UNIVERSITY OF SOUTHAMPTON

A Numerical Model of the Ion Thruster Hollow Cathode Plasma.

Francis T.A. Crawford

A Thesis submitted for the Degree of
Doctor of Philosophy

Astronautics Group
School of Engineering Sciences
Faculty of Engineering and Applied Science

December, 2004

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND APPLIED SCIENCE

SCHOOL OF ENGINEERING SCIENCES (ASTRONAUTICS)

Doctor of Philosophy

A NUMERICAL MODEL OF THE ION THRUSTER HOLLOW CATHODE PLASMA

by Francis T.A. Crawford

Electric propulsion (EP) devices such as gridded and Hall effect ion thrusters (HET) offer significant mass savings compared to chemical rockets. In addition to enhancing existing mission possiblities, it is also true that cutting edge EP systems are an enabling technology for certain high ΔV missions. Although ion and HET thrusters are currently used operationally, there still remains a need for further understanding of the physics associated by the devices in order to enhancing performance and reliability.

The hollow cathode (HC) is a device commonly used in electron bombardment ion thrusters to provide a primary ionising current source, while a secondary cathode is often used as a plume neutraliser. Although HCs are often used, the internal plasma physics is still poorly understood and they exhibit operational modes the cause of which has not been fully explained. This is primarily due to the difficulties associated with recording experimental data from within the very small scale cylindrical cathode.

In order to gain further understanding of the hollow cathode internal physics, a numerical model was developed. Due to the degree of rarefaction and low density, the neutral flow was modelled using a Direct Simulation Monte Carlo (DSMC) model. It was found that the flow was rarefied in the cathode plume, and numerically transitional in Knudsen number within the 'throat' (tip) of the hollow cathode.

Due to the predicted plasma density in the cathode, a Particle-in-Cell (PIC), Monte Carlo Collision (MCC) model was then added to the existing neutral gas model. This was validated against some standard test cases and experimental data for cathode conditions. It was found that the hollow cathode plasma exhibits a dense emitting region extending a short distance inside. The model was then used to characterise the discharge and arguments based on the emitting region hypothesis are extended to explain the results.

Contents

List of Figures ix
List of Tables xi
Acknowledgements xii
Nomenclature xiii
List of Acronyms xv

1 Introduction **1**
1.1 Rocket Propulsion 1
1.1.1 The specific impulse limit in chemical rockets 2
1.2 Electric propulsion 4
1.2.1 Micro-propulsion 7
1.3 Objectives of this research 9
1.4 Thesis structure 10

2 Background and Literature review **11**
2.1 The Hollow Cathode 11
2.1.1 Experimental Cathode Research 13
2.1.2 Analytical 17
2.1.3 Numerical 17
2.2 Approaches to CFD in electric propulsion 19
2.3 Electron Emission 21
2.3.1 Thermo-Field 22
2.3.2 Secondary electron 26
2.3.3 ion wall neutralisation 27
2.4 Mean free path analysis 29
2.4.1 Electron-Neutral 30
2.4.2 Charged particle 32
2.4.3 Neutral-Neutral 35
2.4.4 Ion-Neutral 36

3 Numerical Method **40**
3.1 Direct particle simulation 40
3.1.1 Particle model 40
3.2 Particle distributions 46
3.3 Direct Simulation Monte Carlo 50
3.3.1 DSMC collision sampling procedure 51
3.3.2 Collision Models 52
3.3.3 Producing a free stream flow in the DSMC 55
3.4 Monte Carlo Collisions (MCC) 56

3.4.1 e-A collisions: Event selection . . . 57
3.4.2 e-A collisions: Post collision properties . . . 58
3.4.3 Ion-neutral Collisions . . . 61
3.4.4 Coulomb Collisions . . . 63
3.5 Particle-in-Cell (PIC) Plasma model . . . 65
3.5.1 Field Equations . . . 67
3.5.2 Weighting (particle acceleration) . . . 70
3.5.3 Interpolation . . . 72
3.5.4 Stability . . . 75
3.5.5 Finite difference discretisation . . . 76
3.5.6 Adjustments that enable PIC stability . . . 80

