Boyd's approach does not enforce local charge neutrality. The third approach of note is implemented in the popular open-source DSMC code SPARTA [32, 33], where 'ambipolar ion' particles store both ion and electron velocity components. These ambipolar particles are split apart during collisions and chemical reactions, such that ions and electrons can collide and react independently. At the end of all collisions and chemistry, each electron is re-paired with an ion, and the ambipolar particle is advected according to the ion's velocity. Thus, SPARTA's approach enforces exact charge neutrality. In all of these approaches, electrons are allowed to collide with other particles using their own high velocity components, but the global simulation time step is typically based on the heavy particle collision frequency, not the higher electron collision frequency. Other variations of the ambipolar diffusion approximation may not track electron velocity components or permit collisions between electrons and other species.

The ambipolar diffusion approximation is typically considered valid when the electron Debye length is orders of magnitude smaller than the local average neutral mean free path [34],

$$\lambda_{D,e} \ll \lambda_{mfp,n} \tag{1.5}$$

However, this criterion is not intrinsic to ambipolar diffusion theory on which the ambipolar diffusion approximation is based. Furthermore, the degree to which the ambipolar diffusion approximation captures standard ambipolar diffusion has yet to be formally investigated, particularly for non-continuum, hypersonic flows.

1.3.2 Computational State-of-the-Art in Kinetic Hypersonic Plasma Modeling

To date, comparatively few studies have modeled hypersonic plasmas with a complete electromagnetic or electrostatic description; that is, modeling ions and electrons as separate species (so

as to allow for charge separation) with a two-way coupling to a Maxwell's equations solver. Fewer studies characterize the consequences of using the ambipolar diffusion approximation compared to a full electromagnetic or electrostatic description on quantities of interest.

In Reference [34], Bird coupled a DSMC solver with the ambipolar diffusion approximation to an ambipolar electric field solver. A flowfield was computed, and then the equation for a self-induced electric field formed by a plasma undergoing ambipolar diffusion (known as the Langmuir and Tonks relation) was solved using the electron number densities and temperatures of the DSMC solution. This electric field solution was applied to the flowfield, the flowfield's properties changed due to acceleration and deceleration, and the ambipolar electric field was re-computed. This cycle was repeated until convergence was achieved. For 11-species air (N_2, O_2, NO, N, O, each of their respective ions, and e^-) flow conditions at 93 and 100 km with a freestream velocity of 9.8 km/s, the ambipolar electric field was reported to have little effect on the flow aside from increasing the charged particle densities upstream. The maximum reported electric field value was +50 V/m, located within the shock itself.

Carlson and Hassan developed a similar DSMC approach where ions and electrons were allowed to move independently and were accelerated via an electric field solution that enforces net zero current and exact charge neutrality in each cell [35]. For a stagnation streamline in 11-species air at 0.1 Torr with a freestream velocity of 10 km/s, this resulted in a much greater peak electric field value of approximately +3,800 V/m at the shock compared to the value predicted by the Langmuir and Tonks relation of approximately +2,000 V/m. The electron temperatures also decreased throughout the flowfield, which impacted electron impact ionization processes.

Gallis *et al.* [36] employed particle modeling for heavy species and fluid modeling for electrons with the Coulomb collision approach of Jones *et al.* [37] over a 5° spherically blunted cone flying in a partially ionized helium atmosphere at 10.6 km/s. For a 40% ionized freestream case, the hybrid model produced stronger electric fields that extended 20% further upstream of the vehicle compared to the Langmuir and Tonks relation computed from a standard DSMC solution with the ambipolar diffusion approximation. All electric fields were computed to be on the order of 0.1 to

1.0 V/m. The extended electric field occurs due to the enhanced propagation of plasma species in the presence of a self-consistent electric field.

Farbar and Boyd [38] were the first to directly couple a DSMC code with a Poisson solver so as to obtain self-consistent solutions to the electric field in a hypersonic flowfield. Re-entry conditions of the FIRE II vehicle [39] in 11-species air were simulated at 85 km along a stagnation streamline. Their results showed that the ambipolar diffusion approximation underpredicted stagnation point heat flux at the vehicle surface by between 13% and 16%. The increase in heat flux was primarily attributed to an increase in atomic nitrogen ion flux to the surface of the vehicle due to acceleration via electrostatic fields, contributing both convective and chemical heat flux via wall recombination. Increased convective heat flux was also reported in N_2 and O atoms, since these species gained energy through collisions with the accelerated N^+ ions. The peak electric field value was -130 V/m at the shock, and with resolution of the plasma sheath saw electric fields between +1,700 V/m and +10,000 V/m adjacent to the wall. However, their study employed an artificially reduced plasma density, thus the reported differences with the ambipolar diffusion approximation are potentially underpredicted.

The most recent study on the ambipolar diffusion approximation in kinetic modeling was performed by Kaplan and Oran [40]. DSMC results were compared between simulations with the ambipolar diffusion approximation and without the approximation, allowing ions and electrons to move freely without any electric fields imposed. The latter will be referred to in this thesis as the 'free diffusion approximation'. Under Stardust Sample Return Capsule flight conditions at 71 km and 100 km with 11-species air, they found that the free diffusion approximation resulted in significant charge separation between electrons and ions beyond what would be expected for a typical hypersonic flow.

Although this work primarily concerns kinetic simulation methods and rarefied flows, some recent studies in the CFD literature warrant mention. Blanco and Josyula [41] coupled a CFD solver to an electrostatic Poisson equation solver. For RAM-C II flight conditions at 61 km and Mach 23.9 with 7-species air (N_2, O_2, NO, NO^+, N, O, and e^-) [42, 43], simulations with electrostatic modeling

demonstrated better agreement with flight data than the ambipolar diffusion approximation with respect to electron number density distributions. The computed electric field was 9% greater than that of the Langmuir and Tonks relation, and the mass concentration of N_2 at the wall was increased by 8%. Crucially, the electrostatic modeling simulations predicted an increase of up to 31% heat flux at the wall. However, this study did not resolve the plasma sheath.

Recent work by Parent *et al.* also employed a CFD code coupled to a Poisson solver. While their work did not compare flowfield results with the ambipolar diffusion approximation and with electrostatic modeling, their 11-species air results achieve strong agreement with experimental data. Their work especially highlights the importance of resolving the electron temperature due to heating and cooling of electrons via electric fields [44], the strong coupling of the plasma sheath to the rest of the flowfield [45], and the role of ion mobility models in electron density and temperature calculations [46].

Various works centering around MHD solvers are not substantially applicable to the interests of this work, as although the MHD description solves Maxwell's equations, ideal MHD models electrons and ions as one conductive fluid. Additionally, MHD literature often simulates applied electric and magnetic fields, which adds complexity when trying to understand the fundamentals of a hypersonic plasma. Of note, Hou *et al.* found significant differences in flowfield plasma velocities and temperatures between two-fluid and single-fluid MHD models, which correlate with increasing applied electromagnetic fields, altitude, and freestream velocity [47].

Overall, current literature is in agreement that ambipolar diffusion occurs in the bulk plasma region of a shock layer and charge separation will occur upstream of the shock due to the fast diffusion of electrons once they escape the self-induced electric fields of the plasma. These electric fields generally do not agree with the Langmuir and Tonks relation and result in non-negligible differences in post-shock plasma properties. The most critical of these include the electron temperature, highly influential for prediction of excitation and chemical reaction rate processes, and the surface heat flux to the vehicle, arguably the most important quantity for hypersonic vehicle design. Subsequently, the ambipolar diffusion approximation may be insufficient in capturing these

properties.

Most of the prior studies were performed for a specific set of flow conditions in chemically reacting air. Therefore, their results cannot readily be extrapolated to other flow conditions, vehicle geometries, or atmospheres. Specifically, it is unknown whether the reported differences in quantities of interest originate directly from acceleration and deceleration of charged species through the electric fields and momentum-exchange collisions between charged and neutral species (defined as first-order effects), or from subsequent interactions with particles experiencing first-order effects (defined as second-order effects). It follows that the more species, energy modes, and chemical reactions are included in a flow model, the harder it is to distinguish between first-order and second-order effects in a hypersonic plasma. Furthermore, the vast majority of the hypersonic flowfield modeling community still employs the ambipolar diffusion approximation due to its computational efficiency. Thus, generalized criteria for characterizing a hypersonic plasma of any composition and quantifying the impact of the ambipolar diffusion approximation on first-order plasma effects are needed.

1.4 Scope and Outline of Dissertation

The objective of this dissertation is to characterize first-order plasma diffusion behavior in a hypersonic flowfield and the impact of high-fidelity electrostatic modeling on prediction of quantities of interest for vehicle design. The transitional regime of hypersonic flow is selected for study since it requires consideration of rarefaction in the freestream in analyzing plasma diffusion. Argon gas is selected for analysis throughout this work for several reasons. The primary reason is that modeling a monatomic gas circumvents consideration of rotational and vibrational degrees of freedom and limits the possible binary interactions between particles to collisions, ionization, and CEX. This enables better observation of first-order plasma effects, establishing a baseline for plasma diffusion behavior for which more complex partially ionized gas mixtures can be compared to. Argon is also a common choice of gas for aerospace plasma experiments due to its similar mass to molecular nitrogen ($m_{N_2}/m_{Ar} = 0.701$). Since diffusion processes depend heavily on particle mass, selecting a

monatomic species similar in mass to molecular nitrogen improves generalizability of results to air flows. Finally, future missions to outer planets incentivize study of hypersonic plasma behavior in monatomic flows. For example, the atmospheres of Uranus and Neptune are estimated to comprise about 15 and 19% helium, respectively [48, 49]. Any mission to enter one of these atmospheres will likely require a high entry velocity, which is more likely to generate a strongly ionized hypersonic flow around the spacecraft.

Because hypersonic plasmas are increasingly important for a number of aerospace applications, a large portion of this thesis is dedicated to investigating the advantages and disadvantages of two different kinetic modeling approaches as pertains to the objectives of this work. The conventional kinetic modeling approach is to use a particle-based method where the interactions of real gas particles are represented by computational macroparticles, resulting in a stochastic solution. Effective sampling techniques reduce numerical statistical noise in the solutions, but may still struggle to resolve regions where few macroparticles exist. The second kinetic modeling approach numerically solves multiple gas dynamics equations, resulting in a deterministic solution throughout the flowfield. This avoids the problem of statistical noise, but such approaches often require far more computational resources than particle-based methods. By including these discussions in this dissertation, modelers may reference this work when selecting a computational tool for their studies.

Chapter 2 provides an overview of the deterministic and stochastic kinetic modeling tools used throughout the dissertation research, including the implementation and verification of code enhancements. Chapter 3 studies the physics of charged species diffusion and the validity of the ambipolar diffusion approximation in non-hypersonic partially ionized flows. A set of guidelines for identifying charged species diffusion regime is developed and verified with a weakly ionized expanding gas problem. Chapter 4 analyzes the advantages and disadvantages of using deterministic or stochastic simulation methods in studying partially ionized hypersonic flowfields, particularly in terms of computational cost. A novel one-dimensional stagnation streamline model for deterministic kinetic solvers is derived, verified, and compared with DSMC. Chapter 5 investigates charged species

diffusion in hypersonic flows across a broad range of plasma conditions via a series of stagnation streamline simulations in the transitional flow regime. The charged species diffusion criteria outlined in Chapter 3 are verified for hypersonic flowfields, and the tradeoffs associated with using the ambipolar diffusion approximation are formally quantified. Finally, Chapter 6 summarizes the main findings of this thesis and provides avenues for future work.

Chapter 2

Numerical Methods

In this chapter, a detailed description of the numerical methodology used in this work is provided. First, in Section 2.1, the governing equations for simulating a gas flow with a kinetic approach are given. Next, in Section 2.2, an overview of the Direct Simulation Monte Carlo method (DSMC) is provided. Then, in Section 2.3 an overview of the discrete-velocity method (DVM) is given. Section 2.4 details the numerical scheme behind the Direct Kinetic code: the DVM solver used throughout the majority of this work, and Section 2.5 details MPIC: the coupled DSMC-Particle-in-Cell (PIC) solver used in Chapter 4.

2.1 Governing Equations

Kinetic theory of dilute gas dynamics relates the microscopic evolution of large populations of colliding gas particles to macroscopic properties through integration of collective properties. Here, 'particles' refers to atoms, molecules, ions, and electrons in a gas or plasma. 'Dilute' refers to a gas where the mean spacing between particles is much greater than the particle size, and particles are weakly bound such that their motion can be described as independent of each other [50]. Descriptions of particles are considered in terms of phase space: the combination of physical (or configuration) space and velocity space.

Consider a volume element of configuration space at point $\vec{x}$ with dimensions $dxdydz$ and N identical particles inside. Each molecule has a velocity vector $\vec{v} = (v_x, v_y, v_z)$. Just as dimensions $dxdydz$ define a volume element in configuration space, dimensions $dv_x dv_y dv_z$ define a volume

element in velocity space. The velocity distribution function (VDF), $f(\vec{v})$, for the volume element of space around $\vec{x}$ is then defined as

$$\begin{aligned} N = \int dN &= \iiint_{-\infty}^{\infty} N\hat{f}(\vec{v})dv_x dv_y dv_z \\ &= \iiint_{-\infty}^{\infty} f(\vec{v})dv_x dv_y dv_z \end{aligned} \tag{2.1}$$

where $\hat{f}(\vec{v})$ denotes the VDF normalized to 1.0. If variation of the VDF is considered over the entire configuration space domain, then it is written as $f(\vec{x}, \vec{v})$.

The Boltzmann equation evolves $f(\vec{x}, \vec{v})$ over temporal space. It can be derived through physical conservation arguments [51], or by transforming the Liouville equation via the Bogoliubov-Born-Green-Kirkwood-Yvon (BBGKY) hierarchy [52]. Either derivation requires that only binary collisions occur, and that the velocities of two colliding particles are uncorrelated. For species s, the Boltzmann equation is given by

$$\frac{\partial f_s(\vec{x}, \vec{v}, t)}{\partial t} + \vec{v} \cdot \nabla_x f_s(\vec{x}, \vec{v}, t) + \vec{a} \cdot \nabla_v f_s(\vec{x}, \vec{v}, t) = S(f_s) \tag{2.2}$$

where t describes the time coordinate, $S(f_s)$ is the sum of all source, sink, or collision terms, and $\vec{a}$ is acceleration from all body forces. This work will consider all force contributions from magnetic fields negligible, namely: Earth's magnetic field, which has a maximum value of 60 μT; and the self-induced magnetic field generated by electrical current in the plasma. Gravitational force contributions are also considered negligible. When $S(f_s)$ is zero, Equation 2.2 is referred to as the Vlasov equation.

Macroscopic properties are obtained by evaluating moments of the VDF. For example, number density n, bulk velocity $\vec{u}$, and temperature T are obtained from the zeroth, first, and second moments of the VDF:

$$n(\vec{x},t) = \iiint_{-\infty}^{\infty} f(\vec{x},\vec{v},t)d\vec{v} \tag{2.3}$$

$$\vec{u}(\vec{x},t) = \iiint_{-\infty}^{\infty} \vec{v}\hat{f}(\vec{x},\vec{v},t)d\vec{v} \tag{2.4}$$

$$\epsilon(\vec{x},t) = \iiint_{-\infty}^{\infty} \frac{1}{2}m|\vec{v}|^2\hat{f}(\vec{x},\vec{v},t)d\vec{v} \tag{2.5}$$

$$T(\vec{x},t) = \frac{m}{3k_B}(2\epsilon(\vec{x},t) - u(\vec{x},t)^2 - v(\vec{x},t)^2 - w(\vec{x},t)^2) \tag{2.6}$$

where $\epsilon(\vec{x},t)$ is the mean energy. If $f(\vec{x},\vec{v},t)$ can be described everywhere as a slightly perturbed Maxwellian (known as the Chapman-Enskog, or continuum assumption), then the Navier-Stokes system of equations can be derived from Equations 2.3-2.5. The Maxwellian, or equilibrium, distribution $f_{eq}(\vec{v})$ is defined as

$$f_{eq}(\vec{v},\vec{u},T) = \left(\frac{m}{2\pi k_B T}\right)^{j/2} exp\left(-\frac{m}{2k_B T}|\vec{v}-\vec{u}|^2\right) \tag{2.7}$$

where k_B is the Boltzmann constant, and j is the number of degrees of freedom in velocity space (i.e., the number of components $\vec{v}$ contains).

2.1.1 Non-continuum Behavior

Gas dynamics can be generally categorized according to the non-dimensional global Knudsen number Kn, given by

$$\text{Kn} = \frac{\lambda_{mfp}}{L} \tag{2.8}$$

where λ_{mfp} is the mean free path of the flow (i.e. the average distance traveled by a particle between collisions), and L is the characteristic length scale of the problem. In hypersonic flow analysis, L is typically some property of the vehicle, such as the diameter or nose radius. Figure 2.1 defines the approximate boundaries of each flow regime by its global Knudsen number along with valid governing equations in each regime.

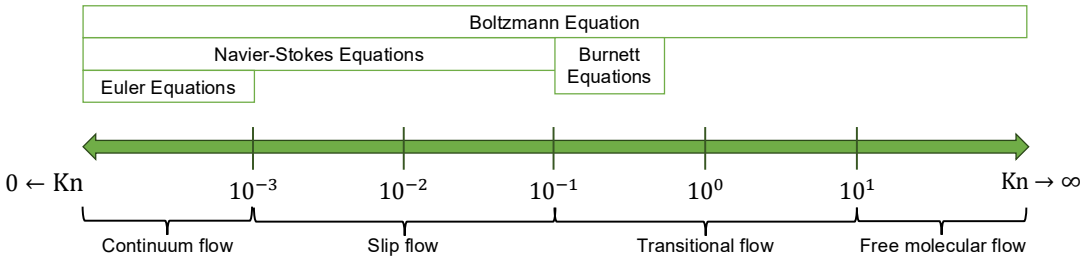

Figure 2.1: Scale of global Knudsen number values, with approximate outlines of flow regimes and their valid governing equations.

As the Knudsen number increases, particles experience fewer collisions relative to the length and time scales of the problem. Since collisions drive a flow towards equilibrium, continuum assumptions break down with increasing Kn and kinetic descriptions are required to accurately capture gas dynamics.

Note that in hypersonics, non-equilibrium flow behavior primarily refers to thermochemical non-equilibrium: where the temperatures of various internal energy modes (e.g. rotational, vibrational, electronic excitation) are not equal to each other across some time or length scale of interest [2]. For example, across a normal shock wave, flow properties change instantaneously across a small distance. Equilibrium will be achieved downstream of the shock front once enough collisions have occurred to drive all temperature modes towards the equilibrium value dictated by the chemical processes. Similarly, in non-continuum flows, collisions drive the VDF further towards the equilibrium state defined by Equation 2.7. It is important to clarify that thermochemical non-equilibrium does not necessarily imply non-continuum VDF behavior; moreover, for the monatomic flows modeled in this work, thermochemical non-equilibrium analysis does not apply since only translational temperature modes exist. For clarity, this work will use 'non-continuum' to define VDF deviation from a Maxwellian or slightly perturbed Maxwellian shape.

In hypersonic flows, the degree of continuum behavior can change throughout the shock layer. To evaluate continuum behavior at a specific point, the gradient-length local Knudsen number [53] can be used,

$$\mathrm{Kn}_{\mathrm{GLL}}(\vec{x}) = \frac{\lambda_{mfp}}{Q}\left|\frac{dQ}{dl}\right| \tag{2.9}$$

where Q is a flow property (typically density or temperature) and dl is the distance between two points. A value of $\mathrm{Kn}_{\mathrm{GLL}}$ above 0.05 is a strong predictor that continuum approaches will fail. Hybrid kinetic-continuum simulation methods often use $\mathrm{Kn}_{\mathrm{GLL}}$ as a metric to evaluate whether to switch between continuum and kinetic methods.

2.2 The Direct Simulation Monte Carlo Method (DSMC)

The most well known kinetic simulation method is Direct Simulation Monte Carlo (DSMC), developed by Graeme Bird [29]. DSMC is a stochastic approach which uses computational macroparticles to model the trajectories of a pre-determined number of real gas particles. The scaling between the number density and the number density of macroparticles is known as the particle weight. As the particle weight approaches 1 (i.e. the number of computational macroparticles approaches the true number of gas particles), it has been demonstrated that solutions of the DSMC method approach the solution of the Boltzmann equation [54]. A full overview of the DSMC method can be found in References [29] and [55].

The DSMC method is founded on the assumption that the movement and collision phases of particle behavior can be decoupled. Thus, the computational time step of a DSMC simulation must be less than the minimum collision time throughout the flowfield. Every time step, the equations of motion are solved for each macroparticle,

$$\frac{d\vec{x}}{dt} = \vec{v}, \frac{d\vec{v}}{dt} = \vec{a} \tag{2.10}$$

though except in rare circumstances, no acceleration forces are considered in DSMC simulations. Integration schemes such as Euler or leapfrog can be used to solve Equation 2.10. Then, within a configuration space grid cell, macroparticles are selected for collisions, chemical reactions, and/or internal energy exchange until the appropriate collision and chemistry rates are achieved. If electromagnetic forces must be considered, a DSMC code can optionally be coupled to a PIC solver [56, 57, 58, 59]. PIC solvers compute the acceleration or deceleration of charged particles through electromagnetic fields using Maxwell's equations (Equation 2.29).

Two DSMC codes will be used throughout this work. The first is ISTAG: a DSMC code which only computes 1D stagnation streamline flows developed by Iain D. Boyd at NASA Ames Research Center [60]. The second is MPIC: a coupled DSMC-PIC solver first developed at the University of Michigan [61].

2.3 The Discrete-Velocity Method (DVM)

The discrete-velocity method (DVM) (also referred in literature as the grid-based kinetic [62] or continuum-kinetic method [63]) numerically solves the Boltzmann equation for a discretized VDF. The first DVM solver was published in 1976, which used the finite difference method to simulate nonlinear evolution of the Vlasov equation [64]. Because the Boltzmann equation is a multidimensional nonlinear hyperbolic partial differential equation (PDE), DVM can employ many of the same numerical advection and stability techniques used for CFD solvers.

In order to simulate realistic hypersonic gas dynamics, the proper heat capacity ratio γ must be captured. For an ideal gas, the classical equipartition theorem defines γ as

$$\gamma = 1 + \frac{2}{j} \tag{2.11}$$

Thus, a monatomic gas with one translational degree of freedom ($j = 1$) has $\gamma = 3$, and a monatomic gas with three translational degrees of freedom ($j = 3$) has $\gamma = 5/3$. Thus, any work meant to be representative of real hypersonic flows should simulate three degrees of freedom in velocity space ($3V$). Note that the collisional DSMC method inherently requires all macroparticles have $3V$, regardless of the configuration space dimensions.

When collisions are modeled in DVM, a collision source term must be solved for on the right side of Equation 2.2. However, the full Boltzmann collision integral is computationally costly to solve and is often unnecessary to compute exactly. Instead, collision models are used depending on the type of collisions and degree of accuracy being resolved. The most common model is the Bhatnagar-Gross-Crook (BGK) model [65], given by

$$S_{BGK} = \frac{1}{\tau(x)}(f_{eq}(x, \vec{v}) - f(x, \vec{v})) \tag{2.12}$$

where τ is the inverse of the collision frequency per configuration space cell and f_{eq} is a Maxwellian distribution initialized to the species temperature and shifted by its bulk velocities. Reference [66] summarizes the properties and tradeoffs of the BGK model as well as its improved variants like the

Shakhov (S-BGK) [67] or ellipsoidal-statistical (ES-BGK) [68] models. Of note, the BGK operator produces a Prandtl number of $Pr = 1$, whereas for a monatomic gas in the hydrodynamic limit, $Pr = \frac{2}{3}$. Examples of this phenomenon impacting flowfield results are presented in Section 4.3.2. In spite of this, the BGK operator is often used due to its simplicity and computational efficiency.

2.4 The Direct Kinetic (DK) Code

The Direct Kinetic (DK) code is the in-house DVM solver used throughout this work. Originally developed at the University of Michigan [69, 70, 71], DK has undergone detailed verification with canonical plasma problems [72, 73, 74, 75, 76], and has been used to model complex engineering systems such as Hall thrusters [77, 78, 79] and hollow cathode devices [80, 81]. This work is the first to use DK to model atmospheric hypersonic flows.

Note that throughout this section and dissertation, the word 'particles' may be used to describe features and results from DK for readability, despite the fact that no physical particles are being simulated. Additionally, note that DK solves for a non-normalized solution to Equation 2.2 such that when integrated the VDF solution sums to the number density of the cell rather than 1.0.

2.4.1 Numerical Scheme

DK is a finite volume solver which satisfies the following key properties:

(1) *Positivity preserving*: All values of the VDF must be positive, since the probability of particles in phase space are bounded between zero and one.

(2) *Mass, momentum, and energy conservation*: Given a closed system, conservation of these properties is essential for accurate results, especially when collisions are implemented.

(3) *At least first order accurate*: The accuracy of the entire solver should at least be first order to reduce numerical error and obtain a converged solution.

A finite volume method is a technique for solving a PDE by solving for average cell (i.e. the volume around a node point on a mesh) properties. These cell-averaged properties are obtained by calculating the flux entering and leaving the cell on all sides of the volume. Consider a simple, one-dimensional (1D) linear advection equation,

$$\frac{\partial q}{\partial t} + u\frac{\partial q}{\partial x} = 0 \tag{2.13}$$

where $q = q(x)$ is the conserved quantity, and u is the characteristic velocity, which is constant for a given point in time. This equation can be discretized as

$$\frac{Q_i^{n+1} - Q_i^n}{\Delta t} = -\frac{1}{\Delta x}(F_{i+1/2} - F_{i-1/2}) \tag{2.14}$$

where Q is $q(x)$ averaged across one cell's length, F is the discretized flux, Δt is the time step, and Δx is the discretized cell size. n denotes the time index, and i denotes the cell number, where $i + 1/2$ and $i - 1/2$ correspond to the right and left cell faces. The flux entering cell i through its left face is identical to the flux leaving cell $i-1$'s right face. Thus, finite volume methods inherently satisfy conservation of mass.

Each cell flux F is represented by a function which must be monotonic (either entirely nonincreasing or nondecreasing). Failure to preserve monotonicity can result in numerical extrema (i.e. undershoots and overshoots) within one cell. These extrema can lead to spurious oscillations in the numerical solution, which may lead to unphysical quantities such as negative density. However, Godunov's theorem [82] states that if a linear numerical scheme is monotonic, it can be at most first-order accurate. Thus, a nonlinear scheme must be used to limit numerical extrema while achieving higher order accuracy.

DK performs advection to second-order accuracy using the monotonic upstream-centered scheme for conservation laws (MUSCL) with a van Leer flux (or slope) limiter adapted for non-uniform grids [83] and a forward Euler time integration scheme. The flux at cell interface $i + 1/2$ is written as

$$F_{i+1/2} = C\left(Q_i + \frac{1.0 - |C|d}{2}(Q_{i+1} - Q_i)\frac{\Psi}{b}\right) \tag{2.15}$$

where C is the CFL number (defined in Equation 2.22), Ψ is the van Leer slope limiter, given by

$$\Psi = \texttt{max}\left[0.0, \texttt{min}(2.0, \Psi')\right] \tag{2.16}$$

$$\Psi' = \texttt{min}\left[2\theta, (1 + |C|d)\frac{\theta - 1}{3} + 1\right] \tag{2.17}$$

$$\theta = \frac{U_i - U_{i-1}}{U_{i+1} - U_i} \tag{2.18}$$

and d and b are constants given by

$$d = \begin{cases} \Delta x_{i+1}/\Delta x_i \ for\ u > 0 \\ \\ 1.0 \text{ for } u < 0 \end{cases} \tag{2.19}$$

$$b = \frac{2\Delta x_i}{\Delta x_i + \Delta x_{i+1}} \tag{2.20}$$

$$\tag{2.21}$$

The combination of these schemes result in a solution that is total variation diminishing. For the Boltzmann equation, referring to Equation 2.13, the VDF in units of particles/($\text{m}^3\cdot$ $(\text{m/s})^3$) is the conserved quantity q and the characteristic velocity is the bulk velocity or acceleration of the cell.

When solving most hyperbolic PDEs with an explicit time integration scheme, the Courant –Friedrichs–Lewy (CFL) condition must be fulfilled. The principle of the CFL condition is that a flow with velocity v_x across a configuration space grid cell of size Δx must have its time step bounded by the following relation for numerical stability:

$$\text{CFL}_x = \left|\frac{v_x \Delta t}{\Delta x}\right| < 1.0 \tag{2.22}$$

In DVM, this condition must be fulfilled for every discrete value of v_x throughout the velocity domain. If charged particles are advected in the presence of an electric field, their time step is

additionally bounded by

$$\mathrm{CFL}_{v_x} = \left| \frac{q_s}{m} \frac{E_x \Delta t}{\Delta v_x} \right| < 1.0 \tag{2.23}$$

where q_s is the species charge, E_x is the electric field in the direction of advection, and Δv_x is the velocity space grid cell size. For lightweight electrons especially, this time step can be extremely small and thus computationally expensive if applied globally. Thus, it is common to implement sub-stepping, where a smaller time step is used to advect specified species, and a larger global time step is used for all other species. The only requirement is that the global time step must be divisible by the sub-step. In DK, Equation 2.23 is checked at every time step, and the sub-step is automatically adjusted if the criterion is not met.

2.4.2 Boundary Conditions

Boundary conditions for solving Equation 2.2 must be considered in terms of configuration space and velocity space. The DK code uses the ghost cell method to apply boundary conditions, where values of the VDF are assigned to two additional cells (known as ghost cells) outside the domain according to the boundary condition desired. For velocity space, the domain range must be sufficiently wide such that the values of the VDF are effectively zero in the last few cells; else, the domain is omitting particles with large velocities, and the flow temperature may be underpredicted. If so, then either a Dirichlet or Neumann boundary condition may be applied to the velocity space ghost cells– both lead to identical solutions as long as the last VDF cells are sufficiently close to zero. For configuration space, the boundary conditions require more consideration. In the following sections, the first physical cell of a domain is labeled i, such that the ghost cells are labeled $i-1$ and $i-2$, and the last physical cell of a domain is labeled n, such that the ghost cells are labeled $n+1$ and $n+2$.

2.4.2.1 Periodic

Periodic refers to two opposing sides of a domain being continuous, such that flow entering one side emerges from the other unchanged. The ghost cell VDFs are set as follows:

$$f_{i-1}(\vec{v}) = f_n(\vec{v})$$

$$f_{i-2}(\vec{v}) = f_{n-1}(\vec{v})$$

$$f_{n+1}(\vec{v}) = f_i(\vec{v})$$

$$f_{n+2}(\vec{v}) = f_{i+1}(\vec{v})$$

2.4.2.2 Inflow

To simulate an inflow boundary condition in the direction $+\hat{x}$, the following VDF is applied to both ghost cells,

$$f_{in}(\vec{v}) = \begin{cases} f_i(\vec{v}) \text{ for } v_x < 0 \\ \\ f_{eq}(\vec{v}, \vec{u}_\infty, T_\infty) \text{ for } v_x > 0 \end{cases} \tag{2.24}$$

where $f_{eq}(\vec{v}, \vec{u}_\infty, T_\infty)$ is initialized according to the desired properties of the inflow. The negative side of the VDF is extrapolated to the zeroth order.

2.4.2.3 Outflow

For an outflow boundary condition in the direction $+\hat{x}$, the following VDF is applied to both ghost cells,

$$f_{out}(\vec{v}) = \begin{cases} 0 \text{ for } v_x < 0 \\ \\ f_n(\vec{v}) \text{ for } v_x > 0 \end{cases} \tag{2.25}$$

This simulates zero back-flow into the domain.

2.4.2.4 Wall

For a wall boundary condition, the accommodation coefficient α determines the fraction of particles that are specularly or diffusely reflected, with $\alpha = 1$ requiring all particles reflect diffusely. For a specular wall on the right hand side of a domain, the following VDF is applied to both ghost cells,

$$f_{spec}(\vec{v}) = \begin{cases} f_n(+v_x, v_y, v_z) \text{ for } v_x < 0 \\ f_n(\vec{v}) \text{ for } v_x > 0 \end{cases} \tag{2.26}$$

such that all particles with velocity v_x are reflected with velocity $-v_x$ and their perpendicular velocity components remain unchanged. For a diffuse wall in the $+\hat{x}$ direction, the following VDF is applied to both ghost cells,

$$f_{diff}(\vec{v}) = \begin{cases} f_{eq}(\vec{v}, \vec{u} = \vec{0}, T_{wall}) \text{ for } v_x < 0 \\ f_n(\vec{v}) \text{ for } v_x > 0 \end{cases} \tag{2.27}$$

where f_{eq} is normalized according to the incoming flux of particles from cell n with velocity $v_x > 0$. To account for α, the following VDF is applied to both ghost cells,

$$f_{wall}(\vec{v}) = \texttt{max}[0, (1-\alpha)f_{spec} + \alpha f_{diff}] \tag{2.28}$$

Chemical reactions with the wall must also be considered. Since at present DK only has monatomic species, only three reactions must be considered: all neutrals will reflect, ions will recombine and reflect from the wall as neutrals, and all electrons will absorb into the wall (i.e., $f_{wall}(\vec{v}) = 0$), as they are assumed to recombine with ions.

2.4.3 Electric Field Modeling

Charged particles interact with electromagnetic fields through the Lorentz force $\vec{F}_{L,s}$, which for non-relativistic particles, is defined as

$$\vec{a} = \frac{\vec{F}_{L,s}}{m_s} = \frac{q_s}{m_s}(\vec{E} + \vec{v}_s \times \vec{B}) \tag{2.29}$$

where $\vec{E}$ and $\vec{B}$ are the electrostatic and magnetic field vectors, respectively. At present, DK does not account for magnetic field effects. DK has four modes for accounting for electrostatic effects. These modes can be accessed with the flag `plasmaPropag`.

The first mode is the 'free diffusion approximation', which assumes any bulk self-induced electric fields are negligible, and DK does not solve the Poisson equation. Ions and electrons retain separate VDFs and macroscopic properties. If Coulomb collisions are turned on with the free diffusion approximation, then strictly speaking, one is accounting for electrostatic effects between particles. If applicable, electrons are sub-stepped.

The second mode invokes the ambipolar diffusion approximation. Ions and electrons are treated as one ambipolar particle with identical macroscopic properties. The ambipolar particle is transported as though it were an ion; that is, the ambipolar particle VDF is advected according to ion properties, and is modified according to ion collisions and reactions. Other particles can still collide and react with electrons (and any rates are calculated with respect to true electron properties, such as mass), but these electrons are assumed to have the macroscopic properties of the ambipolar particles.

For the remaining two modes, the electric field vector is calculated with

$$\vec{E} = -\vec{\nabla}\phi \tag{2.30}$$

where the electric potential ϕ is given by Poisson's equation for electrostatics,

$$\nabla_x^2 \phi = -\frac{\rho_c}{\epsilon_0} \tag{2.31}$$

where ρ_c is the charge density. Using a finite difference second-order discretization with a 3-point stencil, Equation 2.31 can be set up as a linear system, $A\vec{x} = \vec{f}$, where

$$A = \frac{1}{h^2}\begin{bmatrix} \alpha_1 & \alpha_2 & & & & & \\ 1 & -2 & 1 & & & & \\ & 1 & -2 & 1 & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & 1 & -2 & 1 & \\ & & & & 1 & -2 & 1 \\ & & & & & \beta_1 & \beta_2 \end{bmatrix}, \vec{x} = \begin{bmatrix} \phi_0 \\ \phi_1 \\ \phi_2 \\ \vdots \\ \phi_{m-1} \\ \phi_m \\ \phi_{m+1} \end{bmatrix}, \vec{f} = -\frac{e_c}{\epsilon_0}\begin{bmatrix} \alpha_L \\ (n_i + 2n_{ii} - n_e)_1 \\ (n_i + 2n_{ii} - n_e)_2 \\ \vdots \\ (n_i + 2n_{ii} - n_e)_{m-1} \\ (n_i + 2n_{ii} - n_e)_m \\ \beta_R \end{bmatrix} \tag{2.32}$$

where h is the grid width, and α and β are constants relating to the left and right side boundary conditions, respectively, with values listed in Table 2.1 [84].

Table 2.1: Additional matrix and vector entries for Poisson equation boundary conditions.

Boundary condition	α_1	α_2	α_L	β_1	β_2	β_R
Dirichlet (fixed ϕ)	h^2	0	ϕ_L	0	h^2	ϕ_R
Neumann (fixed $-d\phi/dx$)	$-h_L$	h_L	$-d\phi_L/dx$	$-h_R$	h_R	$-d\phi_R/dx$

For a non-uniform grid, if the band values 1, -2, and 1 are labeled a_i, b_i, c_i, then $-1/h^2$ is modified according to the band value it is being multiplied by with Equation 2.33,

$$\frac{1}{h^2} = \begin{cases} 2/(h_L(h_L + h_R)) & \text{for } a_i \\ 1/(h_L h_R) & \text{for } b_i \\ 2/(h_R(h_L + h_R)) & \text{for } c_i \end{cases} \tag{2.33}$$

where h_L is the average of the left adjacent and current cell widths, and h_R is the average of the right adjacent and current cell widths. Periodic Poisson equation boundary conditions are also implemented in DK, where the potential on one side of the domain is set equal to the opposite side.

To solve Equation (2.32), the Portable, Extensible Toolkit for Scientific Computation's (PETSc) linear system solver is used with a LU factorization direct solver as its pre-conditioner [85]. The linear solve is parallelized across all processors. With the potential values, a second order central

difference derivative is used to calculate $\vec{E}$ from Equation 2.30. If a Neumann boundary condition is used, $\vec{E}$ at that boundary is set to the desired $-d\phi/dx$.

The third `plasmaPropag` mode solves for the electric field with the process above, and sub-steps electrons. The Lorentz force is applied by advecting through velocity space in the direction of the electric field. Finally, the fourth mode solves for the electric field and sub-steps both ions and electrons with Δt_e. This mode is recommended when electric fields are particularly strong, where the CFL number in velocity space for ions may exceed 1.0, or when charge separation plays an important role in the flowfield, as ions have better opportunity to catch up with electrons. When electrostatic effects are accounted for, Equation 2.2 is solved with Strang splitting [86] for ions and electrons, such that configuration space and velocity space advection are solved seperately.

2.4.4 Collisions

DK utilizes an explicit–implicit (IMEX) time-stepping scheme where advection is solved explicitly and collisions are solved implicitly [87]. Implicit handling of collisions does not require resolution of mean relaxation time; this is advantageous for stiff collision source terms, which often arise in highly collisional flows. The overall time step is still limited by the CFL condition, however for typical hypersonic flows, this time step is greater than that required for implicit handling of collisions.

DK uses the BGK operator throughout this work. Justification of this choice is addressed in each study. In general, should one attempt to use the methodology of this dissertation for experimental validation or vehicle design, a more rigorous collision operator is recommended. Polyatomic formulation of collision operators is possible to implement [88, 89] but is left for future work.

2.4.4.1 Momentum Exchange (MEX) Collisions

Collision cross sections for elastic momentum exchange (MEX) collisions are modeled with either the hard-sphere (HS) or variable hard-sphere (VHS) model [29]. For the HS model, the collision time $\tau(x)$ for species A colliding with species B is given by

$$\tau_{A-B}^{HS}(x) = \frac{v_{th,A}}{\lambda_{mfp,A-B}^{HS}} \tag{2.34}$$

with

$$\lambda_{mfp,A-B}^{HS} = \frac{1.0}{\sqrt{2}\pi d_{HS}^2 n_B} \tag{2.35}$$

where $d^{HS} = (d_A + d_B)/2$ refers to the hard-sphere diameter. There are multiple definitions of thermal velocity, all of which are within $(1.6 \pm 0.2)\sqrt{k_B T/m}$. In this work, for one dimension in configuration space, one dimension in velocity space (1D1V) flows, v_{th} describes the root mean square of the velocity in any one dimension, given by

$$v_{th}^{1D1V} = \sqrt{\frac{k_B T}{m}} \tag{2.36}$$

and for one dimension in configuration space, three dimensions in velocity space (1D3V) flows, v_{th} describes the most probable speed,

$$v_{th}^{1D3V} = \sqrt{\frac{2k_B T}{m}} \tag{2.37}$$

From the VHS model, the collision time is given by

$$\tau_{A-B}^{VHS}(x) = \frac{\mu_A^{VHS}}{m_A n_B R_{A-B} T_A} \tag{2.38}$$

where R_{A-B} is the specific gas constant of the mixture of species A and B, and μ^{VHS} is the viscosity simulated by the VHS model. μ^{VHS} is given by

$$\mu^{VHS} = \mu_{ref}^{VHS} \left(\frac{T}{T_{ref}} \right)^{\omega} \tag{2.39}$$

where ω is equal to the power-law exponent $\nu + 1/2$ for the species.

For neutral-ion and ion-neutral collisions, it is standard practice to use charge exchange (CEX) cross sections for MEX collisions. To obtain a viscosity value for DK, in the equation for

the standard binary coefficient of viscosity, $\sigma_\mu(g)$ is set equal to the CEX cross section expression from Reference [90],

$$\mu = \frac{\frac{5}{8}\sqrt{2\pi m_r k_B T}}{\left(\frac{m_r}{2k_B T}\right)^4 \int_0^\infty g^7 \sigma_\mu(g) e^{-m_r g^2/(2k_B T)} dg} \tag{2.40}$$

where m_r is the reduced mass, and g is the relative speed. Then, the derived μ is substituted into Equation 2.38. A full derivation for Equations 2.40 and 2.39 can be found in Chapters 5 and 6 of Reference [55], respectively.

2.4.4.2 Charge Exchange (CEX) Collisions

A CEX collision transfers an electron from a slow moving neutral to a fast moving ion, resulting in a fast moving neutral and a slow moving ion, written as

$$A_{slow} + A^+_{fast} \rightarrow A_{fast} + A^+_{slow} \tag{2.41}$$

In hypersonics, CEX collisions serve to equilibrate ions and neutrals. CEX cross sections are obtained from Reference [90]. To perform charge exchange, DK follows the procedure utilized by the open-source plasma physics simulation framework Gkeyll [91, 92],

$$\frac{df_i}{dt} = \sigma_{CEX} V_{CEX} (n_i f_n - n_n f_i), \tag{2.42}$$

$$\frac{df_n}{dt} = -\sigma_{CEX} V_{CEX} (n_i f_n - n_n f_i) \tag{2.43}$$

where V_{CEX} is given by

$$V_{CEX} \equiv \sqrt{\frac{4}{\pi} v^2_{th,i} + \frac{4}{\pi} v^2_{th,n} + v^2_{in}} \tag{2.44}$$

$$v_{in} \equiv |\vec{u}_i - \vec{u}_n| \tag{2.45}$$

2.4.4.3 Coulomb Collisions

A Coulomb collision occurs when charged particles scatter from the influence of other particles' electrostatic fields. These individual Coulomb forces are considered separate from the bulk electrostatic field that develops due to the collective nature of the entire plasma. In DK, collision times for Coulomb collisions are calculated according to the expressions presented in the NRL Plasma Formulary [93], where when both species are near Maxwellian with $T_i \lesssim T_e$ and $Z = 1$, there are only two characteristic collision rates:

$$\nu_e = 2.9 \times 10^{-6} n ln(\Lambda) T_e^{-3/2} \tag{2.46}$$

$$\nu_i = 4.8 \times 10^{-8} n ln(\Lambda) \left(\frac{m_i}{m_p}\right)^{-1/2} T_i^{-3/2} \tag{2.47}$$

where m_p is the mass of a proton. The value of the Coulomb logarithm, $ln(\Lambda)$ is kept constant at 10: a standard practice in plasma physics [27]. The impact of using a constant value versus the full expression is left for future work. However, because collisions with neutrals dominate in this work, the impact of the Coulomb logarithm on flow profiles is likely minimal.

2.4.5 Electron-Impact Ionization (EII)

The reaction for EII is as follows

$$e^- + n \rightarrow 2e^- + i^+ - E_{EII} \tag{2.48}$$

where E_{EII} is the ionization energy of the reaction, and n may be in either ground or an excited state. During EII in DK, the neutral and reactant electron are removed, an ion is generated with the original momentum of the neutral, and two low energy electrons are generated. This assumes the electron is released isotropically in the neutral frame of reference, and the total kinetic energy of the electrons is reduced by the ionization activation energy while the ion retains the momentum and energy of the neutral. The low energy electrons follow a Maxwellian distribution $\hat{f}_{eq,EII}(u_n, T_{iz})$,

where $\hat{f}$ indicates that the VDF is normalized to unity, and (u_n, T_{iz}) indicates the VDF is shifted by the bulk neutral velocity u_n and initialized according to temperature T_{iz}. To get T_{iz}, starting with conservation of energy,

$$\begin{aligned} 2\left(\frac{3}{2}k_B T_{iz}\right) &= \frac{3}{2}k_B T_e - E_{EII} + \frac{m_e u_e^2}{2} - 2\left(\frac{m_e u_n^2}{2}\right) \\ T_{iz} &= \frac{T_e}{2} - \frac{E_{EII}}{3k_B} + \frac{m_e u_e^2}{6k_B} - \frac{m_e u_n^2}{3k_B} \\ &= \frac{T_e}{2} - \frac{T_{EII}}{3} + \frac{m_e u_e^2}{6k_B} - \frac{m_e u_n^2}{3k_B} \end{aligned} \tag{2.49}$$

where the 2's on the first equation line account for two post-reaction electrons [94]. To avoid unphysical results, if T_{iz} is negative, no EII is performed for the cell. The species VDFs are then altered according to

$$\frac{df_e}{dt} = n_n n_e k_{EII}(2\hat{f}_{eq,EII}(u_n, T_{iz}) - \hat{f}_e) \tag{2.50}$$

$$\frac{df_i}{dt} = n_n n_e k_{EII}\hat{f}_n \frac{dv_n^3}{dv_i^3} \tag{2.51}$$

$$\frac{df_n}{dt} = -n_n n_e k_{EII}\hat{f}_n \tag{2.52}$$

2.4.6 Parallel Computation and Flow Diagram

DK is parallelized with the Message-Passing Interface (MPI) across both nodes and processors. At present, DK only has Cartesian grids for configuration space, which are equally divided between processors. Two ghost cells per direction are used to exchange information between processors. A strong scaling analysis of DK's 1D1V and 1D3V dimensionalities is presented in Section 2.4.7.

A general flowchart for DK with both ions and electrons sub-stepped is shown in Figure 2.2. The global timestep is Δt, and the sub-step is Δt_s.

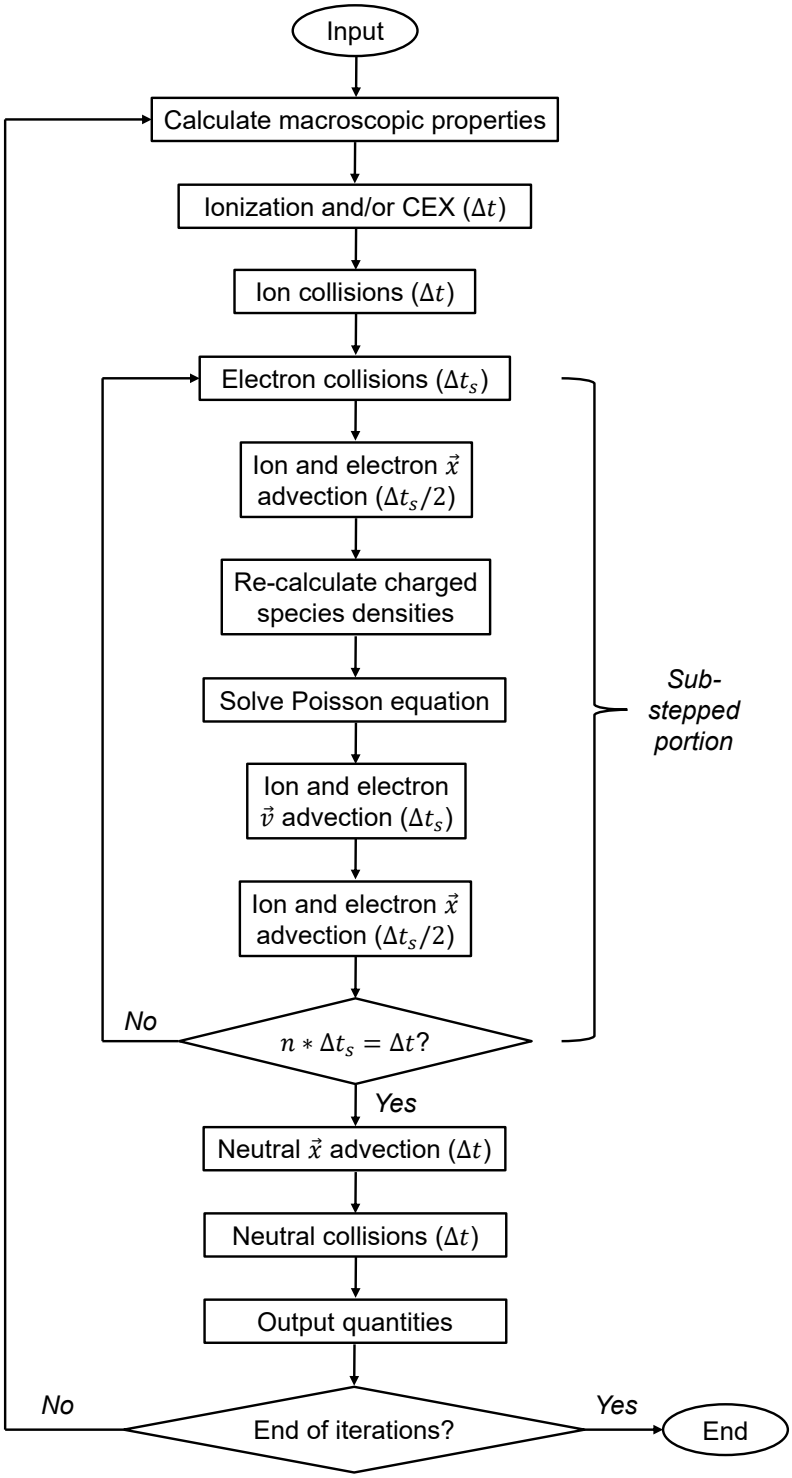

Figure 2.2: Direct Kinetic code flow chart.

2.4.7 Verification of 1D3V: Sod Shock Tube

The classic gas dynamics problem, the Sod shock tube [95], is used to verify DK's 1D3V dimensionality as it tests both advection and collisionality. The problem is initialized as a 1D domain with walls on either side. Maxwellian gases with different properties are initialized on either side of a virtual diaphragm, which is removed at $t = 0$. A shock, contact discontinuity, and rarefaction fan will develop and propagate as the simulation evolves. In this section and throughout this dissertation work, when root mean square error (RMSE) is reported, it is normalized to the range of the data.

Table 2.2: Simulation input parameters for Sod shock tube verification problem.

Parameter	Value
Length of domain, m	0.02
Grid points, configuration space	1000
Diaphragm location, m	0.01
Width of velocity domain, m/s	$\pm$7492
Grid points, velocity space	160
Time step, s	10^{-9}
Total time steps	5000
Driver gas density, m^{-3}	10^{25}
Driver gas temperature, K	1350
Driven gas density, m^{-3}	1.25×10^{24}
Driven gas temperature, m	1080

Table 2.2 lists the parameters used in initializing the DK simulation. Results are compared with exact solutions of the Riemann problem evaluated numerically with Python. Figures 2.3 and 2.4 plot normalized density and velocity profiles of both solutions after 5 μs have elapsed (5,000 time steps), both of which demonstrate excellent agreement with RMSE = 5.99×10^{-2} for density and RMSE = 12.2 for velocity. In the density profile (Figure 2.3), numerical diffusion is visible at the sharp corners of the analytical solution, especially around the contact discontinuity. This is in part due to the van Leer flux-limiting scheme preferring smoothness in regions of steep gradients. Numerical diffusion is also visible at the corners of the velocity profile and along the shock front. The numerical diffusion at the shock front is the largest contributor to the RMSE,

as the analytical solution expects a velocity of zero for five grid points (Figure 2.4c). A numerical oscillation is present in the velocity profile at the contact discontinuity (Figure 2.4b), however its magnitude is less than 2 m/s. Both numerical effects can be minimized with increasing resolution in configuration space. Increasing resolution in velocity space beyond the parameters listed in Table 2.2 does not improve solution quality, however reducing velocity space resolution beyond a point degrades quality of results. Ultimately, for the purposes of verifying the 1D3V dimensionality, these results provide confidence in implemented features.

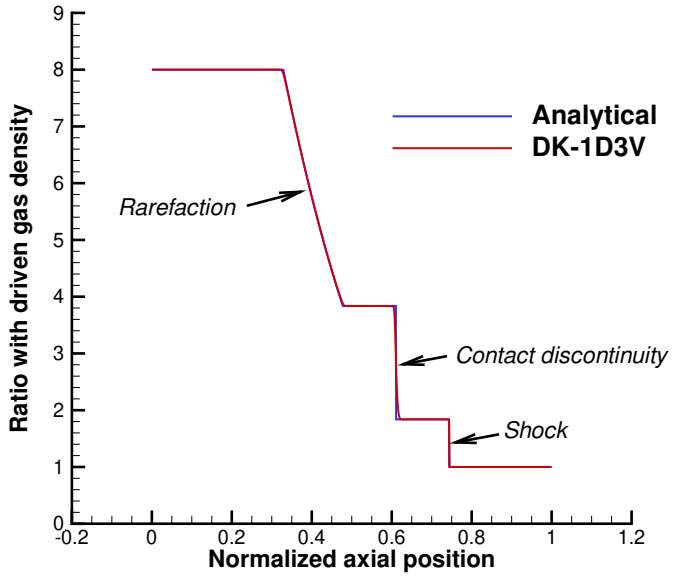

Figure 2.3: Density profile (normalized to the driven gas density) from DK-1D3V and the analytic solution for the Sod shock tube problem at $t = 5$ μs.

The Sod shock tube problem is also used to perform a strong scaling analysis on 1D1V and 1D3V DK. Strong scaling refers to the efficiency of a parallelized code when the problem size is kept fixed and the number of processors (or tasks) is varied. This test is performed on the University of Colorado Boulder's Alpine high-performance computing cluster on AMD EPYC nodes, each of which have 64 processors. When multiple nodes are used, the job is set up such that an equal number of tasks per node are assigned. The `-O3, -ipo` and `-qopenmp` optimization flags

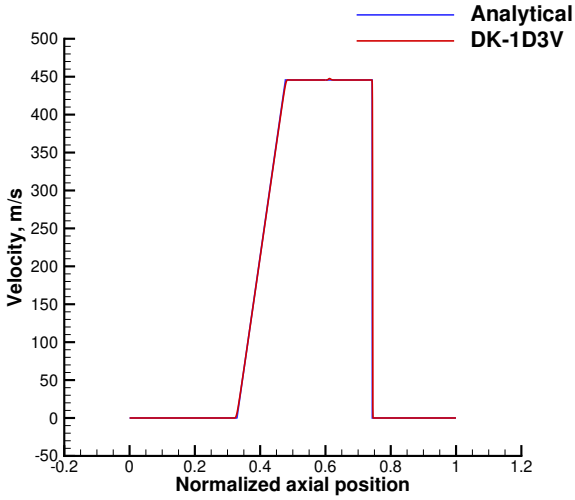

(a) Entire profile.

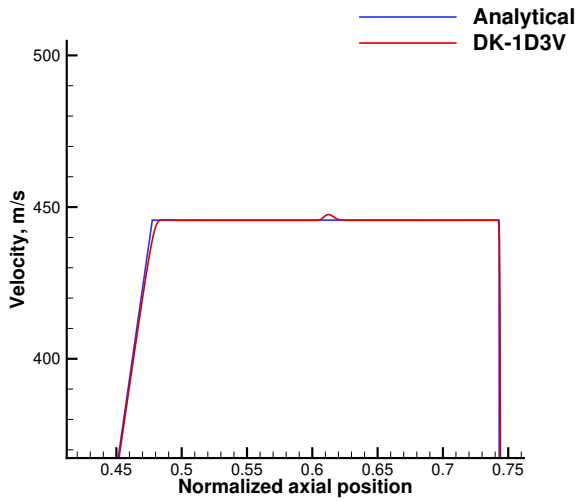

(b) Oscillation in velocity profile at the contact discontinuity.

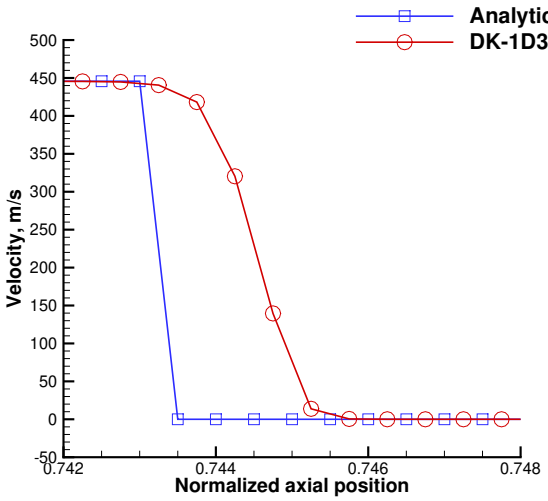

(c) Numerical diffusion in velocity profile at the shock front.

Figure 2.4: Velocity profile from DK-1D3V and the analytic solution for the Sod shock tube problem at $t = 5$ μs.

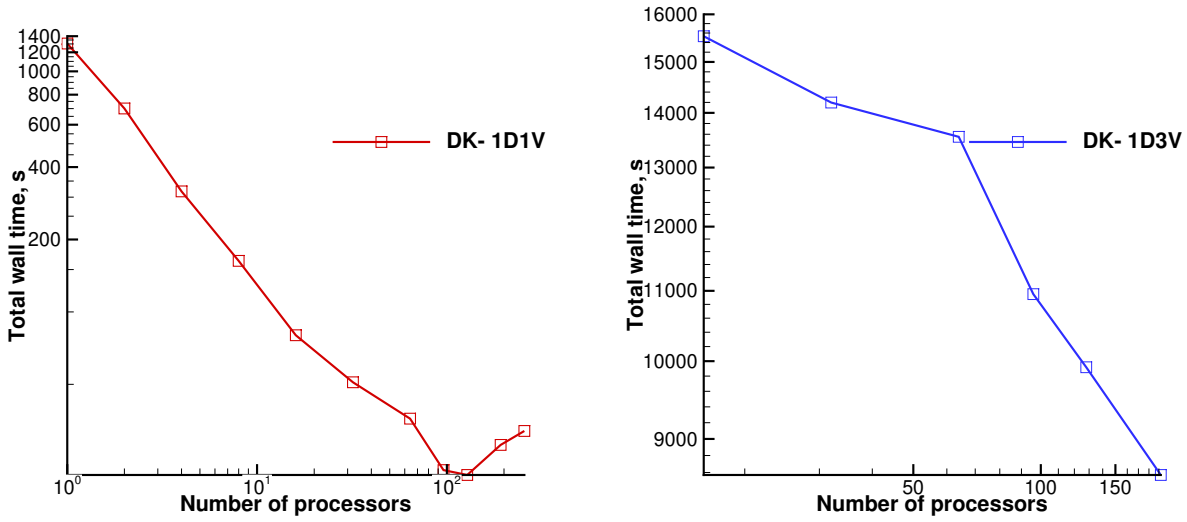

(a) Profile for 1D1V.

(b) Profile for 1D3V.

Figure 2.5: Strong scaling analysis of CPU time (in seconds) versus number of processors.

are used with Intel compilers and the OpenFabrics Interface flag `I_MPI_FABRICS=ofi` is specified during execution*. The simulation parameters of Table 2.2 are also used for this analysis, with two exceptions: in 1D1V, the number of cells in configuration space is increased to 10,000 and the time step is decreased to 10^{-10} (for a total of 50,000 time steps), and in 1D3V the number of cells in configuration space is decreased to 500.

For 1D1V DK, Figure 2.5a plots the total CPU time versus the number of processors (varied from 1 to 256) on a log-log scale. Diminishing returns begin at 96 processors, corresponding to about 104 cells per processor. For 1D3V DK, Figure 2.5b plots the same metric, except the number of processors is varied from 16 to 192†. There is a change in slope at 64 processors, which indicates that DK scales better when inter-node communication is utilized rather than exclusively intra-node.

Ultimately, both results demonstrate that DK scales well in parallel: with an exponential increase in the number of processors, an exponential reduction in CPU time is achieved. However,

*DK's performance degrades when Shared-Memory (SHM) is enabled.

†Below 16 processors, there is insufficent memory to run the simulation. The average number of cells per processor is 2.6 for 192 tasks; thus, 256 processors were not tested as any simulation with less than 2 cells per task is highly impractical.

the increase in computational cost between 1D1V and 1D3V is apparent, as 1D3V takes roughly 3 orders of magnitude more wall time than 1D1V. Regardless of the physics being simulated, 1D3V DK is bottlenecked by the time it takes to cycle through three dimensions in velocity space for each routine in the simulation (e.g., advection, collisions between each pair of species). Thus, for most practical applications, 1D3V DK must run in parallel, and optimally across multiple nodes.

2.5 The MPIC Code

This work also utilizes MONACO–PIC (MPIC), a three-dimensional coupled DSMC-PIC code developed at the University of Michigan [61] and re-factored at the University of Colorado Boulder [96]. Advection of neutral particles and collision dynamics of all species are handled with DSMC, and advection of charged species is handled with PIC. MPIC can optionally run in hybrid particle-fluid mode, such that heavy species are modeled as DSMC-PIC particles while electrons are modeled as a fluid. However, this work will exclusively use an electron particle model. A more comprehensive overview of MPIC in its current state can be found in Reference [96]. The following section will describe the features and code options relevant to this work.

2.5.1 Electron Particle Model

The 'Electron Particle Model' in MPIC simply initializes electrons as negatively charged DSMC-PIC particles and obtains the electrostatic potential by solving Poisson's equation for electrostatics. To do so, MPIC utilizes a compact finite-element Laplacian [97]. In DK, Equation 2.31 is solved as a standard linear matrix system. Because DK parallelizes the Poisson solve and only operates on structured grids, this method of solving the Laplacian is reasonable. However, MPIC can operate on unstructured grids and does not parallelize the Poisson solve across nodes (only among cores on a single node). Therefore, a more general and efficient Laplacian solve is desireable.

Using the divergence theorem on a control volume V with boundary S, the Laplacian can be re-written in discrete-form as

$$\nabla^2\phi = \nabla \cdot (\nabla\phi)$$

$$\int_{V_i} \nabla \cdot (\nabla\phi) dV = \int_S \nabla\phi \cdot \vec{n} dA$$

$$(\nabla^2\phi)_i \approx \frac{1}{V_i} \sum_{i=1}^{n} (\nabla\phi)_i \cdot \vec{n} A_i \tag{2.53}$$

Thus, the discrete Laplacian at cell i can be obtained by summing the gradients normal to each face. In doing so, only the cell neighbors are required for a solution; the entire domain's charge density does not need to be stored in one large matrix. This does not come at the loss of accuracy: compact Laplacian schemes can achieve accuracy at fourth order or higher. Similar derivation steps can be performed to obtain a compact finite-element gradient operator.

Sub-stepping of electrons, ions, or both is implemented in MPIC. There is also the option to invoke the ambipolar diffusion approximation. In MPIC, the ambipolar diffusion approximation propagates ions and electrons with the ion bulk velocity and assumes they have identical positions in configuration space (enforcing exact local charge neutrality). This is implemented by assigning each ion a set of lightweight electron velocity components in addition to its own velocity components. During collisions, each ion is split into two particles: an ion and an electron, the latter of which is assigned the stored electron velocity components. The separated particles are considered for collision pairing independently. Once collisions are completed, each ion is re-paired with an electron and the updated electron velocity components are copied into the ion's storage. The electron particles are then deleted, such that only ions and neutrals remain during advection. This approach is similar to that used in the open-source DSMC solver SPARTA [32]. When using the ambipolar diffusion approximation in MPIC, the outputted electron number density is equal to the ion number density. However, the stored electron velocity components can be used to output an 'ambipolar electron temperature' separate from the ion temperature.

2.5.2 Collisions

MPIC uses the No-Time-Counter (NTC) algorithm [29] to select collision pairs. At every time step and configuration space cell, the maximum number of particle pairs expected to collide is given by

$$N_{coll,max} = \sum_i \sum_j \frac{1}{2} n_i (N_j - \delta_{ij}) [\langle \sigma(d,g) g \rangle]_{max} \Delta t \tag{2.54}$$

where i and j denote each species in the gas mixture, N_j is the number of macroparticles of species j, and $[\langle \sigma(d,g) g \rangle]_{max}$ is the maximum encountered value of $\langle \sigma(d,g) g \rangle$, which is the product of the collision cross section σ and the average relative velocity of the collision pair g (calculated per cell). Only an integer number of particle pairs can be tested. Additionally, the number of pairs must be at least one, and cannot exceed half the number of existing macroparticles. Therefore, the true number of particle pairs to test is*

$$N_{coll,test} = \texttt{floor}[N_{coll,max} + 0.5] \tag{2.55}$$

$$\text{with the following enforced: } 1 \leq N_{coll,test} \leq \texttt{floor}[N/2]$$

After shuffling the order of particles in the cell, $N_{coll,test}$ particle pairs are created, where each particle can only be included in one collision pair. Each pair is then tested for a collision if it meets the following criterion,

$$\texttt{RAND} < P_{coll}$$
$$P_{coll} = \frac{\sigma(d,g) g}{[\langle \sigma(d,g) g \rangle]_{max}} \cdot \frac{N_{coll,max}}{N_{coll,test}} \tag{2.56}$$

In this work, when MPIC is used, momentum exchange during collisions is modeled with VHS. Collisions between neutrals utilize the VHS cross section. Collisions between neutrals and ions

*Note that the MPIC source code uses 0.47 instead of 0.5 in Equation 2.55. This leaves a few extra particles available for three-body reactions.

(both MEX and CEX) utilize the CEX cross section. For neutral-electron and electron-neutral collisions, the cross section is interpolated from data based on the electron energy. In this work, cross section data is retrieved from the IST-Lisbon database via `lxcat.net` [98, 99] and fitted piecewise (see Appendix A). Coulomb collisions between charged particles are not included in MPIC simulations.

2.5.3 Ionization

For a monatomic flow, ionization via a shock wave occurs in two steps. As the incoming flow is neutral, ionization begins with atom-atom impact ionization (AAI). The reaction for AAI is described as a two-step process,

$$\begin{aligned} n + n &\rightarrow n + n^* \\ n + n^* &\rightarrow n + i^+ + e^- - E_{AAI} \end{aligned} \tag{2.57}$$

where n^* represents an excited neutral atom. Once a sufficient population of high energy electrons has been generated, electron-impact ionization (EII) becomes the dominant mechanism for ionization (Equation 2.48). Both atom-atom impact and electron-impact ionization may occur in a 'ladder-climbing' fashion: where an atom undergoing collisions may become excited into any number of electronic excitation states depending on the energy of the collision. The ionization threshold from these excited states is increasingly smaller than the ionization threshold from ground-state, making the probability of ionization higher for excited states. At high entry velocities, ladder-climbing of excitation states becomes the dominant mechanism for ionization rather than direct ionization from ground state. A computational method which resolves each excited state would best capture the ladder-climbing ionization process, though with significant computational expense. Instead, modeling approaches which do not seek detailed study of chemical kinetics use a global ionization rate coefficient dependent on the temperature or energy of the reactants. This ionization rate coefficient accounts for ionization directly from ground state and all excited states.

For EII and AAI, the modified Arrhenius rate equation is used to obtain probabilities of ionization. The modified Arrhenius form is given by

$$k(T)\ [\mathrm{m}^3] = AT^{\eta} exp(-E_a/k_B T) = AT^{\eta} exp(-T_a/T) \tag{2.58}$$

where E_a is the activation energy of the reaction and $T_a = E_a/k_B$. The Arrhenius coefficients used for each reaction can be found in Table A.4.

2.5.3.1 The Total Collision Energy Model

If a particle pair is selected for collision, the pair is additionally tested for chemical reactions. MPIC uses the phenomenological Total Collision Energy (TCE) model [29], which converts the modified Arrhenius rate form to a probability of reaction. Using the VHS model, the probability of a chemical reaction is given by

$$P_{reac} = C_1 \frac{(E_{coll} - E_a)^{\eta+\bar{\zeta}+1/2}}{E_{coll}^{\bar{\zeta}+3/2-\omega}} \tag{2.59}$$

$$C_1 = A \frac{\sqrt{\pi}(1+\delta_{AB})}{2\pi d_{ref}^2} \frac{\Gamma(\bar{\zeta}+5/2-\omega)}{\Gamma(\bar{\zeta}+\eta+3/2)} \sqrt{\frac{m_r}{2k_B T_{ref}}} \frac{T_{ref}^{1-\omega}}{k_B^{\eta-1+\omega}} \tag{2.60}$$

where E_{coll} is the total collision energy (given by the sum of the relative translational energy of the collision pair and the energies of all internal modes of both particles), $\bar{\zeta}$ is the number of internal degrees of freedom of both particles, and δ_{AB} is the Kronecker delta. The rest of this section describes the TCE procedure for a two-body ionization reaction for monatomic gases (i.e., via EII or AAI).

Consider EII, where the particle pair consists of one electron and one neutral. The center-of-mass velocity between the reacting pair is calculated as

$$W_1 = \frac{m_e}{m_e + m_n}\vec{v}_e + \frac{m_n}{m_e + m_n}\vec{v}_n \tag{2.61}$$

The post-collision relative velocity must incorporate the subtraction of activation energy E_a involved in the reaction, and thus is given by

$$g_{rel,1} = \sqrt{\frac{2(E_{coll} - E_a)}{m_r}} \tag{2.62}$$

The standard VHS collision procedure is used to update the velocities of the reacting pair using W_1 and $g_{rel,1}$. In the case of EII, the reacting electron has completed its update, and the neutral must now be ionized into an ion and electron. Functionally, this is accomplished by performing a second collision between the neutral and a new electron with zero velocity. The new center-of-mass velocity is given by

$$W_2 = \frac{m_e}{m_e + m_n}\cancelto{0}{\vec{v}_e'} + \frac{m_n}{m_e + m_n}\vec{v}_n' = \frac{m_n}{m_e + m_n}\vec{v}_n' \tag{2.63}$$

where $\vec{v}_n'$ refers to the updated neutral velocity after the first collision. Similarly, the relative velocity of the collision is given by

$$g_{rel,2} = |\vec{v}_n' - \vec{v}_e| = |\vec{v}_n'| \tag{2.64}$$

The standard VHS collision procedure is again used to update the post-collision velocities of the newly created electron and the neutral, which is then converted into an ion.

2.5.4 Semi-Dynamic Species Weighting Scheme

The presence of trace species in a DSMC flow can lead to an excess of non-trace particles resulting in an inefficient simulation. As discussed in Section 1.1, a hypersonic flow's ionization fraction can range between 10^{-6} to 10^{-1}. Thus, if a constant global particle weight is applied to all species, at absolute minimum a simulation must contain 10 neutral macroparticles for each ion-electron pair.

One way to alleviate this computational cost is to assign different weights to trace particles, known as variable-weight methods. There are two main methods in the literature: the species

weighting scheme (SWS) [29, 100] and the stochastic weighted particle method (SWPM) [101]. SWS assigns a constant particle weight to each chemical species. SWPM allows each particle to have a unique particle weight that can change via any physical processes of interest, such as collisions or surface interactions. SWPM often additionally employs a merging scheme which periodically combines low weight particles. This enables optimization of the number of particles per cell. Either method must carefully preserve mass, momentum, and energy for physically accurate and numerically stable results.

This work implements a semi-dynamic SWS. On creation, each particle of species i is assigned a weight $W_i = w_i f_{num}$, where f_{num} is the global particle weight and w_i is a fraction that modifies the global weight to give a species-specific weight. For example, a non-trace species might have $w_i = 1.0$, while a trace species might have $w_i = 0.001$. Then, through collisions or chemical reactions, a particle may have its own weight reduced. A merging scheme is not included. This setup is well suited for this work as it only involves simulation of three species: Ar, Ar+, and e-. Ar+ and e- are always considered trace species throughout the entire flowfield, thus their weights do not need to be dynamic. Inclusion of ionization requires some Ar particles to have their weight reduced under the circumstances described in Section 2.5.4.3. However, the computational setup under which the semi-dynamic SWS is employed (see Section 4.2.1 and Chapter 5) provides a mechanism for lower weight particles to regularly leave the domain, thus the total number of macroparticles remains controlled.

In this work, w_{Ar} refers to the default neutral weight modifier and $w_p = w_{Ar+} = w_{e-}$ refers to the default plasma species weight modifier.

2.5.4.1 Conservation of Momentum and Energy during Collisions

Collisions are performed using the standard SWS method, a full overview of which is provided in [100]. The details of conservation of momentum and energy during collisions are provided here. In one velocity direction $\hat{x}$, the total linear momentum prior to the collision is given by

$$p_x = W_1 m_1 u_1 + W_2 m_2 u_2 = W_1(m_1 u_1 + \phi m_2 u_2) \tag{2.65}$$

where $\phi = W_2/W_1$ and $W_2 < W_1$. During the collision, the first particle is split into two particles: one with weight W_2 and the other with weight $W_1 - W_2$. The total linear momentum is now written as

$$p_x = W_1[(1-\phi)m_1 u_1 + \phi(m_1 u_1 + m_2 u_2)] \tag{2.66}$$

A standard DSMC collision is performed between the particles of weight W_2. The post-collision linear momentum with post-collision properties indicated by $'$ is now

$$p_x = W_1[(1-\phi)m_1 u_1 + \phi(m_1 u_1' + m_2 u_2')] \tag{2.67}$$

The split particles must then be merged together. The merged particle's velocity is given by

$$u_1'' = (1-\phi)u_1 + \phi u_1' \tag{2.68}$$

such that

$$p_x = W_1(m_1 u_1'' + \phi m_2 u_2') = W_1 m_1 u_1'' + W_2 m_2 u_2' \tag{2.69}$$

Up until this point, these steps explicitly conserve linear momentum in all three directions. However, energy is not conserved. The energy loss per collision is given by

$$\begin{aligned} \Delta E &= W_1 \frac{1}{2} m_1 |\vec{v}_1''|^2 + W_2 \frac{1}{2} m_2 |\vec{v}_2'|^2 \\ &\quad - W_1 \frac{1}{2} m_1 |\vec{v}_1|^2 + W_2 \frac{1}{2} m_2 |\vec{v}_2|^2 \\ &= W_1 \frac{1}{2} m_1 \phi(1-\phi) |\vec{v}_1 - \vec{v}_1'|^2 \end{aligned} \tag{2.70}$$

ΔE can then be added to the total collision energy of another collision between particles of equal weight (ideally, two non-trace particles). In a closed system, this approach conserves energy exactly. In general, variable-weight methods require manual intervention to conserve flowfield properties.

2.5.4.2 The Weight Maximum Filter for Collision Pair Selection

If conventional NTC is used with a variable-weight method, trace species are more likely to be selected as candidates for collision pairs because the ratio between the number of trace macroparticles and the number of non-trace macroparticles no longer reflects the mixture mole fraction. To account for the increased number of trace macroparticles, the weight maximum filter [101, 102] is employed based on the recommendations of Reference [103]. This modifies Equations 2.54 and 2.56 as

$$N_{coll,max,max(w_i)} = \sum_i \sum_j \frac{1}{2} N_i (N_j - \delta_{ij}) f_{num} w_{max} [\langle \sigma(d,g) g \rangle]_{max} \frac{\Delta t}{V} \tag{2.71}$$

$$P_{coll} = \frac{\sigma(d,g) g}{[\langle \sigma(d,g) g \rangle]_{max}} \cdot \frac{N_{coll,max}}{N_{coll,test}} \frac{\texttt{max}(w_i, w_j)}{w_{max}} \tag{2.72}$$

where w_{max} is the maximum weight modifier encountered throughout the simulation. Overall, these modified equations increase the number of attempted collisions but decrease the number of accepted pairs, which has the effect of reproducing the species physical particle collision rates with a slight computational cost increase [103].

2.5.4.3 Reactions

When AAI occurs between two neutrals of equal weight w_n, w_n/w_p ion-electron pairs must be created. Only two post-collision velocities need to be generated, such that all ions share a post-collision velocity and all electrons share a post-collision velocity.

When EII occurs, $w_{e-} \leq w_n$. If $w_{e-} = w_n$, the standard DSMC ionization procedure is performed. If $w_{e-} < w_n$, the neutral's weight modifier is decreased with $w_{n'} = w_n - w_{e-}$. Its velocity components are left as is. An ion is then created with the default ion weight modifer.

If both EII and AAI are enabled, it is possible for AAI to occur between two neutrals of unequal weight, $w_{n,2} < w_{n,1}$. The first collision between the neutrals follows the standard SWS collision procedure from Section 2.5.4.1, and energy loss during the collision must be stored. Then, $w_{n,2}/w_p$ ion-electron pairs are created.

For a catalytic wall, Ar+ particles will reflect off the wall as neutrals. To avoid creating an excessive number of neutrals with $w_n = w_{Ar+}$, w_{Ar}/w_{Ar+} ions must absorb into the wall before one neutral is emitted with w_{Ar}. For the flowfields of interest, $n_n \gg n_p$ at the wall, so the delay before emitting a neutral does not significantly alter results.

2.5.5 Verification Test Cases

This work implemented 1D dimensionality, the electron particle model, and ionization between particle species in MPIC. These aspects of the code are verified with the following test cases. Further integrated testing of 1D MPIC is performed in Appendix B with simulation of a stagnation streamline.

2.5.5.1 Verification of 1D Laplacian Solver: Method of Manufactured Solutions

Although a Laplacian operator has already been implemented and verified in MPIC, this work is the first to implement and utilize 1D dimensionality in the re-factored MPIC. Thus, the Laplacian operator is re-verified with the method of manufactured solutions. Consider a 1D, periodic system where $\phi = \frac{e_c}{\epsilon_0} sin(2x)$. The solution to Poisson's equation for electrostatics is $\nabla^2\phi = -\frac{e_c}{\epsilon_0}(4sin(2x))$. A domain of length π with 100 cells in configuration space, no particles, and Dirichlet boundary conditions on the left and right faces is initialized. Figure 2.6a displays the resulting numerical solution from MPIC alongside the analytical solution for potential. Figure 2.6b displays the electric field, taken as the negative gradient of the potential, to also verify the gradient operator. The RMSE values of the electric potential and field profiles are 2.13 $\times 10^{-12}$ and 2.16 $\times 10^{-11}$, respectively. This demonstrates strong agreement with the analytical solutions and motivates confidence in the 1D MPIC Poisson solver.

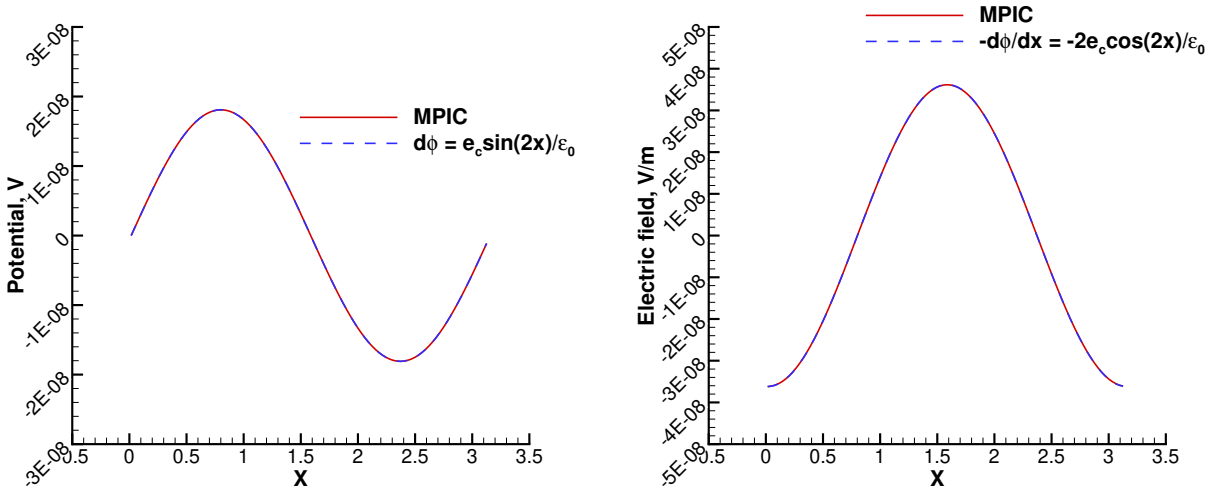

(a) Electric potential profiles.

(b) Electric field profiles.

Figure 2.6: Method of manufactured solutions verification test for MPIC's Poisson solver.

2.5.5.2 Verification of Total Collision Energy Model: 0D Reactor

A zero-dimensional (0D) reactor (also known as a heat bath) is used to verify the TCE model in MPIC. The computational domain is initialized as a 1D enclosed box consisting of one 0.01 m long grid cell with specular wall boundary conditions. The domain is filled with a Maxwellian neutral argon gas with $n_n = 10^{20}$ m^{-3}, $T_n = 100{,}000$ K, and a global particle weight of 5×10^{10}. The properties of the background neutral gas are kept constant to ensure the resulting reaction rate is also kept constant. Only VHS collisions between neutrals are enabled as well as AAI.

In the TCE model, the change in number density of species A due to reactive collisions with species B over time is given by [55]

$$\frac{dn_A}{dt} = -n_A \nu_{A-B} \int_0^{E_{max}} P_{reac}(E_{coll}) f(E_{coll}) d(E_{coll}) \tag{2.73}$$

$$\frac{dn_A}{dt} = -n_A \nu_{A-B} \overline{P}_{reac} \tag{2.74}$$

where the integral expression represents the mean probability of reaction averaged over the distribution function of collision energies found in the flow, $\overline{P}_{reac}$. For a flow at equilibrium, $f(E_{coll})$ is given by

$$f(E_{coll}) = \frac{1}{\Gamma(\zeta_T/2)} \frac{1}{k_B T} \left(\frac{E_{coll}}{k_B T} \right)^{\zeta_T/2-1} e^{-E_{coll}/k_B T} \tag{2.75}$$

where $\zeta_T = 5 - 2\omega$ for a monatomic flow using the VHS model. Thus, to verify the TCE model, the integral expression from Equation 2.73 is computed and compared to $\overline{P}_{reac}$, which is calculated from probability of reaction values outputted from a heat bath simulation.

Figure 2.7 plots the TCE probability of reaction (known as the steric factor) versus E_{coll}/E_a for 2,000 samples. The maximum value of $E_{coll}/E_a = 8.27$ with $P_{reac}(8.27E_a) = 0.939$. The average probability of reaction from this data is $\overline{P}_{reac} = 0.0491$. Using the upper limit of collision energies found in the simulation, $\int_0^{8.27E_a} P_{reac}(E_{coll}) f(E_{coll}) d(E_{coll}) = 0.0502$. This corresponds to 2% error, providing confidence in the validity of MPIC's TCE implementation.

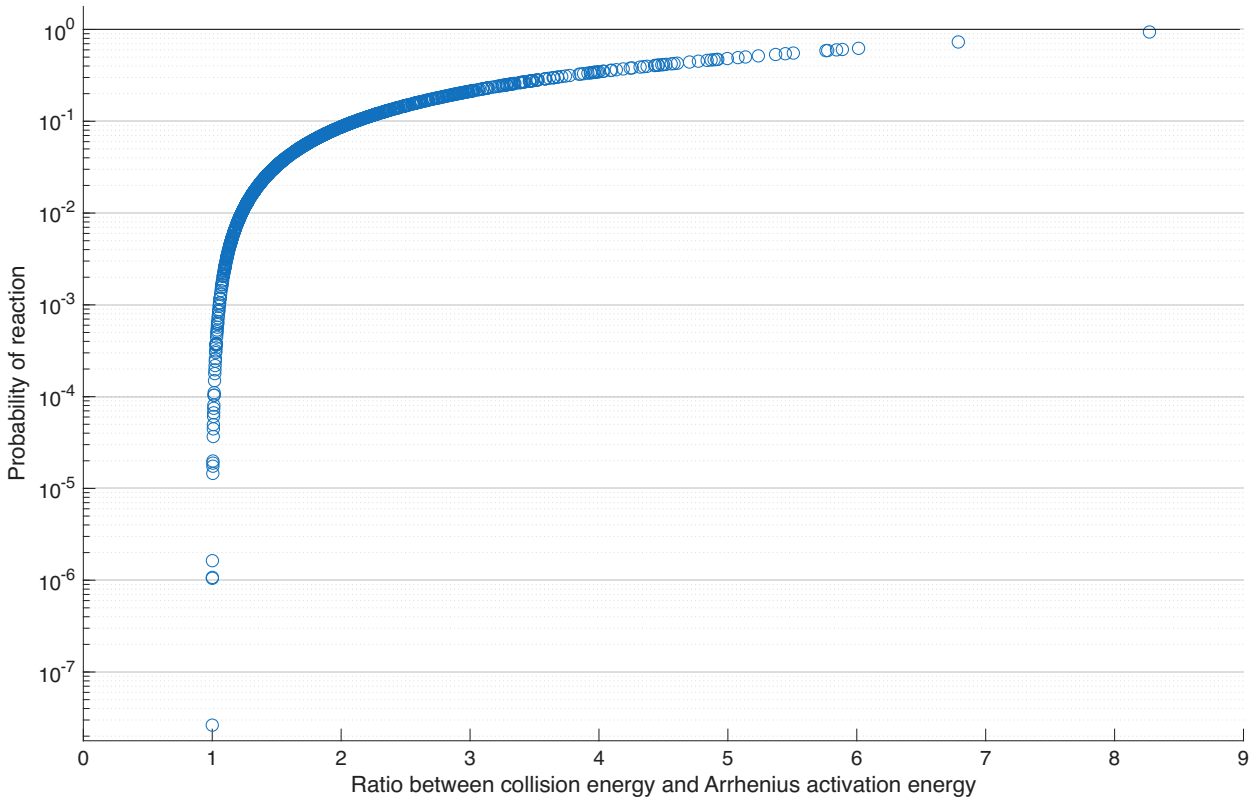

Figure 2.7: Steric factor versus ratio between collision energy and Arrhenius activation energy, E_{coll}/E_a. A line marks the y-axis where $P_{reac} = 1.0$.

Figure 2.7 shows that as E_{coll}/E_a increases, the steric factor continues to increase, with $P_{reac} > 1$ for $E_{coll}/E_a > 8.70$. Probabilities of collision or reaction greater than one are not permitted in DSMC, as they imply more than one collision or reaction should occur. For example, if $P_{reac} = 2.00$, two reactions should occur, but each particle pair can only react once. Strictly speaking, the TCE model should only be used for reaction rates and VHS parameters that satisfy $-(\overline{\zeta} + 1/2) < \eta < -(1 - \omega)$; this bounds Equation 2.59 as $E_{coll} \to E_a$ and as $E_{coll} \to \infty$. However, many commonly used rate expressions fail to meet this criteria, yet can still be used with the TCE model and achieve good agreement with theoretical rate expressions and experimental data [55]. This verification test case used a neutral equilibrium temperature of 100,000 K, which is an upper limit for the expected equilibrium temperatures of the reacting species for the hypersonic flowfields studied throughout this work. Thus, the number of occurrences where $P_{reac} > 1.0$ is expected to be negligible, and the ionization rates are used in this work regardless of the explicit TCE criteria. Alternatively, to keep the steric factor below 1.0 while preserving the ionization rate coefficients, the VHS cross section (Equation 2.60) can be increased a proportionate amount.

2.5.5.3 Verification of Variable Weighting Scheme: 0D Reactor

A 0D reactor is also used to verify the variable weighting scheme. The computational domain is initialized as a 1D enclosed box consisting of one 5 mm long grid cell with specular wall boundary conditions. The domain is filled with a Maxwellian neutral argon gas with $n_n = 1.21 \times 10^{21}$ m^{-3}, $T_n = 500$ K, and $\Delta t = 1.43 \times 10^{-7}$ s.

The first test verifies implementation of the weight maximum filter for collision pair selection (Equation 2.71). Only collisions between argon atoms are enabled. Three cases are run: SWSa, which uses the SWS scheme, a global particle weight of $f_{num} = 10^{15}$, and a species weight of $w_{Ar} = 1.0$; SWSb, which uses the SWS scheme, $f_{num} = 10^{17}$, and $w_{Ar} = 0.01$; and a case with the standard single weight NTC formula (Equation 2.54) with $f_{num} = 10^{15}$ and $w_{Ar} = 1.0$. Each simulation advances 300 time steps. Table 2.3 lists the collision rate (computed using the procedure in Reference [55]) and the number of collision pairs to test at the 300$^{\text{th}}$ time step. The VHS collision

rate is also listed for comparison. To three significant figures, SWSa and SWSb reproduce the VHS collision rate exactly. Although the weight maximum filter formula inherently tests and rejects a greater number of collision pairs compared to the NTC formula, this effect is minimized when all particles have similar weights.

Conservation of momentum and energy are verified by initializing a SWS simulation with $f_{num} = 10^{17}$ where one third of the argon atoms have $w_{Ar,1} = 1.0$ and the rest have $w_{Ar,2} = 0.5$. Conservation of momentum is successfully verified by calculating the total momentum of the 0D reactor before and after the collision procedure. Conservation of energy is verified by plotting the normalized energy of the cell against the number of collisions performed (Figures 2.8a and 2.8b). Each time a collision between two unlike particle weights is performed, energy is lost as the SWS scheme does not inherently conserve energy. If the energy loss ΔE is never added back into the simulation, the total cell energy decreases over time as shown. However, if ΔE is calculated and tracked with Equation 2.70, then added to the collision energy of the next major-to-major particle weight collision, energy is conserved exactly. At any point in time, the total energy loss is less than 0.01% of cell's total energy. So, if $\Delta E \neq 0$ after all collisions are performed (that is, the energy loss carries over to the next time step of the simulation), its impact on the simulation is negligible.

Finally, to verify that energy is conserved when ionization reactions and collisions between neutrals and neutrals, neutrals and ions, and neutrals and electrons are enabled, a SWS simulation is initialized with $f_{num} = 10^{17}$, $w_{Ar} = 1.0$, and $w_{Ar+} = w_{e-} = 0.1$. The temperature of the 0D reactor is also increased to $T_n = 50{,}000$ K. Figure 2.8c plots the normalized cell energy over 1,000 time steps. At first, the domain only contains neutral species, so AAI is the only mechanism to create the smaller weight ions and electrons. Thus, $\Delta E < 0.01\%$ of the cell's total energy for the first 400 time steps of the simulation. As more electrons are created, EII begins, resulting in the creation of neutrals with $w < 1.0$. Between EII, AAI, and collisions between neutrals and particles of unlike weights, ΔE at any given time step can take on a larger value. However, as long as it is added back into the simulation, the energy of the cell is conserved within 0.04% at any given time step. This verifies implementation of the semi-dynamic SWS approach in MPIC.

Table 2.3: Verification of the VHS collision rate for the semi-dynamic species weighting scheme in MPIC.

Parameter	VHS	Single Weight	SWSa	SWSb
Collision rate, 1/s	1.99×10^5	1.98×10^5	1.99×10^5	1.99×10^5
Number of collision pairs to test, $N_{coll,test}$	–	269	273	267

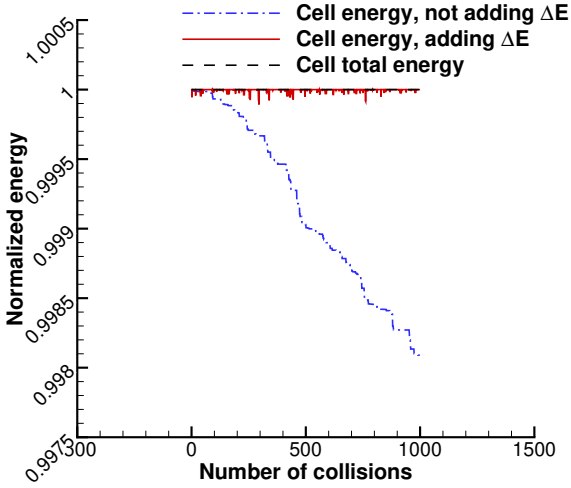

(a) Normalized cell energy versus number of collisions when only collisions between differently weighted particles are enabled.

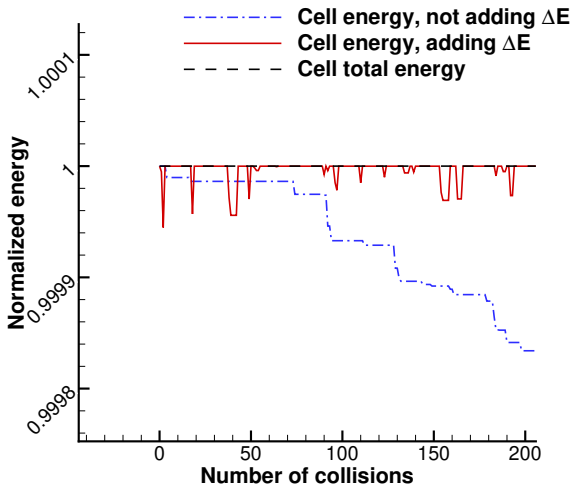

(b) Normalized cell energy versus number of collisions when only collisions between differently weighted particles are enabled, zoomed in on the first few collisions in the simulation.

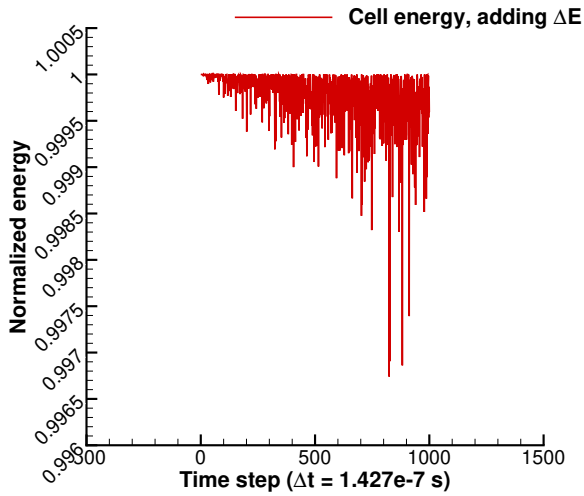

(c) Normalized cell energy versus time step when atom-atom ionization and electron-impact ionization are enabled between differently weighted particles.

Figure 2.8: Relevant plots for verification of the semi-dynamic species weighting scheme in MPIC.

2.6 Discussion of Computational Cost

Computational cost comparisons between particle-based methods and DVM require nuance, and optimal choice of solver will depend on the problem and quantities of interest at hand. Table 2.4 provides an overview of both methods along with key metrics for assessing computational cost. In terms of computational memory, particle-based methods store a mesh in configuration space and N_p macroparticles per cell, each storing at minimum its positions in d dimensions and velocities in j directions. The desired number of macroparticles per cell depends on the quantities being resolved: as a baseline, 20 to 30 particles per cell is considered sufficient. As DVM discretizes full phase space, each configuration space cell contains N_v^j velocity cells multiplied by the number of species s being simulated (e.g., a two-fluid plasma model with ions and electrons has $s = 2$).

Of prime importance is the error associated with each kinetic method. Assuming configuration space grids are sufficiently resolved, for DVM, the global truncation error is the cumulative error associated with numerically solving a differential equation. It can be reduced either by increasing the resolution in velocity space [76] or by increasing the order accuracy of the numerical scheme. For particle-based methods, numerical error stems from statistical noise, which arises when sampling the particle information using averaging techniques. As shown in Table 2.4, numerical statistical noise scales inversely with the square root of the number of macroparticles. In addition to lowering the particle weight, statistical noise can be lowered through ensemble averaging (where the same simulation is run multiple times and results are averaged across runs) or time averaging for steady state simulations (where particle properties are averaged across some time window).

Generally for 1D3V simulations, for obtaining macroscopic properties such as density and temperature, particle-based methods will almost always be more computationally efficient both in terms of computer memory and total wall time. However, because DVM resolves the VDF to equal precision in each phase-space cell, a more just comparison would scale the simulation's number of macroparticles to resolve the VDF to the same degree as a DVM solver. Resolution of VDFs is not necessarily infeasible for particle-based methods, particularly for steady state flows as will be

shown in Section 4.3.2. However, for highly resolved VDFs during unsteady flow simulations, the VDF itself must be ensemble averaged at every bin.

A common argument in favor of DVM is that the process of pairing particles for DSMC collisions becomes increasingly expensive for highly collisional flows, whereas DVM solvers use collision models such as BGK or Fokker-Planck whose computational cost is agnostic to degree of collisionality. However, these collision models can also be coupled to DSMC solvers [104, 105] and thus are not DVM-exclusive advantages. Additionally, there is ongoing debate in the literature over how effective these collision models are in capturing highly non-equilibrium VDF behavior [105, 106] as opposed to standard DSMC collisions.

Finally, in PIC simulations, charged species propagating through an electric field will experience numerical thermalization: an artificial phenomenon where charged macroparticles experience polarization drag and resonantly interact with the fluctuation spectrum, resulting in particles thermalizing towards Maxwellian distributions faster than is physically accurate [107]. This effect can be alleviated with more macroparticles, however, the number of macroparticles required to drive the rate of numerical thermalization below the rate of thermalization due to Coulomb collisions is quite large. As a result, for many PIC studies Coulomb collisions are neglected as numerical thermalization dominates regardless.

Table 2.4: Overview of discrete-velocity and particle-based kinetic simulation methods.

	Discrete-velocity	Particle-based
Solution	VDFs in phase space	Macroparticle trajectories
Differential equation	Hyperbolic PDE	Ordinary differential equation
Fluid description	Eulerian	Lagrangian
Macroscopic properties	Moments of VDFs	Sampling macroparticle information
Collisions	Integral of collision operator	Probabilistic algorithm
Computational memory per physical cell	$s \cdot N_v^j$	$(d+j)N_p$
Numerical error	Global truncation error $\sim \mathcal{O}[(N_v)^{-n}]^*$	Statistical noise $\sim \mathcal{O}[(N_p)^{-1/2}]$

$^*n$ is the order accuracy of the scheme.

2.7 Summary

In this chapter, an overview of the governing equations and numerical approaches for kinetically simulating a hypersonic flow was provided. The DSMC-PIC and DVM methods were described in addition to the codes used throughout this work. Details of collision and ionization models, as well as the argon parameters used in this work, are also included in Appendix A. Verification tests demonstrated successful implementation of the code enhancements developed for this work. A strong scaling analysis compared the scalability and efficiency of 1D1V and 1D3V DVM. In the following chapters, a justification of the choice of solver will be presented for each numerical study.

Chapter 3

Study of Charged Species Diffusion

Before examining charged species diffusion and the efficacy of the ambipolar diffusion approximation in hypersonic flows, it is first necessary to understand how charged species diffuse in general. In this chapter, the physics of charged species diffusion, methods of predicting charged species diffusion regime, and the efficacy of the ambipolar diffusion approximation in non-hypersonic continuum flows are explored. Section 3.1 reviews the various regimes of charged species diffusion including the formal derivation of ambipolar diffusion. Section 3.2 reviews the Phelps criteria: a set of length scale ratios which can be used to identify charged species diffusion regimes in gas discharge vessels. The applicability of the Phelps criteria to other flowfields are tested with DVM and the computational setup described in Section 3.3. Results are presented and discussed in Section 3.3. Guidelines for identifying charged species diffusion regimes are summarized in Section 3.4 and guidelines for using the ambipolar diffusion approximation in non-hypersonic continuum flows are summarized in Section 3.5.

3.1 Regimes of Charged Species Diffusion

Despite its name, the ambipolar diffusion approximation dictates the translational motion of charged species, not diffusion. For example, in CFD, the approximation can be applied in conjunction with any diffusion model of choice. The name and aspects of its implementation are based on the concept of ambipolar diffusion: one of several charged species diffusion regimes for partially ionized flows.

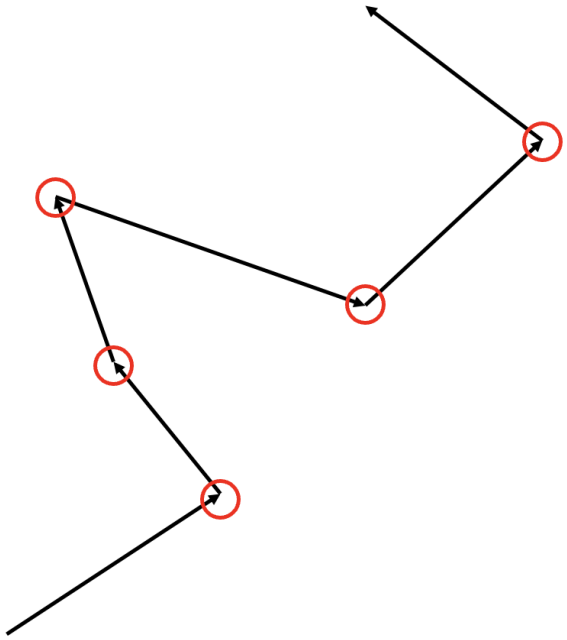

Figure 3.1: Diagram of the "random-walk" trajectory of a particle undergoing elastic and/or inelastic collisions. Each collision event is marked with a red circle.

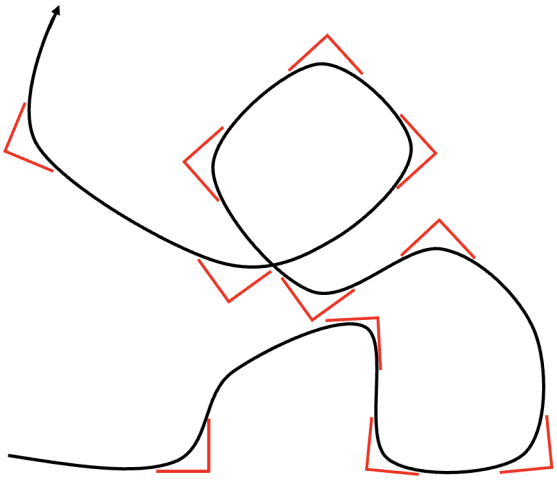

Figure 3.2: Diagram of the trajectory of a charged particle undergoing Coulomb collisions. Each Coulomb collision event is marked with a red right angle symbol.

Diffusion refers to the transport of matter via thermal motion in response to a gradient in concentration. In absence of body forces, neutral particles can only have their trajectory changed via elastic or inelastic collisions. Charged particles may have their trajectory changed via elastic or inelastic collisions, Coulomb collisions, and electromagnetic fields (both applied and self-induced). An important distinction is that a particle undergoing standard elastic and inelastic collisions travels in a random-walk manner with discontinuous derivatives (see Figure 3.1), and, except in rare circumstances, collides with at most two other particles. Conversely, a particle undergoing Coulomb collisions travels along a smooth curve due to contributions of near and long-range interactions with many particles at a time. When the particle's direction has changed by approximately 90°, it is considered to have undergone one Coulomb collision [5] (see Figure 3.2). It is possible for a single Coulomb collision between two particles to result in a scattering angle over 90°, but generally, a Coulomb collision results from accumulative effect of many small-angle deflections [27].

Consider a plasma distributed non-uniformly among a dense background of neutrals, where the plasma particles primarily collide with neutrals rather than each other. The faster, lightweight electrons quickly separate from the ions, leading to a deficit in negative charge. To maintain quasineutrality, the plasma develops an electric field within itself to slow the electrons and accelerate the ions such that they diffuse at the same rate and have equal bulk velocities.

From the fluid equations of motion for a single species plasma, one can derive the self-induced ambipolar electric field $\vec{E}_a$ as well as the mutual diffusion coefficient of ions and electrons, known as the ambipolar diffusion coefficient D_a,

$$\vec{E}_a \equiv \frac{D_i - D_e}{\mu_i + \mu_e} \frac{\nabla n}{n} \tag{3.1}$$

$$D_a \equiv \frac{\mu_i D_e + \mu_e D_i}{\mu_i + \mu_e} \tag{3.2}$$

$$\text{where } \mu_j \equiv \frac{|q_j|}{m_j \nu_{j-n}} \tag{3.3}$$

$$D_j \equiv \frac{k_B T_j}{m_j \nu_{j-n}} \tag{3.4}$$

where the subscript i denotes ion quantities, e denotes electron quantities, and j denotes either species quantities. ν_{j-n} is the collision frequency with neutrals. The mobility and diffusion coefficients are μ_j and D_j, respectively. In the case of isothermal electrons, an alternate expression for the ambipolar electric field was derived by Langmuir and Tonks [108],

$$\vec{E}_{a,LT} = -\frac{k_B T_e}{e} \nabla ln[n_e] \tag{3.5}$$

This expression emphasizes the role of density gradients in inducing ambipolar fields: if the density distribution is uniform, Equation 3.5 is equal to zero. Furthermore, the larger the gradient in plasma number density, the greater the magnitude of the electric field.

If one assumes $\mu_e \gg \mu_i$, the magnitude of D_a is approximately

$$D_a \approx D_i \left(1 + \frac{T_e}{T_i}\right) \tag{3.6}$$

This indicates that the diffusion rate resulting from the self-induced ambipolar electric field is at least equal to the original diffusion rate of the ions. A special case occurs when $T_e = T_i$ and the ambipolar diffusion rate equals the original ion diffusion rate multiplied by a factor of 2.

Several assumptions are made throughout this derivation. First, by starting from the fluid equations of motion, it is implied that the flow is under continuum conditions. How strict this requirement is is explored in Chapter 5. Second, the gas must be sufficiently weakly ionized such that charged particles primarily collide with neutrals rather than each other. This ensures particles diffuse in a random-walk manner consistent with Fick's Law, rather than the collective and often oscillatory motion of strongly and fully ionized plasmas. Finally, and most importantly, the weakly ionized gas must be a 'plasma' according to the criteria presented in Section 1.2.

If a partially ionized gas (where collisions with neutrals dominate over Coulomb collisions) is primarily governed by hydrodynamic forces rather than electrostatic ones, then it is in the free diffusion regime. Free diffusion may occur if the plasma density is sufficiently small, such that the magnitudes of the self-induced electric fields are insufficient in preserving quasineutrality over the

length scales of interest, or if the third criterion of the definition of a plasma is violated.

If a partially ionized gas is effectively collisionless with neutrals and does not exhibit collective plasma behavior, then the gas lies within the free fall regime. Charged particle trajectories are primarily governed by classical mechanics. As an example, charged particles in a vacuum may lie within the free fall regime.

Finally, if the partially ionized gas constitutes a proper plasma but the neutral gas is considered collisionless, then charged species diffusion is described by the ion inertia regime. Here, minimal charge separation occurs such that electrons follow near identical trajectories to the ions.