4 Validation **84**
4.1 Random number generator . . . 84
4.2 Coulomb collision coefficients . . . 86
4.3 DSMC and direct particle transport . . . 87
4.3.1 Supersonic Taylor-Couette flow . . . 88
4.4 Field solver and numerical discretisation . . . 92
4.5 PIC validation . . . 99
4.5.1 Flat plate discharge: Configuration and theory . . . 99
4.5.2 Thermionic to space charge limited emission . . . 100

5 Results: Neutral propellant flow **105**
5.0.3 Boundary Conditions . . . 108
5.1 Numerical Validity . . . 109
5.1.1 Time-step . . . 109
5.1.2 Grid . . . 113
5.2 DSMC cross section model . . . 116
5.3 Reference case: flow structure . . . 121
5.4 Mass Flow rate . . . 124
5.5 Tip radius . . . 128
5.6 Gas . . . 133
5.7 Temperature . . . 136

6 Results: Plasma discharge **140**
6.1 Additional configuration . . . 140
6.2 Initialisation of the discharge . . . 143
6.2.1 Pre-breakdown thermionic discharge . . . 143
6.2.2 Breakdown . . . 145
6.3 Numerical Validity . . . 146
6.3.1 Time-step . . . 146
6.3.2 Mesh . . . 148
6.3.3 Electron Mass Adjustment (value of $\sqrt{f}$) . . . 151
6.4 Discharge structure . . . 155
6.4.1 Axial structure . . . 158
6.4.2 Radial profiles . . . 164
6.4.3 Velocity Distribution Functions . . . 167
6.5 Current and Mass flow rate . . . 171

6.5.1 Current-Voltage . . . 173
6.5.2 Plasma conditions . . . 174
6.6 Tip radius . . . 180
6.7 Anode position . . . 182

7 Discussion **188**
7.1 Neutral Gas . . . 188
7.1.1 Rarefaction . . . 188
7.1.2 Comparison to Experimental Results . . . 190
7.1.3 Cathode geometry . . . 195
7.1.4 Neutral gas parameters . . . 196
7.2 Plasma Discharge . . . 198
7.2.1 Numerical Validity . . . 198
7.2.2 Discharge Structure . . . 199
7.2.3 Comparison to Experiment . . . 199
7.2.4 Cathode Geometry . . . 204
7.2.5 Parallelism . . . 206

8 Conclusions **208**
8.1 General Remarks . . . 208
8.2 Novelty of the Research . . . 209
8.3 Recommendations for Future Work . . . 210

A The Software **214**

B Field solver Algorithms **215**

C Differential Operators in Cylindrical Coordinates **218**

D Experimental setup of Gessini *et al.* **219**

References **220**

List of Figures

1.1 A converging-diverging nozzle. Subscripts $_0$ and $_e$ represent combustion chamber conditions and nozzle exit plane conditions respectively. 3
1.2 Illustration of an ion thruster. Regions denoted I,II and III refer to the coupling plasma, main discharge plasma and ion beam plasma respectively. 7

2.1 The hollow cathode. This a slice through the cathode: it is cylindrical. 12
2.2 Hollow cathode operating modes. ϕ_D is discharge voltage (i.e. the anode voltage when operating in diode configuration), I_d is the discharge current. 15
2.3 Rarefaction limits of various fluid descriptions. 21
2.4 Thermo-Field electron emission. The cathode is at ϕ_c and is located at $x = 0$. The plasma is quasi-neutral for $x > L$ at at potential of ϕ_p. The sheath has shape $\sqrt{1 - e^{-x}}$ and the electric field is shown in the lower plot for this sheath via $E = -d\phi/dx$. The field strength that enhances electron emission is then E_s. 22
2.5 Electron energy distribution for electrons that are successfully emitted. Plotted using equation (2.8) using Tungsten as the material ($\phi = 4.5$ eV) at 1500 K. There is assumed to be no external applied field ($E_s = 0$), so this example is representative of pure thermionic emission. Note: the energy distribution is essentially the tail of a Fermi-Dirac distribution, and additionally that the bulk of the emitted electrons have energies a few times higher that the thermal energy of the emitting surface. 26
2.6 e-Xe collision cross sections vs. electron energy. The data is from Puech and Mizzi [82] 31

3.1 Movement of a particle in the r–θ plane. 42
3.2 Illustration of wall collisions at different thermal accommodation coefficient. 44
3.3 A conic frustum. 48
3.4 Collision geometry of hard sphere molecules (from Bird [9]). $\mathbf{g}$ is the pre-collision relative velocity (notice that the co-ordinate system is aligned with the frame of reference due to $\mathbf{g}$). χ is the resulting deflection angle and $\mathbf{g}^*$ is the post-collision velocity. 53
3.5 Geometric representation of Nanbu's technique for the sampling of an electron-atom collision event. 57

3.6 The Particle-Mesh algorithm applied to plasma physics. Subscripts $_p$ and $_m$ refer to particle and mesh information respectively. Notice that the 'Particle Motion' stage takes place entirely in 'particle space', while the field solution stage is entirely in 'mesh space'. Interpolation and weighting transfer information between mesh and particle. 65
3.7 Illustration of the $\mathbf{v} \times \mathbf{B}$ rotation of velocity in cylindrical coordinates due to azimuthal magnetic field (actually the case $B_\theta > 0$). The method of Boris [13] states that $\Delta\theta = -\omega_c \Delta t$, i.e. $\dot{\theta} = -\omega_c$. 71
3.8 Nearest grid point (NGP) PIC interpolation. Grid points are located at X_j with spacing Δx. The density profile generated by a particle located at x_p is considered. See Figure 2-6a of Birdsall and Langdon [12] 73
3.9 Cloud-in-Cloud (CIC) PIC interpolation. In the same way as the illustration of the NGP scheme (fig. 3.8), mesh points are located at X_j and a test particle is located at x_p. Cell spacing is Δx. For further explanation, see figure 2-6b of Birdsall and Langdon [12] and associated description. 74
3.10 Cartesian grid computational molecule. 77

4.1 The configuration of the Taylor-Couette flow. The concentric shapes represent annular vorticies; this illustrates a stable three-vortex configuration. 88
4.2 Contours of density ratio in a Taylor-Couette flow. Ta = 521637 and Kn = 0.02. This snapshot of the flow field is taken after 30 rotations of the inner cylinder (located at r=1). The aspect ratio of the plot is 2:1, the same as is used in [98] and [9]. 90
4.3 Contours of temperature ratio in a Taylor-Couette flow. Ta = 521637 and Kn = 0.02. This snapshot of the flow field is taken after 30 rotations of the inner cylinder (located at r=1). 91
4.4 Axial-radial velocity vectors due to vorticies in a Taylor-Couette flow. This snapshot of the flow field is taken after 30 rotations of the inner cylinder (located at r=1). 91
4.5 Comparison of analytical and numerical solution of cylindrical Laplace equation. The wall potential on $r = R$ is $V_0 \cos(2\pi z/L)$ 96
4.6 Finite difference truncation error between analytical and numerical solution. 97
4.7 Comparison of various iterative solvers. Shown are the reference Jacobi and BiCG solvers, and the final BiCGStab solver. Note reversed r axis in the top figure. 98
4.8 PIC validation case: configuration. Boundary A is a Neumann condition for the solution of Poisson's equation, and is set to reflect particles at $a_c = 0$ so that there is no radial current at B. 101
4.9 Thermionic and Space charge emission. The anode potential is 50 V and the separation is 2 mm. The plate material has a work function of $\phi_w = 3$ eV. Circular data points are for the discharge simulated using the *real* electron mass ($m_e = 9.1 \times 10^{-31}$ kg) and the crosses represent simulation runs at artificial electron mass of $240000 m_e$, or approximately one Xenon atom. 102