If the charged species dynamics and macroscopic properties are known *a priori*, then the identification of the charged species diffusion regime is straightforward. However, for most applications, simulating a flow with high-fidelity plasma physics modeling just to determine whether resolution of such effects is necessary is impractical. Thus, criteria for predicting the charged species diffusion regime for an arbitrary flow are desirable. In determining such guidelines, one must consider whether the criterion is only properly assessed through full electrostatic modeling. Either electric potential or electric field are clearly unsuitable as they require solution of the Poisson equation. A quantity like electron number density requires more nuance. If EII and electrostatic modeling are included in a simulation, the electric field may impart more energy to electrons, leading to increased EII rates and a higher electron number density. Thus, if the flow is evaluated with a simulation tool that neglects full electrostatic modeling, the electron number density may be underpredicted and the charged species diffusion regime may be incorrectly identified.

It is also clear that the ambipolar diffusion approximation does not precisely represent any of these diffusion regimes, including ambipolar diffusion. The approximation does not account for acceleration of ions through the ambipolar electric field during ambipolar diffusion, but the extent to which this impacts simulation accuracy is case-dependent. Additionally, the criterion of the electron Debye length being orders of magnitude smaller than the neutral mean free path (Equation 1.5), conventionally used to justify the ambipolar diffusion approximation, is not a required assumption for the ambipolar diffusion derivation. $\lambda_{D,e} \ll \lambda_{mfp,n}$ bears a resemblance to the quasineutrality

criterion $\lambda_D \ll L$, however, the neutral mean free path does not define the dimensions of a plasma system. As a counterexample, one could take the limit as $\lambda_{mfp,n} \to \infty$, in which case the plasma is not collisional with neutrals and the charged species are in the ion inertia regime. This motivates evaluation of the ambipolar diffusion approximation and its ability to capture each diffusion regime in a simple test setup prior to assessing its validity for hypersonic flowfields.

3.2 Phelps Criteria for Charged Species Diffusion

In his review paper [109], A.V. Phelps discusses the concept of the diffusion length, Λ [110, 111], and its ratios with the electron Debye length and the ion mean free path as a means of characterizing charged species diffusion in plasma discharge vessels. For a cylindrical and parallel plane geometry, respectively, the diffusion length is defined as

$$\frac{1}{\Lambda^2} \equiv \left(\frac{\pi}{L}\right)^2 + \left(\frac{2.405}{R}\right)^2 \tag{3.7}$$

$$\Lambda \equiv \frac{L}{\pi} \tag{3.8}$$

In the derivation of the diffusion length, it is assumed only a single type of charged particle is present, no external electric fields are present, and the plasma density is approximately zero at the boundaries (i.e. a fully absorbing wall).

Figure 3.3 maps charged particle diffusion regimes according to the ratios between the diffusion length and the ion mean free path $\lambda_{mfp,i}$ (with respect to neutral collisions) and the electron Debye length. The vertical axis, $0 < \Lambda/\lambda_{mfp,i} < \infty$, is a measure of collisionality for the partially ionized gas. As $\Lambda/\lambda_{mfp,i} \to \infty$, the number of collisions increases and diffusion processes dominate charged species mobility. The horizontal axis, $0 < \Lambda/\lambda_{D,e} < \infty$, is effectively a re-statement of the quasineutrality criterion for plasmas. As $\Lambda/\lambda_{D,e} \to 0$, the dimensions of the system $\Lambda \propto L$ approach the magnitude of $\lambda_{D,e}$, and the partially ionized gas no longer behaves like a collective plasma at that length scale.

For collisional rarefied gas flows, the transition from free diffusion ($\Lambda/\lambda_{mfp,i} \to \infty$, $\Lambda/\lambda_{D,e} \to$

0) to ambipolar diffusion ($\Lambda/\lambda_{mfp,i} \to \infty$, $\Lambda/\lambda_{D,e} \to \infty$) is of interest. Phelps quantifies the transition from free to ambipolar diffusion in terms of the effective diffusion coefficient. When compared with the theoretical diffusion coefficients (Equation 3.4), the effective diffusion coefficient describes the change in charged particle loss to absorbing walls in a discharge vessel as the result of ambipolar diffusion. An equivalent quantity is not readily defined nor easily experimentally measured for arbitrary gas flows. Therefore, this study seeks to characterize the transition between free and ambipolar diffusion solely in terms of the diffusion length ratios $\Lambda/\lambda_{mfp,i}$ and $\Lambda/\lambda_{D,e}$.

3.3 Simulation of Weakly Ionized Gas Expansion

In order to investigate charged species diffusion with as few external variables as possible, a simplified weakly ionized gas expansion problem is studied. A uniform background neutral gas is initialized in 1D inside a periodic domain of length L and an ionization fraction is imposed at the center of the domain from $0.25L \leq x \leq 0.75L$. All species are initialized at thermodynamic equilibrium. At $t = 0$ s, the ionized gas is allowed to symmetrically expand into the background neutral gas. In essence, this setup simulates unsteady charged species diffusion under the effect of a large density gradient.

The periodic boundary conditions prevent the charged particles from experiencing wall or inlet effects. Additionally, this relaxes the Phelps criteria assumption that the charged species densities must be zero at the boundaries, enabling assessment of generalizability. Argon is used to model both the neutral and ionized gas. The gas is not allowed to undergo ionization or recombination, such that the amount of ionized gas remains constant throughout the entire simulation. As the setup is in 1D, Equation 3.8 is used to define the diffusion length with L equal to the length of the computational domain.

DK is selected for this study due to its deterministic properties. If a particle-based method were used, the computational cost required to sufficiently reduce numerical noise is impractical for the scale of this study. The unsteady nature of the problem also requires either an extremely small particle weight or ensemble averaging. Significant energy transfer between translational modes

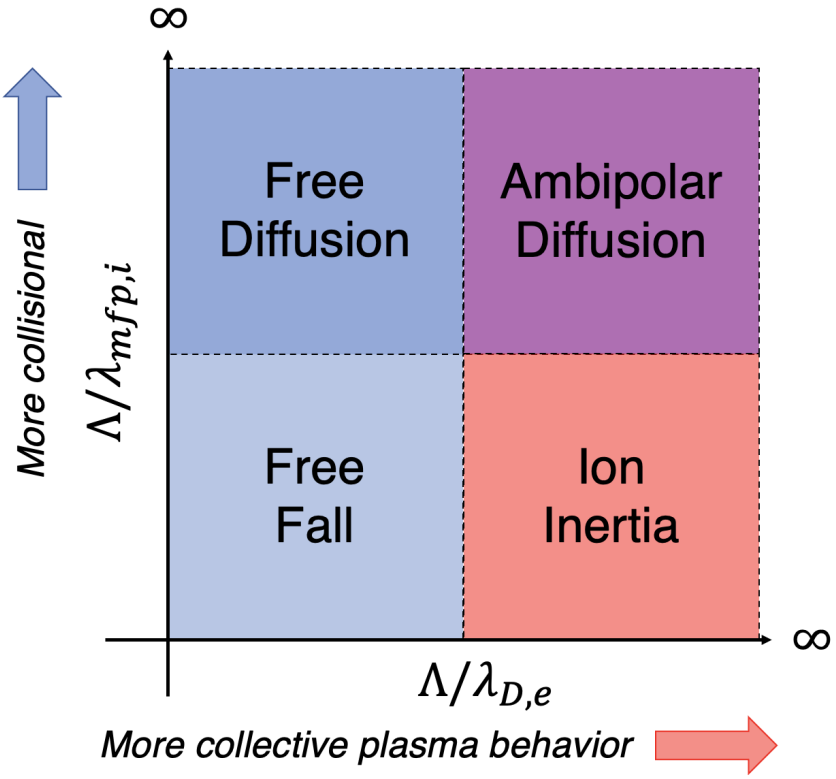

Figure 3.3: Map of charged species diffusion regimes, with scales indicating correlation with diffusion length criteria.

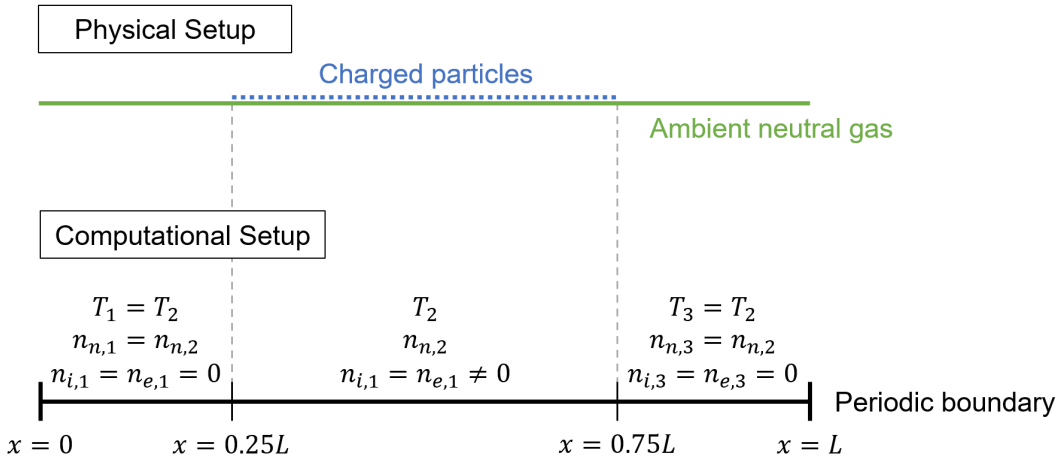

Figure 3.4: Schematic of the physical and computational setup for the weakly ionized gas expansion problem.

is not expected, as the particles diffuse slowly and the electric field remains small for the cases simulated. Thus, 1D1V DVM can be used to reduce computational expense as opposed to 1D3V DVM or 1D3V DSMC.

Both the HS and VHS collision models are tested, and simulations are run with and without Coulomb and CEX collisions. Overall, the conclusions made in each section are invariant of collision model. Though the number density profiles and diffusion rates will be different for HS versus VHS simulations, the diffusion length criteria and conclusions about the ambipolar diffusion approximation hold true for both models. The addition of Coulomb and CEX collisions has a negligible effect on results because their frequencies are small compared to the MEX collision frequency. The following sections discuss results from simulations with the HS model and only MEX collisions between charged particles and neutrals enabled. This represents the closest possible simulation setup to ideal ambipolar diffusion. Initialization parameters used across all simulations are listed in Table 3.1. The range of initial neutral and charged species densities tested are listed in Tables 3.2 and 3.3.

Table 3.1: Initial conditions and simulation parameters used for all cases in the weakly ionized gas expansion problem.

Property	Value
Temperature, K	500
Domain length, m	1.0
Physical space grid points per $\lambda_{mfp,i}$	≥ 21.0
Physical space grid points per $\lambda_{D,e}$	≥ 32.0
Grid points in velocity space	256
Range of ion velocity space, m/s	$\pm$ 16,311
Range of electron velocity space, m/s	$\pm$ 4,404,560

Throughout the following sections, τ_{j-n} refers to the inverse of the collision frequency of a particle of type j colliding with neutrals taken from the center of the domain ($x = 0$ m). Normalized charge separation is defined as

$$\overline{n}_c = \frac{|n_i - n_e|}{n_i} \tag{3.9}$$

Table 3.2: Overview of ratios between diffusion length and ion mean free path tested in weakly ionized gas expansion problem.

$\Lambda/\lambda_{mfp,i}$	Neutral number density, m^{-3}	Ion mean free path, m
0.01	6.12×10^{16}	3.18×10^{1}
0.1	6.12×10^{17}	3.18×10^{0}
1.0	6.12×10^{18}	3.18×10^{-1}
10	6.12×10^{19}	3.18×10^{-2}
100	6.12×10^{20}	3.18×10^{-3}

Table 3.3: Overview of ratios between diffusion length and electron Debye length tested in weakly ionized gas expansion problem.

$\Lambda/\lambda_{D,e}$	Ion/Electron number density, m^{-3}	Debye length, m
0.01	2.35×10^{3}	3.18×10^{1}
0.1	2.35×10^{5}	3.18×10^{0}
1.0	2.35×10^{7}	3.18×10^{-1}
10	2.35×10^{9}	3.18×10^{-2}
100	2.35×10^{11}	3.18×10^{-3}

and is used as a measure of local charge neutrality. For the purposes of this study, quasineutrality is considered achieved when the normalized charge separation value is below 10^{-2} in a given physical space cell. This value is appropriate because for all of the ambipolar diffusion flow cases simulated in this study, a value of $\overline{n}_c = 10^{-2}$ marks the point where the plasma shifts from the ambipolar regime into the transitional diffusion regime. Further justification is provided in Section 3.3.2.

3.3.1 Time Evolution of Ambipolar Diffusion

For ambipolar diffusion, it follows that a finite amount of time must elapse between the start of diffusion and the point at which local charge neutrality is established everywhere. During this time, the ambipolar electric field generates the moment charge neutrality is violated (i.e. when the electrons diffuse away from the ions). The electric field then gradually decays as the plasma neutralizes itself, becoming exactly zero when local charge neutrality is restored everywhere.

Figure 3.5 shows the time evolution of the ion and electron number density profiles throughout the simulation. The diffusion length ratios are $\Lambda/\lambda_{mfp,i} = 100$ and $\Lambda/\lambda_{D,e} = 100$, ensuring the simulation is within the ambipolar diffusion regime (this threshold is further explained in Section 3.3.2). Current time in the simulation is given in terms of seconds and by the normalized number of collision times, given by t/τ_{j-n}. The latter metric demonstrates that the ions have collided a sufficient number of times to study diffusion.

The rate at which local charge neutrality is achieved everywhere occurs on the order of tens of ion collision times and thousands of electron collision times. At the end of the simulation (Figure 3.5c), local charge neutrality is approximately achieved everywhere except the far boundaries of the simulation ($|x/\Lambda| \geq 1.2$). Thus, ambipolar diffusion is only occurring in this simulation at the given point in time for $|x/\Lambda| \leq 1.2$. When the ion and electron number density profiles intersect, large dips form in the charge separation profile. These points are not factored into assessments of quasineutrality, particularly for results in the following sections.

Figure 3.6 shows the magnitude of the electric field calculated from the flow simulation compared with the expected value from the Langmuir and Tonks relation (Equation 3.5). For

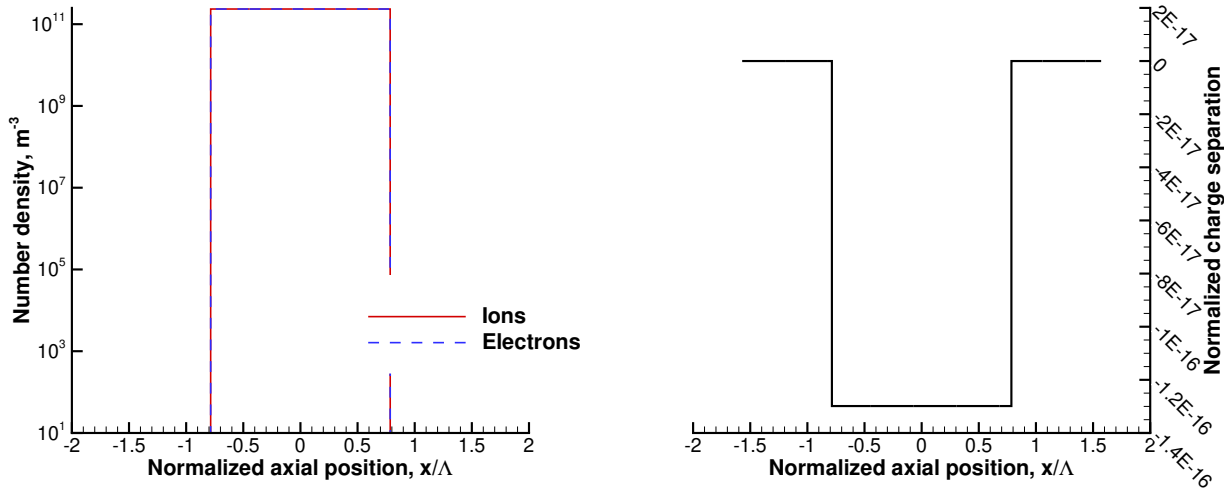

(a) $t = 0$ s (Simulation starting conditions)

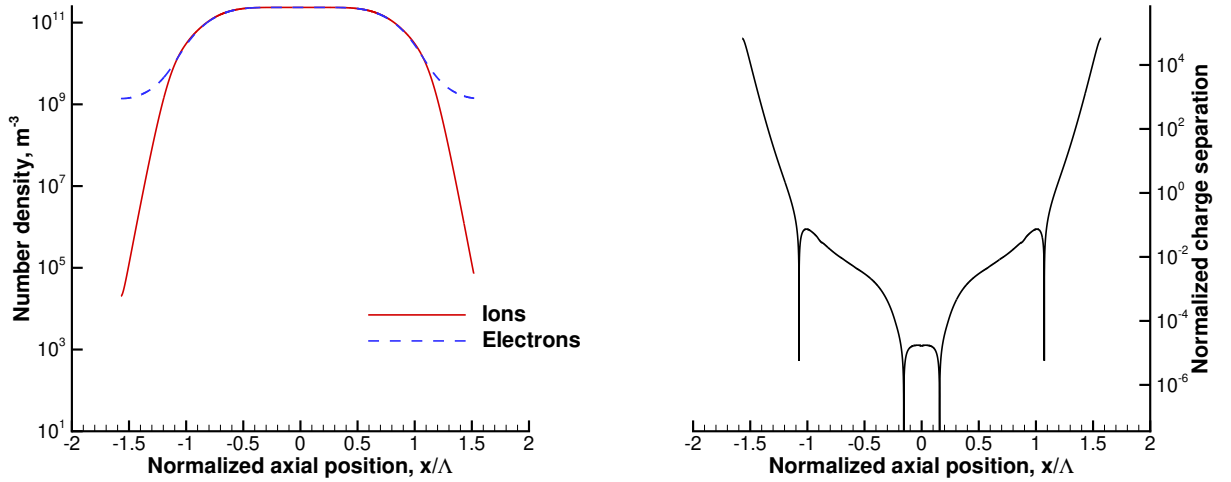

(b) $t = 2 \times 10^{-4}$ s $\approx 16\tau_{i,n} \approx 4283\tau_{e,n}$

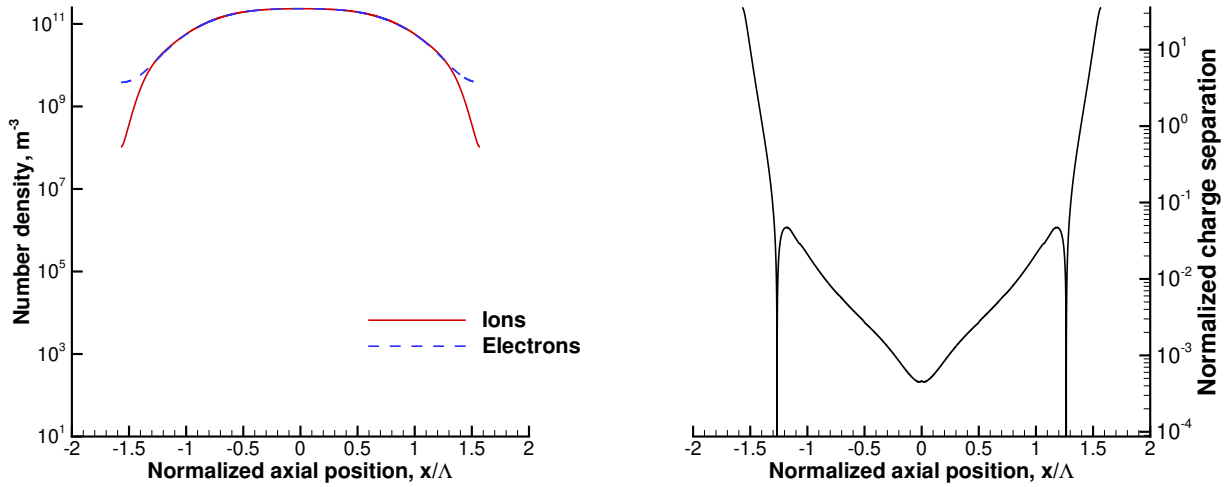

(c) $t = 4 \times 10^{-4}$ s $\approx 32\tau_{i,n} \approx 8566\tau_{e,n}$

Figure 3.5: Number density for ions and electrons (left) and normalized charge separation (right) profiles at various points in time during ambipolar diffusion ($\Lambda/\lambda_{mfp,i} = 100, \Lambda/\lambda_{D,e} = 100$).

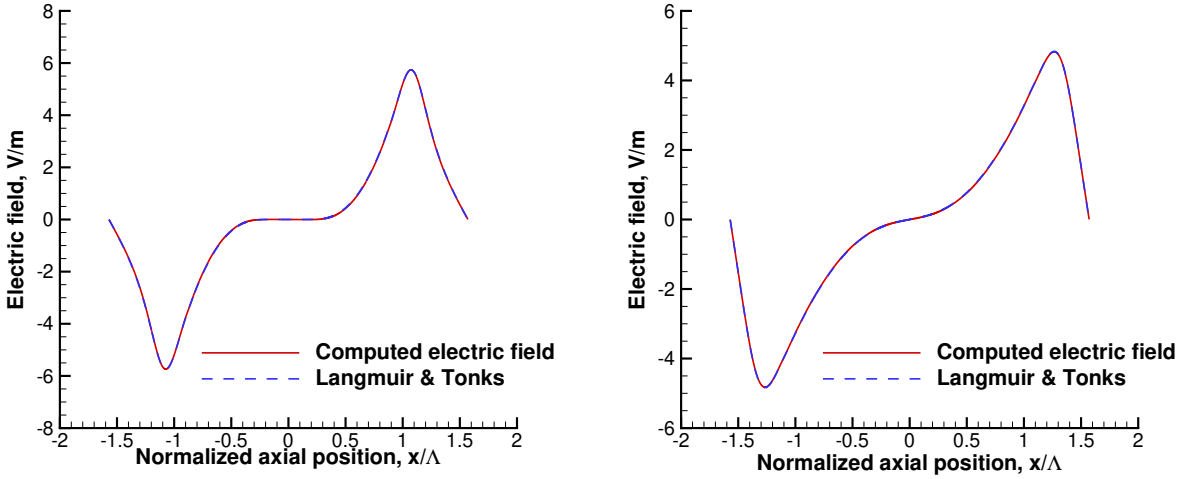

(a) $t = 2 \times 10^{-4}$ s $\approx 16\tau_{i,n} \approx 4283\tau_{e,n}$

(b) $t = 4 \times 10^{-4}$ s $\approx 32\tau_{i,n} \approx 8566\tau_{e,n}$

Figure 3.6: Computed electric field and Langmuir and Tonks relation at various points in time during ambipolar diffusion ($\Lambda/\lambda_{mfp,i} = 100, \Lambda/\lambda_{D,e} = 100$).

visual clarity, the plot of Equation 3.1 is not included, however, there is strong agreement between all three profiles. The RMSE value between the Langmuir and Tonks relation and the DVM-computed electric field at $t = 2 \times 10^{-4}$ s is 0.200%, and the RMSE value between the profiles at $t = 4 \times 10^{-4}$ s is 0.170%. From the halfway point in the simulation to the end, the peak electric field magnitude decays from 5.74 V/m to 4.82 V/m. This demonstrates how the ambipolar field declines in strength as time goes on and charge neutrality is restored. However, because the electron number density profile begins to curve at the center (Figure 3.5c), it forms a nonzero gradient and the region where the electric field is exactly zero shrinks.

Ultimately, this analysis demonstrates that ambipolar diffusion is a time dependent process and that local charge neutrality is not instantaneously achieved by the plasma being in the ambipolar diffusion regime. Simulations with processes that occur on time scales on the order of collision times should be mindful of this time dependence, especially when employing the ambipolar diffusion approximation.

3.3.2 Criteria for Classifying Charged Species Diffusion Regime

In the following section, specific numerical limits for $\Lambda/\lambda_{mfp,i}$ and $\Lambda/\lambda_{D,e}$ are identified in order to classify the charged species diffusion regimes. Additionally, Equation 1.5 is evaluated as a means of predicting ambipolar diffusion.

3.3.2.1 Transition Between Free and Ambipolar Diffusion

To map the transition between free and ambipolar diffusion in terms of the diffusion length ratios, a series of simulations are run which vary the initial values of $\Lambda/\lambda_{mfp,i}$ and $\Lambda/\lambda_{D,e}$ between 0.01 and 100. Each simulation is analyzed in a similar manner to Section 3.3.1. The most collisional of the cases with $\Lambda/\lambda_{mfp,i} = 100$ are discussed here.

For the set of flow conditions simulated, a diffusion length to electron Debye length ratio of $\Lambda/\lambda_{D,e} \geq 100$ is found to be a strong predictor of ambipolar diffusion with macroscopic property profiles similar to those seen in Figures 3.5 and 3.6. When $\Lambda/\lambda_{D,e} \leq 1.0$, charged species diffusion

can be modeled as free diffusion, with Figure 3.7 capturing macroscopic property profiles representative of this regime. At the start of the simulation, the electrons near instantaneously separate from the ions. By the end of the simulation, the electrons have nearly uniform density throughout the domain while the ions lag far behind. The normalized charge separation at the center of the domain is 0.465. This is because the self-induced electric field is not strong enough to restrain the motion of electrons, reaching a peak value of only 0.0431 V/m compared to 4.82 V/m in the ambipolar diffusion case. There is less agreement between the electric field and the Langmuir and Tonks relation with an RMSE value of 3.57%. This is in part because the assumptions under which the Langmuir and Tonks relation is valid are not applicable in the free diffusion regime. It is clear that the ambipolar diffusion approximation would not be valid in this flow case, as it would not capture the significant charge separation between ions and electrons.

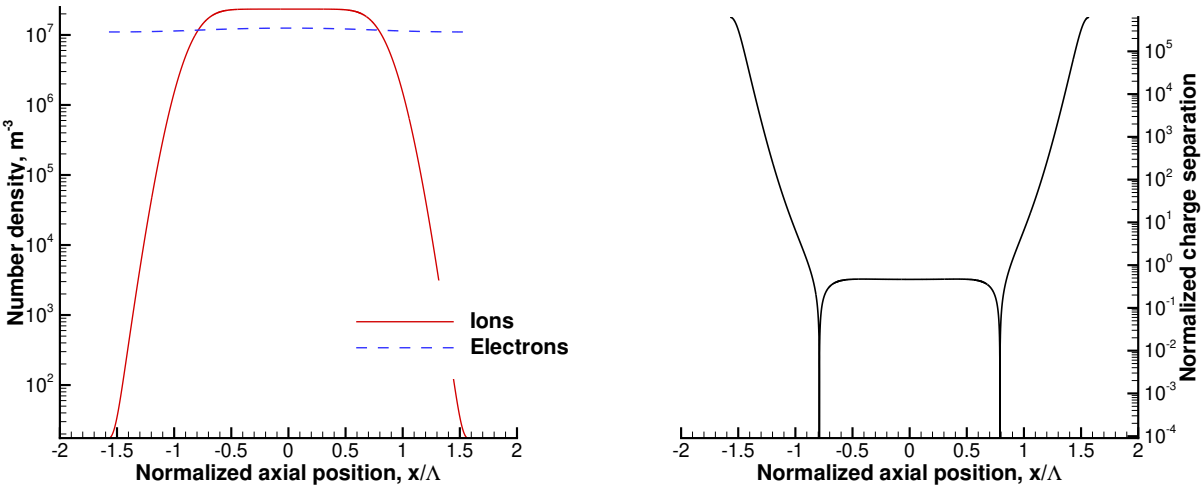

(a) Number density. (b) Normalized charge separation.

Figure 3.7: Number density for ions and electrons (left) and normalized charge separation (right) profiles for free diffusion case at $t = 4 \times 10^{-4}$ s $\approx 32\tau_{i,n} \approx 8566\tau_{e,n}$ ($\Lambda/\lambda_{mfp,i} = 100, \Lambda/\lambda_{D,e} = 1.0$).

A value of $\Lambda/\lambda_{D,e}$ between 1.0 and 100 defines the transitional regime between free and ambipolar diffusion. Figure 3.8 shows macroscopic profiles representative of the regime for $\Lambda/\lambda_{D,e} = 10$. Here, electrons do not separate as far from the ions as they would in the free diffusion regime, with $\overline{n}_c = 0.0873$ at the center of the domain. However, this is still too much charge separation to

call the diffusion ambipolar. The exact value of $\Lambda/\lambda_{D,e}$ for which a flow will enter the ambipolar diffusion regime depends on the collision model used, the types of collisions simulated, and the value of $\Lambda/\lambda_{mfp,i}$, with values ranging from 40 to 80. There is better agreement between the electric field and the Langmuir and Tonks relation than the free diffusion case with an RMSE value of 0.612%. In the transitional regime, direct electric field modeling must be used to properly capture charged species flow dynamics at all points in the flow. It follows that the ambipolar diffusion approximation is inappropriate for this regime for capturing charged species densities. However, if electrons were allowed to diffuse freely, then the simulation would falsely overpredict the degree of charge separation.

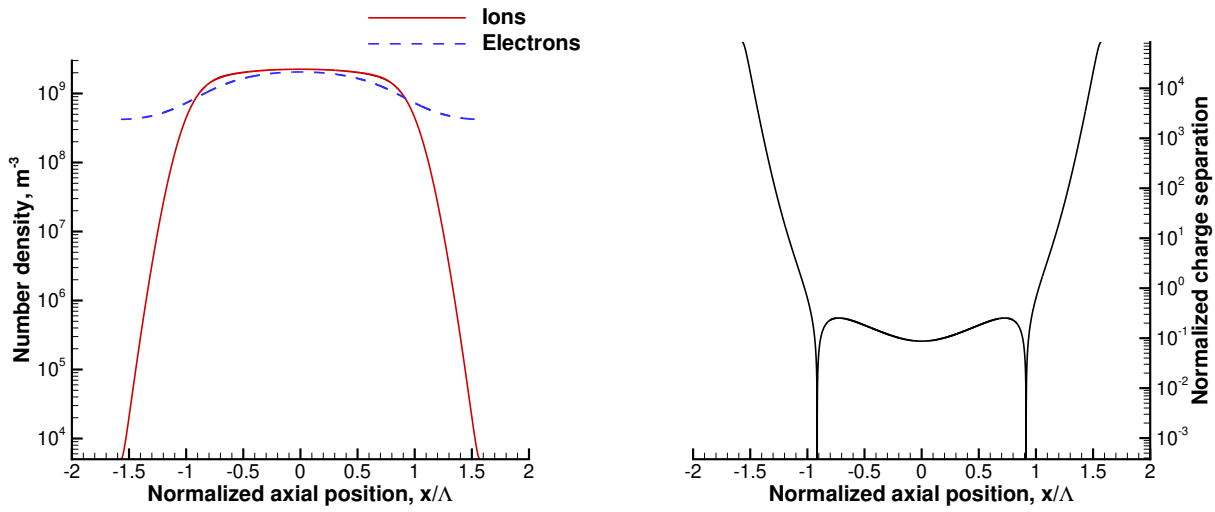

(a) Number density. (b)

Figure 3.8: Number density for ions and electrons (left) and normalized charge separation (right) profiles for transitional diffusion case at $t = 4 \times 10^{-4}$ s $\approx 32\tau_{i,n} \approx 8566\tau_{e,n}$ ($\Lambda/\lambda_{mfp,i} = 100, \Lambda/\lambda_{D,e} = 10$).

The transitional regimes between free diffusion and free fall and between ambipolar diffusion and ion inertia are outside the scope of this work. Free fall and ion inertia diffusion are relevant to applications in the free molecular regime such as the ionosphere (Figure 1.4), which stretches across both low Earth orbit and VLEO. In this study, the correlation between $\Lambda/\lambda_{D,e}$ and the transition between free and ambipolar diffusion is valid for at least $\Lambda/\lambda_{mfp,i} \geq 10$. However, no determinative

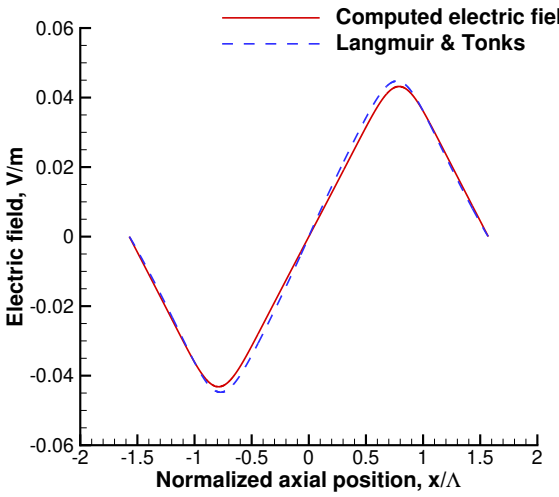

Figure 3.9: Computed electric field and Langmuir and Tonks relation for free diffusion case ($\Lambda/\lambda_{mfp,i} = 100, \Lambda/\lambda_{D,e} = 1.0$).

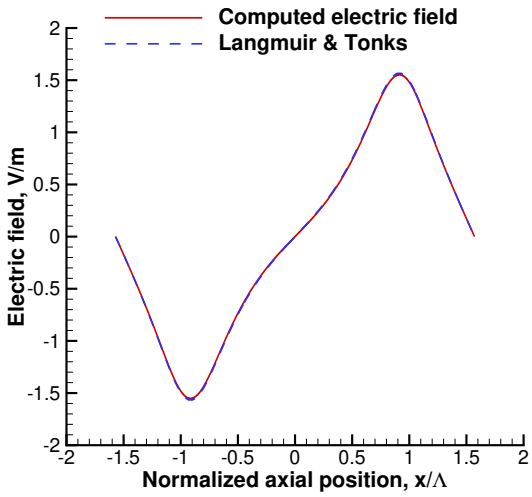

Figure 3.10: Computed electric field and Langmuir and Tonks relation for transitional diffusion case ($\Lambda/\lambda_{mfp,i} = 100, \Lambda/\lambda_{D,e} = 10$).

conclusions can be made for $\Lambda/\lambda_{mfp,i} < 10$. It is likely that, similar to the value of $\Lambda/\lambda_{D,e}$ which marks the transition between free and ambipolar diffusion, the exact value of $\Lambda/\lambda_{mfp,i}$ will vary between flow cases and simulation parameters. Generally, flows with a Knudsen number (Equation 2.8) less than or equal to 0.01 are collisional enough for free diffusion and ambipolar diffusion to occur.

3.3.2.2 Debye Length versus the Mean Free Path

To test the conventional ambipolar diffusion approximation criterion of $\lambda_{D,e} \ll \lambda_{mfp,n}$, a simulation case is set up with a $\Lambda/\lambda_{mfp,i}$ ratio of 10,000 and a $\Lambda/\lambda_{D,e}$ ratio of 100. Use of the HS collision model ensures the neutral mean free path is identical to the ion mean free path. This case violates Equation 1.5 by two orders of magnitude as shown in Figure 3.12. However, with $\Lambda/\lambda_{D,e} = 100$, this case falls within the ambipolar diffusion regime according to the Phelps criteria. The simulation is run for $t = 5 \times 10^{-6}$ s, which translates to $40\tau_{i,n}$ or $4677\tau_{e,n}$.

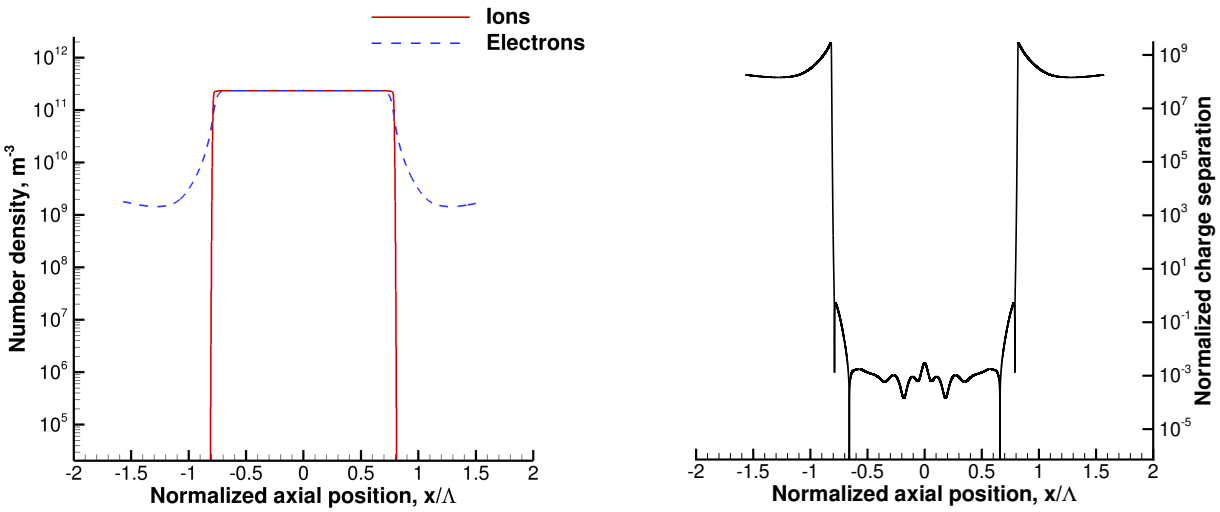

(a) Number density.

(b) Normalized charge separation.

Figure 3.11: Number density for ions and electrons (left) and normalized charge separation (right) profiles for ambipolar diffusion case at $t = 4 \times 10^{-5}$ s $\approx 323\tau_{i,n} \approx 37,418\tau_{e,n}$ ($\Lambda/\lambda_{D,e} = 100, \Lambda/\lambda_{mfp,i} = 10,000$).

Figure 3.11a plots the number densities of the ions and electrons and the charge separation

profile for this flow case. Expansion of ions is much slower compared to the ambipolar diffusion case from in Section 3.3.1 because the neutral number density has increased significantly, thus increasing the number of collisions with neutrals and slowing the diffusion rate of ions. At all points in the flow, $\lambda_{D,e} > \lambda_{mfp,i}$ by at least two orders of magnitude. However, at the points where $\Lambda/\lambda_{D,e} = 100$ ($|x/\Lambda| < 0.789$), local charge neutrality is achieved, indicating ambipolar diffusion is occurring. This demonstrates the Phelps criteria for charged species diffusion is a more accurate predictor of ambipolar diffusion than $\lambda_{D,e} \ll \lambda_{mfp,n}$.

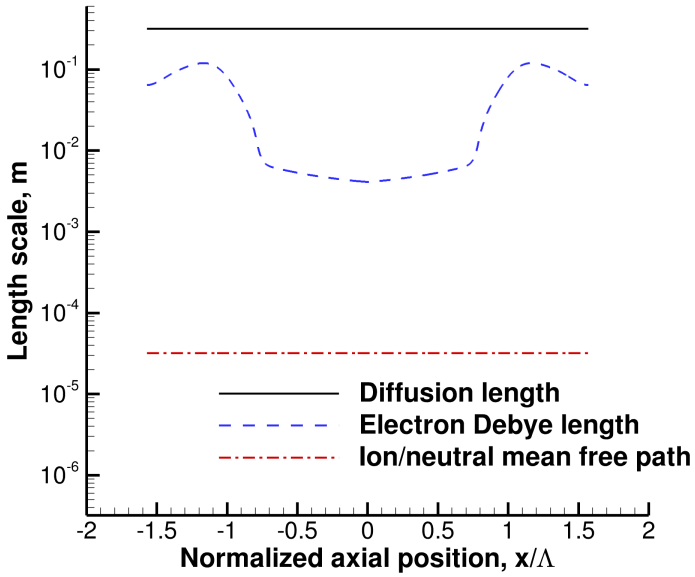

Figure 3.12: Characteristic length scales of simulation, including ion/neutral mean free path, electron Debye length, and diffusion length at $t = 4 \times 10^{-5}$ s $\approx 323\tau_{i,n} \approx 37,418\tau_{e,n}$ ($\Lambda/\lambda_{D,e} = 100, \Lambda/\lambda_{mfp,i} = 10,000$).

3.3.3 Comparison with Ambipolar Diffusion Approximation

To assess the validity of the ambipolar diffusion approximation as it pertains to capturing actual ambipolar diffusion, a flow case is run with the ambipolar diffusion approximation and with full electrostatic modeling. Details on how the ambipolar diffusion approximation is implemented in DK are included in Section 2.4.3. Table 3.4 lists initialization parameters and other flowfield

properties for the test case. Note that in the following plots, the horizontal axis (axial position) is normalized in terms of the Debye length as opposed to the diffusion length. As ions and electrons share the same macroscopic properties when the ambipolar diffusion approximation is invoked, profiles from the ambipolar diffusion approximation simulation are labeled 'ambipolar plasma'.

Table 3.4: Initialization parameters and other flowfield properties for $\Lambda/\lambda_{mfp,i} = 100$, $\Lambda/\lambda_{D,e} = 10,000$ test case.

Property	Value
Neutral number density, m^{-3}	6.12×10^{20}
Ion/Electron number density, m^{-3}	2.35×10^{15}
Ion mean free path, m	3.18×10^{-3}
Electron Debye length, m	3.18×10^{-5}

Figure 3.13a demonstrates that even for a plasma well into the ambipolar diffusion regime ($\Lambda/\lambda_{D,e}$ = 10,000), the ambipolar diffusion approximation does not capture the same number density profiles as full electrostatic modeling. Why this occurs is best understood through the axial velocity profiles (Figure 3.13b). When the ambipolar diffusion approximation is invoked, the axial velocity is effectively zero for $|x/\lambda_{D,e}| <$ 6,400. The velocity increases at the edge of the plasma as particles are diffusing outwards. However, the profile decreases towards zero for $|x/\lambda_{D,e}| >$ 13,000. This indicates that a few ambipolar particles have reached the edges of the computational domain with large axial velocities, but, averaged over the entire VDF, the computed bulk velocity is still small.

Conversely, with full electrostatic modeling and resolution of separate charged species, the ion and electron velocity profiles begin to linearly increase from 6,800 $< |x/\lambda_{D,e}| <$ 15,300, then drop off steeply to the boundaries of the domain. At the beginning of the simulation, the fastest electrons would have escaped from the center of the domain, leading to charge separation that induces the formation of ambipolar electric fields. Once the ambipolar field establishes itself, any electrons leaving the center of the domain will do so at the ambipolar diffusion rate with a proportional number of ions. However, the formation of an electric field will also allow the fastest ions to accelerate to the ends of the domain. Only a small portion of ions will do this (i.e., the

ions initialized with velocities at the tails of the Maxwellian VDF), as the peak ion axial velocity corresponds with a number density of around 10^{10} m^{-3}. In spite of this, the ion acceleration mechanism allows just enough ions to catch up with electrons to allow the domain to achieve local charge neutrality sooner and also ensures the ion and electron bulk velocities match throughout most of the domain.

By multiplying the Argon MEX hard sphere cross section by $1 + T_e/T_i$ (Equation 3.6), the ambipolar diffusion approximation simulation's ion diffusion rate can be artificially increased to match the ambipolar diffusion rate calculated at the end of the full electrostatic simulation. This extends the edge of the ambipolar plasma density profile by about $315\lambda_{D,e}$. This indicates that the differences between the ambipolar diffusion approximation and full electrostatic modeling results outside of the charge neutral region are due to unsteady electric field acceleration effects and not the enhanced ambipolar diffusion rate D_a.

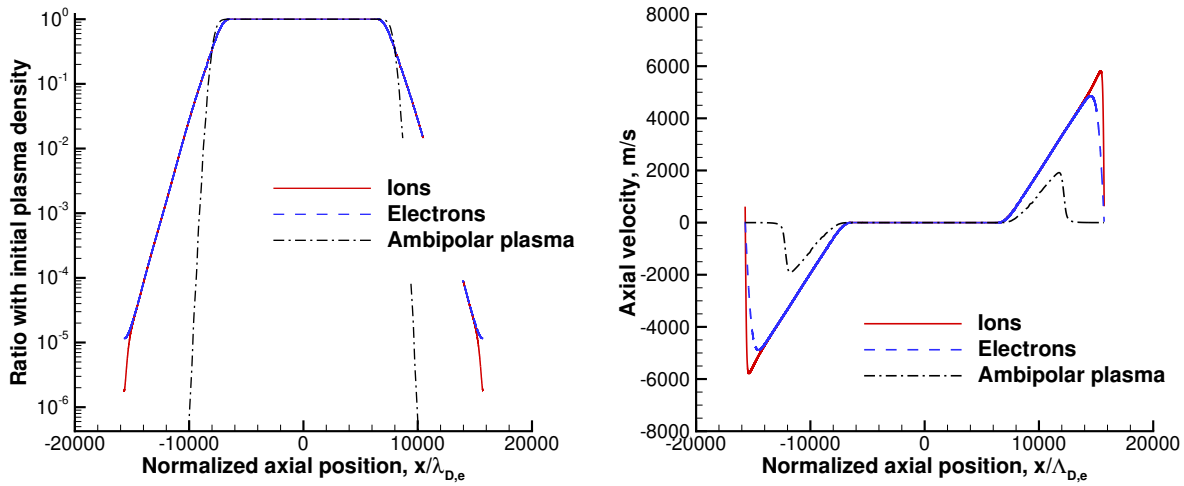

(a) Ratio with initial plasma density.

(b) Axial velocity.

Figure 3.13: Comparison of ratio with initial plasma density (left) and axial velocity (right) profiles for ambipolar diffusion approximation versus direct electric field modeling for ambipolar diffusion case at $t = 5 \times 10^{-5}$ s $\approx 4\tau_{i,n} \approx 755\tau_{e,n}$ ($\Lambda/\lambda_i = 100$, $\Lambda/\lambda_{D,e} = 10,000$).

3.4 Identification of Charged Species Diffusion Regime in Non-Hypersonic Continuum Flows

The Phelps criteria for charged species diffusion, originally developed for plasma discharge vessels, can successfully be used to identify diffusion regime under alternate flow conditions if a suitable diffusion length is chosen. The diffusion length should relate to the maximum length a charged particle can travel before an interaction of interest occurs, such as absorbing into a wall or traversing an entire computational domain. Given a sufficiently collisional plasma (i.e., $\Lambda/\lambda_{mfp,i} \gtrsim 10$), ambipolar diffusion will generally occur for a value of $\Lambda/\lambda_{D,e} \gtrsim 100$. Free diffusion will occur for $\Lambda/\lambda_{D,e} \lesssim 1.0$, and transitional diffusion will occur for $\Lambda/\lambda_{D,e}$ values in between.

These criteria should not be applied without context, particularly at large length scales. For example, a weakly ionized, dense gas with a charged species density on the order of 10^7 m^{-3} will have a Debye length on the order of 10^{-1} m. If $\Lambda \gtrsim \mathcal{O}(10)$ m, $\Lambda/\lambda_{D,e} \gg 100$. This does not necessarily mean the gas has strong electrostatic effects– it simply means the charged species diffusion can be modeled as ambipolar at length scales comparable to Λ because the weakly ionized gas is approximately quasineutral at length scales comparable to Λ. It follows that if one were to observe the flow at a smaller length scale, then Λ should be adjusted appropriately and the charged species diffusion regime should be re-classified.

The ionization fraction of the gas especially should not be used to identify charged species diffusion regime without context. For example, an ionization fraction of 3.84×10^{-6} is tested with $n_i = n_e = 2.35 \times 10^{11}$ m^{-3} and $n_n = 6.12 \times 10^{16}$ m^{-3} as well as $n_i = n_e = 2.35 \times 10^{15}$ m^{-3} and $n_n = 6.12 \times 10^{20}$ m^{-3}. Both gases are classified as weakly ionized. However, the former falls within the ion inertia regime with $\Lambda/\lambda_{mfp,i} = 0.01$ and the latter falls within the ambipolar diffusion regime with $\Lambda/\lambda_{mfp,i} = 100$. All flow cases tested used an ionization fraction of less than 10^{-5}, demonstrating how a range of complex plasma dynamics can occur at low plasma densities and ionization fractions.

If a Poisson solver is employed, agreement between the computed electric field and the theo-

retical ambipolar electric field (Equation 3.1) or the Langmuir and Tonks relation (Equation 3.5) should also not be used to identify charged species diffusion regime. With the exception of highly non-neutral plasmas, the computed electric field will generally agree with theory irrespective of charged species diffusion regime for continuum flows. This is demonstrated with Figures 3.6, 3.9 and 3.10. However, if either Equation 3.1 or 3.5 is computed for a given flow, the maximum of the theoretical electric field profile may be used to gauge whether the charged species density gradients are sufficiently large to warrant detailed consideration of plasma electrostatics. Similarly, a small computed value of the ambipolar electric field does not necessarily mean the flow is in the ambipolar diffusion regime; as in Figure 3.9, a small electric field is also characteristic of the free diffusion regime.

3.5 Guidelines for use of the Ambipolar Diffusion Approximation in Non-Hypersonic Continuum Flows

Based on this study, for a weakly ionized gas flow in the continuum regime, the ambipolar diffusion approximation may be justifiable if the following conditions hold:

(1) That the gas meets the three fundamental criteria for a plasma, as outlined in Section 1.2,

(2) That charged particles primarily collide with neutrals rather than each other, such that charged particles diffuse in a random-walk manner consistent with Fick's Law,

(3) That the ratio between the electron Debye length and the diffusion length scale (proportional to the characteristic length scale of the problem) has a ratio of at minimum $\Lambda/\lambda_{D,e} \sim 100$, and

(4) That sufficient time has elapsed for the self-induced ambipolar electric field to restore charge neutrality. This generally occurs on the order of several ion collision times and thousands of electron collision times.

These criteria do not guarantee simulations performed with the ambipolar diffusion approx-

imation will achieve results identical to full electrostatic field modeling, particularly for unsteady flows. For simulations with $\Lambda/\lambda_{D,e} \approx 100$, even within the region of ambipolar diffusion, the normalized charge separation may be as large as 1%. Whether this is physically significant depends on the application. In the case of free diffusion with $\Lambda/\lambda_{D,e} < 1.0$, the charge separation will exceed 10%. For transitional diffusion with $1.0 < \Lambda/\lambda_{D,e} < 100$, the charge separation may vary between 1% and 10%. Assessment of these criteria require reasonable confidence in the plasma properties *a priori*, especially in the electron number density distributions used to calculate the Debye length throughout the flow. However overall, these criteria represent an improvement from the conventional criterion used to justify the ambipolar diffusion approximation.

The consequences of incorrectly applying the ambipolar diffusion approximation (such as magnitude of percent error) will vary. Because the magnitude of the self-induced ambipolar electric field scales with the gradient in electron density (Equation 3.5), plasmas with higher densities will generate stronger electric fields leading to stronger acceleration or deceleration. Thus, the degree to which the ambipolar diffusion approximation underpredicts charged species diffusion rates scales with plasma density. Workarounds for replicating that enhanced diffusion rate without solving the Poisson equation require further investigation. For the unsteady gas expansion problem studied here, artificially modifying the ion MEX cross sections according to Equation 3.6 does not notably improve agreement between a simulation with the ambipolar diffusion approximation and with full electrostatic modeling.

The conclusions made in this study do not incorporate the possibility of doubly, triply, or negatively charged ions. Namely, if multiple ionic species are present, then the full definition of the Debye length (Equation 1.2) must be used, in which case the criteria presented here may not be valid. This study also did not examine the ambipolar diffusion approximation under the presence of imposed electric or magnetic fields. The presence of a magnetic field can result in anisotropic diffusion and a decoupling of ion and neutral flow velocities [27], in which the ambipolar diffusion approximation may not be valid.

3.6 Summary

The physics of charged species diffusion in non-hypersonic continuum flows was explored through simulation of an expanding weakly ionized gas. Crucially, the conventional criterion for justifying the ambipolar diffusion approximation, $\lambda_{D,e} \ll \lambda_{mfp,n}$, was disproved via theory and numerical simulation. Improved guidelines on identifying charged species diffusion regime and assessing the validity of the ambipolar diffusion approximation were proposed which synthesize fundamental plasma theory, existing guidelines for gas discharge vessels, and the results of this study. The efficacy of these criteria for hypersonic flows is investigated in Chapter 5.

Chapter 4

Assessment of Kinetic Methods for Simulation of Partially Ionized Transitional Hypersonic Flows

In this chapter the strengths and weaknesses of using DVM and DSMC-PIC for studying partially ionized, transitional-regime hypersonic flows are assessed. In Section 4.1, computational requirements for any kinetic simulation method with full electrostatic modeling are discussed. In Section 4.2, a novel 1D3V stagnation streamline model for DVM is derived, verified, and benchmarked with a DSMC equivalent. An analysis of VDFs from both DSMC and DVM is also conducted in Section 4.3.2, revealing VDF features unique to a shock wave in the presence of a solid wall boundary (as opposed to a standard standing shock wave). Finally, in Section 4.4, DK and MPIC are evaluated in terms of suitability for the final objective of this dissertation work: a 1D study of charged species diffusion in transitional regime hypersonic flows.

4.1 Computational Requirements for Kinetic Electrostatic Modeling

Kinetic simulation of partially ionized hypersonic flows with full electrostatic modeling is highly nontrivial, which contributes to its lack of widespread usage in the literature and the popularity of the ambipolar diffusion approximation. The primary challenge is in resolving the enormous difference in length scales between the hypersonic flowfield and the plasma dynamics. Hypersonic vehicles and spacecraft have characteristic length scales on the order of 0.1 m to 1.0 m. The smallest length scale in a typical transitional regime hypersonic flowfield is the collision mean free path, which can range between order 10^{-4} m to 10^{-1} m for flow conditions between 60 km and 100

km on Earth. In rarefied gas dynamics, interactions of interest can occur on scales smaller than or similar to the mean free path (hence the definition of the Knudsen number, Equation 2.8). A computational grid cell much larger than the mean free path will not capture these critical interactions. Subsequently, an estimated maximum ratio of length scales for a kinetic hypersonic flow simulation without electrostatic modeling is of $\sim \mathcal{O}(10^4)$. This is not to say kinetic simulation of hypersonic flowfields without electrostatic modeling is inexpensive; rather, resolution of configuration space tends to be a smaller contributor to total computational cost compared to resolution of collisions, trace species and trace species reactions, excitation states, gas-surface interactions, and other high-fidelity physics effects. Many of these interactions are highly coupled, making hypersonic flow environments some of the most challenging to model.