4.10 Electric potential on axis for a flat plate discharge at two cathode temperatures. 103

5.1 Cathode geometry. The scale for both sub-figures is as for the reference case: $z_t = z_k = r_i = 1$ mm, $z_{ck} = 2$ mm, $r_t = 0.5$ mm, $r_k = 2.5$ mm. The dimensions are essentially that of the T6 cathode. See table 5.1 for further reference dimensions. The points A, B and C are reference sample points where velocity distribution functions are recorded. Taking the bottom left of the upper figure to be at (0,0) mm, A is at (0.5, 0.5) mm, B at (10.5,0.5) mm and C at (15,1) mm. 106
5.2 Neutral gas time-step resolution test: $\dot{m}$ and ν_{DSMC}. These time history plots are taken from the reference $\Delta t = 0.2$ μs case. ν_{DSMC} is a computational value that indicates the number of collisions calculated per second simulation time, *not* the real collision rate in particles per second. 111
5.3 Neutral gas time-step independence test: axial stream velocity. Simulation results at three sample points for $\Delta t = 0.1$, 0.2, 0.4, 0.6, 0.8 and 1 μs. 111
5.4 Neutral gas time-step independence test: density. Simulation results at three sample points for $\Delta t = 0.1$, 0.2, 0.4, 0.6, 0.8 and 1 μs. 112
5.5 Three levels of grid refinement used to evaluate grid independence in the MCC collision routines. 114
5.6 Neutral gas mesh refinement test: axial stream velocity. Simulation results at the three sample points for three levels of mesh density. . . . 114
5.7 Neutral gas mesh refinement: axial stream velocity. Simulation results at the three sample points three levels of mesh density. 115
5.8 Comparison of collision cross section σ for HS and VHS DSMC collision models as a function of collision velocity. 117
5.9 Axial centreline plot of density for the three collision models. 118
5.10 Axial centreline plot of axial stream velocity for the three collision models. 118
5.11 Velocity distribution plots of particles simulated using the hard sphere collision model near the upstream boundary. 119
5.12 Velocity distribution plots of particles simulated using the collisionless model model near the upstream boundary. 120
5.13 Axial centreline plots of number density, temperature, pressure and Mach number for the neutral gas flow in the reference case hollow cathode. 122
5.14 Axial centreline plot of Knudsen number for the hollow cathode under reference case conditions. The characteristic length is assumed to be equal to the tip radius: 700 μm. 123
5.15 Contours of number density for the reference case. 124
5.16 Contours of temperature (K) for the reference case. 124
5.17 Contours of axial stream velocity for the reference case. 125
5.18 Contours of scalar pressure (Pa, log scale) for the reference case. 125
5.19 Contours of Mach number for the reference case. 125
5.20 velocity and density plotted as functions of mass flow rate 126
5.21 Mach number and density plotted as functions of mass flow rate 127

5.22 Knudsen number at the tip (○) and in the plume (×) of the cathode as a function of Xenon mass flow rate. 128
5.23 Axial stream velocity and density at the tip as a function of tip to cathode area ratio. 130
5.24 Axial stream velocity and density near the upstream boundary plotted as a function of tip to cathode area ratio. 131
5.25 Axial stream velocity and number density sampled downstream of the expansion plotted as a function of tip to cathode area ratio. 132
5.26 Mach number as a function of expansion area ratio at three different sample points. Upper plot shows plume Mach number (position C), lower plot shows interior and tip points (positions A and B). See fig. 5.1. 132
5.27 Comparison of gas conditions for three different propellants at reference case configuration. The upper plot shows density, the middle plot temperature and the lower plot scalar pressure. All data are axial centerline values. 134
5.28 Time histories of heat flux between the cathode insert and working gas for three different propellant at a fixed mass flow rate and wall temperature. 135
5.29 Density and Mach number in the cathode plume plotted as a function of cathode temperature. 137
5.30 Scalar pressure and temperature in the cathode plume plotted as a function of cathode temperature. 138
5.31 Exhaust velocity as a function of chamber temperature. The dashed line is numerically evaluated based on chamber temperature, chamber pressure and plume pressure, while the solid line is directly sampled velocity data from the numerical simulation. 139

6.1 Time histories of electron population and anode current for the no-propellant discharge 144
6.2 Radial cut showing plasma potential and electron density at $z = 10$ mm (mid way along the insert axially). The space charge limiting sheath can be seen. 144
6.3 Normalised species densities as a function of simulation time-step, Δt. Values normalised to reference case value. 147
6.4 Normalised species temperature as a function of simulation time-step, Δt. Values normalised to reference case value. 148
6.5 Reference PIC and MCC/DSMC meshes used in the plasma model. . . 149
6.6 Normalised species densities in the tip region (sample point B) as a function of number of PIC mesh cells. Values normalised to reference case value ($N = 10800$). 150
6.7 Normalised species temperatures in the tip region (sample point B) as a function of PIC mesh cells. Values normalised to reference case value ($N = 10800$). 150
6.8 Normalised species densities as a function of electron mass adjustment factor, f. Values normalised to reference case value ($\sqrt{f} = 480$). 152
6.9 Normalised species temperatures as a function of electron mass adjustment factor, f. Values normalised to reference case value ($\sqrt{f} = 480$). . 153