A "complete" simulation of a partially ionized hypersonic flowfield with air would include resolution of the aforementioned flowfield physics effects in addition to detailed plasma dynamics. The smallest length scale to resolve in a plasma physics simulation is the Debye length (Equation 1.2 and/or Equation 1.3). Similar to the mean free path, resolution of the Debye length is required to observe charged species interactions on length scales shorter than the characteristic distance required to shield out charge imbalances from said interactions. Depending on the freestream conditions and chemistry modeling assumptions, transitional hypersonic flow simulations on Earth may observe electron densities between 10^{16} m^{-3} and 10^{20} m^{-3} in the post-shock region [112]. Using an electron temperature estimate of T_e = 5,000 K, $\lambda_{D,e}$ will range between $\mathcal{O}(10^{-5})$ to $\mathcal{O}(10^{-7})$ m in the post-shock region. Thus, the range of length scale ratios for a kinetic hypersonic flow simulation with electrostatic modeling is between $\sim \mathcal{O}(10^4)$ and $\sim \mathcal{O}(10^7)$. Depending on the plasma dynamics, resolution of the electron Debye length with anywhere between 2-16 grid points per $\lambda_{D,e}$ is recommended [76].

The consequences of not resolving the Debye length are much more severe than the consequences of not resolving the mean free path due to strong coupling between the plasma properties and the electric field solver. To electrostatically model a plasma with as few physical assumptions as possible, the Poisson equation (Equation 2.31) must be solved. Equation 2.31 directly uses charge

separation as input, and local charge imbalances are only resolved within a quasineutral plasma at length scales comparable to $\lambda_{D,e}$ (see Section 1.2). Therefore, if $\Delta x > \lambda_{D,e}$, a small charge perturbation between cells may result in a large, unphysical electric field. In turn, the charged species will be accelerated over large distances, exacerbating the charge separation required as input by the Poisson solver. This creates a feedback loop that leads to numerical instability in the plasma. In explicit PIC simulations, this effect is formally known as the grid-instability and leads to unphysical heating of particles when $\Delta x \gg \lambda_{D,e}$ [113]. A similar feedback loop also results if the simulation time step does not resolve the plasma frequency (Equation 1.4), where the plasma frequency characterizes the time required for a plasma to resolve a charge discrepancy. For the plasma densities described above, the plasma frequency will have an order magnitude between 10^{10} s^{-1} and 10^{12} s^{-1}. This requires a time step $\Delta t \leq 10^{-10}$ s for explicit advection methods.

Thus, the computational cost for simulating a partially ionized hypersonic flow in the transitional flow regime with kinetic simulation methods is expected to be very large. A common compromise to make the problem more tractable is artificially increasing the electron mass to decrease the plasma frequency. Typical values include $m_e/m_i = 0.25$ or $m_e/m_i \sim \mathcal{O}(10^{-2})$ [38, 114, 115]. However, this only impacts the size of Δt, not $\lambda_{D,e}$. The most straightforward approach to reduce computational cost in configuration space is reducing dimensionality. A 1D setup is also advantageous for studying the fundamentals of plasma electrostatics and charged species diffusion, as particles can only move in one direction. The contours of a vehicle may result a non-zero curl of the electric field, which may result in the generation of a non-negligible magnetic field necessitating resolution of charged particle gyromotion and ultimately adding complexity to the problem.

DVM still requires discretization of three velocity dimensions to capture real gas effects (see Equation 2.11), however, the computational cost of 3V is not so prohibitive that DVM is infeasible. In fact, DVM has seen increased use in practical gas dynamic simulation in recent years, such as the optimization of two-dimensional vehicle design [116], simulation of a 1D shock in gas mixtures of diatomic molecules with both rotational and vibrational energy states [117], and simulation of supersonic flow around sharp-base geometries [118]. Thus, to assess the suitability of DVM and

DSMC-PIC for studying partially ionized hypersonic flows with electrostatic modeling and minimize computational expense of the final study, a 1D3V computational setup is preferred.

4.2 Development of a One-Dimensional Stagnation Streamline Model for Discrete-Velocity Methods

G.A. Bird developed a theoretical framework for modeling the flowfield along a 1D stagnation streamline of a blunt body moving at hypersonic speed for DSMC [34, 29]. The stagnation streamline is modeled as a constant streamwise (i.e., wall-normal, $\hat{x}$) area flow with an inflow boundary condition on one end and a wall boundary condition on the other. This results in an unsteady shock propagating outwards from the wall. To achieve steady state, once the shock reaches a desired standoff distance, particles are selected and removed according to probabilistic criteria such that mass, momentum, and energy are conserved. Bird's DSMC stagnation streamline model has successfully been used to study many different non-equilibrium phenomena in hypersonic flows, including energy transfer between vibrational and rotational modes [60], shock front radiation [119], dissociation models [120], electron impact ionization [35], and electrostatic effects [35, 38].

Bird's theoretical framework is applicable to any stagnation streamline and thus can form a mathematical basis for a proposed DVM model. Since DVM does not use discrete particles, the probabilistic criteria for removal must be translated into determinstic criteria; that is, the exact VDF of removed particles must be known. The following sections derive this VDF and other relevant equations, then benchmark the new DVM model with its DSMC equivalent both in terms of macroscopic flow properties and microscopic properties like the VDF.

4.2.1 Theoretical Framework

A detailed overview of Bird's DSMC stagnation streamline model and relevant considerations is provided. When the simulation begins, undisturbed freestream particles enter from the inlet side of the computational domain. When the particles reach the solid surface at the other end, they reflect and form an unsteady shock wave which propagates back upstream. The shock wave

propagates until it reaches a specified point along the streamline x_s, where its density ratio with the freestream reaches a specified value. For a monatomic gas, this value should be $\rho/\rho_\infty = 4.0$. Conventionally, x_s is chosen to be the shock standoff distance, known *a priori* from axisymmetric simulation or other literature. Particles are then removed from a downstream region from $x = x_r$ to $x = 0$ m (the wall), such that the removed mass flux equals the inflow mass flux. x_r is a simulation input parameter that can hypothetically take any value. The removal operation is performed over a period of time until steady state stagnation streamline profiles of desired flowfield properties are achieved. Figure 4.1 displays these positions in relation to each other along a one-dimensional axis.

During the removal operation, particles must be chosen for removal in a manner that conserves mass, momentum, and energy. Conservation of mass is represented by

$$d(\rho u) = -\dot{m}_r dx \tag{4.1}$$

where ρ is mass density and $\dot{m}_r$ is the rate of molecular mass removal per unit volume. Conservation of momentum is represented by

$$dp = -d(\rho u^2) - u_r \dot{m}_r dx \tag{4.2}$$

where p is momentum and u_r is the bulk streamwise velocity of the removed particles. Taking the derivative of the first right hand side term in Equation 4.2 gives

$$dp = -\rho u du - u d(\rho u) - \overline{u}_r \dot{m}_r dx \tag{4.3}$$

which reduces to the continuum momentum equation if Equation 4.1 is satisfied and $u_r = u$. In other words, the bulk streamwise velocity of the remaining flow will be unchanged if the VDF of removed particles also has the same computed bulk streamwise velocity. By selecting particles for removal assessment independent of their position within the removal region, the bulk streamwise velocity of the flow is unchanged during removal, and conservation of momentum is satisfied. Lastly, conservation of energy is represented by

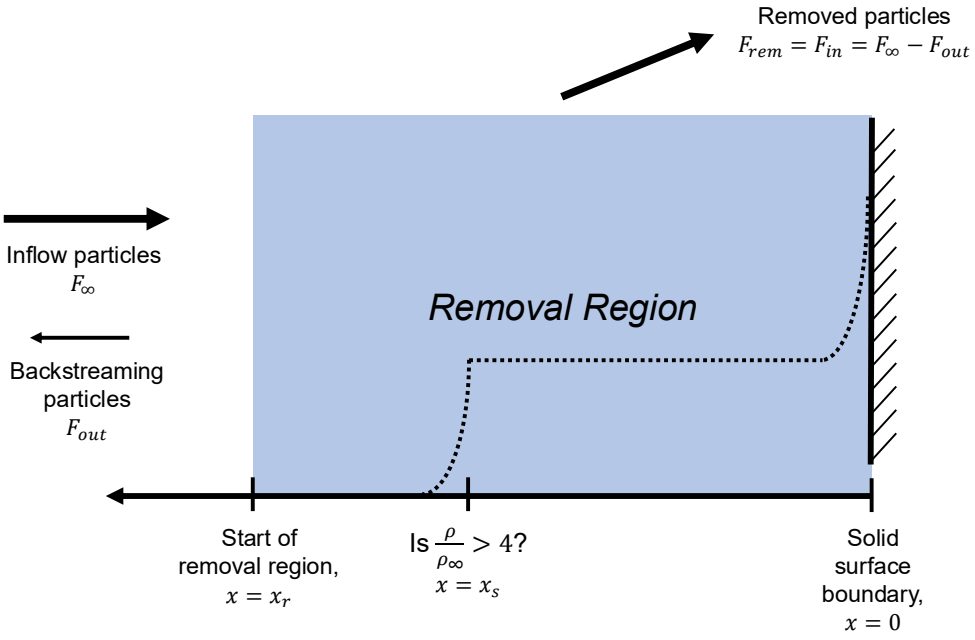

Figure 4.1: Notional schematic for Bird's one-dimensional stagnation streamline scheme [34, 29]. The dashed line represents an example stagnation streamline density profile. The limit of $\rho/\rho_\infty > 4.0$ is valid for a monatomic gas only.

$$d(\rho u h_0) = -\dot{m}_r \overline{e}_{0,r} dx \tag{4.4}$$

where h_0 is the stagnation enthalpy of the flow and $\overline{e}_{0,r}$ is the average specific energy of the removed particles. If particles are removed in a manner that satisfies

$$\overline{e}_{0,r} = h_0 \tag{4.5}$$

then substituting Equation 4.1 into Equation 4.4 results in

$$dh_0 = 0 \tag{4.6}$$

From Equation 4.5, it follows that the average specific energy of the removed particles must exceed that of the flow by $k_B T$ per particle over the average translational flow value of $1.5 k_B T$. Thus, the ratio between the specific energy of the removed particles and the specific energy of the flow will be 5/3. In DSMC (which always requires three velocity translational degrees of freedom), this criterion is met by generating a random number between 0.0 and 1.0 and deleting the selected particle if

$$\texttt{RAND} < \frac{v_y^2 + v_z^2}{(v_y^2 + v_z^2)_{max}} \tag{4.7}$$

$(v_y^2 + v_z^2)_{max}$ is continuously updated if a particle with a larger value is located. Generally, if all species are similar in mass, the same $(v_y^2 + v_z^2)_{max}$ value can be used globally for all species and cells.

If a gas mixture is simulated, then the removed mass flux of each type of atom in the mixture must equal the inflow mass flux of each type of atom. For example, if 50 neutral argon atoms enter the computational domain in one time step, 50 argon atoms must be removed in the same time step. The number of removed argon neutrals versus the number of removed argon ions does not matter as long as either species is selected for removal with the same criteria. This ensures conservation of mass is satisfied regardless of chemical reactions throughout the flowfield.

Since electrons have a much smaller mass than heavy species, they should never be selected for removal using the same $(v_y^2 + v_z^2)_{max}$ as the rest of the flow. When simulating hypersonic flows, electrons are almost never included in the freestream gas composition. Thus, the inflow number of electrons is always zero. However, if only ions are removed with the stagnation streamline scheme, a charge deficit will result. Therefore, if charged species are present, then the number flux of removed electrons should equal the charge number flux of removed ions. For example, if a doubly charged ion is removed, two electrons must be removed. This conserves electrical charge in the flow and effectively functions as a "perpendicular ambipolar diffusion approximation", as it assumes an ion that diffuses away from the stagnation streamline will always bring a number of electrons equal to its charge with it.

In verifying the DVM model, multiple ambiguous elements of the DSMC model are investigated. This includes choice of the removal region starting point x_r, choice of whether the removal scheme should be performed at every time step or only when the density ratio at x_s drops below the specified value, and how well conservation of energy is maintained for non-equilibrium hypersonic flow conditions.

4.2.2 Discrete-Velocity Model Derivation

Assuming a single species flow, the derivation for the DVM stagnation streamline model is as follows. The VDF of removed particles is labeled $f_r(x, \vec{v}) = n_r(x)\hat{f}_r(x, \vec{v})$. $\hat{f}_r(x, \vec{v})$ can be further divided into components (where i, j, and k represent each combination of x, y, and z), each normalized to one,

$$\hat{f}_{r,i}(x, v_i) = \frac{1}{n_r(x)} \int_{-\infty}^{\infty} dv_j \int_{-\infty}^{\infty} f_r(x, \vec{v}) dv_k \tag{4.8}$$

such that

$$f_r(x, \vec{v}) = n_r(x)\hat{f}_{r,x}(x, v_x)\hat{f}_{r,y}(x, v_y)\hat{f}_{r,z}(x, v_z) \tag{4.9}$$

Conservation of particle number for a single species gas requires

$$F_{in} = F_\infty - F_{out} = \alpha_r \int_0^{x_r} dx \int_{-\infty}^{\infty} f_r(x, \vec{v}) d\vec{v} \tag{4.10}$$

where F_{in} is the freestream number flux (F_∞) minus the number flux of particles which diffuse upstream and exit the computational domain (F_{out}). F_{out} is non-negligible to conserve momentum exactly, as discussed in Section 4.3.2. α_r is defined as the global number flux normalization constant, whose value is known once $f_r(x, \vec{v})$ is known.

To ensure conservation of momentum via $u_r = u$, the DVM removal scheme must ensure $\hat{f}_{r,x}(x, v_x) = \hat{f}_x(x, v_x)$. In order for the un-normalized VDF to preserve this relation, the local momentum conservation constant $C_r(x, v_x)$ is defined as

$$C_r(x, v_x) = \hat{f}_x(x, v_x) \left(\int_{-\infty}^{\infty} dv_y \int_{-\infty}^{\infty} f_r(x, \vec{v}) dv_z \right)^{-1} \tag{4.11}$$

This normalization constant can also be calculated once $f_r(x, \vec{v})$ is known.

Since particles with high values of $v_y^2 + v_z^2$ are more likely to be removed, then the magnitude of $\hat{f}_{r,y}(x, v_y)\hat{f}_{r,z}(x, v_z)$ will be lower towards the tails of the velocity domain and larger towards the center. This can be accomplished by defining $\hat{f}_{r,y}(x, v_y)\hat{f}_{r,z}(x, v_z)$ as

$$\hat{f}_{r,y}(x, v_y)\hat{f}_{r,z}(x, v_z) = \frac{v_y^2 + v_z^2}{(v_y^2 + v_z^2)_{max}} \hat{f}_y(x, v_y)\hat{f}_z(x, v_z) \tag{4.12}$$

where $(v_y^2 + v_z^2)_{max}$ is calculated from the maximum range of the DVM velocity domain, and thus is constant throughout the simulation. The expression also indicates that cells with higher temperatures (i.e., wider VDFs) will have more particles removed from them. This holds true even for the DSMC scheme: particles are randomly chosen throughout the removal region to evaluate Equation 4.7 (thus conserving momentum), but whether the particle is actually removed depends on its value of $v_y^2 + v_z^2$ (thus conserving energy according to Equation 4.5).

The final term to define in Equation 4.9 is $n_r(x)$. If steady state is achieved in a DSMC stagnation streamline simulation, then each time step roughly the same number of particles will be

removed from each cell. Thus, $n_r(x)$ must have an analytical form. An *ansatz* for an expression for $n_r(x)$ is proposed as

$$n_r(x) = n(x)\frac{F_{in}}{\int_0^{x_r} n(x)dx}\frac{T(x)}{T_{max}} \tag{4.13}$$

where $n(x)$ is the original number density of the cell and T_{max} is the maximum average translational cell temperature of the removal region during that time step. This expression is verified with numerical experiment in the following section.

Altogether, the final expression for $f_r(x, \vec{v})$ in terms of the original VDF is given as

$$f_r(x, \vec{v}) = \alpha_r \frac{F_{in}}{\int_0^{x_r} n(x)dx}\frac{T(x)}{T_{max}}\frac{v_y^2 + v_z^2}{(v_y^2 + v_z^2)_{max}} C_r(x, v_x) f(x, \vec{v}) \tag{4.14}$$

where $n(x)$ in Equation 4.13 is included in the definition of $f(x, \vec{v})$. To perform the stagnation streamline removal operation in DVM, Equation 4.14 is subtracted from each cell in phase space from x_r to $x = 0$ m. This derivation does not make any assumptions about the shape of the VDF in the flow, and thus is not limited to continuum or non-continuum flow conditions as is characteristic of hypersonic flows. That is, since Equations 4.1–4.6 are defined and verified for a hypersonic stagnation streamline flow around a blunt vehicle, it is not yet known whether either the DSMC or the DVM model can accurately simulate a stagnation streamline in other flow circumstances, such as in fusion reactors.

4.2.3 Verification of Discrete-Velocity Model

Equation 4.13 is first verified with a series of 0D DSMC simulations. Each is initialized with a number of cells filled with particles at equilibrium. Only the DSMC stagnation streamline removal operation is performed: advection, collisions, and chemistry are all neglected, and particles from separate cells do not interact with each other. As such, particle weight is irrelevant. To analyze results, the fraction of particles removed per cell is plotted against a varied parameter of interest. If applicable, the coefficient of determination R^2 is evaluated, where a R^2 value of 1.0 suggests the

linear regression model's values exactly match the observed values from the simulation.

To conserve mass, the fraction of particles removed per cell should be directly proportional to the inflow number flux (equal to the number flux of removed particles). The first numerical experiment compares two cells with the same temperature, $T_1 = T_2 = 500$ K. The first cell contains 2×10^6 particles and the second contains 10^6. This tests the fraction of particles removed from cells with different densities. The inflow number flux varies between 10^2 to 2.5×10^4 particles. This ensures the number of particles removed is small compared to the total number of particles in the simulation and thus the cell temperature of 500 K remains relatively unchanged during removal.

Figure 4.2 demonstrates a clear linear relationship between the inflow number flux and the fraction of particles removed, with a coefficient of determination equal to $R^2 = 0.99989$ (taken from both data sets combined). This verifies that $n_r(x)$ is linearly proportional to F_{in}. Effectively the same fraction of particles is removed from each cell, despite the difference in number density. This suggests under translational equilibrium, as long as the cell temperatures are equal, the particle removal operation will remove an equal fraction of particles from each cell.

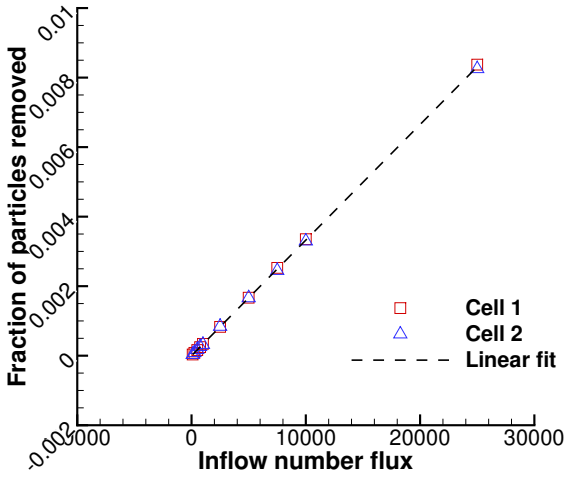

Figure 4.2: Linear relationship between the inflow number flux and the fraction of particles removed. Cell 1 contains 2×10^6 particles, and Cell 2 contains 10^6 particles [106].

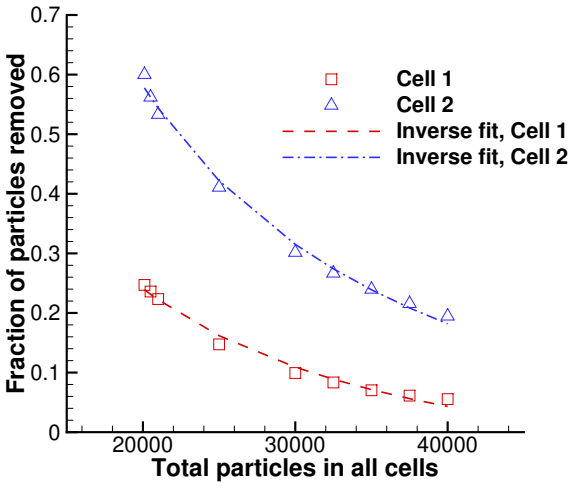

Figure 4.3: Inverse relationship between total number of particles in all cells and the fraction of particles removed. Cell 1 contains 2×10^4 particles at 5,000 K, and Cell 2 contains a varying number of particles at 20,000 K [106].

The second numerical experiment verifies that when the total number of particles in the removal region $\int_0^{x_r} n(x)dx$ increases but the inflow number flux remains the same, the fraction of particles removed per cell will decrease. Two cells are initialized with $T_1 = 5,000$ K, $T_2 = 20{,}000$ K, and $F_{in} = 5 \times 10^3$ particles. The number of particles in Cell 1 remains constant at 2×10^4 particles and the number of particles in Cell 2 is varied between 10^2 and 2×10^4. An inverse relationship results as plotted in Figure 4.3, with coefficients of determination of $R^2 = 0.99989$ and 0.99888 for Cell 1 and 2, respectively. Since Cell 2 has a larger temperature, its fraction of removed particles is always greater than Cell 1. However, as the number of particles in Cell 2 increases, there are more particles available to select, and thus the overall fraction of removed particles per cell is reduced by the same amount. This demonstrates $n_r(x)$ is inversely proportional to $\int_0^{x_r} n(x)dx$.

Equation 4.7 indicates $n_r(x)$ must depend on temperature, as a higher temperature cell will have more particles with $v_y^2 + v_z^2$ near $(v_y^2 + v_z^2)_{max}$. In the third numerical experiment, three cells are simulated with 2×10^6, 5×10^5, and 10^6 particles in each, and $F_{in} = 10^4$ particles. Cell 1 has a temperature of 5,000 K, and Cell 3 has a temperature of 10,000 K. Cell 2's temperature is varied between 100 K to 100,000 K, which results in T_{max} changing depending on the value of T_2. Incidentally, this temperature range covers the range of values typically encountered in hypersonic flow simulations in air. Figure 4.4 plots the fraction of particles removed from each cell as T_2 increases. As expected, the fraction of particles removed from the other cells decreases when more particles are eligible in Cell 2 via a higher initialization temperature. Crucially, the cell with the highest temperature always has the highest fraction of particles removed. This behavior was also observed in the previous numerical experiment. When the temperatures between two cells are identical (i.e., the point where $T_2 = T_1 = 5,000$ K or $T_2 = T_3 = 10,000$ K), the fraction of particles removed from each cell is effectively equal regardless of their number densities. As T_2 continues to increase, the fractions of particles removed from the other cells approach zero. Subsequently, the fraction of particles removed from Cell 2 asymptotically approaches F_{in}/n_2. Ultimately, this indicates $n_r(x)$ is inversely related with T_{max}, and when $T(x) = T_{max}$, the fraction of particles removed in that cell is maximized. Both criteria are fulfilled by making the final term in Equation

4.13 equal to $T(x)/T_{max}$. Because the relationship between $T(x)/T_{max}$ and $n_r(x)$ is nonlinear, a linear regression is not performed for this test.

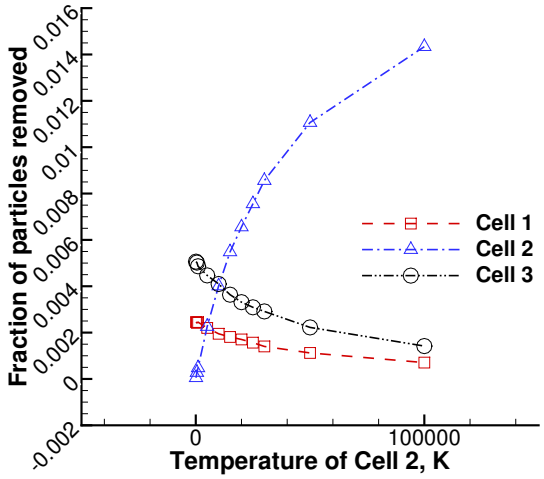

Figure 4.4: Variation of the fraction of particles removed when the temperature of one cell is varied. Cell 1 contains 2×10^6 particles at 5,000 K, Cell 2 contains 5×10^5 particles at varying temperatures, and Cell 3 contains 10^6 particles at 10,000 K [106].

Through these numerical tests, Equation 4.13 is verified piece-wise. Equation 4.14 is verified in totality by implementing the entire DVM stagnation streamline scheme and directly checking for conservation of mass, momentum, and energy. Conservation of mass and momentum for the DVM stagnation streamline model are readily verified by outputting values of inflow number flux, removed number flux, and streamwise bulk velocities prior to and after the removal procedure is performed in one time step.

Conservation of energy requires more consideration. In Bird's original model, the degree to which conservation of energy should be maintained is not well defined. More specifically, Equations 4.2–4.6 are strictly valid if the entire flow is at translational equilibrium, as is the case with the numerical experiments. However, in a kinetic simulation of a rarefied stagnation streamline, the entire removal region may not be at translational equilibrium. Moreover, the region over which the average specific translational energy should be calculated to verify conservation of energy is not well defined. If the average specific translational energy is computed over the entire computational

domain, then its value is dominated by the freestream temperature for an infinitely long domain as opposed to the shock layer. Nevertheless, the DSMC model achieves agreement with axisymmetric simulations [34, 1] and experiment [34]. Thus, for the purposes of verification, the DVM model is only expected to maintain conservation of energy to the same extent as the DSMC equivalent model.

First, conservation of energy is checked in each of the 0D numerical experiments. On a per cell basis, Equation 4.5 is maintained with a percent error ranging from 9.38×10^{-4} to 1.02×10^{-2}. This forms the lower bound of expected error from the DSMC stagnation streamline removal procedure. Then, conservation of energy is assessed in a series of stagnation streamline simulations run with the DVM solver and an in-house DSMC solver [60]. The ratio of the average specific energy of the removed particles to the average specific energy of the flow $\overline{e}_{0,r}/\overline{e}_0$ is calculated and plotted as x_r is moved further upstream. The specific energy per cell is calculated with the following equation,

$$e_0(x) = \frac{1}{2}m\left[\left(\int_{-\infty}^{\infty}(v_x^2 + v_y^2 + v_z^2)\frac{f(x,\vec{v})}{n(x)}d\vec{v}\right) - u(x)^2\right] \tag{4.15}$$

such that the average $\overline{e}_0$ is taken from all values of $e_0(x)$ in the removal region. The range of x_r includes positions upstream and downstream of the shock to assess the dependence of conservation of energy on x_r. This analysis is performed for a Mach 35 flow at 60 km and a Mach 39 flow at 81 km using non-reactive argon to model the gas. Simulation initialization parameters for these flow conditions are included in Tables 4.1 and 4.2.

Ultimately, neither DSMC nor DVM precisely maintains the 5/3 conservation of energy relationship with the same degree of error as the 0D cells at translational equilibrium. Figures 4.5a and 4.5b plot the computed conservation of energy ratio with x_r for the 60 km and 81 km flow cases, respectively. For both flow cases, as x_r is placed further upstream, the conservation of energy ratio decreases. Intuitively, when the removal region is increased in size by including more of the upstream flow, some particles will inevitably be removed with low values of $v_y^2 + v_z^2$. This decreases the average specific energy of the removed particles. Across all flow cases, DVM maintains

conservation of energy with a maximum percent error of 3.60%. DSMC maintains conservation of energy with a maximum percent error of 19.0%, where the points of maximum error occur when x_r is inside the shock layer. Thus, if 19.0% forms the upper bound of expected error from the DSMC stagnation streamline removal procedure, then the DVM model successfully maintains conservation of energy. This concludes the numerical verification of the 1D DVM stagnation streamline model proposed in Section 4.2.2.

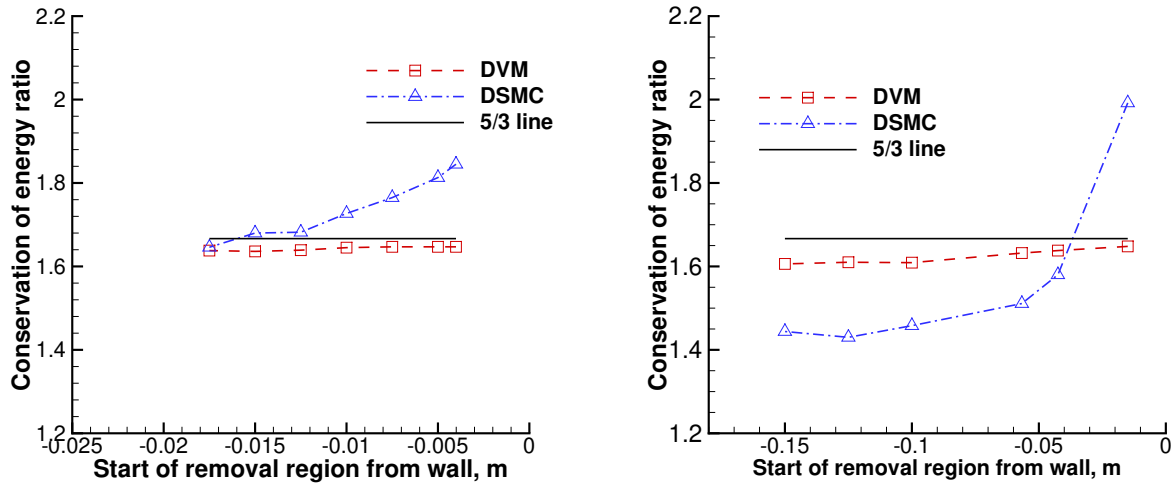

(a) Mach 35 flow case at 60 km.

(b) Mach 39 flow case at 81 km.

Figure 4.5: Maintenance of conservation of energy relation with varying removal region lengths for DVM and DSMC [106].

4.3 Results

To further verify the DVM stagnation streamline model and compare it to its DSMC equivalent, a series of hypersonic flow simulations are run with DSMC and DVM. The flow conditions are loosely based on the re-entry trajectory of the Stardust Sample Return Capsule. Three points in the trajectory are chosen for analysis: 51 km at Mach 25; 60 km at Mach 35; and 81 km at Mach 39 (see Figure 1.3). The first flow case lies within the continuum regime, and is the point of peak pressure on the vehicle forebody. The second flow case is at the boundary of the transitional and continuum flow regimes, and is the point of peak heating on the forebody. The third flow case is

Table 4.1: Freestream and wall input parameters.

Parameter	51 km	60 km	81 km
Freestream number density, molecules/m^3	1.50×10^{22}	4.87×10^{21}	2.64×10^{20}
Freestream temperature, K	253	238	222
Freestream velocity, m/s	7,410	10,350	10,810
Freestream Mach number	25	35	39
Freestream mean free path, m	8.44×10^{-5}	2.55×10^{-4}	4.59×10^{-3}
Wall temperature, K	3300	3750	2000

Table 4.2: Simulation grid and timestep parameters.

Parameter	51 km	60 km	81 km
Domain length, m	0.0200	0.0200	0.200
Points in configuration space	2,800	900	500
Points per dimension of velocity space (DVM only)	144	160	160
Range of velocity space, m/s (DVM only)	$\pm$15,900	$\pm$20,400	$\pm$20,400
Global particle weight (DSMC only)	5.36×10^{14}	5.41×10^{14}	5.28×10^{14}
Start of removal region from wall, DVM, m	0.0150	0.0125	0.100
Start of removal region from wall, DSMC, m	0.0200	0.0175	0.150

in the transitional flow regime. For each flow case, the freestream number density, freestream temperature, and wall temperature directly use the Stardust flow conditions. Due to their difference in specific gas constant, when modeling argon gas instead of air either the Mach number or the freestream velocity can remain consistent with the original Stardust conditions: this study chooses to keep the Mach number consistent and adjusts the freestream velocity accordingly. x_s is selected to yield a shock standoff distance similar to the axisymmetric simulations in References [3] and [121], where the Stardust vehicle forebody consists of a 60 deg half-angle, spherically blunted cone with a nose radius of 0.2202 m.

Tables 4.1 and 4.2 provide the simulation initialization conditions for each solver and each flow case. Configuration space grids are uniformly generated to resolve the local mean free path at all points in the flowfield with at least four grid points. For the DVM simulations, the velocity space grid is uniformly generated across a range that resolves the VDF width of the maximum temperature expected throughout the course of the entire simulation. Specifically, as the unsteady shock propagates upstream (prior to the start of the removal procedure), the temperature increases over time. Once removal begins, the higher energy particles are removed, and the temperature decreases to a more realistic, stable value. For accurate results, the velocity space grid must be wide enough to capture the maximum temperature encountered throughout the entire simulation, including the transient period. The velocity space grid resolution is chosen to sufficiently resolve the finest features of the VDF, as seen in Section 4.3.2. All simulations are run until steady state is achieved.

Section 4.3.1 compares macroscopic flow properties between the DVM and DSMC models. Section 4.3.2 compares the VDFs along the streamline for each of the models and provides insight into VDF features unique to shocks in the presence of a reflecting boundary. For the plotted results, DSMC profiles are time-averaged from the time step when particle removal begins to the end of the simulation. The DVM profiles are taken from the final time step of the simulation.

4.3.1 Comparison of Macroscopic Properties

In this section, the density ratio with the freestream, streamwise bulk velocity, translational temperature, and gradient-length local Knudsen number are plotted as a function of the distance from the wall for each flow case. The profile for $\mathrm{Kn}_{\mathrm{GLL}}$ for the DSMC results is smoothed for readability.

Figure 4.6 plots the macroscopic flow properties for the Mach 25 case at 51 km. The density ratio (Figure 4.6a), streamwise bulk velocity (Figure 4.6b), and translational temperature (Figure 4.6c) profiles agree within 9.13%, 72.0%, and 76.1% RMSE (normalized to the range of the data). The RMSE value for the density ratio profiles is reasonable, especially given DK's use of the BGK operator for a non-equilibrium flow. It is also possible that rigorous iterative adjustment of x_s and x_r may improve agreement, particularly for the shock location. However, this would be highly time consuming without necessarily leading to improved insights into the DVM model. The RMSE of the velocity and temperature profiles is more extreme, primarily due to the very strong gradients of a shock. For example, even if the location of the shock is shifted by a small amount between models, for a freestream temperature of 253 K and a post-shock temperature of 50,000 K, the reported absolute error is 49,747 K. Thus, throughout this section and the next, agreement between models is assessed in terms of overall shape and specific quantities like peak post-shock temperature.

Continuing analysis of the 51 km flow case, the temperature distributions (Figure 4.6c) agree as both models maintain relatively uniform profiles across the shock layer and slightly increase near the wall before tapering off. The percent error between the maximum post-shock temperatures is 0.380%. The DSMC profile contains a small bump at the shock, whereas the DVM profile does not. Since translational temperature is defined in terms of a Maxwellian distribution, but the VDFs in the vicinity of a shock are non-Maxwellian (see Figure 4.10a), temperature is a poorly defined quantity in this environment. That said, the bump appearing in only the DSMC profile may be caused by inherent differences in collision operators between DSMC and DVM. Section 4.3.2 expands further on this result and possible explanation. As for the gradient-length local Knudsen

number (Figure 4.6d), the point where $\mathrm{Kn}_{\mathrm{GLL}}$ exceeds 0.05 for each profile coincides precisely with the shock standoff distance. Both codes predict similar degrees of continuum and non-continuum behavior for the same locations in the flowfield.

Figure 4.7 presents the macroscopic flow properties for the Mach 35 case at 60 km. The density ratio (Figure 4.7a), streamwise bulk velocity (Figure 4.7b), and gradient-length local Knudsen number (Figure 4.7d) profiles show similar levels of agreement to the 51 km case. As 60 km is on the lower altitude end of the transitional flow regime, the temperature profiles for both codes are slightly rounded in the post-shock region as opposed to a uniform flat temperature. The DSMC profile still peaks at the shock front whereas the DVM profile does not. The percent error between the maximum post-shock temperatures is 2.41%. This is potentially due to the BGK operator, which is discussed in Section 4.3.2.

Figure 4.8 presents the macroscopic flow properties for the Mach 39 case at 81 km. As the most non-equilibrium flow case, this case is most important for evaluation of suitability between DSMC and DVM. The DVM scheme yields a slightly wider shock layer than DSMC (Figure 4.8a) which causes all other profiles to be spatially offset from DSMC. The least amount of agreement is demonstrated in the temperature profiles (Figure 4.8c), with 3.16% error between maximum post-shock temperatures. Of particular note, the DSMC temperature relaxes to the freestream value more rapidly than the DVM models. In the literature, this is attributed to the inability of the BGK operator to produce the correct Prandtl number [122]. This is supported by examining the flow's VDFs upstream of the shock in Section 4.3.2. There is a small discontinuity in the DVM $\mathrm{Kn}_{\mathrm{GLL}}$ profile at the start of the removal region (see Figure 4.8d); this is caused by the first order derivative of $\mathrm{Kn}_{\mathrm{GLL}}$ and does not impede interpretation of results. It does, however, demonstrate that x_r may require placement further upstream if derivatives are required for quantites of interest.

Both DVM and DSMC are sensitive to the chosen start of the removal region x_r. Decision criteria for choice of x_r are not well defined in the literature, though most studies place x_r after the shock standoff distance. However, in this study both DVM and DSMC require x_r placement upstream of the shock; otherwise discontinuities appear at x_r in all macroscopic property profiles.

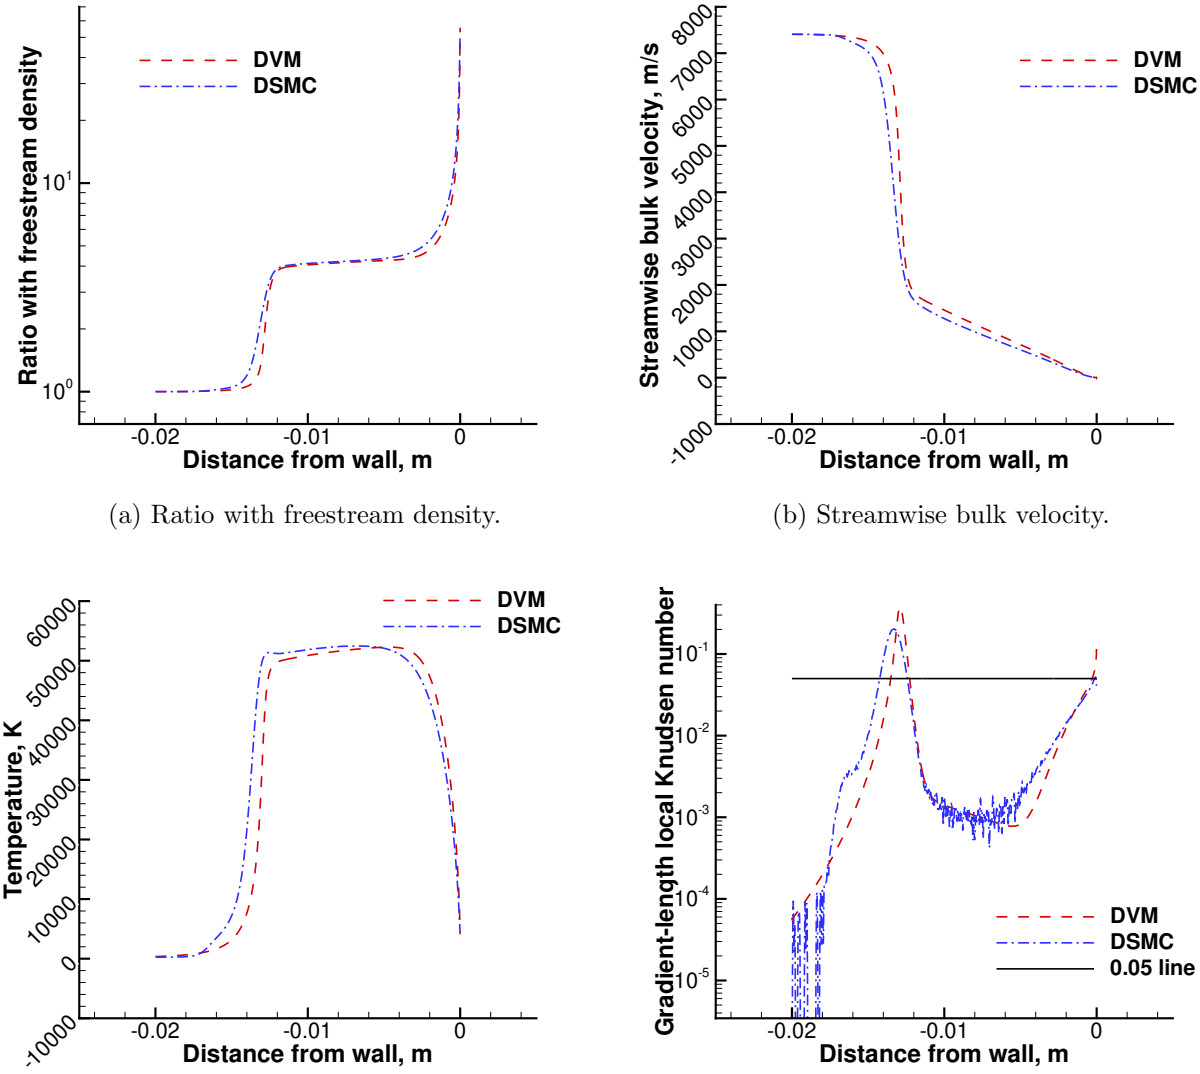

(a) Ratio with freestream density.

(b) Streamwise bulk velocity.

(c) Temperature.

(d) Gradient-length local Knudsen number.

Figure 4.6: DVM and DSMC macroscopic property profiles for the Mach 25 flow at 51 km [106].

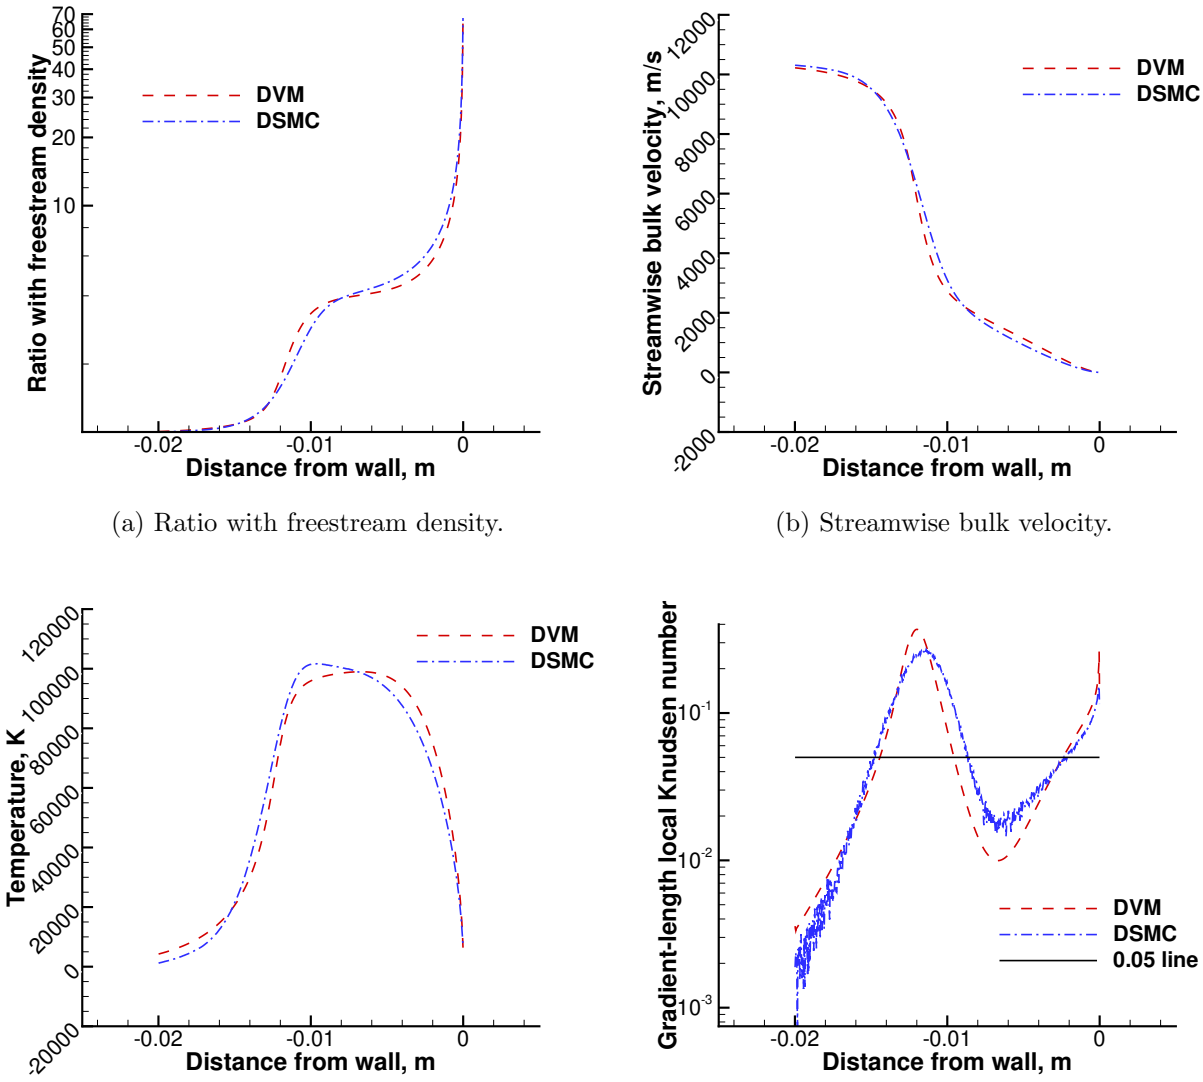

(a) Ratio with freestream density.

(b) Streamwise bulk velocity.

(c) Temperature.

(d) Gradient-length local Knudsen number.

Figure 4.7: DVM and DSMC macroscopic property profiles for the Mach 35 flow at 60 km [106].

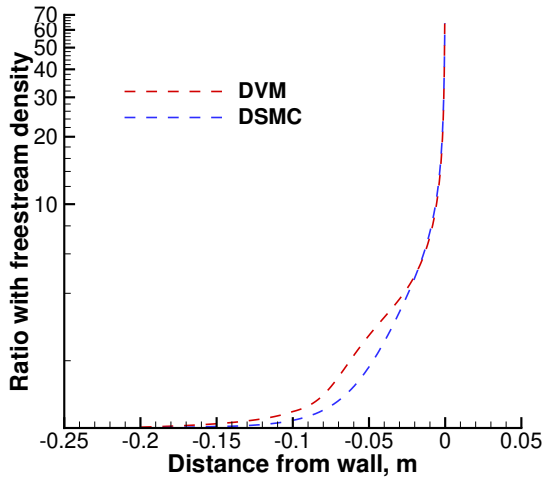

(a) Ratio with freestream density.

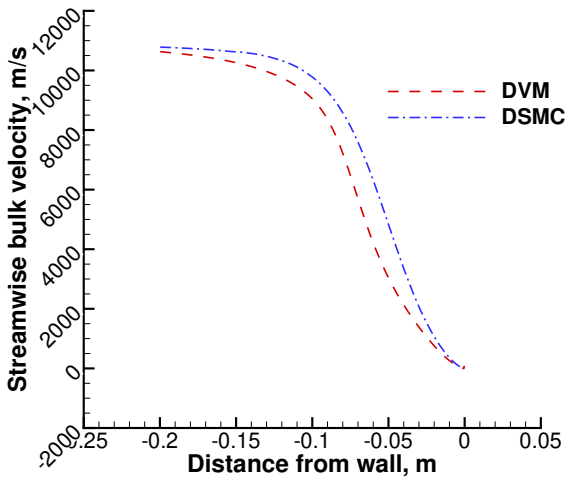

(b) Streamwise bulk velocity.

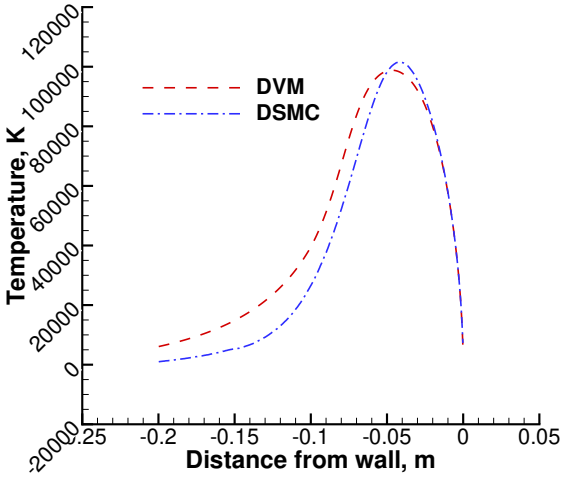

(c) Temperature.

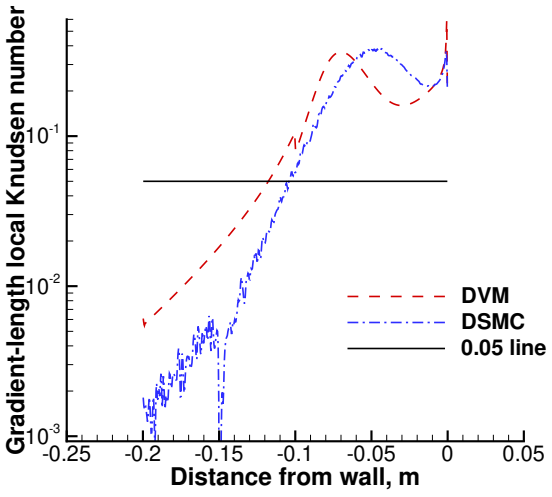

(d) Gradient-length local Knudsen number.

Figure 4.8: DVM and DSMC macroscopic property profiles for the Mach 39 flow at 81 km [106].

Figure 4.9 provides an example of this effect in the density profiles of DVM and DSMC for the Mach 39 flow at 81 km. While the overall shock shape is maintained, the smoothness of both profiles is clearly disrupted at $x = x_r = -0.0425$ m compared to Figure 4.8a when $x_r = -0.1$ m. Upstream of the shock, flow properties are more uniform and each cell in the removal region will have a smaller fraction of removed particles due to the lower translational energy. Thus, if x_r is placed upstream, the discontinuity is less likely to appear because removal is performed over a larger region and the flowfield density distribution is less disturbed. Trivially, x_r should not be placed infinitely far upstream: Figure 4.5 verifies that the average energy of removed particles decreases as x_r moves further upstream. This decrease must at least be bounded, since cold freestream particles with low values of $v_y^2 + v_z^2$ are less likely to be chosen for removal. x_r should also not be placed too far upstream for computational efficiency: a larger removal region results in more particles being evaluated but ultimately rejected in DSMC, and a larger removal region in DVM requires Equation 4.14 to be manually subtracted from more VDF cells.

Conversely, if x_r is located within the shock standoff distance from the wall, the density flowfield is suddenly disturbed. Depending on how flowfield properties are outputted, the perturbation may not smooth out in time before quantities are saved. The presence of discontinuities when x_r is placed within the shock was not reported in previous work using Bird's original DSMC stagnation streamline method. Explanations relating to the stochastic nature of DSMC (e.g., time-averaging obscuring the discontinuity; a smaller particle weight enabling more uniform selection of particles), do not hold up under scrutiny since the discontinuity phenomenon also occurs in deterministic DVM. The most likely explanation relates to whether the particle removal procedure is performed at each time step. In the literature, it is unclear whether the particle removal operation should be executed at every time step or only at time steps where the density ratio at the specified point x_s exceeds the value specified [1]. If a discontinuity results after particle removal, but the simulation is allowed to advance several time steps without removal due to the stablized density profile, the discontinuity may smooth out during that period.

There are advantages and disadvantages to particle removal at each time step or only when

the density ratio at x_s exceeds the specified value. Removing particles at each time step circumvents the calculation of density at x_s each time step. The flowfield is also guaranteed to reach steady state eventually, since mass is conserved every time step (neglecting electron mass). At the same time, performing the removal operation every time step is also computationally expensive, particularly for parallelized implementations (see Appendix B). Additionally, if particle removal begins while the chemical composition of the flow is still in a transitory state (e.g., particle removal begins before the flow has the opportunity to ionize), an incorrect flowfield may result. However, by allowing the simulation to advance several steps in between removal operations, the shock may oscillate slightly around x_s. This effect is minimized with increasing configuration space resolution around x_s.

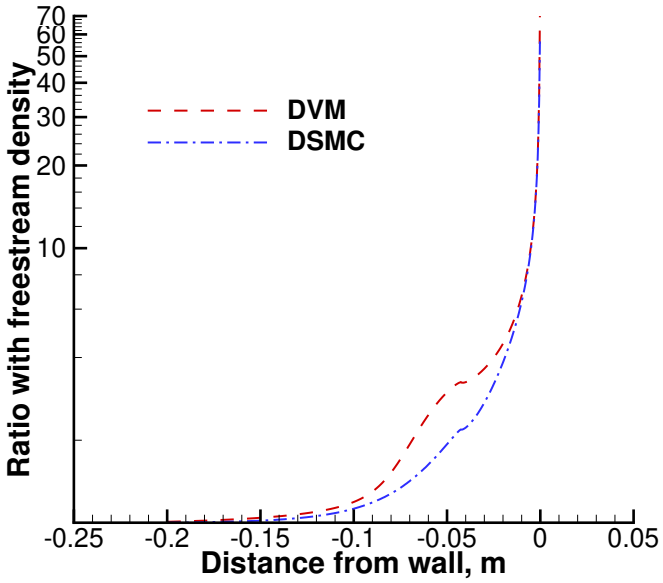

Figure 4.9: Discontinuities in density profiles when removal region begins inside the shock layer for DVM and DSMC for the Mach 39 flow at 81 km [106].

For either DSMC or DVM, concrete guidelines for choosing x_r are desireable. In the best case scenario, experimental data or results from multidimensional simulations are available. Then, users may iteratively adjust x_r and x_s to maximize agreement with the flowfield of interest. In absence of such data, x_r should always be the minimum distance from the wall that produces

Table 4.3: Locations of velocity distribution function slices, in meters, with respect to the wall.

Altitude	Solver	Upstream ($\rho/\rho_\infty = 1.25$)	Shock ($\rho/\rho_\infty = 2$)	Standoff ($\rho/\rho_\infty = 4$)
51 km	DVM	1.32×10^{-2}	1.28×10^{-2}	1.11×10^{-2}
	DSMC	1.32×10^{-2}	1.12×10^{-2}	7.62×10^{-3}
60 km	DVM	1.31×10^{-2}	1.17×10^{-2}	7.03×10^{-3}
	DSMC	1.32×10^{-2}	1.12×10^{-2}	7.62×10^{-3}
81 km	DVM	9.06×10^{-2}	6.24×10^{-2}	2.72×10^{-2}
	DSMC	7.56×10^{-2}	4.74×10^{-2}	2.43×10^{-2}

smooth macroscopic property profiles during steady state. For this specific study, smooth profiles are achieved when x_r is located upstream of the shock standoff distance and removal is performed at every time step. The DSMC model is generally more sensitive to placement of x_r, where smooth profiles are achieved when x_r is placed at an upstream location where the temperature gradient is close to zero. Occasionally, this requires the removal region to encompass the entire domain such as in the 51 km flow case. The DVM model is generally less restrictive with respect to the choice of x_r. Smooth profiles are observed when x_r is located at an upstream location where either the density or the temperature gradient just begins to increase. For the configuration space grids of this work, selection of x_r is found to be insensitive to changes below $\Delta x = 0.005$ m in either code. As previously discussed, x_r should not be placed too far upstream to ensure optimal conservation of energy and reduce the computational expense of the removal scheme.

4.3.2 Analysis of the Velocity Distribution Function Along a Stagnation Streamline

A recent paper on Bhatnagar–Gross–Krook methods [105] emphasized the importance of analyzing VDFs in addition to flowfield macroscopic properties when evaluating kinetic simulation options for study of non-equilibrium flows. Additionally, since the DVM stagnation streamline scheme entails direct VDF manipulation via Equation 4.14, ensuring the DSMC and DVM VDFs agree is necessary for verification of the DVM stagnation streamline scheme.

The structure of a VDF in the vicinity of a normal shock wave is well known and has been verified numerically and validated experimentally many times in the literature [123, 124, 125, 126,

127]. Upstream and downstream of the shock, the VDFs are Maxwellian. Within the shock, the streamwise, or parallel, component of the VDF can be decomposed as a superposition of two Maxwellian distributions: a fast, low temperature component centered at the pre-shock velocity and a slower, high temperature component centered at the post-shock velocity. This is known as the bimodal Maxwellian structure, and it occurs because in order for the VDF to transition from the high velocity, cold upstream state to the low velocity, hot downstream state, the particles must experience a sufficient number of collisions while traveling across the shock. For example, at the very center of the shock, some particles will have collided many times and follow the slow VDF distribution, while some particles may not have collided at all and will follow the original fast VDF distribution. The bimodal Maxwellian structure also exists in the perpendicular VDFs, where the VDFs remain symmetric about zero regardless of their location along the streamline if the flow exclusively propagates in parallel.

A non-equilibrium stagnation streamline differs from the normal shock setup in several respects. Firstly, the flowfield velocity and temperature do not remain constant after the shock. For a no-slip wall boundary condition, the flowfield velocity will decrease towards zero as it approaches the wall. For an isothermal wall boundary condition, the temperature profile will approach the wall temperature. Thus, the slow, high temperature component of the VDF will decrease in width and shift towards zero as the flow approaches the wall. Second, under non-continuum flow conditions, it is possible for many particles to pass through the shock without colliding at all. Evidence of the fast Maxwellian VDF component may exist as far downstream as the wall itself. Finally, if the stagnation streamline is observed along a line of sight to a multidimensional vehicle, the VDFs may develop other features (e.g., the development of a bulk velocity in the $\hat{y}$ direction). However, these considerations are irrelevant for a 1D flowfield.

Three points along the stagnation streamline are selected for comparison: just upstream of the shock, at the center of the shock, and at the shock standoff distance. Since the DSMC and DVM density profiles are slightly offset from one another in all flow cases, it is appropriate to assess VDFs from cells with equivalent densities rather than identical configuration space locations.

Table 4.3 provides the density ratios and corresponding configuration space locations for the VDF slices analyzed from each flow case and solver. When plotted, each VDF is normalized by its cell number density. Only the streamwise VDFs are discussed, as the perpendicular VDF information only supports the streamwise conclusions without revealing additional insights.

Figures 4.10–4.12 present profiles of the streamwise VDFs from DVM and DSMC. Consistent with the normal shock case, the VDFs upstream of and within the shock display the bimodal Maxwellian structure. Not shown, in the highly non-continuum 81 km flow case, the pre-shock component is present throughout the shock layer up until the wall. Throughout all three flow cases, a third distinct VDF component also emerges centered around a large negative velocity. This feature is most readily observed in the DVM upstream VDFs, where a Maxwellian distribution at -2510 m/s at 51 km, -3600 m/s at 60 km, and -4460 m/s at 81 km are observed. The third Maxwellian component is also visible in the VDFs within the shock at 60 km and 81 km. The prominence and shape of the third component do not significantly change with increasing configuration space resolution for either DSMC or DVM or with increasing velocity space resolution for DVM.

Labeled as the "trimodal Maxwellian", for a 1D computational domain, this VDF structure was first observed in Reference [128] and is also observed in the results of Reference [105]. Both studies used DSMC to simulate a normal shock with a moving specularly reflecting wall to model a piston. In Reference [128], the origin of the trimodal Maxwellian strucutre was attributed to particles significantly decelerated by collisions at the shock front. However, if collisions were the cause, similar features would appear in conventional normal shock flowfields. Additionally, collisions drive a VDF towards a Maxwellian centered around the bulk flow velocity, which should not be negative for a shock. The center of the third Maxwellian component in References [128] and [105] and in this study shifts further negative as one moves upstream. The height of the third component also decreases in magnitude, indicating fewer particles follow that distribution. This provides evidence of three main trajectory options for a particle traveling through a non-equilibrium hypersonic flowfield. The first is a particle that enters the shock layer and collides many times, following a slow, high temperature distribution. The second is a particle that enters the shock layer but does not

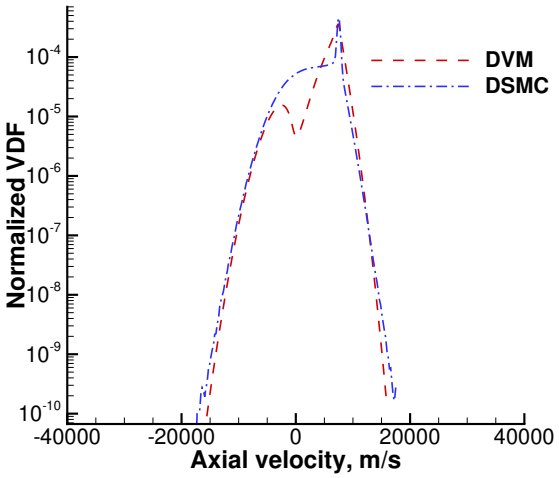

(a) Upstream of the shock.

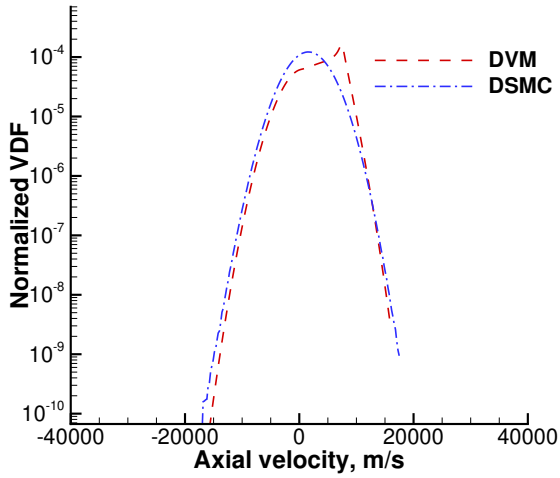

(b) Within the shock.

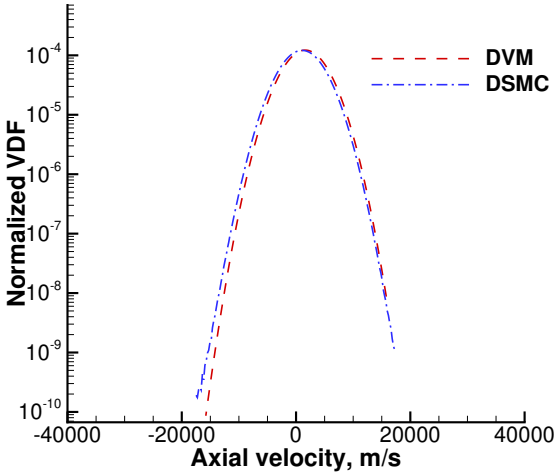

(c) At the shock standoff distance.

Figure 4.10: DVM and DSMC streamwise VDF slices for the Mach 25 flow at 51 km [106].

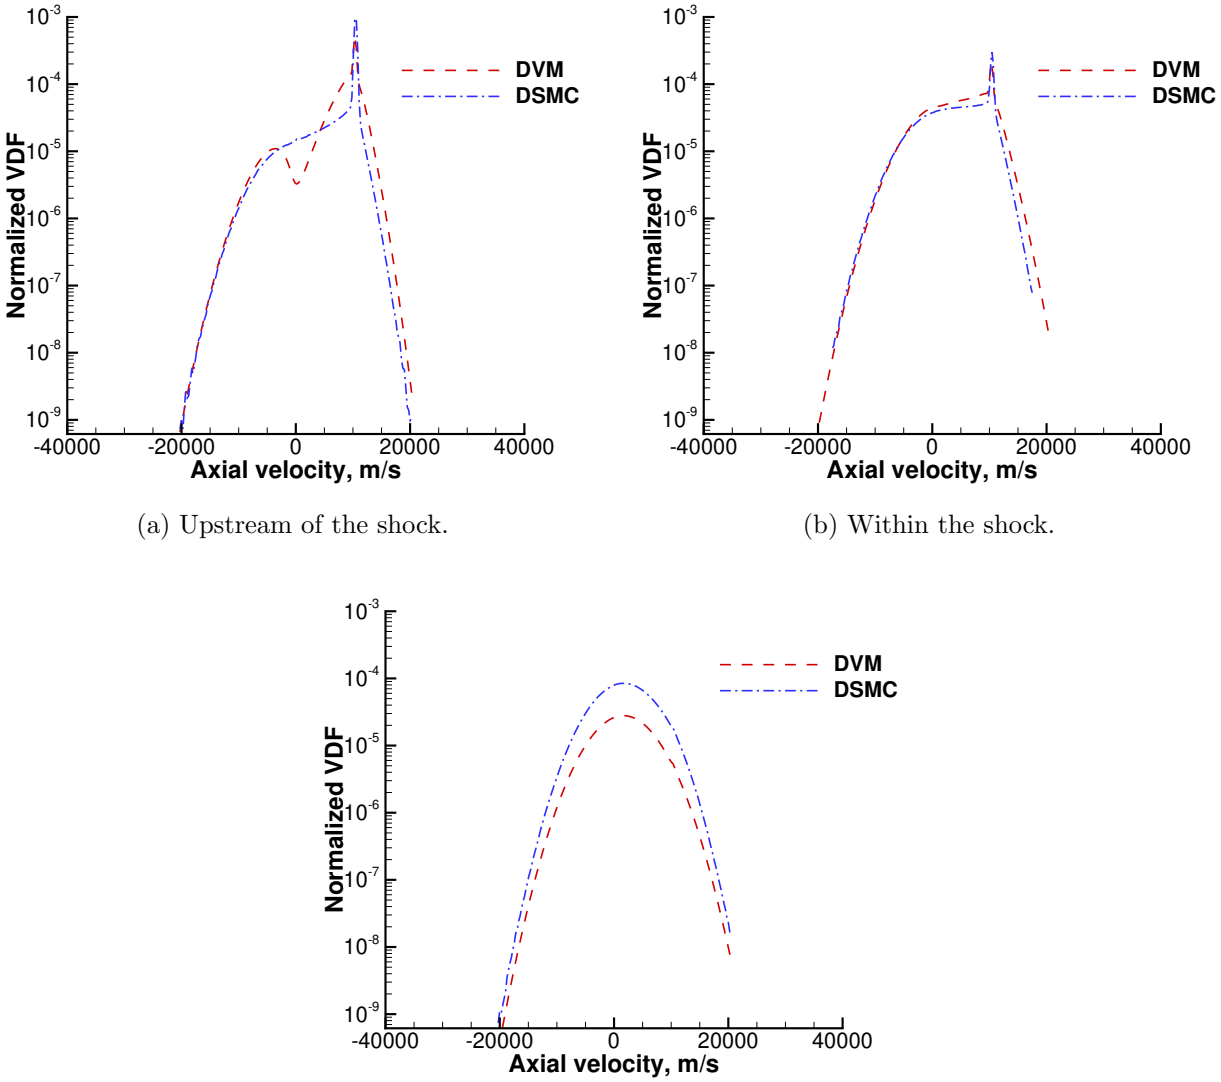

(a) Upstream of the shock.

(b) Within the shock.

(c) At the shock standoff distance.

Figure 4.11: DVM and DSMC streamwise VDF slices for the Mach 35 flow at 60 km [106].

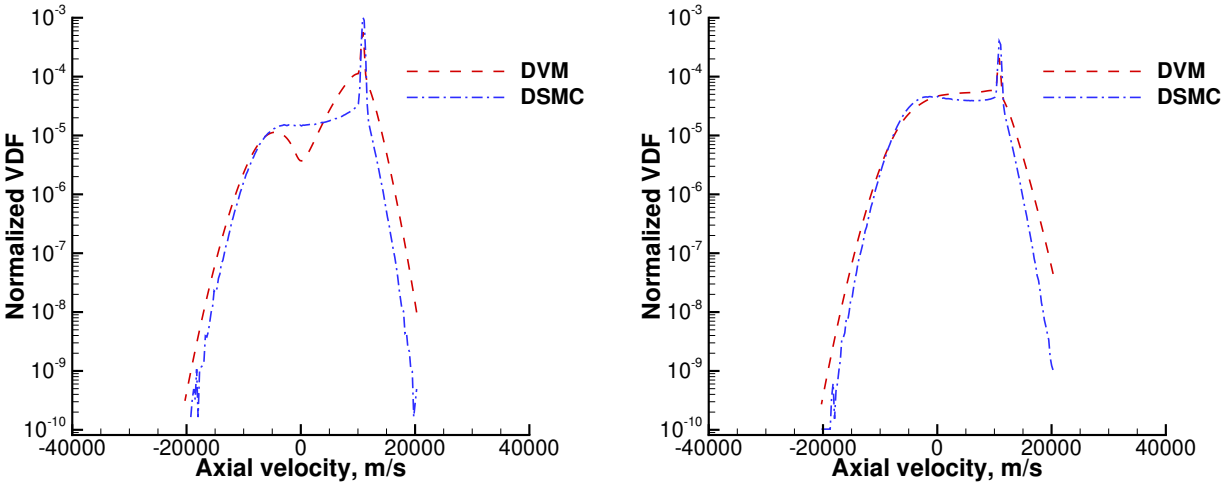

(a) Upstream of the shock.

(b) Within the shock.

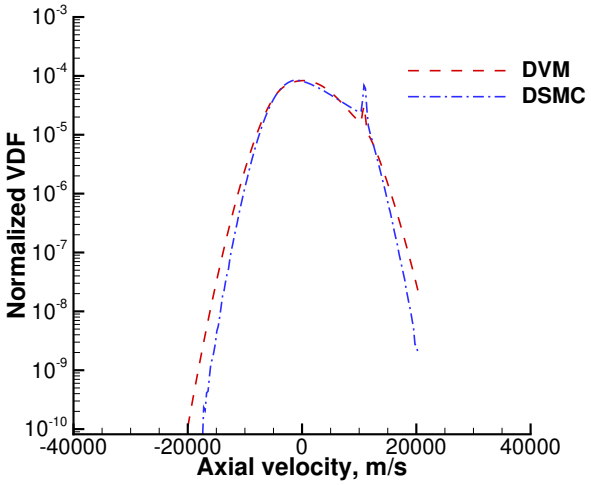

(c) At the shock standoff distance.

Figure 4.12: DVM and DSMC streamwise VDF slices for the Mach 39 flow at 81 km [106].

collide a sufficient number of times to thermalize to the bulk flowfield properties, thus retaining its fast, low temperature characteristics from the freestream. Since the trimodal Maxwellian structure is unique to VDFs in the presence of a reflecting boundary, the last trajectory option must involve the wall itself, and consists of a particle which reflects off the wall and propagates all the way upstream. Since the third Maxwellian is centered at a large negative velocity, it follows that these particles must have a high amount of translational energy even after reflecting from the isothermal wall. Furthermore, these high energy backstreaming particles naturally decline in population as one moves upstream, since the further a particle travels, the more likely it has experienced at least one thermalizing collision.

Backstreaming particle behavior is intuitively expected in free molecular flow, such as along the stagnation streamline of a low Earth orbit satellite. However, this study highlights the existence of the trimodal VDF structure at altitudes as low as 51 km. Possible secondary-effects from these high energy backstreaming particles warrants further consideration, particularly for chemically reacting or multidimensional flows. For example, References [129] and [130] simulated chemically reacting air around the Stardust spacecraft in the rarefied regime with DSMC. Highly non-Maxwellian VDFs were observed in the vicinity of the shock with multiple peaks at negative velocities. The additional complexity in the VDF structure led to differences in electron-impact excitation rates and atomic nitrogen excited state number densities between the quasi-steady-state assumption (QSS) and DSMC-derived calculations. Subsequently, there were differences in computed radiative heat flux between QSS and DSMC. Thus, the trimodal Maxwellian VDF structure is not limited to 1D flows, and may have subtle but non-negligible impact on higher order flowfield quantities relevant to vehicle design.

In terms of benchmarking VDFs for the DVM stagnation streamline model, the DVM-BGK and DSMC simulations achieve levels of agreement comparable to previous work [105, 131] given the use of the BGK operator. In particular, the DVM simulations generally observe a more defined third Maxwellian component than DSMC, which may have an impact on temperature calculations. For example, in the 81 km flow case, the third Maxwellian component is present in the VDFs

as far upstream as the first configuration space cell. Although the magnitude of the peak is small, indicating few particles follow this distribution relative to the freestream distribution, it still factors into the overall width of the VDF, increasing the calculated temperature upstream. This is consistent with the work of Reference [105], which compared VDFs for non-equilibrium flows between DSMC simulations with standard stochastic collision routines, DSMC simulations with the Shakov (S-BGK) [67] collision operator, and DSMC simulations with the ellipsoidal–statistical (ES-BGK) [68] collision operator. For the piston-induced normal shock wave flow case, the ES-BGK VDFs displayed a more defined third Maxwellian component and a larger streamwise temperature upstream of the shock compared to standard DSMC and DSMC with the S-BGK model. This suggests that the discrepancies in VDF profiles between DSMC and DVM in this work can be attributed to the choice of collision operator rather than the DVM stagnation streamline removal method itself. However, it also highlights the inaccuracies of the BGK model in simulating non-equilibrium hypersonic flows.

4.4 Evaluation of the Discrete-Velocity Method for Electrostatic Hypersonic Flowfield Simulation

Development of the 1D stagnation streamline model for DVM enables computational cost assessment of DK appropriate for the partially ionized hypersonic flowfield problem. In this study, the DK simulations are run on anywhere between 40 to 240 processors with total wall times on the order of hours or days. The largest DK simulation, the 51 km flow case at Mach 25, requires 672 total core-hours for a computational grid of 2,800 points in configuration space. Comparatively, the ISTAG DSMC simulations are run on a local computer without parallelization and total wall times on the order of minutes or hours. The largest ISTAG simulation in this work (51 km at Mach 25) requires 10 total core-hours. Thus, for this specific study, the DSMC simulations are always more computationally efficient, even for resolution of higher order quantities such as the particle VDF.

One consideration unique to DVM simulations of partially ionized hypersonic flows is velocity

space resolution for the charged species. In DK, each species is resolved with its own VDF. For the neutral-only flows of this work, a modest velocity space grid is used with 160 grid points per direction, corresponding to a bin size of 255 m/s. According to Reference [76], if applied to a plasma simulation, this degree of resolution would be sufficient to resolve the number densities, bulk velocities, electric potential, and electric field with normalized RMSE on the order of 0.100%. However, to observe the same RMSE in profiles of local charge separation, critical to analyzing charged species diffusion as discussed in Chapter 3, at least 640 grid points are recommended. 640 cells for each of the three dimensions of velocity space brings the VDF resolution to 262,144,000 velocity space grid cells *per configuration space grid cell* and *per charged species*. Since advection and the electric field only exist in the axial direction for a 1D problem, an argument can be made for reducing the computational grid resolution in the perpendicular directions. This brings the number of VDF grid points to $640 \times 8 \times 8 = 40{,}960$ points per configuration space grid cell per charged species. Unfortunately, this number is still intractable for the expected number of configuration space grid cells required to resolve the Debye length both in terms of simulation wall time and required computer memory. Various methods of further reducing computational cost for DVM without compromising flowfield accuracy were explored throughout the development of the stagnation streamline model, the most notable of which is the Parallel-Kinetic-Perpendicular-Moment (PKPM) model [63, 132, 133, 134]. Details of this approach, why it was ultimately not chosen for the final study of this dissertation, and avenues for enhancement of the PKPM model for study of collisional flows are included in Chapter 6.3.4. Ultimately, if the DVM method is used, a compromise would inevitably have to be made in resolution of local charge separation, detracting from any claims made about the nature of charged species diffusion in hypersonic flows.

There are also physics-motivated reasons for why DVM is ultimately not chosen for the final study of this dissertation work. The first is that the results of Section 4.3 highlight the disadvantages of the BGK collision operator in simulation of non-equilibrium flows. However, upgrading DK's collision operator to either S-BGK or ES-BGK would increase the computational cost of the simulation as both require calculation of higher order terms such as heat flux for S-BGK

or stress tensor for ES-BGK. Since resolution of charged species upstream of a stagnation streamline shock has not been studied in detail in the literature, accurate simulation of collisions in the vicinity of the shock (especially for backstreaming neutrals) is necessary. Finally, referring to Section 2.4.5 and Equations 2.50–2.52, a bulk ionization rate k_{EII} is applied to the entire VDF. Therefore, the VDF bins at $v_x = v_{max}$ and $v_x = 0$ m/s see an equal amount of ionization. This is acceptable for the fusion-based applications for which the ionization model was derived (where the plasma temperatures are on the order of 10^7 K and above); however, non-continuum hypersonic flowfield modeling requires a preference for ionization in the tails of the VDF, such that high energy particles are more likely to ionize than low energy particles at the center of the VDF. The TCE model of DSMC is ultimately more accurate in simulating ionization processes relevant to hypersonic flows.

DSMC-PIC methods still have many computational cost limitations, even with enhancements such as variable particle weighting schemes, and physical accuracy limitations, such as numerical thermalization (Section 2.6). The primary concern of coupling a DSMC solver to a PIC solver is the number of macroparticles required to simulate all flowfield effects. As the plasma density increases, the number of configuration space grid points increases in order to resolve the Debye length. For a non-collisional PIC simulation, statistical significance in results depends on the total number of macroparticles per Debye length, not per cell. However, for collisional DSMC, at least 20-30 macroparticles *per cell* are required to accurately model collisions so that the pair selection procedure can sample from a sufficient number of particles [55]. If trace species are present, assuming identical weights, many more macroparticles per cell are required to resolve the trace species collision rates. Using the semi-variables SWS does not alleviate this cost entirely. For example, if a flowfield ionization fraction of 10^{-6} is desired, the minimum possible plasma weight modifier would be 10^{-6}. Thus, for one AAI reaction, 10^6 macroparticles would be generated in one cell with the ambipolar diffusion approximation, and 2×10^6 macroparticles would be generated if ions and electrons are modeled as independent species. A fully dynamic variable particle weighting scheme such as SWPM [101] would reduce this cost by setting $w_p = w_n$ at the time of AAI, then reduce w_p as necessary over the course of the simulation. However, fully dynamic weighting schemes

require consideration of their own limitations and complexities, such as error introduced via particle merging and resolution of low probability collision events [135, 136].

4.5 Summary

Computational cost considerations when using either DVM or DSMC to simulate partially ionized, transitional regime hypersonic flows with electrostatic modeling were evaluated throughout this chapter in detail. A 1D stagnation streamline model for DVM was developed and benchmarked with its DSMC equivalent. The model was used to study VDFs in the vicinity of the shock, revealing features potentially relevant to hypersonic vehicle design and study of charged species diffusion upstream of a shock layer. However, even with the reduced dimensionality enabled by the novel model, DVM proved too computationally costly if proper VDF resolution of higher order quantities is desired. Additionally, certain aspects of the DK code specifically are ill-suited for studying non-equilibrium hypersonic flows. Although DSMC-PIC solvers have their own set of computational cost and physical accuracy considerations, these are ultimately deemed more manageable than those of DVM for the remaining objectives of this dissertation.

Chapter 5

Study of Plasma Diffusion in Transitional Hypersonic Flows

In this chapter, the physics of charged species diffusion, methods of predicting charged species diffusion regime, and the efficacy of the ambipolar diffusion approximation in transitional hypersonic flows are investigated. Section 5.1 reviews the computational setup and initialization conditions of the hypersonic flowfields, analyzes flowfields with the ambipolar diffusion approximation, and determines an appropriate length scale for calculation of the diffusion length. Section 5.2 investigates the dependence of hypersonic plasma properties on domain length and choice of Poisson equation boundary condition at the inflow boundary. Section 5.3 investigates the free and transitional plasma diffusion regimes in hypersonic flows with the ambipolar diffusion approximation and with full electrostatic modeling. Section 5.4 discusses the general characteristics and first-order effects of a hypersonic plasma identified across all flow cases. Section 5.5 verifies that the criteria for identification of charged species diffusion regimes proposed in Section 3.5 can be used to classify hypersonic plasmas and discusses other useful quantities for plasma characterization. Finally, Section 5.6 evaluates the efficacy of the ambipolar diffusion approximation in simulation of hypersonic flows.

5.1 Simulation of Partially Ionized Hypersonic Flow

The 1D DSMC stagnation streamline model described in Section 4.2.1 is used to simulate a partially ionized hypersonic flowfield. Details of its implementation and benchmarking in MPIC (including how to parallelize the scheme) are included in Appendix B.

For the reasons outlined in Section 1.4, argon is used to model the gases of this section. Momentum exchange collisions between neutrals and neutrals, neutrals and ions, and neutrals and electrons are modeled using the VHS cross section, the Sakabe and Izawa CEX cross section [90], and the fitted IST-Lisbon cross section data [98, 99], respectively. The VHS model is used to model translational energy exchange. Electron-impact ionization and atom-atom ionization are modeled, but the CEX reaction is not. This is because CEX is a second-order plasma effect: ions may gain energy through acceleration via electric fields, thus increasing the CEX rate. Baseline parameters for all interactions are found in Appendix A. Excited states are not explicitly modeled, however, the chosen ionization rates account for the 'ladder-climbing' of excited states an atom will experience before ionizing– characteristic of super-orbital hypersonic flows in air (see Section 2.5.3). The electron mass is artificially increased to $m_e = 10^{-2} \cdot m_i$ to reduce the computational cost of resolving fast electron motion while capturing a sufficiently large difference in mass between ions and electrons. Use of the true electron mass will lead to differences in ionization rates and density distributions throughout the flowfield. Since the focus of this work is on general flow behavior and scaling arguments rather than specific flow conditions or spacecraft, modification of the electron mass does not detract from the conclusions made in this study. This is further justified after discussing all results in Section 5.5.

Whether to include Coulomb collisions in kinetic simulation of partially ionized rarefied flows is subject to debate. The characteristic cross section of a Coulomb collision is much larger than a typical VHS cross section; thus, implementing true Coulomb collisions can be computationally expensive for particle-based methods as many particle pairs must be made across configuration space cells. Furthermore, Coulomb collisions drive the particle VDFs towards Maxwellians, but their impact may be negligible compared to collisions with neutrals or numerical thermalization [107]. Finally, the derivation of ambipolar diffusion assumes charged particles primarily collide with neutrals; so, omission of Coulomb collisions allows for observation of 'ideal' ambipolar diffusion. Calculation of the Coulomb collision rates for several cases is included in Section 5.4 to demonstrate that their impact is negligible compared to MEX collisions with neutrals, and thus Coulomb

collisions are not modeled in this study.

The configuration space grids are uniformly generated such that the local mean free path is always resolved with at least two points and that the *minimum* Debye length of the entire flowfield is resolved with at least 2 points. Although uniform grids are inefficient for resolving the electron Debye length, they ensure that the plasma is always well-resolved upstream of the bulk plasma region regardless of evolving properties. Specific resolution for each flow case is included in tables for each section. The global particle weight is chosen to yield at least 20 neutral particles per cell in the freestream. The species weight modifier for Ar^+ and e^- is selected to yield at least 100 particles per cell in the bulk plasma region. Each simulation is run to steady state, and the final macroscopic property profiles are time averaged over 10^6 time steps during the steady state period.

When electrostatic modeling is used, the simulation starts from an ambipolar diffusion approximation solution already at steady state, runs until the new flowfield with electrostatic modeling achieves steady state, and then is time averaged over 10^6 time steps. Physically, this is sufficiently representative of a hypersonic flowfield. A plasma generated through re-entry flow conditions begins to form at altitudes much higher than the freestream conditions simulated here. Technically speaking, a hypersonic re-entry flowfield at 60 km or 81 km should already be ionized at the start of the simulation. However, the ionization fractions of the flowfields are not known *a priori*, and the freestream gas should remain fully neutral. Thus, the closest approximation to reality is to start from an ambipolar diffusion approximation solution. For MPIC's implementation of the ambipolar diffusion approximation, it is straightforward to split each ambipolar particle and allow the ion and electron to move independently for the rest of the simulation.

The following plasma properties are discussed throughout this section. The first is the maximum distance from the wall for which statistically significant collisional plasma properties are obtained. As discussed in Section 4.4, to properly resolve plasma collision rates there must be a sufficient number of plasma particles per cell available for pair selection and per Debye length for accurate results. A smaller plasma weight will increase this maximum distance; thus, this quantity is reflective of modeling resolution and domain length rather than flowfield characteristics. Beyond

the 'maximum resolved distance', time averaging over the steady state period allows for some interpretation of plasma properties. However overall, there is less confidence in the plasma diffusion behavior than in the well-resolved region. For example, high velocity ambipolar macroparticles may travel further upstream if they are never selected for collision and consequently plasma density may be overpredicted. At the same time, the plasma density distributions remain constant over the steady state period. This indicates that even if the collision rates or charged particle interactions are underresolved far upstream of the shock layer, the plasma will not diffuse to the edge of the domain by default if the simulation is allowed to continue infinitely. This work defines the 'maximum resolved distance' as the point where there is at least one charged particle per cell and at least 10 charged particles per Debye length. Although the latter metric is not required for ambipolar diffusion approximation solutions (since the Poisson equation is never solved), all simulations for a flow case are always simulated with the same configuration space grids.

The second quantity is the shortest electron Debye length present in the flow. When the ambipolar diffusion approximation is invoked, the Debye length is computed from the ambipolar plasma density and the ambipolar electron temperature. The range of plasma densities and ionization fractions (with respect to the ions) within maximum resolved distance are also listed to further characterize the plasma. Finally, the convective and chemical stagnation point heat fluxes [137] by species are reported for each flow case. In these simulations, only Ar^+ contributes to chemical heat flux as it neutralizes at the wall. Of note, during an ion-wall neutralization reaction, MPIC adds the energy of the ion before the collision to the total ion convective heat flux, but subtracts the energy of the reflected neutral from the total neutral convective heat flux. Thus, the total convective heat flux from all species is calculated correctly, but the ion convective heat flux is inflated if interpreted incorrectly. The ambipolar electrons' contribution to heat flux is also reported by using the ambipolar particles' electron velocity components to calculate kinetic energy.

Table 5.1: Simulation parameters for each case in survey of flow conditions with the ambipolar diffusion approximation.

	60 km		81 km	
Parameter	Baseline	Reduced	Baseline	Reduced
Domain length, m	0.03	0.03	0.30	0.30
Points in configuration space	1,350	1,350	750	750
Reduced density factor, F_r	1.0	10^8	1.0	10^8
Global weight	5.41×10^{15}	5.41×10^7	5.28×10^{15}	5.28×10^7
Plasma weight modifier	0.05	0.05	0.05	0.05
Time step, s	10^{-10}	10^{-10}	10^{-9}	10^{-9}
Shock standoff distance from wall, m	0.003	0.003	0.0243	0.0243
Start of removal region from wall, m	0.0175	0.0175	0.150	0.150

5.1.1 Survey of Flow Conditions with the Ambipolar Diffusion Approximation

The 60 km and 81 km flow cases from Table 4.1 are run again with the ionization reactions and collisions described above enabled and the ambipolar diffusion approximation on. Simulation initialization parameters are included in Table 5.1. The domain lengths are intentionally much larger than the shock layer in order to observe how far upstream the plasma macroparticles will propagate. Two simulation setups are considered. The first uses the default parameters for collision cross sections found in Appendix A; these are labeled the 'baseline' simulations. The second artificially reduces the freestream number density while scaling the collision cross sections to reproduce the same mean free paths and ionization fractions α_p as the baseline simulations; labeled the 'reduced' simulations. Details of the approach are included in Section 5.1.3. Various length scales relevant to hypersonic flows are also reviewed to determine an appropriate diffusion length Λ for applying the charged species diffusion regime criteria presented in Section 3.4.

Figure 5.1 plots the species number densities, ionization fraction, temperatures, electron Debye length, and ratio of hypersonic flowfield diffusion length $\Lambda_h/\lambda_{D,e}$ (defined and elaborated on for each figure in Section 5.1.2) for the Mach 35 flow at 60 km. The post-shock region is strongly ionized, with α_p ranging from 2.47×10^{-3} to 3.79×10^{-1} within the region of statistically significant plasma. The impact of enabling ionization on the neutral flowfield properties is most evident in the temperature profile which appears far rounder than the neutral-only profile in Figure 4.7c.

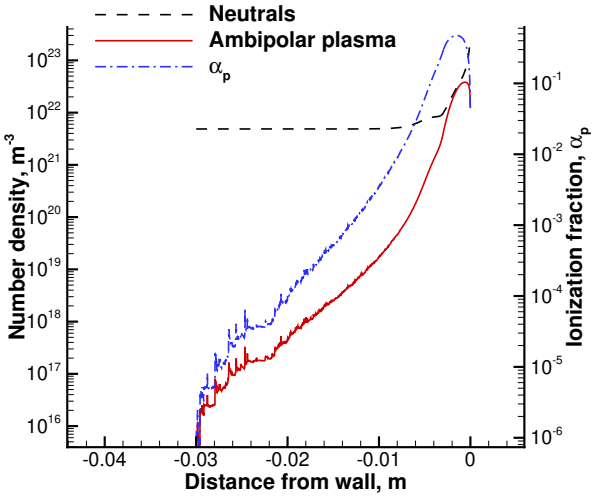

(a) Number densities for neutrals and ambipolar plasma (left axis) and ionization fraction (right axis).

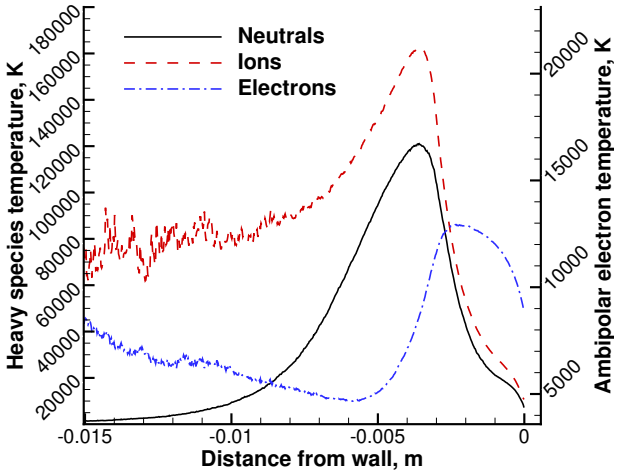

(b) Temperature for neutrals and ambipolar ions (left axis) and ambipolar electrons (right axis).

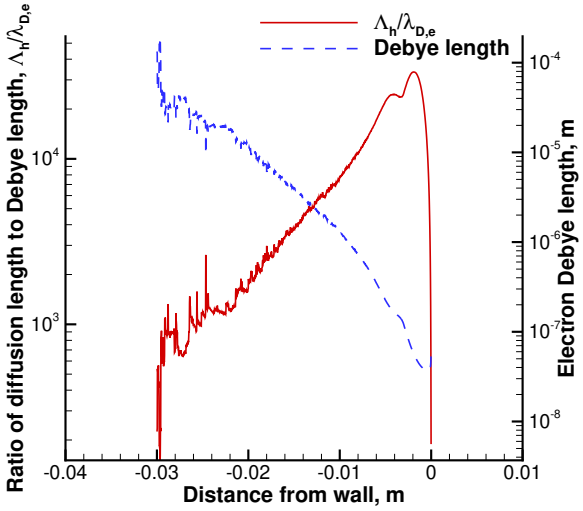

(c) Electron Debye length (left axis), computed from the ambipolar plasma density and the ambipolar electron temperature, and the ratio of the hypersonic diffusion length to the Debye length (right axis).

Figure 5.1: Baseline flowfield results for the Mach 35 flow at 60 km with the ambipolar diffusion approximation on.

Table 5.2: Quantities of interest from the baseline hypersonic flow survey (ambipolar diffusion approximation on).

Quantity	60 km	81 km
Center of the shock from wall, m	-0.00386	-0.0302
Plasma maximum resolved distance, m	-0.0109	-0.178
Range of ionization fraction	[2.47 $\times 10^{-3}$, 3.79 $\times 10^{-1}$]	[2.46 $\times 10^{-3}$, 4.36 $\times 10^{-1}$]
Range of plasma density, m^{-3}	[1.22 $\times 10^{19}$, 3.82 $\times 10^{22}$]	[6.60$\times 10^{17}$, 5.56 $\times 10^{21}$]
Minimum electron Debye length, m	3.73$\times 10^{-8}$	1.72 $\times 10^{-7}$

The plasma reaches its peak ionization fraction downstream of the shock at $x = -0.00153$ m, but achieves its peak ion density at $x = -0.000564$ m. The ion and neutral temperatures follow similar trends to each other within the shock layer, with the ions being roughly 10,000 K to 20,000 K hotter in the post-shock region (Figure 5.1b), but deviate upstream. This is consistent with expectations, as only high velocity ions produced inside the shock layer have enough kinetic energy to diffuse upstream. As discussed in Section 4.3.1, temperature as a metric is only strictly defined for a flow at equilibrium; thus, temperature profiles within the shock itself describe the general width of the VDF rather than the actual average energy of the particles. The electron temperature is considerably lower than the other two species, with its temperature ranging from 9,090 K to 12,800 K throughout the shock layer, and gradually increasing upstream.

The electron Debye length decreases throughout the flowfield as one moves towards the wall, except within the plasma sheath, where the ambipolar plasma density drops and thus the electron Debye length increases rapidly. Of particular interest is the plasma within and upstream of the shock. The center of the shock (i.e., $\rho/\rho_\infty = 2.0$ for a monatomic gas) is located at $x = -0.00386$ m. The plasma is well resolved up to $x = -0.0109$ m, and remains strongly ionized both upstream and downstream of the shock with $\alpha_p > 10^{-3}$. This somewhat contradicts the current understanding of the structure of a partially ionized hypersonic shock layer (Figure 1.5), where it is generally assumed the plasma upstream of the shock is diffuse and negligible. Granted, this is a strongly ionized flow case at Mach 35. From the formal definition of a plasma (Section 1.2), the plasma parameter at $x = -0.0109$ m is $N_D = 196$, $\omega_{p,e}\tau_{e-n} = 13,100$, and $\lambda_{D,e} = 1.57 \times 10^{-6}$ m. A generic

characteristic length scale L for a hypersonic flowfield is on the order of 0.01 m to 1.0 m, in which case the quasineutrality criterion $\lambda_{D,e} \ll L$ holds. Therefore, the partially ionized gas upstream of the shock is non-negligible and should be treated as a proper plasma.

Figure 5.2 plots the species number densities, ionization fraction, temperatures, electron Debye length, and ratio of hypersonic flowfield diffusion length for the Mach 39 flow at 81 km. The region of statistically significant plasma is also strongly ionized throughout. The neutral flowfield properties do not significantly change in terms of general shape with respect to the neutral-only flow (Figure 4.8). The plasma reaches its peak ionization fraction at $x = -0.0127$ m and its peak plasma density at $x = -0.00628$ m. The ion and neutral temperatures follow similar trends to the 60 km flow case, with the ions being roughly 9,000 K to 13,000 K hotter in the post-shock region (Figure 5.2b). The electron temperature ranges from 8,700 K to 13,400 K within the shock layer. The electron temperature upstream of the shock is similar in magnitude to the 60 km flow case.

5.1.2 The Hypersonic Diffusion Length

To apply the Phelps criteria for charged species diffusion, an appropriate diffusion length (Equation 3.8) must be chosen. In Chapter 3, Λ is calculated from the length of the entire computational domain. A similar approach could be applied here, but in the case of an infinitely long domain, Λ will always be much greater than $\lambda_{D,e}$ and the criterion becomes trivial if observing local behavior. Λ is meant to be representative of the typical distance a plasma species will traverse. The majority of plasma species form inside the shock layer and will not travel the full length of the computational domain. Furthermore, using the width of the computational domain in one direction is less sensible for a multi-dimensional simulation. Another option for a 1D setup is the width of the shock layer, measured from the shock standoff distance to the wall. This may be suitable for certain applications, such as continuum flows where the shock flowfield properties have sharper gradients. However, in rarefied flows the shock itself can be quite wide, and as shown in Figure 5.2a the plasma can extend far upstream of the shock. Using the width of the shock layer is also less sensible for application to multi-dimensional geometries; for example, a plasma formed along

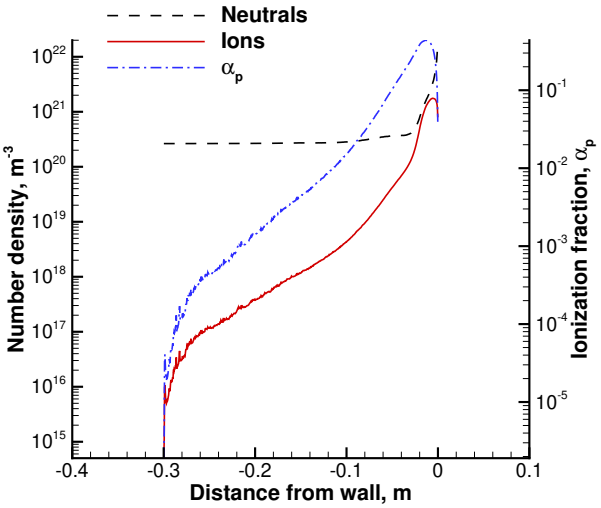

(a) Number densities for neutrals and ambipolar plasma (left axis) and ionization fraction (right axis).

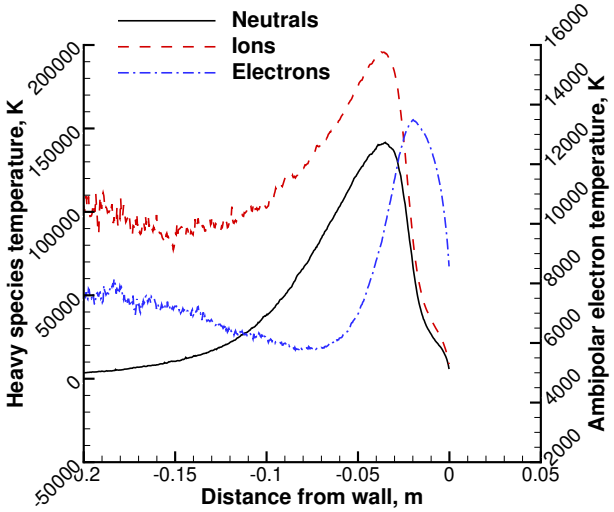

(b) Temperature for neutrals and ambipolar ions (left axis) and ambipolar electrons (right axis).

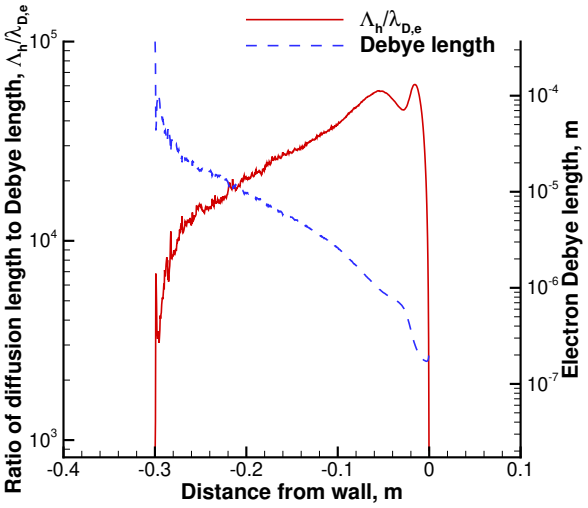

(c) Electron Debye length (left axis), computed from the ambipolar plasma density and the ambipolar electron temperature, and the ratio of the hypersonic diffusion length to the Debye length (right axis).

Figure 5.2: Baseline flowfield results for the Mach 39 flow at 81 km with the ambipolar diffusion approximation on.

the leeward side of a hypersonic vehicle should not use the shock standoff distance at the front of the vehicle to represent its diffusion length.

This work defines the local diffusion length Λ_h for a 1D parallel plane within a hypersonic flowfield as

$$\Lambda_h(x) = x - x_w \tag{5.1}$$

where for a length scale calculated at location x, $x - x_w$ is the shortest distance to the wall or vehicle surface. The Phelps diffusion ratio of $\Lambda_h/\lambda_{D,e} = \Lambda_h(x)/\lambda_{D,e}(x)$ is then calculated with respect to the local Debye length. Since many hypersonic simulation solid surface boundary conditions are at least partially catalyctic or absorbing to charged species, this definition partially fulfills the Phelps's criteria's derivation assumption of the plasma density being zero at the boundaries. Equation 5.1 aptly accounts for the range of length scales present in a hypersonic flowfield. Within the bulk plasma region where diffusion is expected to be ambipolar, Λ_h is similar in magnitude to the shock layer width and thus is representative of the typical distance a charged species will travel to the wall. A crucial length scale where the plasma is not charge neutral is the plasma sheath adjacent to the surface of the vehicle. Since plasma sheaths have a width on the order of several $\lambda_{D,e}$, $\Lambda_h/\lambda_{D,e}$ should have $\sim \mathcal{O}(1)$ within a sheath. The criteria presented in Section 3.4 characterizes this regime as transitional diffusion and would correctly recommend full electrostatic modeling. Finally, in the limit of $x \to \infty$, the plasma formed by a hypersonic vehicle will have its diffusion approximated as ambipolar. This highlights a subtle point in understanding plasma quasineutrality in hypersonics: within a hypersonic shock layer, there may be many fluctuations in charge separation and changes in diffusion regime. But, to an observer sufficiently far from the vehicle, the plasma can be treated as quasineutral since these fluctuations are negligible compared to the scale of Λ_h. Finally, the form of Equation 5.1 can readily be tested for use in multi-dimensional geometries by creating a line of sight to the vehicle. Validation of this approach is outside the scope of this work, but recommendations for exploration of this idea are included in Section 6.3.1.

Specific limits for $\Lambda_h/\lambda_{D,e}$ relevant to hypersonic flows are explored and verified in Section 5.5. In the analyses preceding that section, the limits of Section 3.4 will be used for general discussion. Returning to the flow survey, both flow cases generate plasmas with $\Lambda_h/\lambda_{D,e} > 10{,}000$ in the bulk plasma region, indicating strong collective behavior and quasineutrality. Within the plasma sheath for each flow case, $\Lambda_h/\lambda_{D,e}$ drops rapidly. The configuration space grids in this survey do not resolve the Debye length, thus the value of $\Lambda_h/\lambda_{D,e}$ adjacent to the wall is overpredicted in the plots. However, it is clear from the trends that with more grid resolution adjacent to the wall, $\Lambda_h/\lambda_{D,e}$ will decrease into the transitional diffusion regime with $1.0 < \Lambda_h/\lambda_{D,e} < 100$.

5.1.3 Artificial Density Reduction Method

Both re-entry flow cases produce a strongly ionized shock layer that likely falls entirely within the ambipolar diffusion regime. It would be preferable to study a diverse range of realistic flow cases, particularly 'weakly' ionized plasmas with $n_p \lesssim \mathcal{O}(10^{13})$ m^{-3}. However, as discussed in Section 4.4, even with the current semi-dynamic SWS scheme the macroparticle count would quickly become intractable. This work is primarily concerned with studying plasma diffusion of all regimes in hypersonic flowfields, developing criteria for predicting said regimes, and assessing the strengths and weaknesses of the ambipolar diffusion approximation. Simulating realistic flow conditions with argon is useful for gauging potential consequences for vehicle design. However, realistic flow conditions are not inherently necessary for testing prediction criteria or evaluating the ambipolar diffusion approximation, especially if scaling arguments can be made.

To explore a wider range of $\Lambda/\lambda_{D,e}$ in hypersonic flowfields and reduce the computational cost per simulation, a method of artificially reducing the freestream densities of the flowfield is employed [38, 138]. For each of the baseline flow cases in Table 5.1, the freestream number density is reduced by the amount required to bring the maximum plasma density down to a specified magnitude. The argon collision cross sections are then increased to yield the same mean free path as the baseline flow; this multiplication factor is labeled F_r. This approach ensures the profiles of axial velocity, temperature, Kn, α_p, and $\Lambda/\lambda_{mfp,i}$ across the flowfield remain similar in magnitude to the baseline

flow case when the ambipolar diffusion approximation is invoked; verification of which is included in this section. If electrostatic modeling is used, these properties may change due to the impact of the electric field on the charged species.

The artificial density reduction method does not preserve the ratio of $\Lambda/\lambda_{D,e}$. $\lambda_{D,e}$ is inversely proportional to the square root of the plasma density (Equation 1.3) and will increase if the freestream density is decreased while the ionization fraction is maintained. Subsequently, the artificial density reduction method will substantially decrease the computational cost of a plasma simulation which resolves the Debye length. It also enables study of hypersonic flows with realistic shock layer widths in the free or transitional charged species diffusion regimes if $\Lambda/\lambda_{D,e}$ is reduced a sufficient amount. However, care must be taken so that the ratio of $\Lambda/\lambda_{D,e}$ is not decreased below the ambipolar limit if observation of ambipolar diffusion is desired. Furthermore, since the magnitude of the electric field scales with the magnitude of charge separation, any reported differences in macroscopic flow properties between a flowfield obtained with electrostatic modeling and without may be underestimated.

To demonstrate its successful implementation in this work, both flow cases are re-run with $F_r = 10^8$. Several flowfield properties from the Mach 35 flow at 60 km simulations are plotted and analyzed to quantify agreement. Figure 5.3a plots the flowfield density ratio with the freestream as well as the ionization fraction across the flowfield. The flowfield density ratio with the freestream density profiles agree within 0.51% RMSE, and within the shared region of statistical significance, the ionization fraction profiles agree within 2.65% RMSE. For neutral temperature and electron temperature (Figure 5.3b), the profiles agree within 1.24% and 3.95% RMSE, respectively. This error is dominated by the fact that the profiles are spatially shifted from one another by a maximum of 3.07%, but overall, collision dynamics and subsequent ionization with the increased cross sections are being simulated correctly. Therefore, despite the decrease in true number density, the overall shock layer is still representative of super-orbital velocity conditions. As expected, in Figure 5.3c, the ratio of hypersonic diffusion length to electron Debye length profiles do not agree. The reduced density profile is $\sqrt{F_r} = \sqrt{10^8}$ orders of magnitude smaller than the baseline profile, consistent with

the definition of the electron Debye length (Equation 1.3). The significance of the $\Lambda_h/\lambda_{D,e}$ profile for the reduced density flow case will be further discussed in Section 5.3.2.