6.10 Neutral downstream mass flux time history for the reference and minimum (case C) electron mass acceleration factor. 154
6.11 Contours of electron number density for the reference case discharge. Note log scale colourmap. 155
6.12 Contours of ion number density for the reference case discharge. Note log scale colourmap. 156
6.13 Contours of neutral number density for the reference case discharge. Note log scale colourmap. 156
6.14 Contours of electric potential (V) for the reference case discharge. . . . 157
6.15 Contours of current density |j| for the reference case discharge. 158
6.16 Contours of azimuthal induced magnetic field B_θ for the reference case discharge. 158
6.17 Axial centreline number density of the three species for the reference case discharge 159
6.18 Axial centreline temperature of the three species for the reference case discharge, see note in text regarding calculation of temperature. The sudden drop on the downstream value represents the anode: no temperature is recorded here. The drop in temperature is not part of the modelled physics; it is a plotting anomaly. 161
6.19 Axial centreline mean velocity in the axial direction of the three species for the reference case discharge. The sudden drop on the downstream value in the upper plot represents the anode: no data is recorded here. The drop in temperature is not part of the modelled physics; it is a plotting anomaly 162
6.20 Axial centreline plots of current density — $e(n_i u_i - n_e u_e)$, neutral scalar pressure and plasma potential for the reference case discharge 163
6.21 electron density: Radial cuts at various points axially (see individual plot title) showing electron density, reference case discharge. 165
6.22 ion density: Radial cuts at various points axially (see individual plot title) showing ion density, reference case discharge. 166
6.23 Plasma potential: radial cuts at various points axially (see individual plot title) showing plasma potential, reference case discharge. Cathode (at $z = 1$ mm, plots 1-3 and $z = 0.7$ mm, plot 4) is at 0 V. 167
6.24 Velocity distribution plots of electrons in the cathode interior, at $z = 5$ mm. 169
6.25 Velocity distribution plots of electrons in the cathode plume, at $z = 12$ mm, $r = 0.2$ mm. 169
6.26 Velocity distribution plots of Xenon ions in the cathode interior. 170
6.27 Velocity distribution plots of Xenon ions in the cathode plume. 171
6.28 Time history of anode current and voltage for the reference 20 A discharge case, Xe gas at 1 mg s^{-1}. 173
6.29 Current-voltage behaviour of the cathode using Xenon at two different mass flow rates ('plume' and 'spot' conditions). 174
6.30 Cathode internal (i.e. reference position A) plasma density as a function of current for two mass flow rates, 1 mg s^{-1} (−) and 3.29 mg s^{-1} (−−). The data are normalised to the 10 A, 1 mg s^{-1} case for each species. 175

6.31 Cathode internal (i.e. reference position A) species temperature as a function of current for two mass flow rates, 1 mg s^{-1} (−) and 3.29 mg s^{-1} (−−). The data are normalised to the 10 A, 1 mg s^{-1} case for each species. 177
6.32 Cathode plume plasma density as a function of current for two mass flow rates, 1 mg s^{-1} (−) and 3.29 mg s^{-1} (−−). The data are normalised to the 10 A, 1 mg s^{-1} case for each species. 178
6.33 Cathode plume species temperature as a function of current for two mass flow rates, 1 mg s^{-1} (−) and 3.29 mg s^{-1} (−−). The data are normalised to the 10 A, 1 mg s^{-1} case for each species. 179
6.34 Interior species density as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm. 181
6.35 Plume species temperature as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm. 181
6.36 Anode voltage as a function of anode position at a fixed current. 183
6.37 Interior species density as a function of anode position. Anode position of 0 represents reference case. 184
6.38 Plume species density as a function of anode position. Anode position of 0 represents reference case. 184
6.39 Plume species temperature ratio as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm. 185
6.40 Plume species pressure ratio as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm. 186
6.41 Plume species axial stream velocity ratio as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm. . . . 186