Table 5.3 compares the stagnation point heat flux by species for the baseline flow cases and the artificially reduced density flow cases with the ambipolar diffusion approximation on. Each of the reported heat flux values for the reduced density simulations is scaled by F_r. The heat fluxes from the baseline simulations and the reduced density simulations agree within 3.15% on a species-by-species basis, and within 1.72% for total stagnation point heat flux. This verifies that the artificial reduced density method can reasonably simulate the order of magnitude of heat fluxes from a full scale density simulation.

For the Mach 35 flow at 60 km, the peak ionization fraction in the bulk plasma region is ∼48%, therefore the ambipolar plasma species contributing ∼34% to the total stagnation point heat flux is roughly proportional to their mass fraction. Electrons consistently contribute the least to the total heat flux, in part due to their lightweight mass as well as their low temperature relative to heavy species.

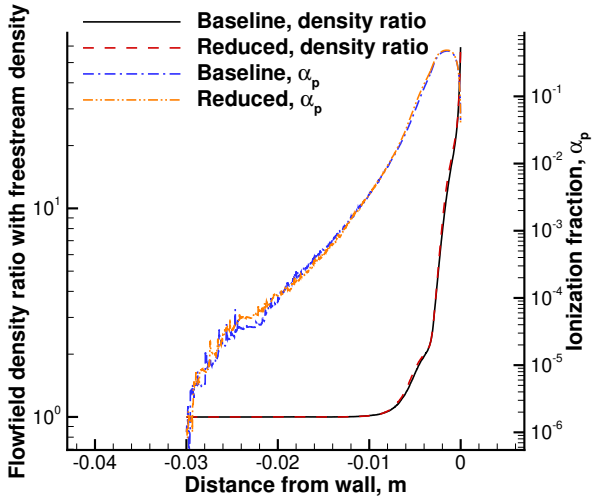

(a) Ratio of neutral density with freestream density (left axis) and ionization fraction (right axis).

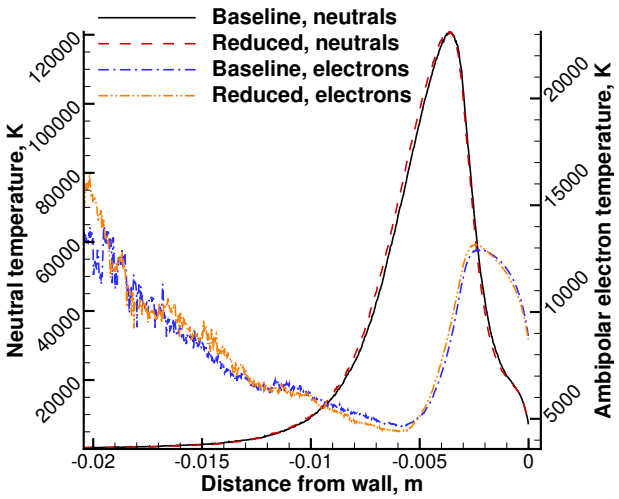

(b) Temperature for neutrals (left axis) and ambipolar electrons (right axis).

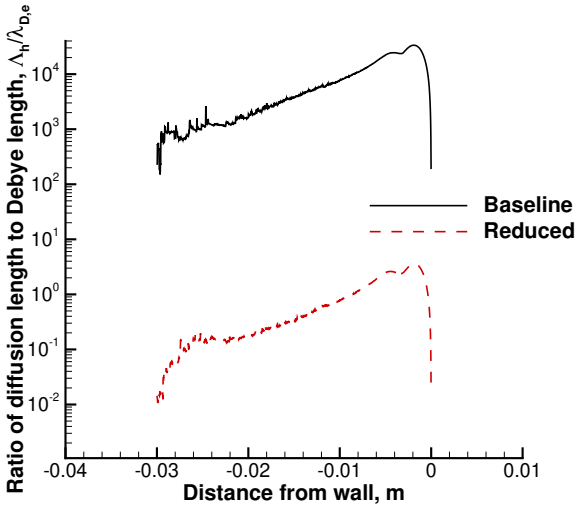

(c) Ratio of the hypersonic diffusion length to the Debye length.

Figure 5.3: Comparison between flowfield properties of Mach 35, 60 km flow case with baseline parameters and with an artificially reduced density of $F_r = 10^8$ (ambipolar diffusion approximation on).

Table 5.3: Stagnation point heat flux by species for simulations with baseline parameters and with an artificially reduced density of $F_r = 10^8$ (ambipolar diffusion approximation on). Error is calculated with respect to the baseline data.

	60 km			81 km		
Stagnation point heat flux, W/cm^2	Baseline	Reduced $\times F_r$	Error	Baseline	Reduced $\times F_r$	Error
Ar, convective	5,770	5,930	+2.77%	331	337	+1.81%
Ar^+, convective	708	690	−2.54%	25.7	26.0	+1.17%
Ar^+, chemical	2,220	2,150	−3.15%	100	102	+2.00%
Ambipolar e^-, convective	161	156	−3.10%	7.12	7.23	+1.50%
Total heat flux, plasma	3,090	3,000	−2.91%	133	135	+1.50%
Percentage of total, plasma	34.4%	33.5%	−0.900%	28.7%	28.6%	−0.100%
Total heat flux, all species	8,860	8,930	+0.790%	464	472	+1.72%

5.2 Boundary Conditions for the Poisson Equation

Generally, simulation solutions for a hypersonic flowfield without electrostatic modeling will be invariant to the size of a domain as long as the domain encompasses the entire shock layer. For example, if the computational domains of the flowfields of Section 5.1.1 are doubled in size, the solution does not change since all perturbations from the freestream are captured within the initial domain lengths. The same is not necessarily true for solutions to the Poisson equation depending on the boundary conditions implemented.

Physically, infinitely upstream from a hypersonic vehicle the electric field will be zero as the flowfield comprises neutral gas. The vehicle itself may be treated as an electrically insulating surface, and the plasma sheath will bias the potential at the wall until the net current to the wall is zero. This is known as a floating wall potential, and relative to the freestream potential of 0 V, the floating potential of the vehicle will be negative.

To represent this environment in a computational domain, several combinations of boundary conditions have been used throughout the literature. Most studies in the literature implement a Neumann ($(d\phi/dx)_{in} = 0$ V/m) condition at the inflow boundary and a Dirichlet ($\phi_w = \phi_{const}$) condition at the wall boundary with $\phi_w = 0$ V. Theoretically, this setup replicates physical reality because solutions to the Poisson equation are gauge invariant. That is, if the wall potential is fixed to zero and the inflow potential rises to $\phi_{in} = 2$ V, the flowfield solution should be the same if the potential solution is shifted, such as fixing $\phi_{in} = 0$ V and allowing ϕ_w to drop to -2 V. Blanco and Josyula [41] used a Neumann condition at the inflow boundary and a Robin condition ($(d\phi/dx)_w = 0$ V/m and $\phi_w = 0$ V) at the wall boundary. The authors justified the zero gradient at the wall due to the fact that the plasma sheath was unresolved in the simulations. In any case, for a converged solution to the Poisson equation, at least one boundary must be grounded with a constant potential.

Computationally, implementing a true floating potential at the wall as a self-adjusting Dirichlet condition is highly nontrivial. As the plasma potential is adjusted, the plasma will oscillate in

response which perturbs the current at the wall. This leads to a feedback loop of plasma instability, and the simulation may never reach steady state. In simulations of plasma sheaths this effect scales with the length of the computational domain L since the plasma wave can travel further upstream from the wall, where $L > 50\lambda_{D,e}$ is considered a long computational domain. Clearly, hypersonic flowfields with $L \gtrsim \mathcal{O}(0.1)$ m far exceed this recommended maximum length. Methods of implementing a floating potential boundary condition exist for local simulations of plasma sheaths [114] and were investigated for use in this work, but ultimately did not lead to stable plasma simulations. Thus, development of a novel floating wall potential condition for hypersonic flow simulation is outside the scope of this work. Numerical verification of gauge invariance between simulations with $(d\phi/dx)_{in} = 0$ V/m and $\phi_w = \phi_{const}$ and simulations with $(d\phi/dx = 0)_{in} = 0$ V/m, $\phi_{in} = 0$ V, and $\phi_w = \phi(t)$ via a floating wall potential condition is recommended for future work. Exact distributions of electric potential may be relevant for study of space-charge limited sheath effects.

Two configurations of Poisson equation boundary conditions are investigated in this section, each using a Dirichlet boundary condition at the wall of $\phi_w = 0$ V. The first is a Neumann boundary condition at the inflow, standard in the literature. The second is a Dirichlet condition at the inflow with $\phi_{in} = 0$ V. The Mach 39 flow at 81 km with $F_r = 10^8$ (Table 5.7) is used for analysis as the resulting electric fields are relatively small in magnitude but still non-negligible. Thus, if solutions to the Poisson equation are sensitive to the varied parameters in a small electric field case, it follows that those parameters are even more important for obtaining correct solutions as the strength of the ambipolar field increases. Note that the plotted electric potential and field solutions are time-averaged over the steady state period, but the actual Poisson solver in MPIC always calculates a new solution each time step with respect to the instantaneous charge separation throughout the flowfield.

5.2.1 Neumann Condition

Figure 5.4 compares solutions to the Poisson equation when a Neumann condition is applied to the inflow boundary for a domain length of $L = 0.5$ m, 0.3 m, and 0.2 m. The grid cell size is

preserved across all cases. When the simulation restarts from an ambipolar diffusion approximation solution, similar to the expanding weakly ionized gas flow of Chapter 3, electrons quickly diffuse upstream and away from the ions. The resulting self-induced ambipolar electric field is negative for most of the shock layer and upstream of the shock in order to accelerate ions upstream and decelerate electrons. A positive, rapidly increasing electric field also forms adjacent to the wall, indicative of sheath behavior. The electric field profiles (Figure 5.4a) strongly agree within the statistically significant plasma region with a maximum of 3.98% RMSE between all flow cases. The potential profiles (Figure 5.4b) demonstrate less agreement: as the length of the domain increases, the potential relative to the grounding wall increases. This aligns with expectations in a numerical sense: a Neumann boundary condition enforces $(d\phi/dx)_{in} = 0$ V/m exactly and only at the boundary. If the computational domain increases in size, the solver may relax to a solution that approaches zero gradient as gradually as possible. Hypothetically, even if the charge separation is identically zero throughout the domain upstream of the shock, this only guarantees $d^2\phi/dx^2 = 0$ V/m^2, not necessarily $\phi = 0$ V. The electric potential must always be interpreted as a quantity relative to the grounding point, which via the initialization of the Poisson solver is the solid wall with $\phi_w = 0$ V. The most important quantity for modeling purposes is the electric field since it accelerates charged particles, which does not demonstrate strong spatial dependence when a Neumann condition is applied to the inflow.

5.2.2 Dirichlet Condition

Figure 5.5 compares flowfield solutions when a Dirichlet condition is applied to the inflow boundary for a domain length of $L = 0.3$ and $L = 0.2$ m. The Neumann profile for $L = 0.3$ m is also plotted for reference. Within the shock layer, the electric field profiles are similar in shape between the Neumann and the Dirichlet inflow conditions. However, the Dirichlet profiles have nonzero electric fields at the inflow boundary. This is numerically correct because the charge density at the boundary is nonzero and negative for all flow cases, but, from a modeling perspective this is not representative of a hypersonic flowfield. Furthermore, the solution is far more dependent on

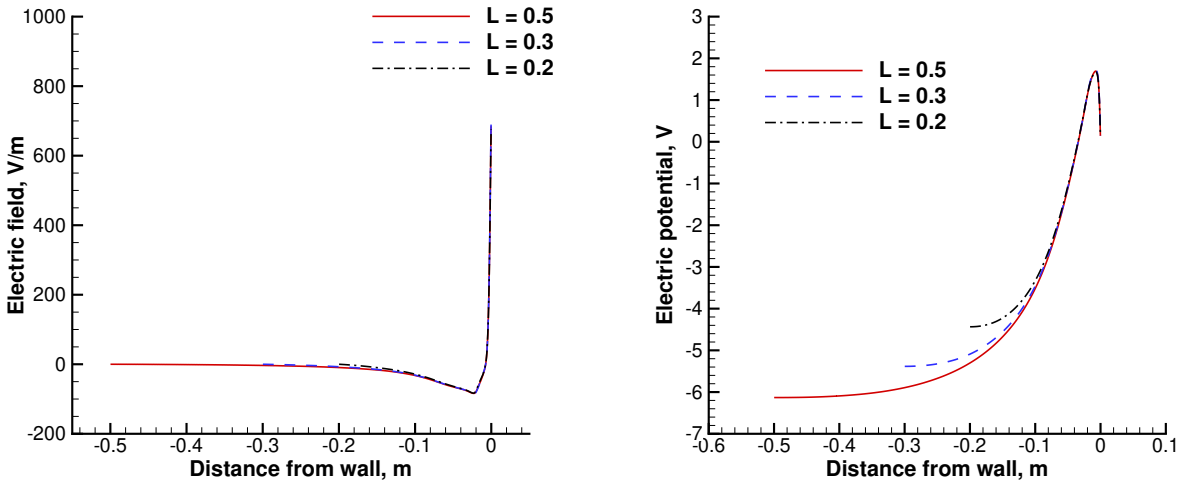

(a) Electric field.

(b) Electric potential.

Figure 5.4: Variation in Poisson equation solutions when the domain length is varied for a Neumann boundary condition at the inflow (Mach 39 flow at 81 km, $F_r = 10^8$).

the length of the domain since the potential curve must always curve back to zero at the inflow boundary. Although the Neumann profiles are also spatially dependent, when the length of the domain is increased, the electric fields upstream of the shock are closer to zero than the Dirichlet electric fields.

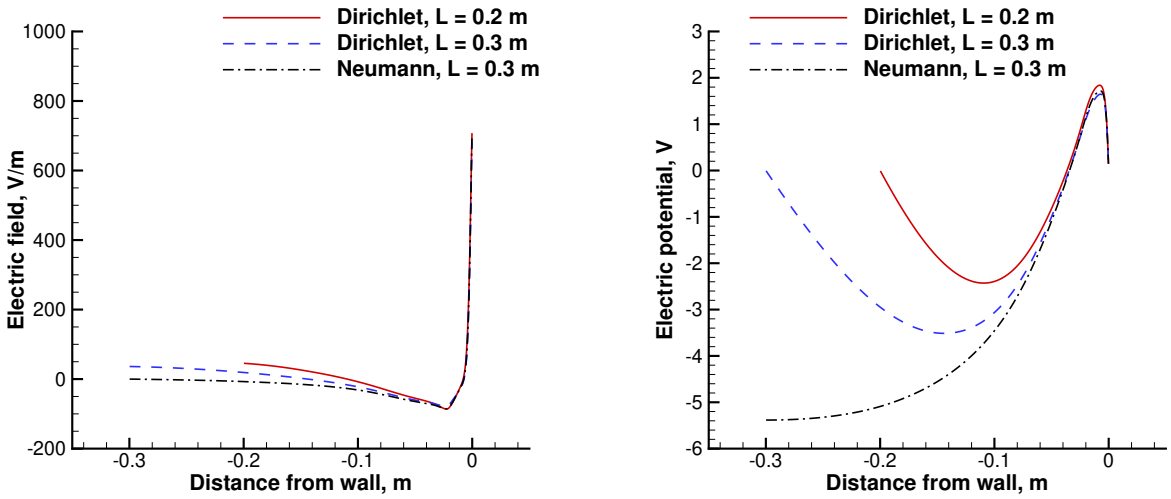

(a) Electric field.

(b) Electric potential.

Figure 5.5: Variation in Poisson equation solutions when the domain length is varied for a Dirichlet boundary condition at the inflow (Mach 39 flow at 81 km, $F_r = 10^8$).

Despite the differences in boundary conditions, the ion and electron number density profiles demonstrate strong agreement, with a maximum difference of 2.34% RMSE between all solutions in the region of statistical significance, and maximum 0.725% RMSE between the Neumann solutions. The electron temperature profiles demonstrate similar levels of agreement. This is in part because the magnitude of the electric field is small throughout the flowfield, thus any variation in Poisson solution will be small. Thus, the electric field and potential profiles within the bulk plasma region are not strongly dependent on the inflow boundary condition as long as the boundary is sufficiently far upstream. As will be shown in Section 5.5, despite the small magnitude of the electric field, flowfield solutions will still be different than those obtained with the ambipolar diffusion approximation.

From this analysis, the Neumann condition is recommended for use at the inflow boundary, consistent with existing literature. Additionally, care should be taken to extend the domain length sufficiently far upstream so as to capture the gradual return of the electric field to zero. Otherwise, the electric field profile compresses in width, resulting in sharper gradients that may impact charged species acceleration. Ideally, this would be the length required for the plasma density to return to exactly zero; practically, this length should encompass at least the width of the statistically significant plasma.

A Robin condition at the inflow boundary is not tested in this analysis due to limitations of the MPIC code. It can be inferred that the resulting electric potential profiles from a Robin boundary condition would look similar to the Dirichlet profiles, but with a gradual sloping of the potential curve towards zero at the inflow such that $(d\phi/dx)_{in} = 0$ V/m. This would still result in a spatially dependent solution, where as the length of the domain increases, the distance over which the electric field is nonzero also increases. Additionally, fixing both $\phi_{in} = 0$ V and $\phi_w = 0$ V implies the two boundaries have zero electric potential relative to each other. This is not correct physically since in reality, the wall itself becomes biased due to the presence of the plasma. A Robin condition at the inflow boundary would be required to capture a potential difference between the wall and the freestream if a floating potential condition at the wall is successfully developed for hypersonic flowfield-scale simulation.

5.3 Charged Species Diffusion Regimes across Transitional Hypersonic Flows

Strictly speaking, a self-induced electric field will always have some effect on the charged particles inside it– what classifies as a negligible effect depends on the application. At the same time, at a sufficiently large length scale, a partially ionized collisional gas of any Debye length may be re-classified as a plasma and into the ambipolar diffusion regime. Thus, this work characterizes plasma diffusion regimes of hypersonic flows relative to typical vehicle length scales of $\mathcal{O}(0.01)$ m to $\mathcal{O}(1.0)$ m as well as properties observed in the flowfields themselves.

In the following section, each flow case is run with the ambipolar diffusion approximation

and with full electrostatic modeling (with $(d\phi/dx)_{in} = 0$ V/m, $\phi_w = 0$ V, and modeling of ions and electrons as separate species). For certain flow cases, a simulation is also run with the 'free diffusion approximation', defined in Section 1.3.2. Normalized charge separation $\overline{n}_c$ is calculated using Equation 3.9. As in Section 3.3, $\overline{n}_c < 10^{-2}$ is used to assess whether a flow is quasineutral. For the flows of that study, $\overline{n}_c < 10^{-2}$ marked the point where a flow transitions into the ambipolar diffusion regime. Since that study is conducted with the deterministic discrete-velocity method with highly resolved configuration and velocity space grids and the true electron mass, it indicates that $\overline{n}_c = 10^{-2}$ is a reasonable value for numerically assessing whether a flow is quasineutral.

When the Langmuir and Tonks relation (Equation 3.5) is calculated from the electrostatic solution, it uses the electron properties; when it is calculated from the ambipolar diffusion approximation solution, it uses the ambipolar plasma density and the ambipolar electron temperature.

5.3.1 Free Diffusion

This work defines the free plasma diffusion regime in hypersonic flows as having no portion of the shock layer undergo ambipolar diffusion, evidenced by the values of $\omega_{p,e}\tau_{e-n}$, $\Lambda_h/\lambda_{D,e}$, and $\overline{n}_c > 0.1$. The free diffusion regime is most relevant for weakly ionized hypersonic flows with $\alpha_p < 10^{-4}$, such as on the wake side of a hypersonic vehicle at high altitudes [139].

In the following flow cases, $\lambda_{D,e} > \lambda_{mfp,n}$. However, this is reflective of the low density of the partially ionized gas rather than the validity of Equation 5.1 in predicting ambipolar diffusion. Simulation initialization parameters are included in Table 5.4, and quantities relevant to analysis are included in Tables 5.5 and 5.6.

5.3.1.1 Free Diffusion at 60 km

Figure 5.6 presents flowfield results with the ambipolar diffusion approximation and with full electrostatic modeling for the Mach 35 flow at 60 km when $F_r = 10^{10}$. Select results from the free diffusion approximation solution are also shown. The charge separation throughout the shock

Table 5.4: Simulation parameters for observation of free diffusion in hypersonic flowfields.

Parameter	60 km	81 km
Domain length, m	0.03	0.3
Uniform grid width Δx, m	2.22×10^{-5}	4.00×10^{-4}
Points per minimum $\lambda_{D,e}$	≥ 954	≥ 53.6
Reduced density factor, F_r	10^{10}	10^{10}
Global weight	5.41×10^{5}	5.28×10^{5}
Plasma weight modifier	0.05	0.05
Time step, s	10^{-10}	10^{-9}

layer (Figure 5.6b) remains above 30%*; thus the flowfield is not quasineutral. This is because the partially ionized gas violates the third criterion for the definition of a plasma: $\omega_{p,e}\tau_{e-n} = 0.455$ at the point of peak electrostatic electron density, and $\omega_{p,e}\tau_{e-n} = 0.199$ at the point of peak ambipolar plasma density. The profile of $\Lambda_h/\lambda_{D,e}$ from both the electrostatic and ambipolar diffusion approximation solutions also confirms this, with both achieving $(\Lambda_h/\lambda_{D,e})_{max} < 1.0$ within the shock layer. This classifies the partially ionized gas as in the free diffusion regime. The plasma parameter N_D is not illustrative by itself for predicting the plasma diffusion regime. For this flow case, the second criterion for the definition of a plasma ($N_D \gg 1$) holds, but all three criteria must be satisfied for the gas to be considered a plasma (and thus, exhibit sufficient collective behavior to support ambipolar diffusion).

Despite being in the free diffusion regime, the self-induced electric field has a non-negligible impact on flowfield properties. In Figure 5.6a, if the electric field had no impact on the flow, then the electrostatic electron number density profile would align with the free diffusion approximation solution. Instead, the electrostatic electron number density is at least an order of magnitude lower than the free diffusion approximation electron number density throughout the flowfield. Figure 5.6c plots the electron temperatures, where both the free diffusion approximation and the ambipolar diffusion approximation overpredict the electron temperature throughout the shock layer compared

*At one point slightly upstream of the shock center, the charge separation decreases below 30%; this is an intersection of density profiles and is not indicative of quasineutrality or ambipolar diffusion at that point. This also occurs for the Mach 39 flow case at 81 km with $F_r = 10^{-10}$.

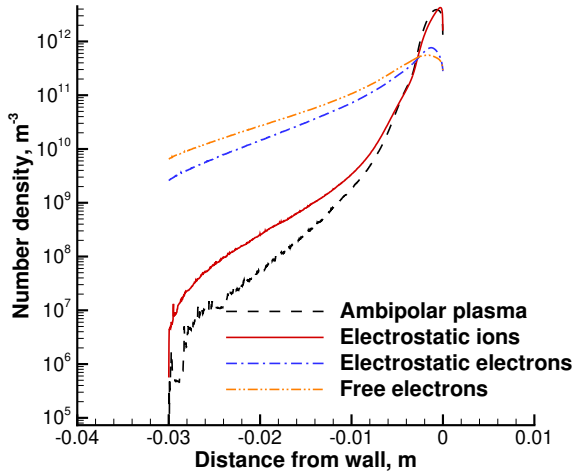

(a) Charged species densities.

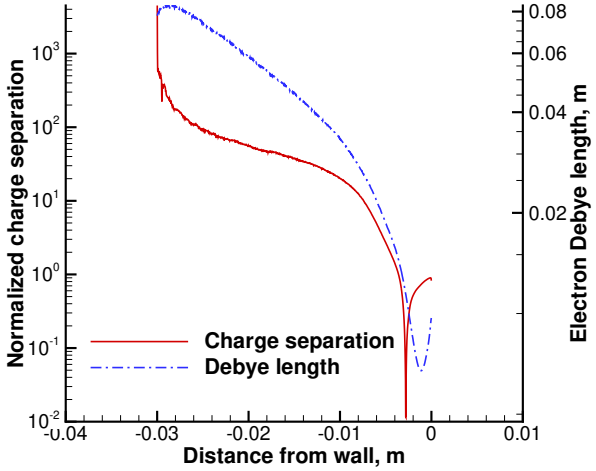

(b) Normalized charge separation from electrostatic simulation (left axis) and electron Debye length (right axis) from the electrostatic solution. The black dashed line separates the domain where the charge separation switches from negative to positive.

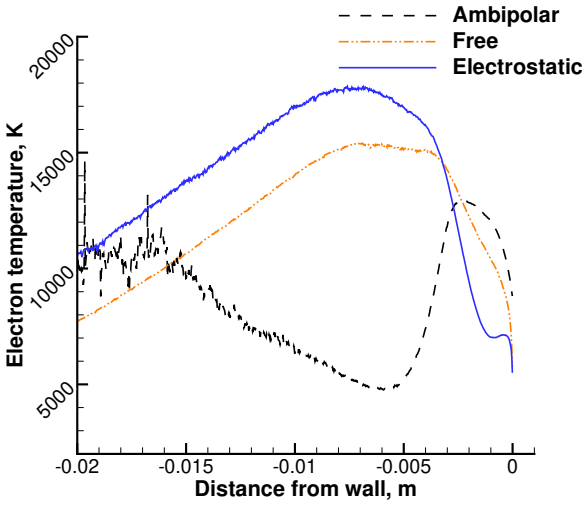

(c) Electron temperature.

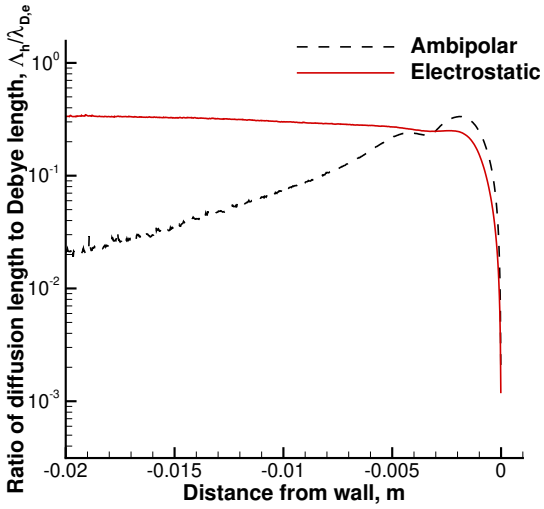

(d) Ratio of the hypersonic diffusion length to the electron Debye length.

Figure 5.6: Flowfields with the ambipolar diffusion approximation and with full electrostatic modeling (Mach 35 flow, 60 km, $F_r = 10^{10}$, free diffusion).

to the electrostatic solution. Specifically, the ambipolar diffusion approximation overpredicts the electron temperature by upwards of 39.6%. This can *only* be due to the self-induced electric field, which has a peak value of -24.2 V/m in the shock layer (Figure 5.7). The negative electric field decelerates the electrostatic electrons, decreases their kinetic energy, and uses that energy to maintain itself. Upstream of the shock, the electrostatic electron temperature is significantly higher than the other two models because only high energy electrons can escape the electric field.

Conversely, by forcing electrons to follow the ion density distribution via the ambipolar diffusion approximation, the ambipolar electrons collide with high temperature neutrals in the shock layer, but cannot carry that increased energy out of the shock layer unless ions propagate upstream. Upstream of the shock, the ambipolar diffusion approximation electron temperature solution decreases because the gas is more rarefied and fewer collisions with high velocity neutrals occur. In that sense, the ambipolar electron temperature follows a standard temperature curve throughout the shock layer (that is, similar to the ion and neutral temperature profiles of Figures 5.1b and 5.2b), where the temperature increases across the shock from the freestream and decreases towards the wall. The free diffusion approximation electrons are more evenly distributed throughout the domain and thus do not experience any strong gradients in temperature, except at the wall. In terms of overall shape, the free diffusion approximation electron temperature aligns more with the ambipolar diffusion approximation than the electrostatic solution. Finally, the peak neutral temperature throughout the flowfield decreases by 13.0%. This supports the hypothesis that mobile electrons transport energy out of the shock layer: either by diffusing upstream or by quickly absorbing into the wall.

The Poisson solver electric field does not agree with either of the Langmuir and Tonks solutions, which overpredict the bulk region peak electric field value by a factor of 20 (electrostatic) or 86 (ambipolar diffusion approximation). Note that the Langmuir and Tonks relation always overpredicts the electric fields immediately adjacent to the wall because of the sharp electron density gradient; this is unphysical and can be neglected in analysis because the Langmuir and Tonks relation does not apply to plasma sheaths.

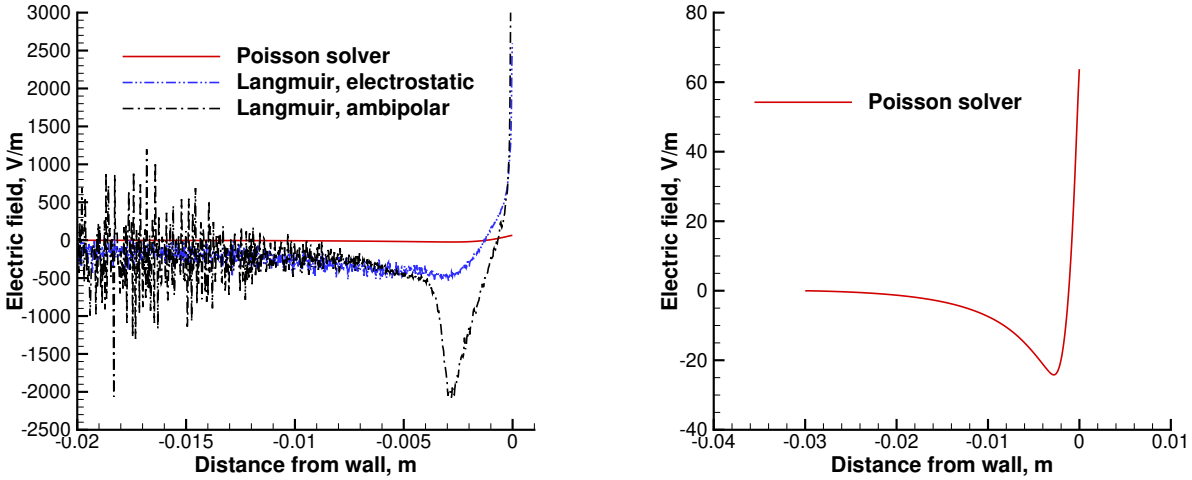

(a) Electric field.

(b) Electric field from only the Poisson solver.

Figure 5.7: Comparison between electric field profiles from the Poisson solver and the Langmuir and Tonks relation (Mach 35 flow, 60 km, $F_r = 10^{10}$, free diffusion).

The influence of the self-induced electric field on electron energy is also evident from their convective heat flux, which increases by 349% compared to the ambipolar diffusion approximation solution (Table 5.6). At first, this is counter-intuitive, since the electrostatic electrons have a lower temperature throughout the shock layer and the ambipolar plasma density is higher in the shock layer. However, when electrons are allowed to move independently, they will always reach the wall faster than ions due to their mass. This contributes to a larger number flux of electrons to the wall, thus leading to a larger electron convective heat flux. A plasma sheath forms whose strong, positive electric fields have two main consequences: ions are accelerated towards the wall, resulting in higher ion convective and chemical heat fluxes; and, electrons repelled by the sheath fields lose energy from deceleration, resulting in a lower electron temperature throughout the shock layer compared to the ambipolar diffusion approximation. Evidence of the sheath cooling effect and its strong coupling to the electron temperature throughout a hypersonic flowfield in the continuum regime with CFD was presented in Reference [44]. Thus, this phenomenon is not unique to either fluid or kinetic modeling. As for the other stagnation point heat flux quantities, the electrostatic solution reports an increase in total heat flux contribution from plasma species of 38.0% and an increase in total heat flux from all species of 15.2% relative to the ambipolar diffusion approximation. The increase in total heat flux of 15.2% may be considered non-negligible; and is similar in magnitude to other studies' reported increases in vehicle surface heat flux when electrostatic modeling is used [38, 41].

5.3.1.2 Free Diffusion at 81 km

The Mach 35 flow at 60 km with $F_r = 10^8$ ultimately fails to be categorized as a proper plasma due to its high collisionality with neutrals with $\omega_{p,e}\tau_{e-n} < 1.0$, leading more electrons to diffuse away and quasineutrality to be violated. Thus, investigation of free diffusion when the freestream gas is more rarefied warrants investigation. Figure 5.9 presents flowfield results with the ambipolar diffusion approximation and with full electrostatic modeling for the Mach 39 flow at 81 km when $F_r = 10^{10}$. Throughout the flowfield, the charge separation remains above 15% and thus the partially ionized shock layer is also not quasineutral. In fact, the electrostatic electrons diffuse

out so far that the entire computational domain is statistically significant for them– although not to the same extent as the free diffusion approximation electrons, indicating the electron motion is still being restrained by the self-induced electric fields. The ambipolar diffusion approximation overpredicts the electron temperature by upwards of 135%. Additionally, the ambipolar diffusion approximation ambipolar plasma density profile extends out 10% further than the electrostatic ion density profile, and the approximation underpredicts the peak ion density by 20.4%. This is caused by the sheath's electric field (Figure 5.8), which switches from negative to positive at $x = -0.001$ m, directly overlapping with the region where the electrostatic ion density exceeds the ambipolar ion density. Unlike the 60 km case, where more collisions occur to slow ion motion, the plasma sheath attracts ions towards the wall relatively unobstructed, resulting in the higher density and an increase in the ion chemical heat flux of 10.9% relative to the ambipolar diffusion approximation. This is unique to the free diffusion regime in highly rarefied flows, where it is possible to have charged species densities so small that both the Debye length and the sheath's length are similar in scale to the shock layer itself, and collisions with neutrals can strongly impact sheath dynamics since $\lambda_{mfp,i} < \lambda_{D,e}$.

Outside of the sheath, the Poisson solver's electric field is small in magnitude and relatively flat, reaching a peak value of only -8.77 V/m. However, the electron cooling effect is still present in the electron temperature profiles (Figure 5.9c), supported by the fact that the free diffusion approximation electrons have a higher temperature throughout the shock layer once more. The Langmuir and Tonks relation completely fails to predict the electric field by 757%, as expected for a strongly non-continuum flow at 81 km.

Both the ambipolar and electrostatic profiles of $\Lambda_h/\lambda_{D,e}$ predict a value less than 1.0 inside and outside of the shock layer, consistent with the limits identified in Section 3.4. Additionally, $\omega_{p,e}\tau_{e-n} \lesssim \mathcal{O}(1)$ for both models throughout the flowfield, indicating that collisions with neutrals dominate over the partially ionized gas's ability to respond to charge disturbances.

For this flow case, the total contribution to heat flux from plasma species increases by 25.4% when electrostatic modeling is used, similar in magnitude to the 60 km flow case. However, the

overall difference in total stagnation point heat flux remains small at +4.76%. Additionally, the peak neutral temperature decreases only by 7.41% compared to 13.0% in the 60 km case. This indicates a dependence between first-order plasma effects and the degree of rarefaction in the atmosphere.

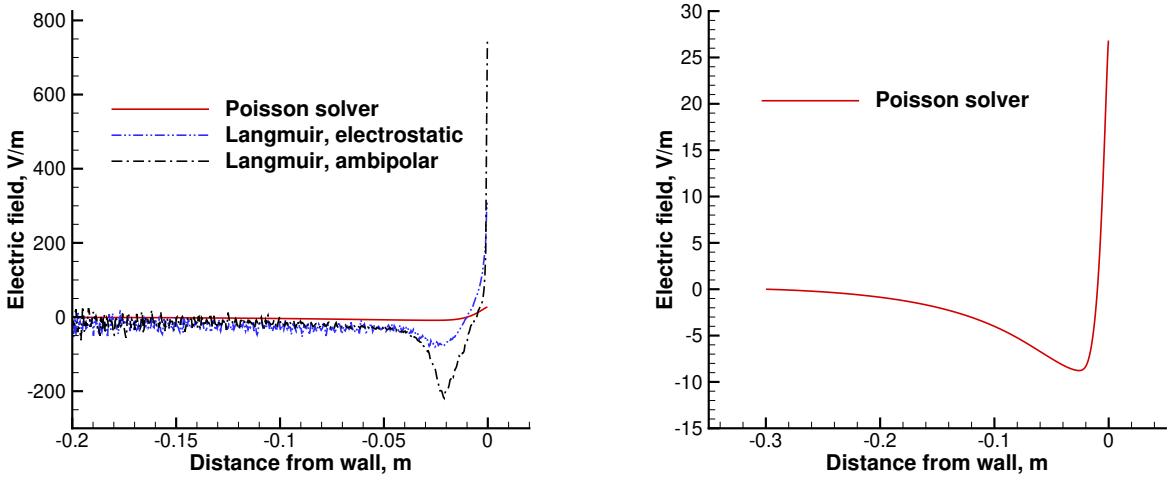

(a) Electric field.

(b) Electric field from only the Poisson solver.

Figure 5.8: Comparison between electric field profiles from the Poisson solver and the Langmuir and Tonks relation (Mach 39 flow, 81 km, $F_r = 10^{10}$, free diffusion).

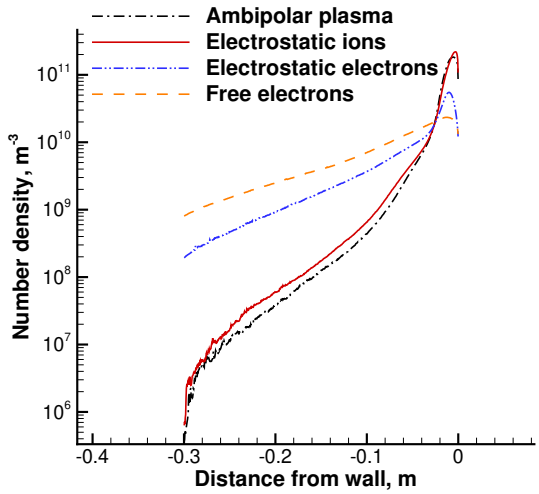

(a) Charged species densities.

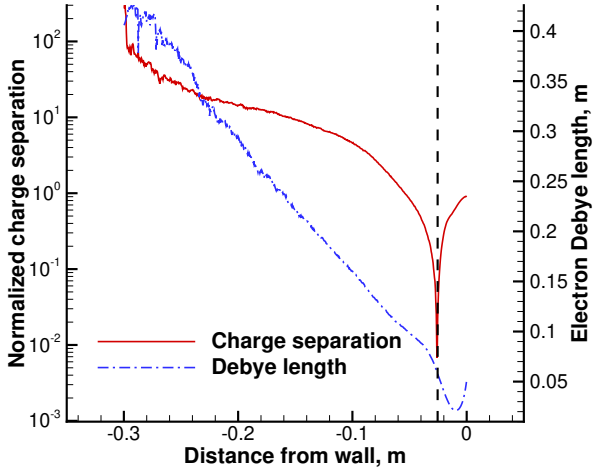

(b) Normalized charge separation from electrostatic simulation (left axis) and electron Debye length (right axis) from the electrostatic solution. The black dashed line separates the domain where the charge separation switches from negative to positive.

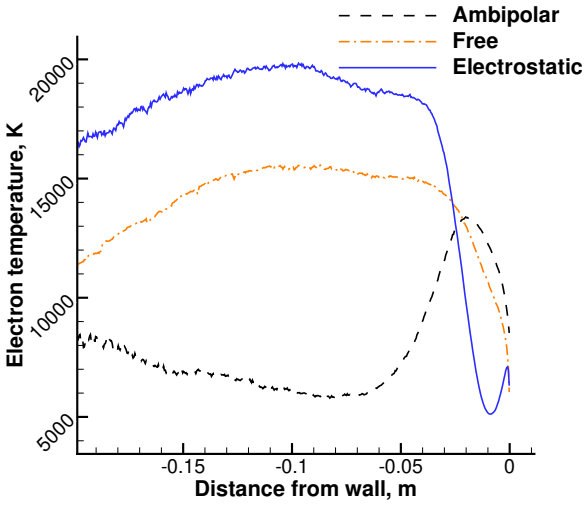

(c) Electron temperature.

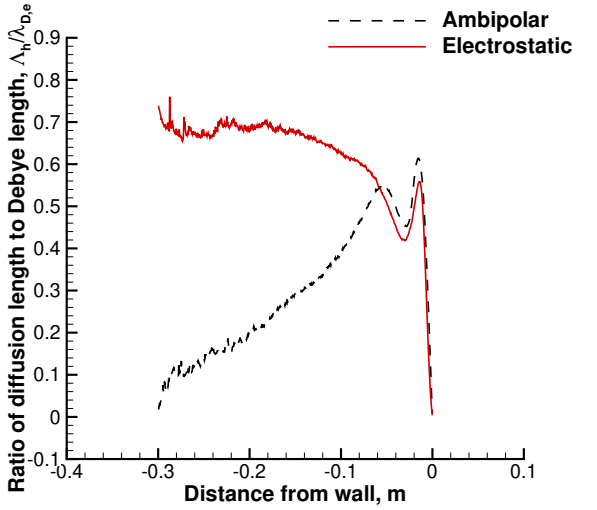

(d) Ratio of the hypersonic diffusion length to the electron Debye length.

Figure 5.9: Flowfields with the ambipolar diffusion approximation and with full electrostatic modeling (Mach 39 flow, 81 km, $F_r = 10^{10}$, free diffusion).

Table 5.5: Partially ionized gas properties of interest for simulations of free diffusion in hypersonic flows ($F_r = 10^{10}$).

	60 km		81 km	
Plasma property	Ambipolar	Electrostatic	Ambipolar	Electrostatic
Ions maximum resolved distance, m	-0.0110	-0.0133	-0.177	-0.196
Electrons maximum resolved distance, m	-0.0110	-0.0300	-0.177	-0.300
Maximum $\Lambda_h/\lambda_{D,e}$ in shock layer	0.353	0.234	0.614	0.558
N_D at peak electron density	9.39×10^5	8.80×10^5	2.22×10^6	3.85×10^6
N_D at maximum resolved distance, electrons	5.94×10^6	5.95×10^6	1.21×10^8	5.86×10^7
$\omega_{p,e}\tau_{e-n}$ at peak electron density	0.199	0.455	1.35	3.88
$\omega_{p,e}\tau_{e-n}$ at maximum resolved distance, electrons	0.141	0.592	0.320	0.660

Table 5.6: Stagnation point heat flux by species for simulations of free plasma diffusion regime in hypersonic flows. Each heat flux value is multiplied by $F_r = 10^{10}$. The change is calculated relative to the ambipolar diffusion approximation value.

	60 km			81 km		
Heat flux, W/cm^2	Ambipolar	Electrostatic	Change	Ambipolar	Electrostatic	Change
Ar, convective	5,714	5,910	+3.43%	328	316	−3.66%
Ar^+, convective	671	846	+26.1%	25.4	26.1	+2.76%
Ar^+, chemical	2,124	2,540	+19.6%	101	115	+13.9%
e^-, convective	153	687	+349%	7.17	27.0	+277%
Total heat flux, plasma	2,950	4,070	+38.0%	134	168	+25.4%
Percentage of total, plasma	34.0%	40.8%	+6.80%	28.9%	34.7%	+5.80%
Total heat flux, all species	8,660	9,980	+15.2%	462	484	+4.76%

5.3.2 Transitional Diffusion

This work defines the transitional plasma diffusion regime for hypersonic flows as having some portion of the post-shock region undergo ambipolar diffusion and/or some portion of the post-shock region has a normalized charge separation of $\overline{n}_c \leq 0.1$, or 10%. Across and upstream of the shock, a majority of partially ionized gas undergoes free diffusion. Simulation initialization parameters are listed in Table 5.7, and quantities relevant to analysis are included in Tables 5.8 and 5.9.

Table 5.7: Simulation parameters for observation of transitional diffusion in hypersonic flowfields.

Parameter	60 km	81 km
Domain length, m	0.030	0.300
Uniform grid width Δx, m	2.22×10^{-5}	4.00×10^{-4}
Points per minimum $\lambda_{D,e}$	≥ 15.7	≥ 3.95
Reduced density factor, F_r	10^8	10^8
Global weight	5.41×10^7	5.28×10^7
Plasma weight modifier	0.05	0.05
Time step, s	10^{-10}	10^{-9}

5.3.2.1 Transitional Diffusion at 60 km

Figure 5.10 presents flowfield results with the ambipolar diffusion approximation and with full electrostatic modeling for the Mach 35 flow at 60 km when $F_r = 10^8$ (increasing the plasma density by a factor of 100 from the free diffusion case). Excluding the plasma sheath, the post-shock plasma has normalized charge separation values between 8.73 and 16.3%. For the third criterion of the definition of a plasma, $\omega_{p,e}\tau_{e-n} = 8.86$ at the point of peak electrostatic electron density and $\omega_{p,e}\tau_{e-n} = 2.01$ at the point of peak ambipolar plasma density. The profile of $\Lambda_h/\lambda_{D,e}$ (Figure 5.10d) from the electrostatic and ambipolar diffusion approximation solutions is above 1.0 throughout the shock layer. Together, these quantities indicate the partially ionized gas is in the transitional flow regime.

As in the free diffusion flow cases, the contribution to heat flux from plasma species increases when electrostatic modeling is used. Of interest are the ion heat fluxes: where the ion convective

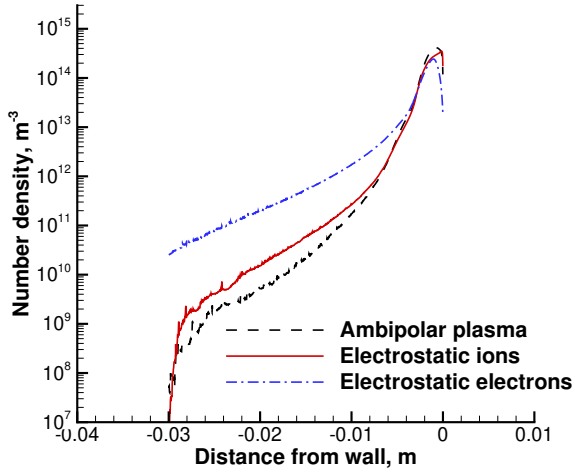

(a) Charged species densities.

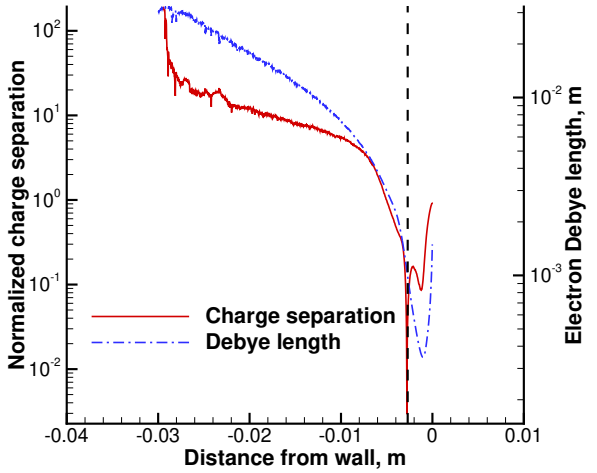

(b) Normalized charge separation from electrostatic simulation (left axis) and electron Debye length (right axis) from the electrostatic solution. The black dashed line separates the domain where the charge separation switches from negative to positive.

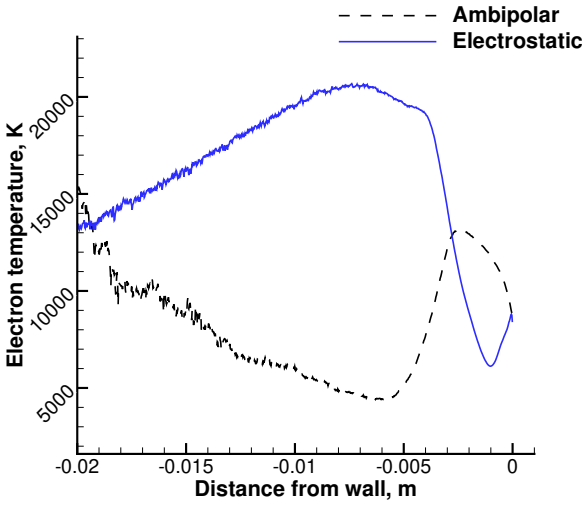

(c) Electron temperature.

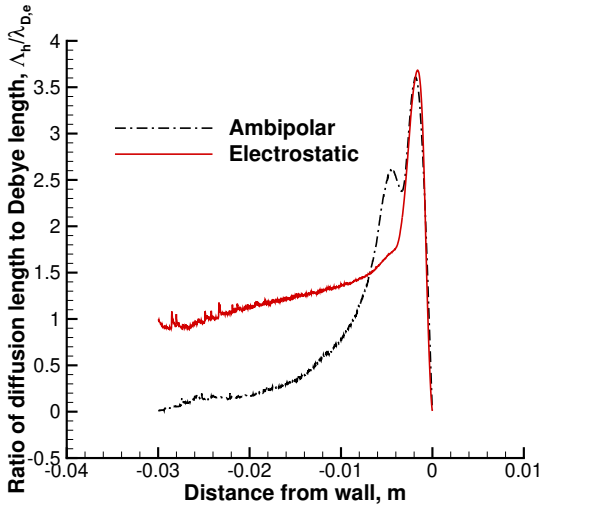

(d) Ratio of the hypersonic diffusion length to the electron Debye length.

Figure 5.10: Flowfields with the ambipolar diffusion approximation and with full electrostatic modeling (Mach 35 flow, 60 km, $F_r = 10^8$, transitional diffusion).

heat flux decreases by 11.6%, but the ion chemical heat flux increases by 27.0%. This indicates more ions are hitting the wall overall, as chemical heat flux is proportional to the number flux of ions. The electric field switches from negative to positive at $x = -0.00101$ m, which corresponds with the point on the electrostatic ion number density curve (Figure 5.10a) where the density increases sharply before dropping immediately adjacent to the wall. Thus, the increase in ion chemical heat flux is due to acceleration via the plasma sheath.

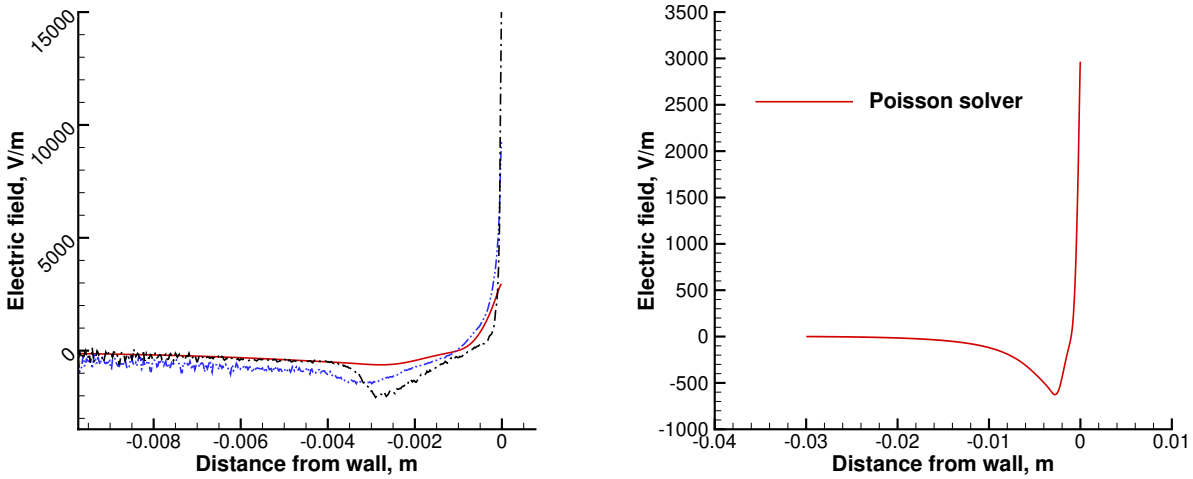

(a) Electric field.

(b) Electric field from only the Poisson solver.

Figure 5.11: Comparison between electric field profiles from the Poisson solver and the Langmuir and Tonks relation (Mach 35 flow, 60 km, $F_r = 10^8$, transitional diffusion).

Meanwhile, the decrease in convective heat flux is caused by fewer high energy ions hitting the wall, as high energy ions are more likely to propagate back upstream and accelerate via the negative electric fields rather than stay in the shock layer. The total increase in stagnation point heat flux from all species between the ambipolar diffusion approximation and the electrostatic solutions is only 5.60%: this is because the neutral convective heat flux decreases by 7.59%, and as the dominant species, neutrals still contribute the most to the heat flux. The neutral convective heat flux decreases because more electrons remain within the shock layer when stronger self-induced electric fields form. Thus, more energy transfer can occur from high-velocity neutrals to electrons via

collisions. This is supported by the fact that the peak neutral temperature decreases by 7.34% with electrostatic modeling, and that the electrostatic electron temperature increases overall from the free diffusion flow case. This trend in electron temperature continues into the ambipolar diffusion case of Section 5.3.3.1.

As in the free diffusion case, the Poisson solver electric field does not agree with either of the Langmuir and Tonks solutions, which overpredict the electric field by at least a factor of 2.0 (electrostatic) or 3.2 (ambipolar diffusion approximation).

5.3.2.2 Transitional Diffusion at 81 km

Figure 5.13 presents flowfield results with the ambipolar diffusion approximation and with full electrostatic modeling for the Mach 39 flow at 81 km when $F_r = 10^8$. From $x = -0.0376$ m to $x = -0.00757$ m, the plasma achieves a normalized charge separation value below 10^{-1}. This comprises the entire post-shock region, where $\Lambda_h/\lambda_{D,e}$ ranges between 4.00 and 6.61 for either model. Thus, the entire shock layer is in the transitional diffusion regime. The ambipolar diffusion approximation overpredicts the peak post-shock plasma density by 19.4%, and also overpredicts the electron temperature due to the effects of electron cooling. Specifically, at the electrostatic electron temperature minimum of 6,830 K, the ambipolar diffusion approximation overpredicts the electron temperature by upwards of 43.3%. This corresponds with the point where the electric field exceeds +389 V/m at $x = -0.00186$ m. The neutral macroscopic properties do not significantly change, with the peak neutral temperature decreasing by 3.25%. At the maximum resolved distance for electrons of $x = -0.238$ m, $\omega_{p,e}\tau_{e-n} = 99.0$, $N_D = 2.16 \times 10^6$, $\Lambda_h/\lambda_{D,e} = 2.58$, and $\overline{n}_c = 1.71$. Thus, although the partially ionized gas does not meet the criteria for a plasma due to not fulfilling the quasineutrality condition, there is a significant number of electrons upstream capable of quickly responding to charge disturbances as indicated by the high value of $\omega_{p,e}\tau_{e-n}$.

For the Mach 35 flow at 60 km, the Langmuir and Tonks electric field profiles have much more defined peaks located between the center of the shock and the shock standoff distance. At Mach 39 km and 81 km, the Langmuir and Tonks relation calculated with respect to the electrostatic

solution has a much smoother profile that is more consistent with the Poisson solver electric field. However, the magnitude of the peak electric field is still overpredicted by 78.8%. The Langmuir and Tonks relation calculated with respect to the ambipolar diffusion approximation displays a sharp peak located 16.8% closer to the wall than the Poisson solver's electric field.

The differences in total stagnation point heat flux are smaller for the transitional diffusion flow cases compared to the free diffusion flow cases at both 60 km and 81 km. This indicates that with increasing strength of the self-induced electric fields, cooling effects such as electron sheath cooling dominate over heating effects such as the acceleration of ions via the plasma sheath in terms of impact on stagnation point heat flux.

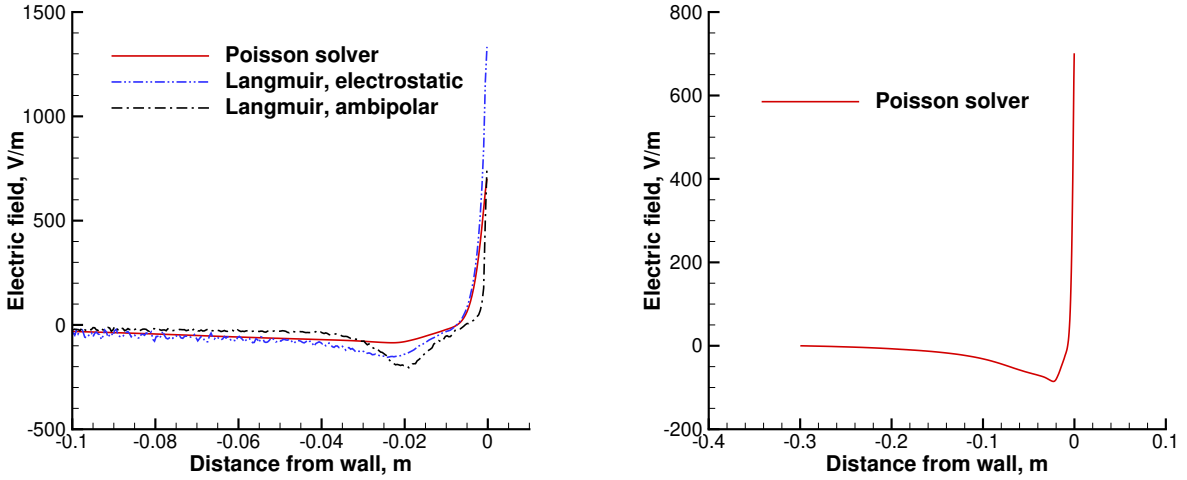

(a) Electric field in region of statistical significance.

(b) Electric field, zoomed in region of peak values.

Figure 5.12: Comparison between electric field profiles from the Poisson solver and the Langmuir and Tonks relation (Mach 39 flow, 81 km, $F_r = 10^8$, transitional diffusion).

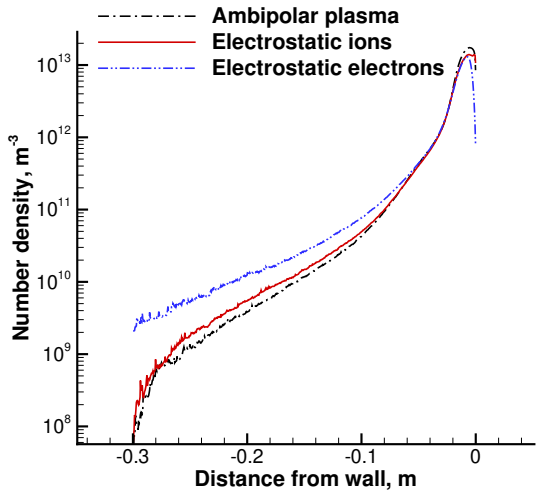

(a) Charged species densities.

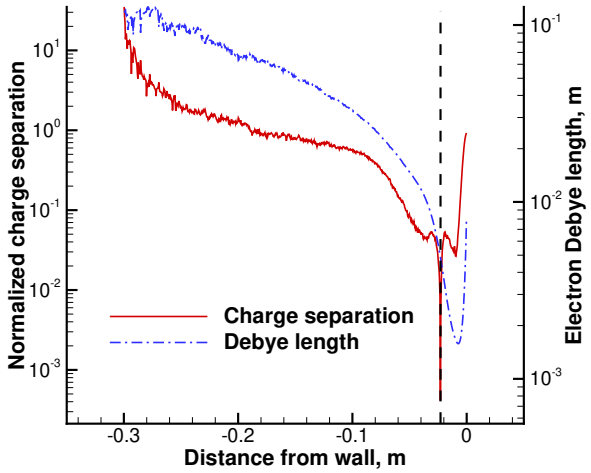

(b) Normalized charge separation from electrostatic simulation (left axis) and electron Debye length (right axis) from the electrostatic solution. The black dashed line separates the domain where the charge separation switches from negative to positive.

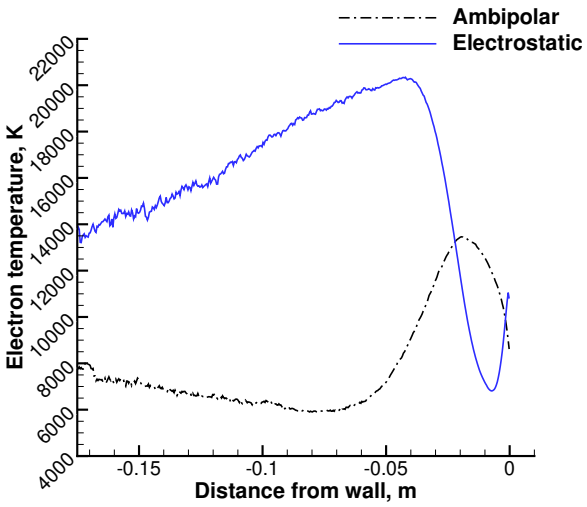

(c) Temperature for ions (left axis) and electrons (right axis).

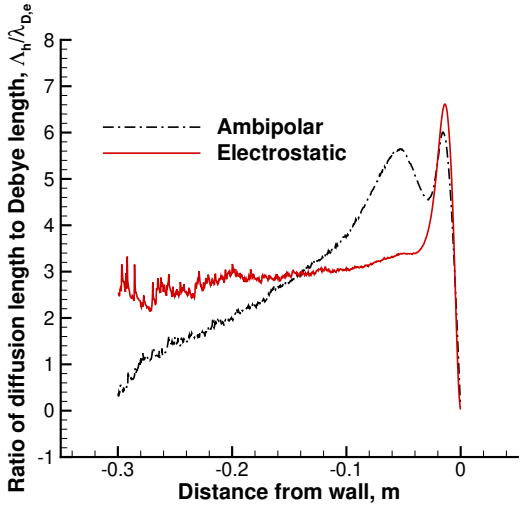

(d) Ratio of the hypersonic diffusion length to the electron Debye length.

Figure 5.13: Flowfields with the ambipolar diffusion approximation and with full electrostatic modeling (Mach 39 flow, 81 km, $F_r = 10^8$, transitional diffusion).

Table 5.8: Plasma properties of interest for simulations of transitional plasma diffusion in hypersonic flows ($F_r = 10^8$).

	60 km		81 km	
Plasma property	Ambipolar	Electrostatic	Ambipolar	Electrostatic
Ions maximum resolved distance, m	-0.0108	-0.0124	-0.175	-0.191
Electrons maximum resolved distance, m	-0.0108	-0.0232	-0.175	-0.238
Maximum $\Lambda_h/\lambda_{D,e}$ in shock layer	3.60	3.68	5.00	6.61
N_D at peak electron density	8.17×10^4	4.25×10^4	4.07×10^5	2.18×10^5
N_D at maximum resolved distance, electrons	1.94×10^6	4.63×10^6	1.16×10^6	2.16×10^6
$\omega_{p,e}\tau_{e-n}$ at peak electron density	2.01	8.86	7.35	17.4
$\omega_{p,e}\tau_{e-n}$ at maximum resolved distance, electrons	1.23	0.405	9.00	99.1

Table 5.9: Stagnation point heat flux by species for simulations of transitional plasma diffusion regime in hypersonic flows. Each heat flux value is multiplied by $F_r = 10^8$. The change is calculated relative to the ambipolar diffusion approximation value.

	60 km			81 km		
Heat flux, W/cm^2	Ambipolar	Electrostatic	Change	Ambipolar	Electrostatic	Change
Ar, convective	5,930	5,480	-7.59%	337	301	-10.7%
Ar^+, convective	690	610	-11.6%	26.0	18.3	-29.6%
Ar^+, chemical	2,150	2,730	$+27.0\%$	102	122	$+19.6\%$
e^-, convective	156	614	$+294\%$	7.23	24.4	$+237\%$
Total heat flux, plasma	3,000	3,950	$+31.7\%$	135	165	$+17.8\%$
Percentage of total, plasma	33.5%	41.9%	$+8.40\%$	28.7%	35.3%	$+6.60\%$
Total heat flux, all species	8,930	9,430	$+5.60\%$	472	466	-1.27%

5.3.3 Ambipolar Diffusion

This work defines the ambipolar diffusion regime for hypersonic flows as having a majority of the post-shock region undergo ambipolar diffusion with charge separation on the order of 1% or lower. Simulation initialization parameters are included in Table 5.10, and quantities relevant to analysis are included in Tables 5.11 and 5.12.

Table 5.10: Simulation parameters for observation of ambipolar diffusion in hypersonic flowfields.

Parameter	60 km	81 km
Domain length, m	0.030	0.200
Uniform grid width Δx, m	1.59×10^{-5}	1.26×10^{-4}
Points per minimum $\lambda_{D,e}$	≥ 2.69	≥ 3.49
Reduced density factor, F_r	10^6	10^7
Global weight	3.87×10^9	1.66×10^8
Plasma weight modifier	0.05	0.10
Time step, s	10^{-10}	10^{-9}

5.3.3.1 Ambipolar Diffusion at 60 km

Figure 5.14 presents flowfield results with the ambipolar diffusion approximation and with full electrostatic modeling for the Mach 35 flow at 60 km when $F_r = 10^6$. Of note, the post-shock reduced plasma density is of $\mathcal{O}(10^{16})$ m^{-3}, which is a realistic post-shock plasma density for certain hypersonic flow conditions [16]. Excluding the plasma sheath, the normalized charge separation is below 1% for the entire post-shock region. The charge separation exceeds 1% at $x = -0.00453$ m, corresponding with $\Lambda_h/\lambda_{D,e} = 11.1$ from the electrostatic solution and $\Lambda_h/\lambda_{D,e} = 25.4$ from the ambipolar diffusion approximation solution. This indicates that ambipolar diffusion is occurring throughout the post-shock region. Nevertheless, electrostatic modeling changes many of the shock layer properties: the peak plasma density decreases by 35.0%, and the overall shape of the neutral density profile becomes more smooth (Figure 5.15). The peak neutral temperature also decreases by 4.79%, as collisions with diffusing electrons decreases the total kinetic energy of the shock layer. Electron cooling is also observed throughout the shock layer, where at the minimum

electron temperature, the ambipolar diffusion approximation overpredicts the electron temperature by 18.9%. However, at $x = -0.000329$ m, the electrostatic electron temperature *exceeds* that from the ambipolar diffusion approximation. This corresponds with the point where the electric field exceeds +1740 V/m; any electrons capable of pushing past this field must have increasingly higher kinetic energy, resulting in the temperature increase. Upstream of the shock, the electrons diffuse away from ions as the plasma enters the free diffusion regime, though this happens more gradually than the free diffusion case. When the normalized charge separation exceeds 0.1, $\Lambda_h/\lambda_{D,e} = 8.48$ from the electrostatic solution and $\Lambda_h/\lambda_{D,e} = 19.3$ from the ambipolar diffusion approximation solution.

Although the post-shock region is undergoing ambipolar diffusion, the Langmuir and Tonks relations from either model do not agree with the Poisson solver solution. Specifically, both overpredict the maximum magnitude of the electric field by at least 51%, and the Langmuir and Tonks relation from the ambipolar diffusion approximation solution shifts the location of the peak value towards the wall by 9.07%.

For species-specific stagnation point heat flux quantities, all electrostatic heat flux changes relative to the ambipolar diffusion approximation demonstrate similar trends to the transitional diffusion case. This indicates that the logic applied towards analyzing the transitional diffusion partially ionized flowfield also applies to the ambipolar diffusion flowfield. Most importantly, the difference in total stagnation point heat flux is below ±5%. This is due to the increased strength of the self-induced electric field, proportional to the overall plasma density. The peak electric field value in the bulk plasma region increases from −626 V/m to −1,750 V/m from the transitional diffusion to the ambipolar diffusion flow case, and the peak sheath electric field increases from +2,950 V/m to +24,700 V/m. The increased strength of the plasma sheath's field repels more electrons, resulting in an increase in electron convective heat flux of only 167% compared to 294% in the transitional diffusion case. Additionally, the plasma sheath is shorter in width due to the decreased Debye length, thus fewer ions are attracted towards the wall. Finally, the change in neutral convective heat flux increases from −7.59% in the transitional diffusion case to −10.9% in

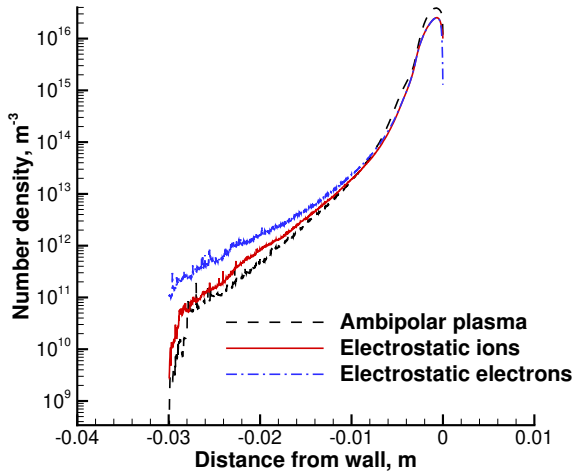

(a) Charged species densities.

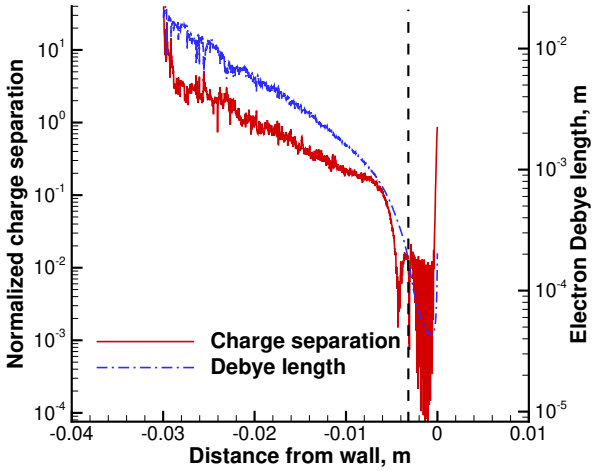

(b) Normalized charge separation from electrostatic simulation (left axis) and electron Debye length (right axis) from the electrostatic solution. The black dashed line separates the domain where the charge separation switches from negative to positive.

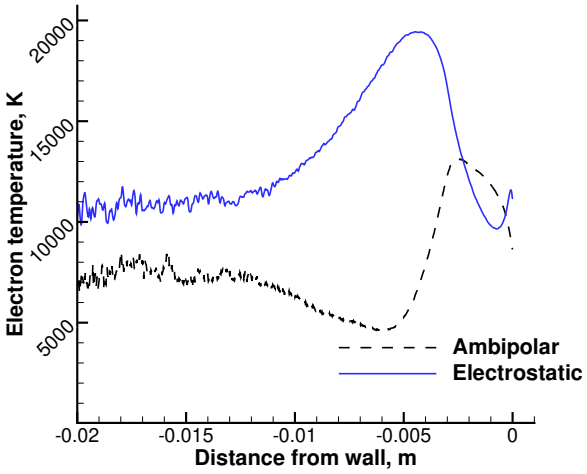

(c) Electron temperature.

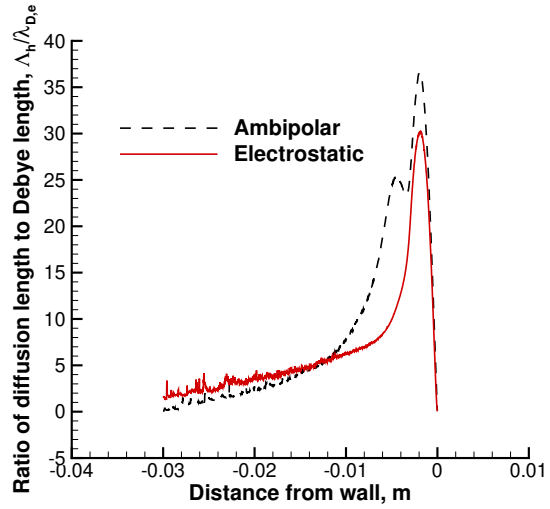

(d) Ratio of the hypersonic diffusion length to the electron Debye length.

Figure 5.14: Flowfields with the ambipolar diffusion approximation and with full electrostatic modeling (Mach 35 flow, 60 km, $F_r = 10^6$, transitional diffusion).

the ambipolar flow case. The artificial reduced density method preserves the species collision rates between flow cases of different F_r. Therefore, if the bulk plasma electric fields are stronger, but the degree of collisionality is the same, then as kinetic energy is transferred to electrons via collisions more of that energy is lost via deceleration to sustain a stronger electric field. Since the stronger electric fields maintain a lower degree of charge separation, the electrons stay in the shock layer longer, resulting in a higher electron temperature.

5.3.3.2 Ambipolar Diffusion at 81 km

Figure 5.18 presents flowfield results with the ambipolar diffusion approximation and with full electrostatic modeling for the Mach 39 flow at 81 km when $F_r = 10^7$. In terms of $\overline{n}_c$ (Figure 5.18b) and $\Lambda_h/\lambda_{D,e}$ (Figure 5.18d), the flowfield meets the criteria to identify ambipolar diffusion. In the normalized charge separation profile at 60 km with $F_r = 10^6$ (Figure 5.10b), the point where the sign switches from negative to positive roughly corresponds with the shock standoff distance, and the charge separation rapidly increases beyond this point. Whereas at 81 km with $F_r = 10^7$, a greater portion of the upstream region has charge separation below 1%. The charge separation exceeds 1% at $x = -0.0687$ m, corresponding with a value of $\Lambda_h/\lambda_{D,e} = 10.4$ for the electrostatic solution and 21.0 from the ambipolar diffusion approximation solution. This indicates that between the 60 km and 81 km flow cases, the value of $\Lambda_h/\lambda_{D,e}$ where the flow enters the ambipolar diffusion regime is fairly consistent at $\Lambda_h/\lambda_{D,e} \approx 10$ for the electrostatic solution and $\Lambda_h/\lambda_{D,e} \approx 20$ for the ambipolar diffusion approximation solution.

The flowfield demonstrates similar trends in charged species number densities, electron temperature, and species-specific stagnation point heat fluxes as the 60 km case with $F_r = 10^6$. The peak plasma density in the post-shock region decreases by 31.1% between the ambipolar diffusion approximation and electrostatic solutions, and the peak neutral temperature decreases by 4.71%. At the minimum post-shock electron temperature for the electrostatic solution, the ambipolar diffusion approximation overpredicts the electron temperature by 49.6%.

As with all flow cases in this study, the Poisson solver electric field does not demonstrate

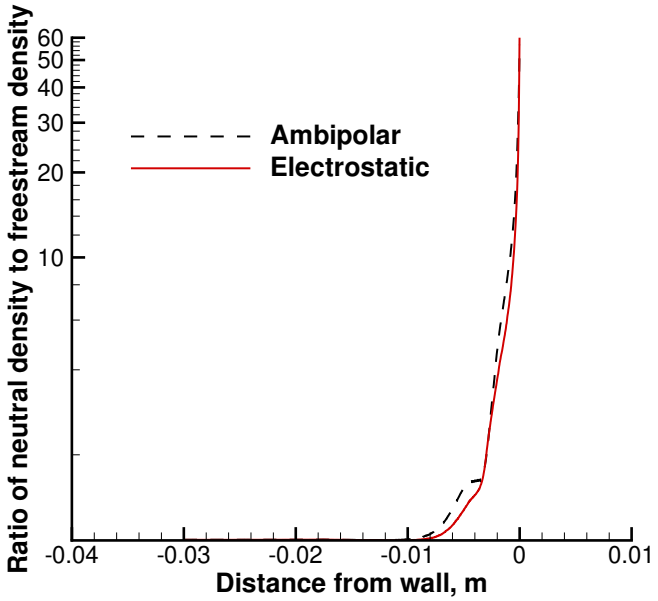

Figure 5.15: Ratio of neutral density with freestream density.

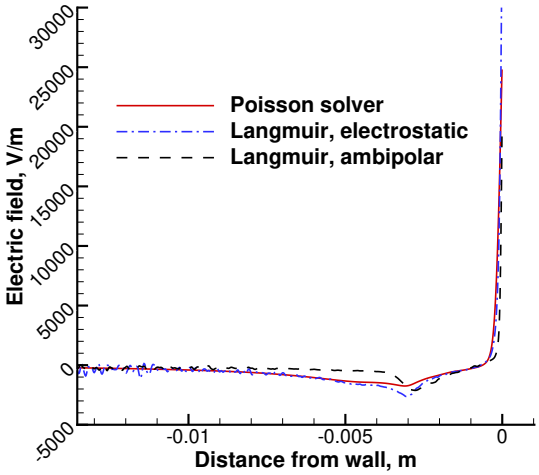

(a) Electric field.

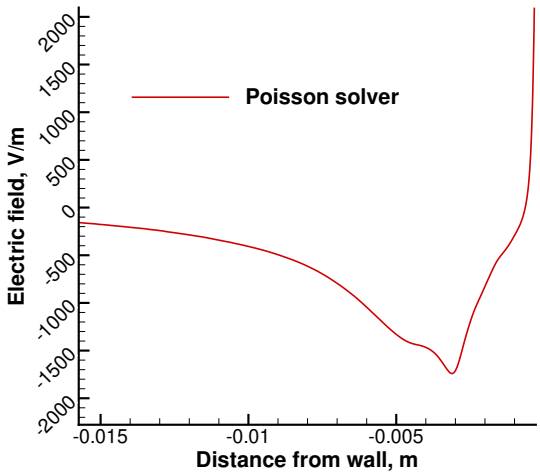

(b) Electric field from only the Poisson solver, zoomed in on bulk plasma region.

Figure 5.16: Comparison between electric field profiles from the Poisson solver and the Langmuir and Tonks relation (Mach 35 flow, 60 km, $F_r = 10^6$, ambipolar diffusion).

agreement with the Langmuir and Tonks relation (Figure 5.17). The peak bulk plasma electric field value is −115 V/m from the Poisson solver, −164 V/m from the electrostatic Langmuir and Tonks relation (shifted 7.14% closer to the wall), and −205 V/m from the ambipolar diffusion approximation Langmuir and Tonks relation (shifted 17.4% closer to the wall). From the results of Section 3.3, strong agreement is expected between the Poisson solver and the Langmuir and Tonks relation if ambipolar diffusion is occurring. Since ambipolar diffusion theory begins with continuum fluid equations, it is possible that 'true' ambipolar diffusion is not occurring at all in these transitional flows. Hence, this work defines hypersonic ambipolar diffusion in terms of charge separation rather than conventional plasma physics theory.

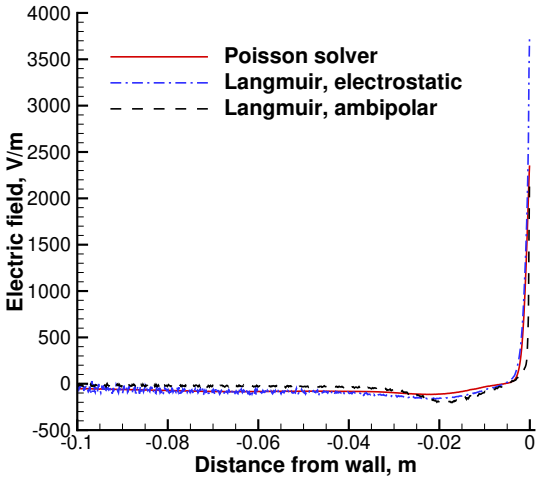

(a) Electric field.

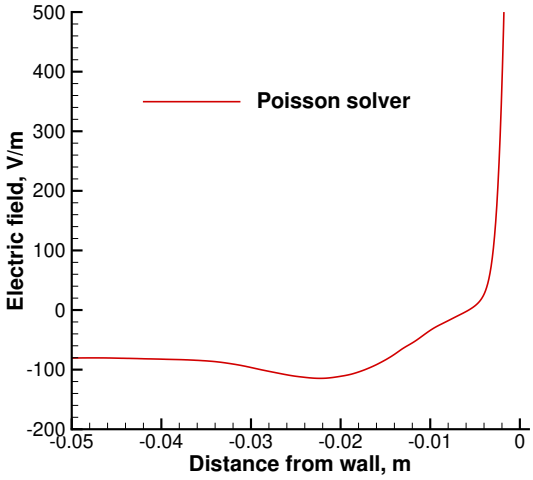

(b) Electric field from only the Poisson solver, zoomed in on bulk plasma region.

Figure 5.17: Comparison between electric field profiles from the Poisson solver and the Langmuir and Tonks relation (Mach 39 flow, 81 km, $F_r = 10^7$, ambipolar diffusion).

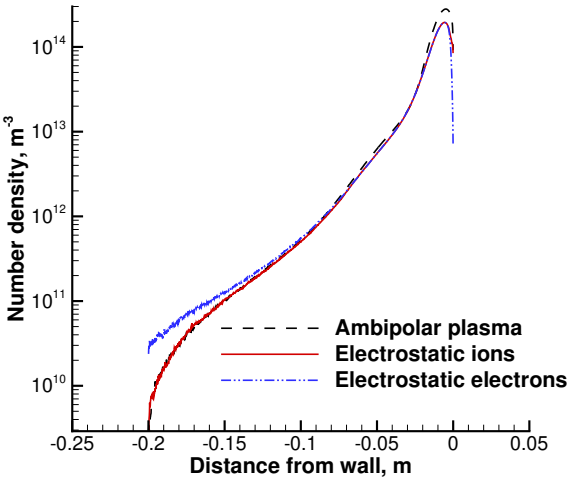

(a) Charged species densities.

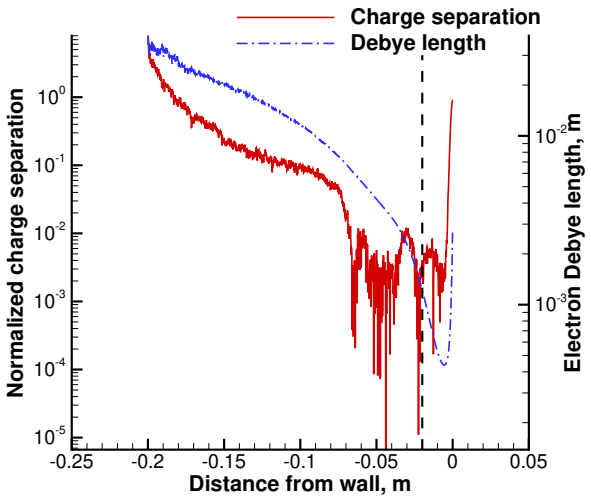

(b) Normalized charge separation from electrostatic simulation (left axis) and electron Debye length (right axis) from the electrostatic solution. The black dashed line separates the domain where the charge separation switches from negative to positive.

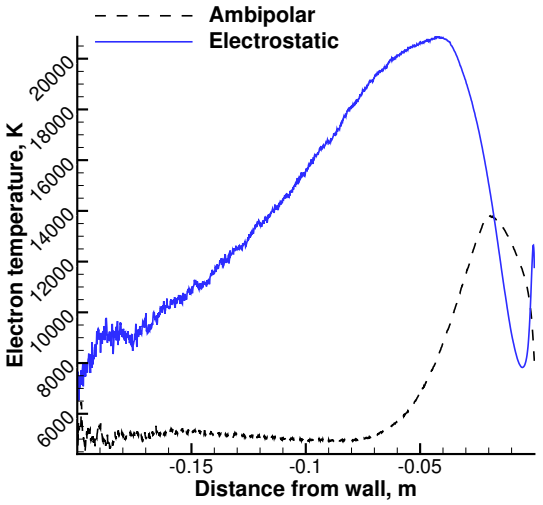

(c) Temperature for ions (left axis) and electrons (right axis).

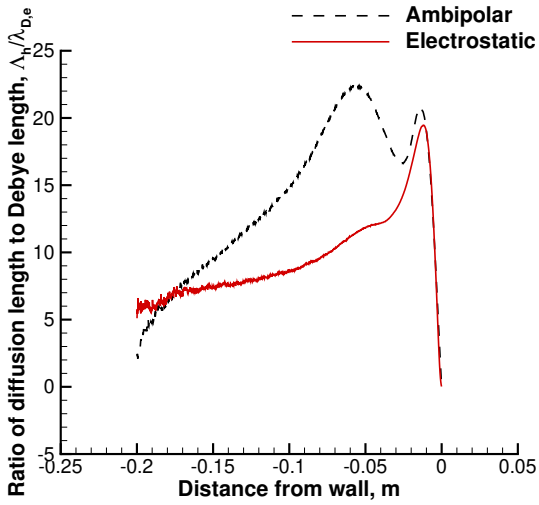

(d) Ratio of the hypersonic diffusion length to the electron Debye length.

Figure 5.18: Flowfields with the ambipolar diffusion approximation and with full electrostatic modeling (Mach 39 flow, 81 km, $F_r = 10^7$, ambipolar diffusion).

Table 5.11: Plasma properties of interest for simulations of ambipolar plasma diffusion in hypersonic flows.

	60 km, $F_r = 10^6$		81 km, $F_r = 10^7$	
Plasma property	Ambipolar	Electrostatic	Ambipolar	Electrostatic
Ions maximum resolved distance, m	-0.0108	-0.0112	-0.101	-0.106
Electrons maximum resolved distance, m	-0.0108	-0.0112	-0.101	-0.137
Maximum $\Lambda_h/\lambda_{D,e}$ in shock layer	36.5	30.1	20.2	19.5
N_D at peak electron density	8.35×10^3	8.36×10^3	1.03×10^5	6.86×10^4
N_D at maximum resolved distance, electrons	2.19×10^5	4.13×10^5	6.79×10^5	4.41×10^6
$\omega_{p,e}\tau_{e-n}$ at peak electron density	18.9	28.0	33.1	81.4
$\omega_{p,e}\tau_{e-n}$ at maximum resolved distance, electrons	11.4	4.89	7.04	7.95

Table 5.12: Stagnation point heat flux by species for simulations of transitional plasma diffusion regime in hypersonic flows. Each heat flux value is multiplied by F_r. The change is calculated relative to the ambipolar diffusion approximation value.

	60 km, $F_r = 10^6$			81 km, $F_r = 10^7$		
Heat flux, W/cm^2	Ambipolar	Electrostatic	Change	Ambipolar	Electrostatic	Change
Ar, convective	5,690	5,070	-10.9%	290	273	-5.86%
Ar^+, convective	650	447	-31.2%	24.3	16.6	-31.7%
Ar^+, chemical	2,090	2,310	$+10.5\%$	127	145	$+14.2\%$
e^-, convective	149	398	$+167\%$	8.40	26.2	$+212\%$
Total heat flux, plasma	2,890	3,160	$+9.34\%$	160	188	$+17.5\%$
Percentage of total, plasma	33.7%	38.4%	$+4.70\%$	35.5%	40.8%	$+5.30\%$
Total heat flux, all species	8,580	8,230	-4.08%	450	461	$+2.44\%$

5.4 General Characteristics of a One-Dimensional Hypersonic Plasma

Despite the vast differences in plasma properties and diffusion regimes, several trends exist across all flow cases. An electric field, negative in sign, always forms within the shock layer. It reaches its peak value between the center of the shock and the shock standoff distance, and increases slowly as it approaches either the inflow boundary or the wall. Within a sufficient distance of the wall, the electric field switches sign as it couples to the fields of the plasma sheath adjacent to the solid boundary surface. The electric field profile does not agree with the Langmuir and Tonks relation (Equation 3.5), and in many cases, the Langmuir and Tonks relation overpredicts the magnitude of the self-induced electric field by at least 50%. This is due to the highly non-continuum nature of a hypersonic shock layer: where even if a portion of the shock layer is in the continuum regime according to Kn_{GLL}, this is not sufficiently continuum for the Langmuir and Tonks relation to apply. The Langmuir and Tonks relation should especially not be used for hypersonic plasma sheath analysis, as it significantly overpredicts the electric fields within the sheath. Thus, approaches that accelerate charged particles according to calculated profiles of the Langmuir and Tonks relation are recommended only for high-level assessment of how electrostatic modeling may impact a hypersonic flowfield.

The first-order hypersonic plasma effects observed in this study include the following:

(1) 'Electron cooling', first identified in Reference [44], where electrons lose kinetic energy via deceleration by the self-induced electric fields. In an attempt to maintain a quasineutral shock layer, electrons are repelled either by positive plasma sheath electric fields (to prevent electrons from absorbing at the wall), or by the negative bulk plasma electric fields (to prevent electrons from diffusing upstream and out of the shock layer). This work demonstrates that electron cooling can occur for any charged species diffusion regime and in kinetic modeling of rarefied hypersonic flows. That is, electron cooling is not unique to continuum flows or high plasma densities.

(2) 'Energy loss', where energy transfer via collisions with independently moving charged

species (particularly mobile electrons) that diffuse out of the shock layer (either upstream or to the wall) results in a slightly lower neutral temeprature throughout the shock layer and a decreased neutral convective heat flux at the wall. Depending on the plasma conditions, this may result in a net decrease in total convective heat flux at the wall.

(3) 'Ion acceleration', where ions are accelerated out of the shock layer due to the negative bulk plasma electric fields or towards the wall when the electric field switches sign in order to couple to the plasma sheath's fields. If ions are allowed to recombine at the wall, the ions will see an overall increase in heat flux. By explicitly separating the two, this study identifies that for flows in the transitional or ambipolar diffusion regimes, the ion convective heat flux will decrease, but the chemical heat flux will increase by a greater amount. This indicates that more ions are recombining at the wall with smaller kinetic energies. The smaller kinetic energy stems from the fact that the more time an ion spends in the shock layer, the more kinetic energy it loses to the negative bulk electric fields. Eventually, the ions reach the positive side of the electric field, which accelerates them towards the wall. The distance between this point and the wall is generally short– thus ions do not gain significant kinetic energy in this region and the reported convective heat flux is lower.

(4) 'Collisional plasma sheath dynamics', where if $\lambda_{D,e} > \lambda_{mfp,i}$ or $\lambda_{mfp,e}$, the plasma sheath with positive electric fields may occupy a sizable portion of the post-shock region. This can result in electrons within that region having a higher temperature, since only highly energetic electrons can resist the sheath electric fields. A large collisional plasma sheath may also facilitate increased ion acceleration towards the wall, especially in the free diffusion regime.

These first-order effects provide a baseline framework to analyze more complex, chemically reacting flows. For example, in 11-species air modeling, the first-order effect of acceleration of N^+ ions may be lessened through the second-order effect of collisions between N^+ and N_2 and N^+ and O, resulting in larger contributions towards total heat flux from N_2 and O than N^+ when

electrostatic modeling is used [38]. Or, flowfields which are sensitive to reactions involving charged species are likely to experience more second-order effects when electrostatic modeling is used.

The impact of first or second-order plasma effects on vehicle surface heat flux, especially species-specific quantities, is highly relevant when engineering new thermal protection materials. For example, if a material is intentionally designed to reduce heat flux contributions from ions, capturing plasma effects via electrostatic modeling may be important.

For the Mach 35 flow case at 60 km with $F_r = 10^6$, using $n_e = 10^{16}$ m^{-3} and $T_e = 10{,}000$ K, an estimate for the electron Coulomb collision frequency (Equation 2.46) is 6.68 $\times 10^4$ s^{-1}. The plasma frequency for those parameters is 2.09 $\times 10^8$ s^{-1}, and with $n_\infty \times F_r = 4.87 \times 10^{21}$ m^{-3}, ν_{e-n} = 1.37 $\times 10^6$ s^{-1}. For the Mach 39 flow case at 81 km with $F_r = 10^8$, using $n_e = 10^{13}$ m^{-3} and T_e = 6,900 K, an estimate for the electron Coulomb collision frequency is 1.16 $\times 10^2$ s^{-1}, $\omega_{p,e} = 6.61 \times 10^6$ s^{-1}, and with $n_\infty \times F_r = 2.64 \times 10^{20}$ m^{-3}, $\nu_{e-n} = 3.63 \times 10^4$ s^{-1}. Thus, electron collisions with neutrals dominate over Coulomb collisions for the strongly ionized flow cases observed in this study, especially within the shock layer where $n_n > n_\infty$. Similarly negligible rates are calculated for ion Coulomb collision frequencies. Thus, neglection of Coulomb collisions is not expected to strongly impact the conclusions made in this study, as in Chapter 3. It follows that the more 'weakly' ionized a flow is, the smaller the impact of Coulomb collisions. In general, it is unlikely that large scattering angle Coulomb collisions would occur in a hypersonic flowfield because the characteristic mean free path of collisions with neutrals is shorter than the characteristic Coulomb mean free path (see Figure 3.2). If Coulomb collisions are non-negligible in a hypersonic flow, then they likely will distribute energy more evenly between all flow species, as energy transfer can occur between charged species, not just between ions or electrons and neutrals. However, to isolate the impact of Coulomb collisions, an extremely high number of macroparticles would be required to suppress the rate of numerical thermalization below the rate of thermalization via Coulomb collisions.

5.5 Guidelines for Characterization of Plasma Diffusion in Hypersonic Flows

If the diffusion length is taken to be the shortest distance between a point in the flowfield and the vehicle surface (Equation 5.1), then the Phelps criteria for charged species diffusion can successfully be applied to characterize plasma diffusion in 1D transitional hypersonic flows. The limits identified in Section 3.4 translate well, where $\Lambda_h/\lambda_{D,e} < 1.0$ is a strong predictor of the free diffusion regime with charge separation on the order of 10% or more, and $\Lambda_h/\lambda_{D,e} > 100$ is a strong predictor of the ambipolar diffusion regime with charge separation of 1% or less. The transitional plasma diffusion regime is defined as having charge separation between 1% and ∼10%. An appropriate lower bound of $\Lambda_h/\lambda_{D,e}$ is 1.0, but the upper limit will vary based on the flow case. In Section 3.3.2.1, the value of $\Lambda/\lambda_{D,e}$ for which a flow enters the ambipolar diffusion regime is found to vary between 40 and 80 depending on the types of collisions modeled and the value of $\Lambda/\lambda_{mfp,i}$ throughout the flow. For the 60 km and 81 km cases of this study, $\Lambda_h/\lambda_{D,e} \gtrsim 10$ and $\Lambda_h/\lambda_{D,e} \gtrsim 20$ mark the transition into the ambipolar diffusion regime for the electrostatic and ambipolar diffusion approximation models, respectively. Since the limits are consistent across flow cases using the same numerical model, the artificial reduced density method may be useful in efficiently identifying the transitional regime upper limit for a given model. That is, one may run an artificially reduced density flow case, determine a value for $\Lambda_h/\lambda_{D,e}$ where $\overline{n}_c < 0.01$, then apply that limit to a full density flowfield to accurately characterize plasma diffusion regime.

Incidentally, the limits of Section 3.4 are identified with continuum plasma simulations that model the true electron mass of $m_e = 1.37 \times 10^{-5} \cdot m_i$, whereas the simulations here use $m_e = 10^{-2} \cdot m_i$. This highlights two key conclusions: the first is that the increased electron mass does not prevent the full range of plasma diffusion regimes from being observed; the second is that the hypersonic diffusion length definition in Equation 5.1 serves as an appropriate length scale for characterizing diffusion regime and assessing quasineutrality. However, for improved prediction of plasma diffusion regime, all three criteria for the definition of a plasma should be considered. The plasma parameter N_D is not particularly illustrative for the flow cases of this study, as $N_D \gg 1$

holds for all flow cases throughout this study. Although the formal definition of a plasma only requires $\omega\tau > 1$, this work finds that $\omega_{p,e}\tau_{e-n} < 10$ is a strong predictor of free diffusion and charge separation over 10% in transitional hypersonic flows.

All flow cases simulated in this study are 'strongly' ionized in the post-shock region, despite the large variation in plasma densities and diffusion regimes, due to the artificially reduced freestream densities. This illustrates that ionization fraction by itself is not useful in predicting charged species diffusion regime unless the total density of the flowfield is known. Ionization fraction is illustrative in predicting how much quantities like stagnation point heat flux may be impacted by electrostatic modeling, as a larger ionization fraction means a greater portion of the flowfield is impacted by plasma effects.

The hypothetical hypersonic flow conditions for which the impacts of self-induced electric fields can be neglected entirely requires low charged species densities, such that the magnitude of the self-induced electric fields are even smaller than those reported in Section 5.3.1, and a highly rarefied freestream, such that electrons cannot quickly diffuse away from ions and energy transfer via collisions is minimized. Thus, the partially ionized hypersonic flow must be in the free fall regime. This could occur for re-entry flow conditions in vacuum or near-vacuum conditions, such as Earth's exosphere. However, these conditions are generally not relevant for re-entry spacecraft vehicle design and thus are not explored in this study.

5.6 Guidelines for use of the Ambipolar Diffusion Approximation in Hypersonic Flows

When the ambipolar diffusion approximation is used, this work recommends always resolving separate electron velocity components in order to better resolve electron temperature and use realistic electron kinetic energies in collisions and reactions. Additionally, the electron collision frequency should always be resolved. If it is employed in this manner, then the ambipolar diffusion approximation is highly useful for predicting plasma diffusion regime in a hypersonic flowfield. Post-shock values of $\Lambda_h/\lambda_{D,e}$, N_D, and $\omega_{p,e}\tau_{e-n}$ computed from simulations invoking the ambipolar

diffusion approximation agree with the values computed from electrostatic simulations within an order of magnitude. Throughout all flow cases, the ambipolar diffusion approximation profile of $\Lambda_h/\lambda_{D,e}$ exhibits two peaks, whereas the electrostatic profile only exhibits one. Therefore, if one is estimating the electrostatic profile of $\Lambda_h/\lambda_{D,e}$ from an ambipolar diffusion approximation solution, the second peak should be disregarded. The slope of the left side of the peak adjacent to the wall is fairly consistent between the ambipolar diffusion approximation and electrostatic solutions (see Figures 5.13d and 5.14d as examples). Therefore, if the left side of the peak's profile is extrapolated to $\Lambda_h/\lambda_{D,e} \approx 5$, the ambipolar diffusion approximation profile of $\Lambda_h/\lambda_{D,e}$ may more accurately predict the transition into the free diffusion regime.

For hypersonic plasmas undergoing ambipolar diffusion, the ambipolar plasma density distribution generally agrees with the electrostatic distributions for either ions or electrons within 35%. The ambipolar diffusion approximation fails to accurately predict electron temperature due to the electron cooling effect, which is present in every flowfield of this study regardless of freestream conditions or charged species diffusion regime. Prediction of electron temperature is most relevant for super-orbital re-entry flow conditions in an Earth atmosphere, where EII is the dominant source of electrons [19, 20]. Finally, for a monatomic flow with EII, AAI, collisions between neutrals, collisions between ions and neutrals, and collisions between electrons and neutrals, the ambipolar diffusion approximation's total stagnation point heat flux agrees with the electrostatic solution within $\pm 5\%$. This indicates that when differences in total surface heat flux on the order of 14–30% are reported in the literature for chemically reacting air [38, 41], it is primarily due to second-order effects rather than the first-order effects. This work is the first to report potential *decreases* in total vehicle surface heat flux when electrostatic modeling is used due to the role of electrons in transporting kinetic energy out of the shock layer via diffusion upstream. However, this effect is likely very sensitive to flow conditions and included chemistry.

5.7 Summary

The physics of charged species diffusion in transitional flow regime hypersonic flows was explored through simulation of a partially ionized one-dimensional stagnation streamline. A characteristic diffusion length scale for hypersonic flows was identified and used to define the free diffusion, transitional diffusion, and ambipolar diffusion charged species regimes for hypersonic flows in conjunction with the formal definition of a plasma. Improved guidelines on using the ambipolar diffusion approximation in hypersonic flowfield modeling were proposed which synthesize fundamental plasma theory, the hypersonic plasma behavior observed in this work, and the practical engineering approach required to apply flowfield modeling results towards hypersonic vehicle design.

Chapter 6

Conclusion

In the final chapter, Section 6.1 reviews each chapter of this dissertation and highlights notable insights for the field. A list of contributions to the state-of-the-art is provided and discussed in the context of rarefied gas dynamics and hypersonics research in Section 6.2. Finally, avenues for future work are discussed in Section 6.3.

6.1 Summary of Dissertation

Chapter 1 motivated the investigation of plasma diffusion in transitional hypersonic flows in terms of advancing fundamental understanding of hypersonic gas dynamics and enabling development of advanced aerospace technologies. Definitions for a plasma, the structure of a partially ionized shock layer, and the ambipolar diffusion approximation were provided. A review was provided of the current literature on electrostatic modeling in hypersonic flow simulation was presented.

Chapter 2 described the governing equations of kinetic simulation of hypersonic flows, the Direct Simulation Monte Carlo (DSMC) method, the discrete-velocity method (DVM), and the two kinetic solvers used throughout this work. Verification of the new physical models and code enhancements added to the solvers to accomplish the objectives of this work was presented.

Chapter 3 began with a review of the various charged species diffusion regimes in partially ionized flows as well as the formal derivation of ambipolar diffusion. The Phelps criteria for identifying charged species diffusion regimes in gas discharge vessels was discussed. The generalizability of Phelps's criteria was investigated through the simulation of a weakly ionized expanding gas flow

with a DVM solver. The conventional criterion for evaluating the validity of the ambipolar diffusion approximation was disproved, and it was shown that the ambipolar diffusion approximation does not accurately capture plasma density distributions or bulk velocities in continuum, non-hypersonic, unsteady flows.

Chapter 4 evaluated the suitability of DSMC and DVM for achieving the final objectives of this dissertation work. A one-dimensional stagnation streamline model for DVM was derived, verified, and compared with DSMC. An analysis of the velocity distribution functions along a stagnation streamline was performed, revealing features unique to a shock in the presence of a reflecting solid surface boundary. Several advancements were made in the DVM method itself and understanding its broader applicability to rarefied gas dynamics and hypersonic flow simulation research.

Chapter 5 investigated the impact of electrostatic modeling in partially ionized hypersonic flows. Appropriate length scales for applying the Phelps criteria in hypersonic flow simulation were identified. A series of hypersonic flow simulations was performed with electrostatic modeling and with the ambipolar diffusion approximation for a broad range of freestream and plasma diffusion conditions. Several first-order plasma effects with non-negligible impact on flowfield properties were identified. The efficacy of the ambipolar diffusion approximation was quantified in terms of impact on stagnation point heat flux, charged species density distributions, and electron temperature. It was demonstrated that the Phelps criteria for charged species diffusion can successfully be applied to hypersonic flow simulation with the proposed length scales.

6.2 Research Contributions

A list of the original contributions of this dissertation to the state of the art in computational partially ionized plasma modeling and hypersonic flowfield prediction is given here. These contributions are contained in References [106, 134, 140, 141].

(1) Phelps' criteria for identification of charged species diffusion regimes were derived in the

context of gas discharge vessels: a closed system with fully absorbing walls (ensuring zero plasma density at the boundaries) and a continuum, steady state plasma. This work expanded on Phelps' original criteria and applied them to two vastly different flowfield configurations:

(a) An unsteady, weakly ionized expanding gas flow with nonzero plasma densities at the boundaries. Many low temperature plasma applications involve unsteady weakly ionized gases, including electric propulsion and plume-surface interactions. Thus, insights made with this simple setup are applicable to other areas of research.

(b) A one-dimensional, monatomic hypersonic flow with only one fully absorbing wall boundary. This setup allowed for investigation of hypersonic charged species diffusion dynamics independent of perpendicular transport and extensive chemistry, improving the generalizability of this work's results.

The criteria proposed in this dissertation enable prediction of charged species diffusion regime and expected degree of charge separation prior to electrostatic modeling. That is, the criteria can be reasonably evaluated from a simulation with the ambipolar diffusion approximation and then used to make an informed decision on whether full electrostatic modeling is required to accurately capture flowfield quantities of interest.

(2) A number of 'first-order' plasma effects were identified in hypersonic plasmas across all diffusion regimes. These effects are only observable with full electrostatic modeling, and include electron cooling, energy loss in the shock layer, ion acceleration, and collisional plasma sheath dynamics. These criteria enable proper identification of the impacts of electrostatic modeling on hypersonic flows versus 'second-order' effects, such as changes in chemistry rates based on changes in plasma density or temperature distributions with electrostatic modeling.

(3) Guidelines for use of the ambipolar diffusion approximation in kinetic simulation of rar-

efied flows were provided. Specifically, the conventional criterion of the electron Debye length being orders of magnitude less than the neutral mean free path was disproved as a robust means of flowfield assessment, in terms both of the validity of the ambipolar diffusion approximation and the prediction of actual ambipolar diffusion. With respect to first-order plasma effects when a hypersonic plasma is undergoing ambipolar diffusion, the consequences of using the ambipolar diffusion approximation in hypersonic flow modeling include incorrect prediction of charged species densities, overprediction of the electron temperature in the shock layer, and inaccurate calculation of vehicle surface heat flux quantities, especially on a species-by-species basis.

(4) The discrete-velocity method has several advantages over DSMC for numerical investigation of rarefied gas flows, particularly for simulation of unsteady flows or capturing instantaneous VDFs of a flowfield. The novel one-dimensional stagnation streamline model developed for this work adds to the tools available to the hypersonics modeling community. Additionally, many clarifications for the original DSMC stagnation streamline model were made, particularly the width of the particle removal region used. Ultimately, if DVM is employed in practical rarefied gas dynamics research, this work strongly recommends use of either the S-BGK or ES-BGK model.

(5) The 'trimodal Maxwellian' VDF shape was observed in both DSMC and DVM simulations of a one-dimensional stagnation streamline and retroactively identified in studies throughout the literature. The presence of backstreaming particles within a shock may impact chemistry rates throughout the flowfield, with subsequent effects on quantities including predicted radiative heat flux.

6.3 Future Work

The insights into fundamental hypersonic plasma behavior and kinetic simulation methods made throughout this dissertation work provide many avenues for future work, including but not

limited to those highlighted in the following sections.

6.3.1 Experimental Validation and Continuum Investigation of First-Order Hypersonic Plasma Effects

This work highlights the importance of studying hypersonic plasmas under a range of plasma diffusion regimes. Experimental investigation would provide clarification on the actual impact of first-order plasma effects on realistic hypersonic flowfields. An ideal experiment would obtain measurements of charge separation at various locations in the facility (especially locations where free or transitional diffusion is expected), electron and neutral species temperatures, and total heat flux to a solid surface or target. Such an experiment could use any gas composition where ionized species are present, though the use of monatomic gases would facilitate comparison with computational results.

The presence of first-order plasma effects and the applicability of the expanded Phelps criteria in continuum flows warrant investigation. The value of $\Lambda_h/\lambda_{D,e}$ for which a hypersonic plasma enters the ambipolar diffusion regime increased between the 81 km and the 60 km flow cases due to the increased collisionality in the freestream. It follows that this limit would increase further as one enters the continuum regime of hypersonic flows. As most spacecraft experience peak heating and pressure in the continuum portion of a re-entry trajectory, quantification of the tradeoffs in using the ambipolar diffusion approximation in CFD will be highly relevant for practical vehicle design. If such a study is considered, this work highly recommends use of a two-fluid plasma model, which would allow for observation of charge separation between ions and electrons, enabling proper assessment of the free and transitional plasma diffusion regimes.

In multidimensional setups, a new definition of the local hypersonic diffusion length may be needed. For example, Equation 5.1 may need to be normalized to a property of the vehicle in axisymmetric flows. However, once the appropriate local diffusion length is defined, then the limits identified in this work should apply.