7.1 Comparison of the experimental data of Fearn and Patterson for hollow cathode upstream pressure to numerical results for the same cathode geometry. Numerical data is plotted for a range of insert temperatures and experimental data is shown for 5 and 10 A discharge sweeps. In the simulation, the pressure monotonically increases with cathode temperature, so the larger pressure data points are at the higher temperatures. 191
7.2 Axial centreline plot of axial stream velocity. 193
7.3 Comparison of numerical, experimental and analytical results for the thrust produced by cold gas flow in the hollow cathode. 194
7.4 Radial magnetic field profile in the plume. In the simulation, the radial cut is made at $z = 14$ mm. Plot (b) is from Patterson and Fearn [77] . 202
7.5 Comparison of experimental data of Rudwan [88] to simulation results for integrated axial electron temperature. Simulation data is given as average axial temperature and compensated average axial temperature, where the plume temperature is weighted to account for the experimental configuration. 204

List of Tables

1.1 Magnitude of ΔV required for some typical space missions. Adapted from Fearn [32]. 1

1.2 Key features and performance of current electric propulsion devices. Adapted and condensed from Martinez-Sanchez and Pollard [58] 6

2.1 Experimental data describing plasma properties of the hollow cathode. This table was originally complied by Edwards [31] and has been adapted and extended to include among others the results of the original author. Cathodes with names such as `UK-25` refer to the those that accompany RAE/DERA/Qinetiq gridded ion thrusters where the number refers to the main discharge chamber diameter in centimetres. `SERT` refers to the US 'Space Electric Rocket Test' program that evolved to provide the NASA DS-1 ion thruster. 14

2.2 Analytical cathode models. All are 1D models that assume some emitting plasma region and model the three species flow using 1D approximations. 18

2.3 Numerical modelling of the hollow cathode. The MHD plasma simulations of Murray were never completed, but initial results can be found in Murray et al. [65]. 19

2.4 Fluid models. 20

2.5 Collision types, cross sections and mean free paths of the various collisions that can occur in the maximum density region of the hollow cathode discharge. The right-hand column denotes those collision types that are included in the numerical model; the adjacent column shows the preceding section number that contains the discussion of that collision type. 38

3.1 Extracting macroscopic fluid properties from particle data. m is molecular mass of species a. Further properties can be extrapolated from these using gas laws, etc. 45

3.2 VHS parameters for the inert gases. Diameters are referenced to a temperature of 273 K, properties are given for standard conditions (101,325 Pa and 0°C). Degrees of freedom is 3 in all cases since all the elements are insert monatomic species: the ratio of the specific heats is then $\gamma = (\xi + 2)/\xi = (3 + 2)/3 = 1\frac{2}{3}$. Radon, the most massive of the insert gas species, is omitted for two reasons: (a) there is no gas dynamic data available and (b) it is not available in sufficient quantities to be a viable propellant for the hollow cathode. 54

4.1 Tabulation of $A(\tau)$ between 0.01 and 4 generated by the test function within the code. This matches the data of Nanbu [69] 86

5.1 Neutral flow: reference case parameters. Lengths are in mm. 107

6.1 Discharge: reference case parameters. 142

7.1 Plasma properties reproduced for comparison with this research. Some of this table was originally complied by Edwards [31] and has been adapted and extended to include several other, more recent results. Cathodes with names such as `UK-25` refer to the those that accompany RAE/DERA/Qinetiq gridded ion thrusters where the number refers to the main discharge chamber diameter in centimetres. `SERT` refers to the US 'Space Electric Rocket Test' program that evolved to provide the NASA DS-1 ion thruster. 200

Acknowledgements

First, I would like to thank my supervisor, Prof. Steve Gabriel for all of his help and advice, and a seemingly endless flow of ideas and questions. I would also like to thank those who where involved in experiments on hollow cathodes at Southampton, Ismat Rudwan, Sabrina Pottinger and Paolo Gessini, all of whom were very helpful to me during the period of this work. Many thanks go to Kevin Meyer of AM & PM research for inviting me to present this research at a fluid dynamics conference.

I would like to thank everyone in the Astronautics group would I knew over the years and who helped make the time so enjoyable. In no particular order: Scott Walker, Mark Paine, Steffen Dransfelt, David Simonin, Neil Williams, Duncan Goultry, Gavin Gittins.