6.3.2 Sensitivity of Air Chemistry to Electrostatic Modeling

The increase in uncertainty quantification work in hypersonics modeling literature emphasizes the inherent sensitivity of quantities of interest to modeled chemistry rates, notably as relates to an air atmosphere. The first-order plasma effects described in this work present an important source of structural uncertainty, defined as how well a model (i.e., a simulation with the ambipolar diffusion approximation) accurately captures the true system (i.e., a hypersonic flowfield with non-negligible electric fields). However, the number of simulations required to produce a sufficient sample size for an uncertainty quantification study is large, and may be considered unfeasible with kinetic simulation methods. A hypothetical future study could run a single, highly-resolved electrostatic DSMC-PIC simulation for a freestream condition of interest in the continuum or near-continuum regimes (i.e., a flow condition where DSMC and CFD agree). Then, the electric fields could be exported and imposed onto a CFD simulation (either as a full scale flowfield simulation or a smaller scale flowfield, such as a standing shock). This would not equate to a full electrostatic CFD simulation, but it would provide a close estimate of the electric fields generated. A more efficient sensitivity analysis could then be conducted which would characterize the impact of electric fields on hypersonic non-equilibrium chemistry rates.

6.3.3 Hypersonic Plasma Sheath Simulation

This dissertation work concludes that, like most hypersonic gas dynamics phenomena, all aspects of a hypersonic plasma are highly coupled to one another. Additionally, collisional plasma sheath dynamics are most relevant for weakly ionized flows in the free or transitional diffusion regime, as the Debye length and therefore plasma sheath may be large relative to the shock layer. As a result, if analyzing hypersonic plasma sheaths for purposes such as ETC or other space-charge limited applications, the entire flowfield may need to be resolved for accurate results as the negative bulk plasma electric fields are highly coupled to the positive plasma sheath electric fields. This strong coupling plays an important role in electron cooling, charged species diffusion out of the

shock layer, and ion acceleration towards the vehicle surface when the plasma sheath is large.

6.3.4 Improvements to DVM for Efficient Rarefied Gas Dynamics Simulation

Approaches including adaptive mesh refinement for velocity space grids [142], use of GPU architectures, and combination of species into one VDF may alleviate many computational costs commonly associated with DVM. One approach, deemed out-of-scope for this work, was the Parallel-Kinetic-Perpendicular-Moment (PKPM) method [63, 132, 133]. In absence of anisotropic effects, a one-dimensional DVM simulation in a Cartesian coordinate system will have identical VDFs in the perpendicular directions. In essence, the PKPM model evolves the axial and perpendicular VDFs of a flow independently. The axial VDF is evolved directly with the Boltzmann equation as though it were a 1D1V simulation, allowing for full retention of non-continuum behavior. The VDF in the perpendicular direction is given a spectral representation. A single spectral expansion term models the perpendicular VDFs as Maxwellian. When used to simulate the one-dimensional stagnation streamline configuration of Chapter 4, moderate agreement was achieved between full 1D3V DVM, DSMC, and DVM-PKPM with a single expansion term. Stronger agreement would be achieved if more spectral terms are included, however, this would require an extension of the higher order PKPM model to account for collision operators. At present, the higher order PKPM model derivation only applies to collisionless flows.

Bibliography

[1] E. D. Farbar, Kinetic Simulation of Rarefied and Weakly Ionized Hypersonic Flow Fields. PhD thesis, University of Michigan, 2010.

[2] J. D. Anderson, Hypersonic and High Temperature Gas Dynamics. New York: McGraw-Hill, 1989.

[3] D. Olynick, Y.-K. Chen, and M. E. Tauber, "Aerothermodynamics of the Stardust Sample Return Capsule," Journal of Spacecraft and Rockets, vol. 36, pp. 442–462, May 1999.

[4] D. A. Kontinos and M. J. Wright, "Introduction: Atmospheric Entry of the Stardust Sample Return Capsule," Journal of Spacecraft and Rockets, vol. 47, pp. 865–867, Nov. 2010.

[5] D. A. Frank-Kamenetskii, Plasma: The Fourth State of Matter. London: The Macmillan Press, LTD, 1972.

[6] S. Hu, J. Tan, X. Li, Q. Ye, and Z. Liu, "A review of studies on the infrared remote sensing characteristics of hypersonic cruise vehicles," in AOPC 2020: Optical Sensing and Imaging Technology (D. Liu, X. Luo, Y. Jiang, and J. Lu, eds.), (Beijing, China), p. 254, SPIE, Nov. 2020.

[7] X.-X. Song, S. Zhang, J.-F. Zhang, and Y. Zheng, "Study on the Variation of Gas Radiation Characteristics of Hypersonic Reentry Vehicles," AIAA Journal, vol. 61, pp. 5277–5287, Dec. 2023.

[8] S. Yuan, L. Shi, Y. Zhai, B. Yao, F. Li, and Y. Du, "An unsupervised classification method of flight states for hypersonic targets based on hyperspectral features," Chinese Journal of Aeronautics, vol. 36, pp. 434–446, May 2023.

[9] E. L. Resler and W. R. Sears, "The Prospects for Magneto-Aerodynamics," Journal of the Aerospace Sciences, vol. 25, pp. 235–245, Apr. 1958.

[10] V. Bityurin, V. A. Zeigarnik, and A. L. Kuranov, "On a perspective of MHD technology in aerospace applications," in 27th Plasma Dynamics and Lasers Conference, American Institute of Aeronautics and Astronautics Inc, AIAA, 1996.

[11] M. Kim and I. D. Boyd, "Effectiveness of a magnetohydrodynamics system for mars entry," Journal of Spacecraft and Rockets, vol. 49, no. 6, pp. 1141–1149, 2012.

[12] H. K. Ali, D. N. Nguyen, and R. D. Braun, "A Conceptual Modeling Approach for Performance of Magnetohydrodynamic Drag During Planetary Entry," Journal of Spacecraft and Rockets, vol. 62, pp. 176–185, Jan. 2025.

[13] H. Alkandry, K. Hanquist, and I. D. Boyd, "Conceptual Analysis of Electron Transpiration Cooling for the Leading Edges of Hypersonic Vehicles," in 11th AIAA/ASME Joint Thermophysics and Heat Transfer Conference, (Atlanta, GA), American Institute of Aeronautics and Astronautics, June 2014.

[14] L. Uribarri and E. H. Allen, "Electron Transpiration Cooling for Hot Aerospace Surfaces," in 20th AIAA International Space Planes and Hypersonic Systems and Technologies Conference, (Glasgow, Scotland), American Institute of Aeronautics and Astronautics, July 2015.

[15] D. Vatansever, N. Nuwal, and D. A. Levin, "Numerical investigations of spatiotemporal dynamics of space-charge limited collisional sheaths," Physics of Plasmas, vol. 31, p. 093508, Sept. 2024.

[16] K. Y. Monroe and I. D. Boyd, "Circuit Analysis of an Electron Transpiration Cooling System for Hypersonic Leading Edges," AIAA SCITECH 2025 Forum, 2025.

[17] G. V. Candler, "Rate Effects in Hypersonic Flows," Annual Review of Fluid Mechanics, vol. 51, pp. 379–402, Jan. 2019.

[18] T. B. Petrova, G. M. Petrov, J. R. Maxwell, and J. R. Penano, "A quasi-one-dimensional model for stagnation streamline plasma properties at hypersonic speeds," Physics of Plasmas, vol. 32, no. 073509, 2025.

[19] A. Lemal, C. M. Jacobs, M.-Y. Perrin, C. O. Laux, P. Tran, and E. Raynaud, "Prediction of Nonequilibrium Air Plasma Radiation Behind a Shock Wave," Journal of Thermophysics and Heat Transfer, vol. 30, pp. 197–210, Jan. 2016.

[20] T. K. West, C. O. Johnston, and S. Hosder, "Uncertainty and Sensitivity Analysis of Afterbody Radiative Heating Predictions for Earth Entry," Journal of Thermophysics and Heat Transfer, vol. 31, pp. 294–306, Apr. 2017.

[21] T. B. Petrova, G. M. Petrov, and J. R. Peñano, "Transport properties of high Mach number hypersonic air plasmas," Plasma Sources Science and Technology, vol. 33, p. 115008, Nov. 2024.

[22] G. M. Petrov, T. B. Petrova, and J. R. Peñano, "Time-dependent electron Boltzmann equation for hypersonic plasmas," Physics of Plasmas, vol. 31, July 2024.

[23] A. Lemal, C. M. Jacobs, M.-Y. Perrin, C. O. Laux, P. Tran, and E. Raynaud, "Air Collisional-Radiative Modeling with Heavy-Particle Impact Excitation Processes," Journal of Thermophysics and Heat Transfer, vol. 30, pp. 226–239, Jan. 2016.

[24] T. T. Aiken and I. D. Boyd, "Collisional-radiative modeling of shock-heated nitrogen mixtures," Journal of Applied Physics, vol. 135, Mar. 2024.

[25] T. T. Aiken and I. D. Boyd, "Sensitivity Analysis of Ionization in Two-Temperature Models of Hypersonic Air Flows," Journal of Thermophysics and Heat Transfer, vol. 38, pp. 478–490, Oct. 2024.

[26] N. Carter, T. T. Aiken, and I. D. Boyd, "Sensitivity Analysis and Uncertainty Quantification of Plasma Formation in Nonequilibrium Hypersonic Airflows," in AIAA SCITECH 2025 Forum, (Orlando, FL), American Institute of Aeronautics and Astronautics, Jan. 2025.

[27] F. F. Chen, Introduction To Plasma Physics and Controlled Fusion. Springer International Publishing Switzerland, 2016.

[28] E. Josyula and W. F. Bailey, "Governing Equations for Weakly Ionized Plasma Flowfields of Aerospace Vehicles," Journal of Spacecraft and Rockets, vol. 40, pp. 845–857, Nov. 2003.

[29] G. A. Bird, Molecular Gas Dynamics and the Direct Simulation of Gas Flows. Oxford: Clarendon Press, 1994.

[30] G. A. Bird, "Low Density Aerothermodynamics," in AIAA Paper 85-0994, 1985.

[31] I. D. Boyd, "Monte Carlo Simulation of Nonequilibrium Flow in a Low-Power Hydrogen Arcjet," Physics of Fluids, vol. 9, no. 10, pp. 3086–3095, 1997.

[32] S. Plimpton, S. Moore, A. Borner, A. Stagg, T. Koehler, J. Torcyznski, and M. Gallis, "Direct Simulation Monte Carlo on Petaflop Supercomputers and Beyond," Physics of Fluids, vol. 31, no. 086101, 2019.

[33] M. Gallis, S. Plimpton, and S. Moore, "SPARTA Documentation Web Page," 2024.

[34] G. A. Bird, "Direct Simulation of Typical AOTV Entry Flows," in AIAA Paper 86-1310, (Boston), June 1986.

[35] A. B. Carlson and H. A. Hassan, "Direct simulation of re-entry flows with ionization," Journal of Thermophysics and Heat Transfer, vol. 6, pp. 400–404, July 1992.

[36] M. Gallis, R. Prasad, and J. Harvey, "The effect of plasmas on the aerodynamic performance of vehicles," in 29th AIAA, Plasmadynamics and Lasers Conference, (Albuquerque,NM,U.S.A.), American Institute of Aeronautics and Astronautics, June 1998.

[37] M. E. Jones, D. S. Lemons, R. J. Mason, V. A. Thomas, and D. Winske, "A Grid-Based Coulomb Collision Model for PIC Codes," Journal of Computational Physics, vol. 123, pp. 169–181, Jan. 1996.

[38] E. D. Farbar and I. D. Boyd, "Modeling of the plasma generated in a rarefied hypersonic shock layer," Physics of Fluids, vol. 22, no. 10, 2010.

[39] E. S. Cornette, "Forebody temperatures and calorimeter heating rates measured during project FIRE II reentry at 11.35 kilometers per second," NASA Technical Report TM X-1305, NASA Langley Research Center, 1966.

[40] C. R. Kaplan and E. S. Oran, "Analysis of the ambipolar diffusion approximation in hypersonic ionizing reentry flows," Journal of Thermophysics and Heat Transfer, vol. 36, no. 3, pp. 1–11, 2022.

[41] A. Blanco and E. Josyula, "Numerical Modeling of Hypersonic Weakly Ionized External Flowfields with Poisson's Equation," AIAA Journal, vol. 58, no. 8, pp. 3464–3475, 2020.

[42] G. V. Candler and R. W. MacCormack, "Computation of weakly ionized hypersonic flows in thermochemical nonequilibrium," Journal of Thermophysics and Heat Transfer, vol. 5, pp. 266–273, July 1991.

[43] F. Grasso and G. Capano, "Modeling of ionizing hypersonic flows in nonequilibrium," Journal of Spacecraft and Rockets, vol. 32, pp. 217–224, Mar. 1995.

[44] B. Parent, "Electron heating and cooling in hypersonic flows," Physics of Fluids, vol. 33, p. 046105, Apr. 2021.

[45] B. Parent, P. Thoguluva Rajendran, and A. Omprakas, "Electron losses in hypersonic flows," Physics of Fluids, vol. 34, p. 016110, Jan. 2022.

[46] F. M. Rodriguez Fuentes and B. Parent, "Impact of ion mobility on electron density and temperature in hypersonic flows," Physics of Fluids, vol. 37, p. 013609, Jan. 2025.

[47] Y. Hou, Y. Feng, and X. Zheng, "Study of Two-Fluid Effects in Weakly Ionized Plasma Under Electromagnetic Fields," AIAA Journal, vol. 62, no. 4, pp. 1354–1365, 2024.

[48] B. Conrath, D. Gautier, R. Hanel, G. Lindal, and A. Marten, "The helium abundance of Uranus from Voyager measurements," Journal of Geophysical Research: Space Physics, vol. 92, pp. 15003–15010, Dec. 1987.

[49] K. H. Baines, "Neptune: Atmosphere," in Encyclopedia of Planetary Science, pp. 523–532, Dordrecht: Springer Netherlands, 1997.

[50] T. Bodineau, I. Gallagher, L. Saint-Raymond, and S. Simonella, "On the dynamics of dilute gases," European Mathematical Society Magazine, pp. 13–22, Apr. 2023.

[51] W. G. Vincenti and C. H. Kruger, Introduction to Physical Gas Dynamics. Krieger Publishing Company, 1965.

[52] N. A. Krall and A. W. Trivelpiece, Principles of Plasma Physics. McGraw-Hill, 1973.

[53] I. D. Boyd, G. Chen, and G. V. Candler, "Predicting failure of the continuum fluid equations in transitional hypersonic flows," Physics of Fluids, vol. 7, no. 1, pp. 210–219, 1995.

[54] W. Wagner, "A Convergence Proof for Bird's Direct Simulation Monte Carlo Method for the Boltzmann Equation," Journal of Statistical Physics, vol. 66, no. 3, pp. 1011–1044, 1992.

[55] I. D. Boyd and T. E. Schwartzentruber, Nonequilibrium Gas Dynamics and Molecular Simulation. Cambridge: Cambridge University Press, 2017. Publisher: Cambridge University Press Place: Cambridge.

[56] R. Hockney and J. Eastwood, Computer Simulation Using Particles. London: McGraw-Hill, 1981.

[57] J. M. Dawson, "Particle Simulation of Plasmas," Reviews of Modern Physics, vol. 55, no. 403, 1983.

[58] C. A. Birdsall and A. B. Langdon, Plasma Physics via Computer Simulation. New York: McGraw-Hill, 1985.

[59] J. Villasenor and O. Buneman, "Rigorous charge conservation for local electromagnetic field solvers," Computer Physics Communications, vol. 69, no. 2-3, pp. 306–316, 1992.

[60] I. D. Boyd, "Rotational and vibrational nonequilibrium effects in rarefied hypersonic flow," Journal of Thermophysics and Heat Transfer, vol. 4, no. 4, pp. 478–484, 1990.

[61] C. Cai, Theoretical and Numerical Studies of Plume Flows in Vacuum Chambers. PhD thesis, University of Michigan, Ann Arbor, 2005.

[62] C. Cui and J. Wang, "Grid-based Vlasov simulation of collisionless plasma expansion," Physics of Plasmas, vol. 28, no. 9, 2021.

[63] P. Cagas, A. Hakim, J. Juno, and B. Srinivasan, "Continuum Kinetic and Multi-Fluid Simulations of Classical Sheaths," Physics of Plasmas, vol. 24, Oct. 2017.

[64] C. Z. Cheng and G. Knorr, "The integration of the vlasov equation in configuration space," Journal of Computational Physics, vol. 22, no. 3, pp. 330–351, 1976.

[65] P. L. Bhatnagar, E. P. Gross, and M. Krook, "A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems," Physical Review, vol. 94, pp. 511–525, May 1954.

[66] Y. Zheng and H. Struchtrup, "Ellipsoidal statistical Bhatnagar–Gross–Krook model with velocity-dependent collision frequency," Physics of Fluids, vol. 17, p. 127103, Dec. 2005.

[67] E. Shakhov, "Generalization of the Krook Kinetic Relaxation Equation," Fluid Dynamics, vol. 3, no. 5, pp. 142–145, 1968.

[68] L. H. Holway, Jr., "Kinetic Theory of Shock Structure Using an Ellipsoidal Distribution Function," in Rarefied Gas Dynamics, Volume 1, vol. 1, p. 193, 1965.

[69] K. Hara, Development of Grid-Based Direct Kinetic Method and Hybrid Kinetic-Continuum Modeling of Hall Thruster Discharge Plasmas. PhD thesis, University of Michigan, 2015.

[70] A. L. Raisanen, A Two-dimensional Hybrid-Direct Kinetic Model of a Hall Thruster. PhD thesis, University of Michigan, 2020.

[71] A. R. Vazsonyi, Deterministic-Kinetic Computational Analyses of Expansion Flows and Current-Carrying Plasmas. PhD thesis, University of Michigan, 2021.

[72] K. Hara, T. Chapman, J. W. Banks, S. Brunner, I. Joseph, R. L. Berger, and I. D. Boyd, "Quantitative study of the trapped particle bunching instability in Langmuir waves," Physics of Plasmas, vol. 22, no. 2, 2015.

[73] K. Hara, I. Barth, E. Kaminski, I. Y. Dodin, and N. J. Fisch, "Kinetic simulations of ladder climbing by electron plasma waves," Physical Review E, vol. 95, no. 5, pp. 1–7, 2017.

[74] K. Hara and K. Hanquist, "Test cases for grid-based direct kinetic modeling of plasma flows," Plasma Sources Science and Technology, vol. 27, no. 6, 2018.

[75] A. L. Raisanen, K. Hara, and I. D. Boyd, "Two-dimensional hybrid-direct kinetic simulation of a Hall thruster discharge plasma," Physics of Plasmas, vol. 26, no. 12, 2019.

[76] W. H. R. Chan and I. D. Boyd, "Grid-point requirements for direct kinetic simulation of weakly collisional plasma plume expansion," Journal of Computational Physics, vol. 475, p. 111861, Feb. 2023.

[77] K. Hara, I. D. Boyd, and V. I. Kolobov, "One-dimensional hybrid-direct kinetic simulation of the discharge plasma in a Hall thruster," Physics of Plasmas, vol. 19, no. 11, 2012.

[78] K. Hara, M. J. Sekerak, I. D. Boyd, and A. D. Gallimore, "Mode transition of a Hall thruster discharge plasma," Journal of Applied Physics, vol. 115, no. 20, 2014.

[79] W. H. R. Chan, K. Hara, and I. D. Boyd, "Effects of multi-dimensionality and energy exchange on electrostatic current-driven plasma instabilities and turbulence," Journal of Plasma Physics, vol. 90, no. 2, pp. 1–25, 2024. ISBN: 0022377824.

[80] K. Hara and C. Treece, "Ion kinetics and nonlinear saturation of current-driven instabilities relevant to hollow cathode plasmas," Plasma Sources Science and Technology, vol. 28, no. 5, 2019.

[81] A. R. Vazsonyi, K. Hara, and I. D. Boyd, "Non-monotonic double layers and electron two-stream instabilities resulting from intermittent ion acoustic wave growth," Physics of Plasmas, vol. 27, no. 11, 2020.

[82] P. K. Sweby, "Godunov Methods," Tech. Rep. Numerical Analysis Report 7/99, The University of Reading, New York, 2001.

[83] X. Zeng, "General approach to enhance slope limiters in muscl schemes on nonuniform rectilinear grids," SIAM Journal on Scientific Computing, vol. 38, no. 2, pp. A789–A813, 2016.

[84] R. J. LeVeque, Finite Difference Methods for Ordinary and Partial Differential Equations. Philadelphia: Society for Industrial and Applied Mathematics, 2007.

[85] S. Balay, S. Abhyankar, M. F. Adams, S. Benson, J. Brown, P. Brune, K. Buschelman, E. M. Constantinescu, Dalcin, A. Dener, V. Eijkhout, J. Faibussowitsch, W. D. Gropp, V. Hapla, T. Isaac, P. Jolivet, D. Karpeev, D. Kaushik, M. G. Knepley, F. Kong, S. Kruger, D. A. May, L. C. McInnes, R. T. Mills, L. Munson, T. Munson, J. E. Roman, K. Rupp, P. Sanan, J. Sarich, B. F. Smith, S. Zampini, H. Zhang, H. Zhang, and J. Zhang, "PETSc Web page," 2024.

[86] G. Strang, "On the Construction and Comparison of Difference Schemes," SIAM Journal on Numerical Analysis, vol. 5, pp. 506–517, Sept. 1968.

[87] S. Pieraccini and G. Puppo, "Implicit-explicit schemes for BGK kinetic equations," Journal of Scientific Computing, vol. 32, pp. 1–28, July 2007.

[88] M. Pirner, "A review on BGK models for gas mixtures of mono and polyatomic molecules," Fluids, vol. 6, no. 11, pp. 1–17, 2021.

[89] V. Djordjić, G. Oblapenko, M. Pavić-Čolić, and M. Torrilhon, "Boltzmann collision operator for polyatomic gases in agreement with experimental data and DSMC method," Continuum Mechanics and Thermodynamics, vol. 35, pp. 103–119, Jan. 2023.

[90] S. Sakabe and Y. Izawa, "Simple formula for the cross sections of resonant charge transfer between atoms and their positive ions at low impact velocity," Physical Review A3, vol. 45, no. 3, pp. 2086–2089, 1992.

[91] A. Hakim, "Neutral Models in Gkeyll."

[92] E. T. Meier and U. Shumlak, "A general nonlinear fluid model for reacting plasma-neutral mixtures," Physics of Plasmas, vol. 19, no. 7, 2012.

[93] A. S. Richardson, "NRL Plasma Formulary," tech. rep., Naval Research Laboratory, Washington, D.C., 2019.

[94] C. Wersal and P. Ricci, "A first-principles self-consistent model of plasma turbulence and kinetic neutral dynamics in the tokamak scrape-off layer," Nuclear Fusion, vol. 55, Nov. 2015.

[95] G. A. Sod, "A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws," Journal of Computational Physics, vol. 27, no. 1, pp. 1–31, 1978.

[96] C. P. Lipscomb, Simulation of Hall Thruster Plume Expansion in Vacuum Chamber and Space Environments. PhD thesis, University of Colorado Boulder, Boulder, 2025.

[97] J. Hammel, "Development of an Unstructured 3-D Direct Simulation Monte Carlo/Particle-In-Cell Code and the Simulation of Microthruster Flows," Master's thesis, Worcester Polytechnic Institute, 2002.

[98] "IST-Lisbon database." www.lxcat.net, retrieved on April 23, 2024.

[99] L. L. Alves, "The IST-LISBON database on LXCat," Journal of Physics: Conference Series, vol. 565, no. 1, 2014.

[100] I. D. Boyd, "Conservative species weighting scheme for the direct simulation Monte Carlo method," Journal of Thermophysics and Heat Transfer, vol. 10, pp. 579–585, Oct. 1996.

[101] S. Rjasanow and W. Wagner, Stochastic Numerics for the Boltzmann Equation, vol. 37 of Springer Series in Computational Mathematics. Berlin/Heidelberg: Springer-Verlag, 2005.

[102] A. Y. Hong and M. A. Gallis, "An improved stochastic weighted particle method for boundary driven flows," Journal of Computational Physics, vol. 517, p. 113353, Nov. 2024.

[103] V. Charton, E. Yamaoka, T. Morimoto, and K. Kinefuchi, "Species Weighting Scheme in Direct Simulation Monte Carlo Applied to Jet Plume Simulations," 2025.

[104] S. Kim and E. Jun, "A stochastic Fokker–Planck–Master model for diatomic rarefied gas flows," Journal of Computational Physics, vol. 506, p. 112940, June 2024.

[105] W. Park, S. Kim, M. Pfeiffer, and E. Jun, "Evaluation of stochastic particle Bhatnagar–Gross–Krook methods with a focus on velocity distribution function," Physics of Fluids, vol. 36, p. 027113, Feb. 2024.

[106] M. Petrusky and I. D. Boyd, "Discrete-Velocity Simulation of One-Dimensional Hypersonic Stagnation Streamline," Journal of Thermophysics and Heat Transfer, pp. 1–13, Mar. 2026.

[107] S. Jubin, A. T. Powis, W. Villafana, D. Sydorenko, S. Rauf, A. V. Khrabrov, S. Sarwar, and I. D. Kaganovich, "Numerical thermalization in 2D PIC simulations: Practical estimates for low-temperature plasma simulations," Physics of Plasmas, vol. 31, Feb. 2024.

[108] L. Tonks and I. Langmuir, "A General Theory of the Plasma of an Arc," Physical Review, vol. 34, pp. 876–922, Sept. 1929.

[109] A. V. Phelps, "The Diffusion of Charged Particles in Collisional Plasmas: Free and Ambipolar Diffusion at Low and Moderate Pressures," Journal of Research of the National Institute of Standards and Technology, vol. 95, no. 4, pp. 407–431, 1990.

[110] W. P. Allis, "Motions of Ions and Electrons," tech. rep., Massachusetts Institute of Technology, 1956. Publication Title: Technical Report 299.

[111] E. W. McDaniel, Collision Phenomena in Ionized Gases. New York: John Wiley and Sons Inc., 1964.

[112] K. Y. Monroe, M. P. Georgin, and I. D. Boyd, "Electron transpiration circuits for hypersonic leading edges," Journal of Applied Physics, vol. 138, p. 163304, Oct. 2025.

[113] A. Langdon, "Effects of the spatial grid in simulation plasmas," Journal of Computational Physics, vol. 6, pp. 247–267, Oct. 1970.

[114] L. Jolivet and J. F. Roussel, "Numerical modeling of plasma sheath phenomena in the presence of secondary electron emission," IEEE Transactions on Plasma Science, vol. 30, no. 1, pp. 318–326, 2002.

[115] D. Vatansever, N. Nuwal, and D. A. Levin, "A Parametric Study for Kinetic Modeling of Emissive Sheaths Using Electrostatic Particle-in-Cell Method," in AIAA AVIATION 2022 Forum, (Chicago, IL & Virtual), American Institute of Aeronautics and Astronautics, June 2022.

[116] B. Evans and S. Walton, "Aerodynamic optimisation of a hypersonic reentry vehicle based on solution of the Boltzmann–BGK equation and evolutionary optimisation," Applied Mathematical Modelling, vol. 52, pp. 215–240, Dec. 2017.

[117] P. Clarke, P. Varghese, and D. Goldstein, "A low noise discrete velocity method for the Boltzmann equation with quantized rotational and vibrational energy," Journal of Computational Physics, vol. 352, pp. 326–340, 2018.

[118] D. Toussaint, L. B. Cedex, and H. Noubel, "Influence of Rarefaction Degree and Aft-Body Geometry on Supersonic Flows," AIAA Journal, pp. 1–21, 2024.

[119] A. Berghausen, J. Taylor, and H. Hassan, "Direct Simulation of Shock Front Radiation in Air," Journal of Thermophysics and Heat Transfer, vol. 10, no. 3, 1996.

[120] I. J. Wysong and D. C. Wadsworth, "Effect of DSMC dissociation models on a stagnation streamline flowfield," in 32nd Thermophysics Conference, 1997.

[121] I. D. Boyd, K. A. Trumble, and M. J. Wright, "Modeling of Stardust Entry at High Altitude, Part 1: Flowfield Analysis," Journal of Spacecraft and Rockets, vol. 47, pp. 708–717, Sept. 2010.

[122] M. Pfeiffer, "Particle-based fluid dynamics: Comparison of different Bhatnagar-Gross-Krook models and the direct simulation Monte Carlo method for hypersonic flows," Physics of Fluids, vol. 30, p. 106106, Oct. 2018.

[123] H. M. Mott-Smith, "The solution of the Boltzmann equation for a shock wave," Physical Review, vol. 82, no. 6, pp. 885–892, 1951.

[124] G. A. Bird, "Aspects of the structure of strong shock waves," Physics of Fluids, vol. 13, no. 5, pp. 1172–1177, 1970.

[125] H. Alsmeyer, "Density profiles in argon and nitrogen shock waves measured by the absorption of an electron beam," Journal of Fluid Mechanics, vol. 74, no. 3, pp. 497–513, 1976. ISBN: 0022112076001.

[126] G. Pham-Van-Diep, D. Erwin, and E. P. Muntz, "Nonequilibrium Molecular Motion in a Hypersonic Shock Wave," Science, vol. 245, pp. 624–626, Aug. 1989.

[127] S. Mazouffre, P. Vankan, R. Engeln, and D. C. Schram, "Behavior of the H atom velocity distribution function within the shock wave of a hydrogen plasma jet," Physical Review E, vol. 64, p. 066405, Nov. 2001.

[128] G. A. Bird, "The velocity distribution function within a shock wave," Journal of Fluid Mechanics, vol. 30, no. 3, pp. 479–487, 1967.

[129] Z. Li, T. Ozawa, I. Sohn, and D. A. Levin, "Modeling of electronic excitation and radiation in non-continuum hypersonic reentry flows," Physics of Fluids, vol. 23, p. 066102, June 2011.

[130] Z. Li, I. Sohn, and D. A. Levin, "Effects of Non-Maxwellian Distributions on Shock-Layer Radiation from Hypersonic Reentry Flows," Journal of Thermophysics and Heat Transfer, vol. 27, pp. 183–187, Jan. 2013.

[131] Q.-H. Sun, C.-P. Cai, and W. Gao, "On the validity of the Boltzmann-BGK model through relaxation evaluation," Acta Mechanica Sinica, vol. 30, pp. 133–143, Apr. 2014.

[132] P. Cagas, J. Juno, A. Hakim, A. LaJoie, F. Chu, S. Langendorf, and B. Srinivasan, "An investigation of shock formation versus shock mitigation of colliding plasma jets," Physics of Plasmas, vol. 30, p. 053903, May 2023.

[133] J. Juno, A. Hakim, and J. M. TenBarge, "A parallel-kinetic-perpendicular-moment model for magnetised plasmas," Journal of Plasma Physics, vol. 91, p. E129, Oct. 2025.

[134] M. Petrusky and I. D. Boyd, "A Novel Stagnation Streamline Model for Discrete-Velocity Simulation of Hypersonic Flows," in AIAA AVIATION FORUM AND ASCEND 2025, (Las Vegas, Nevada), American Institute of Aeronautics and Astronautics, July 2025.

[135] S. Rjasanow, T. Schreiber, and W. Wagner, "Reduction of the Number of Particles in the Stochastic Weighted Particle Method for the Boltzmann Equation," Journal of Computational Physics, vol. 145, pp. 382–405, Sept. 1998.

[136] G. Oblapenko, D. Goldstein, P. Varghese, and C. Moore, "Hedging direct simulation Monte Carlo bets via event splitting," Journal of Computational Physics, vol. 466, p. 111390, Oct. 2022.

[137] Y. Ko and E. Jun, "Radiative equilibrium boundary condition and correlation analysis on catalytic surfaces in DSMC," International Journal of Heat and Mass Transfer, vol. 255, p. 127693, Feb. 2026.

[138] E. Kawamura, A. J. Lichtenberg, M. A. Lieberman, and J. P. Verboncoeur, "Double layer formation in a two-region electronegative plasma," Physics of Plasmas, vol. 16, no. 12, 2009.

[139] D. P. Weaver and B. D. Shizgal, eds., Rarefied Gas Dynamics: Theory and Simulations. Washington DC: American Institute of Aeronautics and Astronautics, Jan. 1994.

[140] M. Petrusky and I. D. Boyd, "Evaluation of the Ambipolar Diffusion Approximation Using an Eulerian Boltzmann-Poisson-BGK Solver," in Rarefied Gas Dynamics (M. Grabe, G. Oblapenko, and M. Torrilhon, eds.), pp. 669–677, Cham: Springer Nature Switzerland, 2026. Series Title: Springer Aerospace Technology.

[141] M. Petrusky and I. D. Boyd, "Characterizing Plasma Diffusion in Transitional Hypersonic Flows," In preparation, 2026.

[142] C. Baranger, J. Claudel, N. Hérouard, and L. Mieussens, "Locally refined discrete velocity grids for stationary rarefied flow simulations," Journal of Computational Physics, vol. 257, no. PA, pp. 572–593, 2014. arXiv: 1304.5611.

[143] J. Annaloro, V. Morel, A. Bultel, and P. Omaly, "Global rate coefficients for ionization and recombination of carbon, nitrogen, oxygen, and argon," Physics of Plasmas, vol. 19, July 2012.

[144] L. C. Pitchford, L. L. Alves, K. Bartschat, S. F. Biagi, M. C. Bordage, A. V. Phelps, C. M. Ferreira, G. J. M. Hagelaar, W. L. Morgan, S. Pancheshnyi, V. Puech, A. Stauffer, and O. Zatsarinny, "Comparisons of sets of electron–neutral scattering cross sections and swarm parameters in noble gases: I. Argon," Journal of Physics D: Applied Physics, vol. 46, p. 334001, Aug. 2013.

[145] H. W. Drawin and F. Emard, "Atom-Atom Excitation and Ionization in Shock Waves of the Noble Gases," Physics Letters, vol. 43, no. 4, 1973.

[146] A. J. Kelly, "Atom—Atom Ionization Cross Sections of the Noble Gases—Argon, Krypton, and Xenon," The Journal of Chemical Physics, vol. 45, pp. 1723–1732, Sept. 1966.

Appendix A

Argon Parameters

The following appendix provides the details and parameters for the collision and ionization models used throughout this dissertation work.

Table A.1: Collision cross section models used for each reaction pair.

Collision pair	Model
Ar + Ar	VHS
Ar + Ar^+	CEX (Sakabe and Izawa [90])
Ar + e^-	VHS or fitted cross-section data (IST-Lison Database [98, 99])

Table A.2: Collision pair parameters used in the variable hard-sphere model.

Collision pair	ω	Reference temperature, K
Ar + Ar	0.81	273 [29]
Ar + Ar^+	0.81	273
Ar + e^-	0.70	273 [38]

For argon, the expression for the CEX cross section from Reference [90] is given as follows

$$\sigma_{CEX}(g)\ [\mathrm{cm}^2] = (1.81 \times 10^{-14} - 2.12 \times 10^{-15}\log_{10}(g))(I/I_0)^{-1.5} \tag{A.1}$$

where I is the first ionization potential of argon (15.7596 eV) and I_0 is the ionization potential of hydrogen (13.6 eV).

For neutral-electron collisions, elastic momentum-transfer cross section data is obtained from the IST-Lisbon database via `lxcat.net` [98, 99] and fitted piecewise:

Table A.3: Species parameters used in the variable hard-sphere and hard-sphere models.

Species	d_{ref}^{HS}, m	d_{ref}^{VHS}, m
Ar	3.4×10^{-10}	4.17×10^{-10} [29]
Ar^+	3.4×10^{-10}	4.17×10^{-10}
e^-	1×10^{-10}	1×10^{-10} [38]

$$\sigma_{e-n}(E_e)\ [\mathrm{m}^2] = \begin{cases} exp[-50.309 - 1.816ln(E_e) - 0.131(ln(E_e))^2] & \text{for } E_e \leq 0.25 \\ exp[-45.875 + 1.336ln(E_e) - 0.143(ln(E_e))^2] & \text{for } 0.25 < E_e \leq 15 \\ 1.681 \times 10^{-18} E_e^{-0.995} & \text{for } E_e > 15 \end{cases} \tag{A.2}$$

where the electron energy E_e is in units of eV. Figure A.1 plots the piecewise fit alongside the IST-Lisbon data. Because the data is nonlinear, a single test cannot accurately characterize whether the fit is "good". The fitted parameter uncertainties are reported in Table A.5 and are relatively small in magnitude, indicating the fitted parameters are well constrained. Using the fitted parameter uncertainties, a reduced chi-squared goodness-of-fit test is performed for each of the three regions, yielding statistics of $\chi_\nu^2 = 0.98, 1.12, 0.91$, respectively. In this context, $\chi_\nu^2 \approx 1$ indicates the residual scatter is statistically consistent with the uncertainty predicted by the parameter covariance, and that the model is not over or underfitting the data. However, a proper chi-squared goodness-of-fit test would use the error of the IST-Lisbon data. The study did not report individual uncertainties for each of the data points, instead reporting 20-30% error with experiment [144]. Applying 25% error across the entire dataset yields statistics of $\chi_\nu^2 = 1.84$, 0.19, and 0.059. In this context, $\chi_\nu^2 \gg 1$ indicates a poor model fit or underestimated uncertainties, and $\chi_\nu^2 <$ indicates the model is

Table A.4: Baseline reaction rate coefficients used in the Total Collision Energy chemistry model.

Reaction	Rate coefficient, m^3/s
$Ar + Ar \rightarrow Ar + Ar^+ + e^-$	$8.9169 \times 10^{-23} T^{1.1935} exp(-129450/T)$
$e^- + Ar \rightarrow Ar^+ + 2e^-$	$1.23 \times 10^{-19} T^{1.511} exp(-141480/T)$ [143]

overfitting the data, either from overestimated error variance or improper noise fitting. Therefore for region 1, $\chi^2_\nu = 1.84$ indicates the model has clear but modest discrepancy with the IST-Lisbon data. This is evident from Figure A.2, where the piecewise fit slightly underpredicts the data. However, the large value of χ^2_ν can also be from underprediction of error: if instead the upper bound of 30% is used, the test yields $\chi^2_\nu = 1.28$, which is within acceptable range of $\chi^2_\nu \approx 1$. As for regions 2 and 3, the $\chi^2_\nu < 1$ mostly reflects that the error is overestimated in these regions. Taking all of the tests performed into account, we conclude that the piecewise fit of the cross section data is sufficiently accurate for the purposes of this study.

Table A.5: Best-fit parameters and associated 1σ uncertainties for each region of the piecewise model of the IST-Lisbon elastic momentum-transfer cross section data for electron-neutral collisions.

Region	**Parameter**	**Value**	1σ **uncertainty**
Region 1	a_1	-50.3088	± 0.167
Region 1	b_1	-1.8159	± 0.050
Region 1	c_1	-0.13144	± 0.0039
Region 2	a_2	-45.8754	± 0.283
Region 2	b_2	$+1.3359$	± 0.094
Region 2	c_2	-0.14306	± 0.0068
Region 3	a_3	-40.9270	± 0.061
Region 3	p_3	-0.99460	± 0.0063

For atom-atom impact ionization, the original expression from Reference [145] decomposes k_{AAI} into separate rate coefficients for direct ionization from ground state ($L_0(T_A)$) and ionization from all possible excited energy states ($G_{0R}(T_A)$), such that

$$k_{AAI}(T_A) = L_0(T_A) + 2G_{0R}(T_A) \tag{A.3}$$

G_{0R} is the rate coefficient of excitation via atom-atom collisions into the first resonance level, which when doubled approximately equals the rate coefficient for excitation into all excited levels. Eq. (A.3) assumes that all excited particles are eventually ionized in a given time step. L_0 and G_{0R} are given by

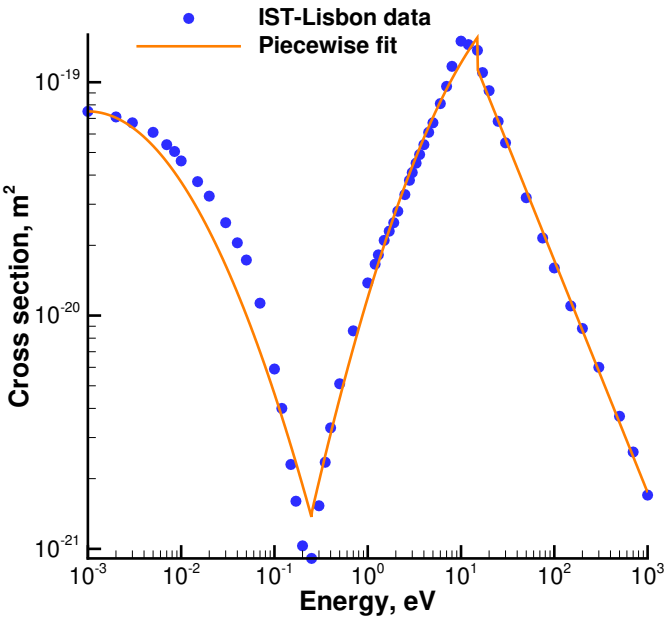

Figure A.1: Elastic momentum-transfer cross section data for electron-neutral collisions from the IST-Lisbon database.

$$L_0(T_A) = 64\pi a_0^2 \left(\frac{E_1^H}{E_0}\right)^2 \sqrt{\frac{k_B T_A}{\pi m_A}} \xi_0^2 \frac{m_e m_A}{m_H(m_e + m_A)} \Psi_A \left(\frac{E_0}{k_B T_A}\right) \tag{A.4}$$

and

$$G_{0R}(T_A) = 64\pi a_0^2 \left(\frac{E_1^H}{E_0}\right)^2 \sqrt{\frac{k_B T_A}{\pi m_A}} \xi_0^2 f_{0R} \frac{m_e m_A}{m_H(m_e + m_A)} \Psi_A \left(\frac{E_{0R}}{k_B T_A}\right) \tag{A.5}$$

where a_0, m_H, and E_1^H are the first Bohr radius, mass, and first ionization energy of atomic hydrogen. For argon, $E_0 = 15.76$ eV, $E_{0R} = 11.6$ eV, $\xi_0 = 6$, and $f_{0R} = 0.32$. $\Psi_A(\chi)$ is a function given by

$$\Psi_A(\chi) = \frac{1 + 2/\chi}{1 + \left(\frac{2m_e}{m_e + m_A}\chi\right)^2} \cdot e^{-\chi} \tag{A.6}$$

where the $e^{-\chi}$ term is not included in [145], but makes Eq. (A.3) more consistent with the expression presented in Reference [146]¶.

Using a non-linear least squares fit to the modified Arrhenius rate form, $A_{AAI} = 8.9169 \times 10^{-23}$ m^3/s/K, $\eta_{AAI} = 1.1935$, and $T_{AAI} = 129,450$ K for argon. The coefficient of determination for the fit is $R^2 = 1$, with a maximum fit error of 7.44% at $T = 5000$ K. Figure A.2 shows the original Kelly and Drawin formulae alongside the modified Arrhenius fit.

¶Personal communication, Timothy T. Aiken, June 11, 2025.

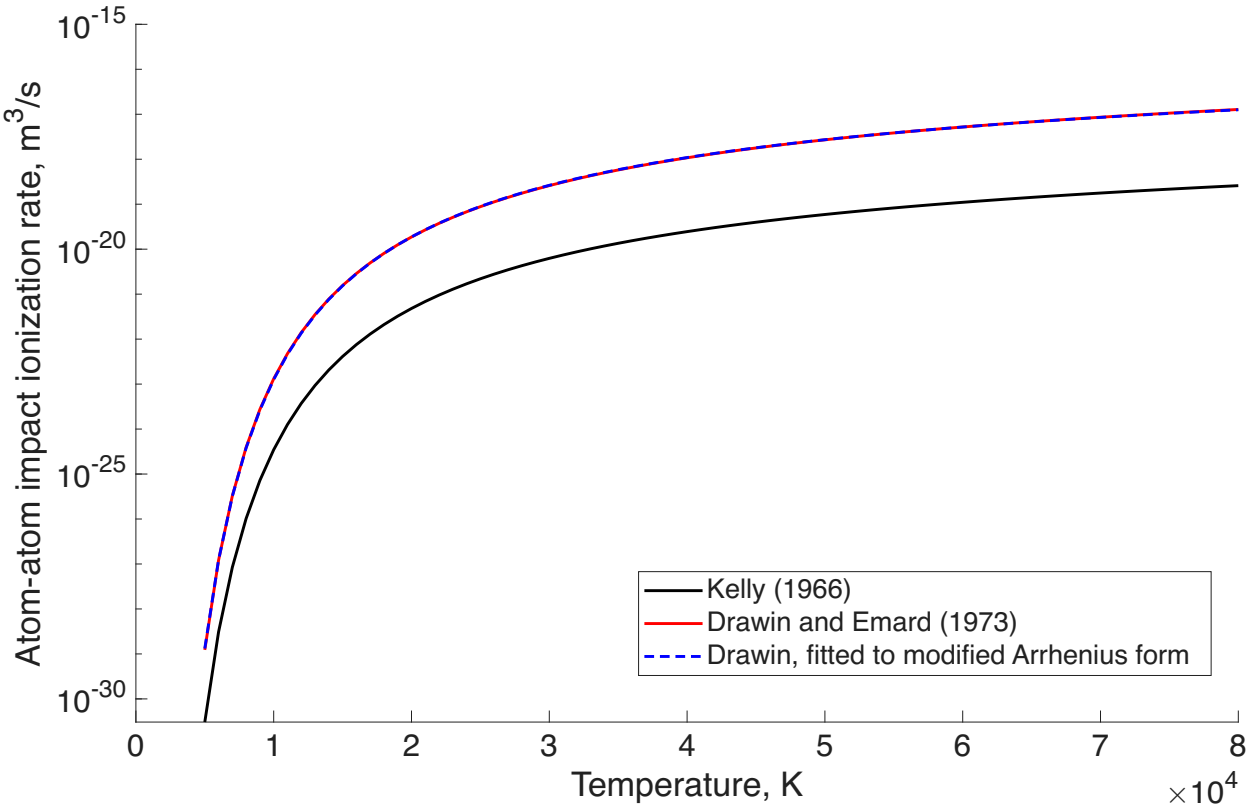

Figure A.2: Atom-atom impact ionization rates from Kelly [146], Drawin and Emard [145] and modified Arrhenius fit of Drawin and Emard data.

Appendix B

Parallelized DSMC Stagnation Streamline Scheme

The following section benchmarks MPIC's parallelized implementation of the 1D DSMC stagnation streamline model of Section 4.2.1 with the ISTAG code. Conservation of mass, momentum, and energy is verified in MPIC for the stagnation streamline scheme, but, for the sake of brevity, these tests are not included here. The Mach 39 flow at 81 km is used for testing with simulation initialization parameters located in Tables 4.1 and 4.2. Pseudo-code for the parallel scheme is also provided.

Three different MPIC modes are benchmarked. The first is 2D MPIC, where a configuration space grid of 500 cells in the $\hat{x}$ direction and one cell in the $\hat{y}$ direction is initialized. Symmetry boundary conditions are applied for the two $\hat{y}$ boundaries. The second is 1D MPIC using the standard NTC collision pair selection method. The third mode is 1D MPIC with the semi-dynamic SWS pair selection method. For this third mode only, 90% of the inflow comprises neutral argon with a weight modifier of $w_1 = 1.0$, and the other 10% comprises neutral argon with a weight modifier of $w_2 = 0.01$. Since the only difference between these particle types is their weight, the final stagnation streamline flowfield should demonstrate strong agreement with the other profiles if the semi-dynamic SWS is implemented correctly. The only difference between 1D and 2D MPIC simulations in this section is that 2D MPIC performs a wall collision when particles encounter a $\hat{y}$ boundary. The $\hat{y}$ boundaries do not exist in 1D MPIC, and thus no wall collisions occur in the $\hat{y}$ direction, reducing computational cost.

Table B.1 lists the RMSE error taken with respect to ISTAG for each flowfield property and

each MPIC solver option, normalized to range of the ISTAG data and written as a percent. For all properties, all solvers have less than 1% RMSE with the ISTAG data. Figure B.1 displays the ratio of the neutral density with the freestream density for all four codes. All profiles demonstrate strong agreement in relation to each other within the context of the 1D stagnation streamline scheme. Of note, 2D MPIC's density at the wall is considerably lower than the other codes, reaching a ratio of only 36.51. ISTAG's density ratio is 61.72, default 1D MPIC's ratio is 57.61, and semi-dynamic SWS MPIC's ratio is 56.05. This result motivates the use of a true one-dimensional MPIC in Chapter 5. Figures B.2 and B.3 display the profiles of bulk velocity and temperature. Good agreement is achieved again for the 1D stagnation streamline scheme, however, it is clear 1D MPIC with the semi-dynamic SWS exhibits the most error with respect to ISTAG's profile. This is primarily due to the slight difference in numerical form between Equations 2.71 and 2.72 of the weight maximum filter and Equations 2.55 and 2.56 of the standard NTC method when SWS uses different particle weights. As discussed in Section 4.3.1, the DSMC 1D stagnation streamline model is highly sensitive to placement of x_r. It follows that with iterative adjustment of x_r in the semi-dynamic SWS MPIC setup, stronger agreement could be achieved between the profiles.

Overall, these results demonstrate that the 1D stagnation streamline scheme is successfully implemented in MPIC.

Table B.1: Root mean square error, normalized with respect to the range of the ISTAG data, for each flowfield property in the verification of the 1D stagnation streamline scheme in MPIC.

MPIC solver option	MPIC-2D	MPIC-1D	MPIC-1D, SWS
Ratio with freestream density	0.0299%	0.0444%	0.1934%
Streamwise bulk velocity	0.0424%	0.0371%	0.1394%
Temperature	0.0978%	0.0635%	0.1836%

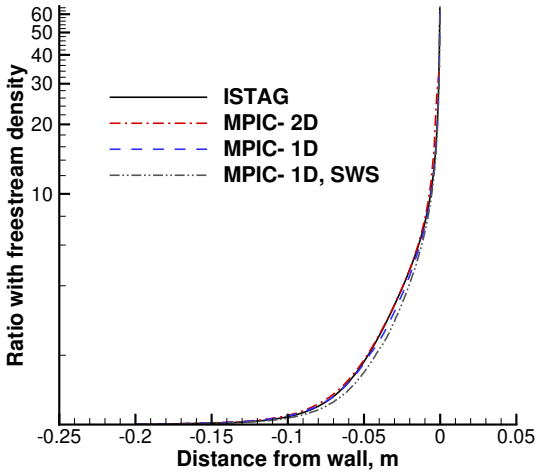

Figure B.1: ISTAG and MPIC ratio with freestream density profiles for Mach 39 flow at 81 km.

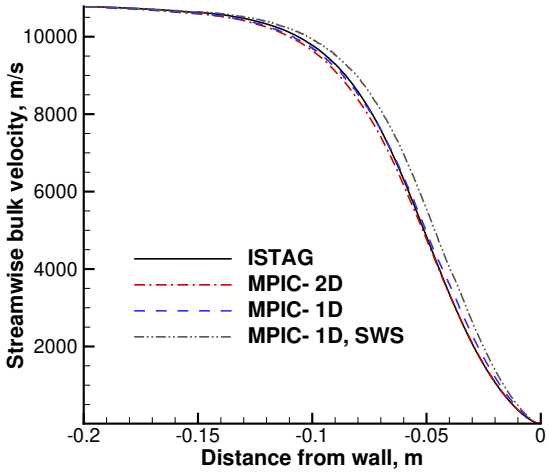

Figure B.2: ISTAG and MPIC streamwise bulk velocity profiles for Mach 39 flow at 81 km.

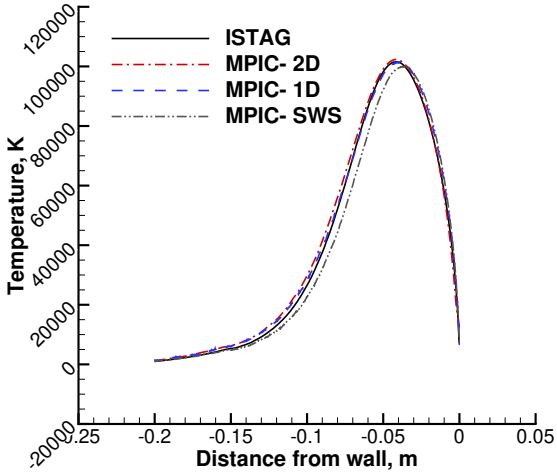

Figure B.3: ISTAG and MPIC temperature profiles for Mach 39 flow at 81 km.

```
Data: F_in ≥ 0
while F_in > N_eligible,global do
    for each cell do
        N_eligible,cell = 0;
        for each particle in cell do
            evaluate particle for removal;
            if particle is removable then
                N_eligible,cell++;
            end
        end
    end
    if rank = 0 then
        collect global array of N_eligible,cell;
        N_eligible,global = Σ_cells N_eligible,cell;
        broadcast N_eligible,global to all ranks;
    end
    else
        send local array of N_eligible,cell to rank 0;
    end
end
if rank 0 then
    N_removed,global = 0;
    while N_removed,global < F_in do
        generate random number i between 1 and N_eligible,global;
        find cell with ith particle;
        N_remove,cell++;
        N_removed,global++;
    end
    broadcast array of N_remove,cell to all ranks;
end
for each cell do
    for each particle in cell do
        if particle is removable then
            delete particle;
            N_remove,cell−;
        end
        if N_remove,cell = 0 then
            break;
        end
    end
end
```

Algorithm 1: Parallelized algorithm for DSMC stagnation streamline method. MPI operations are indicated with the LaTeX `teletype` environment.