Finally, I would like to thank Lizzy for supporting me throughout the time of this research. You were patient, trusting and supportive at all times. Thank you.

Nomenclature

Δt	timestep, s
Δx	mesh spacing, m
Λ	ratio of Coulomb collisions impact parameters
α	macroscopic recombination rate, eq. (2.31)
β	$\sqrt{m/(2kT)}$, eq. (3.44)
γ	ratio of specific heats, eq. (1.5)
δ_{se}	secondary electron emission yield
δ_n	ion-wall neutralisation probability
ϵ_0	permittivity of free space, F m^{-1}
ϵ	electron emission energy, eV
λ	mean free path, m
λ_D	Debye length, m
μ	viscosity coefficient, N s m^{-2}; also
μ	for particles A and B, the reduced mass, $m_A m_B/(m_A + m_B)$
μ_0	permeability of free space, N A^{-2}
ν	collision rate, s^{-1}
ξ	molecular degrees of freedom
ρ	charge density, C m^{-3}
τ_{pe}	electron plasma period, s
τ_{pi}	ion plasma period, s
τ	mean time between collisions: ν^{-1} s; also
τ	Coulomb collision scattering paramter, eq. (3.81)
ϕ	electric potential, V
χ	scattering angle
ω	viscosity index, see table 3.2
ω_c	cyclotron frequency, Hz
ω_{pe}	electron plasma frequency, Hz
ω_{pi}	ion plasma frequency, Hz
ω_r	Linear solver residual tolerance

A_0	Richardson's constant, 1.2×10^6 A m^{-2}, eq.(2.5)
$\mathbf{B}$	magnetic field, T
$\mathbf{E}$	electric field, V m^{-1}
I_{sp}	specific impulse, s, eq. (1.4)
Kn	Knudsen number
Pr	Prandtl number
Re	Reynolds number
Ta	Taylor number, eq. (4.13)
R	uniformly distributed random number of range $[0, 1]$
$\mathbf{R}$	uniformly distributed random vector with $\|\mathbf{R}\| = 1$
Sc	Schmidt number
$S(x)$	PIC interpolation/weighting shape factor
a	speed of sound, m s^{-1}, eq. (1.8)
a_c	thermal accomodation coefficient, eq. (3.11)
b	Coulomb collision impact parameter, $\approx \lambda_D$, m
b_0	90° scattering Coulomb collision impact parameter, m, (2.19)
c	speed of light in a vacuum, 3×10^8 m s^{-1}
d	molecular diameter, m
e	elementary charge, 1.6×10^{-19} C
$\hat{\mathbf{r}}, \hat{\theta}, \hat{\mathbf{z}}$	unit vectors
f	electron mass adjustment factor (fluxes hence altered by $\sqrt{f}$)
$f(v)$	particle velocity distribution fn., m^{-3}
$\mathbf{g}$	Relative velocity between two particles, m s^{-1}
$g(v)$	particle velocity magnitude distribution fn., m^{-3}
g_0	acceleration due to gravity at Earth's surface, 9.81 m s^{-2}
$\mathbf{h}$	post-collision scattering velocity, m s^{-1}
h	Planck constant; $\hbar = h/2\pi$
$\mathbf{j}$	current density, C m^{-2} s^{-2}
k	Boltzmann's constant, 1.38065×10^{-23}J K^{-1}
n	number density, m^{-3}
s	molecular speed ratio, m s^{-1}, eq. (3.47)
$\mathbf{u}$	stream velocity, m s^{-1}
$\mathbf{v}$	thermal velocity, m s^{-1}
v_{ex}	rocket exhaust velocity, m s^{-1}, eq. (1.2)
v_{th}	thermal velocity (most probable velocity in a Maxwellian), β^{-1} m s^{-1}
$\mathbf{x}$	particle position, m

List of Acronyms

BDS	backward difference scheme
CDS	central difference scheme
CEX	charge exchange (collision)
CFD	computational fluid dynamics
CG	conjugate gradient
CR	collisional-radiative
DSMC	direct simulation Monte Carlo
ECR	electron cylcotron resonance
EP	electric propulsion
FEEP	field effect electric propulsion
FFT	fast Fourier transform
FD	finite difference
FDS	forward difference scheme
HET	Hall effect thruster
LAM	local area multicomputer
LTE	local thermodynamic equilibrium
MCC	Monte Carlo collision
MHD	magneto-hydro-dynamics
MPD	magneto-plasma-dynamic (thruster)
MPI	message passing interface
NSSK	north south station keeping
PIC	particle in cell
PPT	pulsed plasma thruster
RIT	radio-frequency ion thruster

Chapter 1

Introduction

This chapter provides general background material covering spacecraft propulsion, in particular electric propulsion. The origin of the fundamental limitation in the performance of chemical rockets is discussed in detail here because it coincidentally introduces some useful terminology for later sections.

The following sections cover the objectives of the research presented and the structure of its presentation in the ensuing chapters.

1.1 Rocket Propulsion

Space missions are usually characterised by the magnitude of the velocity increment, ΔV, required to complete them. This velocity change is a fixed parameter for a given mission: regardless of the spacecraft mass, the velocity change must be completed. Typical values of ΔV for a variety of missions are shown in table 1.1. It can be shown that the amount of ΔV that a spacecraft can obtain is a function of the available fuel mass and the propulsion exhaust velocity. This can be derived from the rocket equation:

$$M\frac{dv}{dt} = -v_{ex}\frac{dM}{dt} \tag{1.1}$$

Mission	ΔV (km s^{-1})
Escape from 480 km altitude orbit	3.15–7.59
Earth orbit to Mars orbit and return	14
Earth orbit to Jupiter orbit and return	64
North-South Station keeping for 1 year	41-51 m s^{-1}

Table 1.1: Magnitude of ΔV required for some typical space missions. Adapted from Fearn [32].

where M refers to spacecraft mass and v_{ex} is the rocket exhaust velocity. The solution of this equation is

$$\Delta V = v_{ex} \ln \left(\frac{M_0}{M_t} \right), \tag{1.2}$$

where M_0 is the initial spacecraft mass and M_t is the mass at time t. If ΔM is the mass of propellant consumed, then the result becomes

$$\Delta V = v_{ex} \ln \left(\frac{M_0}{M_0 - \Delta M} \right). \tag{1.3}$$

This result is important because it tells us that the total ΔV that is available is proportional to the rocket exhaust velocity, v_{ex}. Such a result may similarly be derived by considering the standard metric of propulsion performance, specific impulse, which is defined as

$$\mathrm{I}_{sp} = \frac{T}{\dot{m} g_0}, \tag{1.4}$$

where T is thrust, $\dot{m}$ is propellant flow rate (kg/s) and g_0 is the gravitational acceleration. If we take $T = \dot{m} v_{ex}$, then the expression for specific impulse becomes $\mathrm{I}_{sp} = v_{ex}/g_0$. In other words, the performance of any propulsive device that relies on the expelling of propellant is fundamentally related to the absolute exhaust velocity of the propellant. The reason why electric propulsion (EP) receives so much attention is due to the fact that EP thrusters can provide v_{ex} far in excess of conventional chemical rockets. Table 1.2 shows the characteristics of a wide range of electric propulsion devices. In the table, row five shows values for I_{sp} which, for the more advanced EP rockets, are of the order of thousands. Compare this to the fact that the maximum I_{sp} for the most advanced chemical rocket — liquid bi-propellant — is around 430.

1.1.1 The specific impulse limit in chemical rockets

Why do chemical rockets fall so short of electric rockets? In order to answer this question, it is worth briefly examining some basic compressible nozzle flow theory. In fact, as will become clear, the hollow cathode acts in a gas dynamic sense in the same way as a converging-diverging nozzle, so the following analysis will be useful later.

Consider figure 1.1. This shows an example of a converging-diverging nozzle. The minimum area or *throat* point is shown with a dotted line. Subscripts e refer to exit conditions. There are various expressions that relate the pressure, temperature and Mach number ratios for such a nozzle that are derived from elementary isentropic compressible flow theory. From these expressions, it is possible to understand the fundamental limit on chemical rockets. For further derivation of the following, see Anderson [2] or Hill and Peterson [43]. For the quasi 1D nozzle in figure 1.1, the

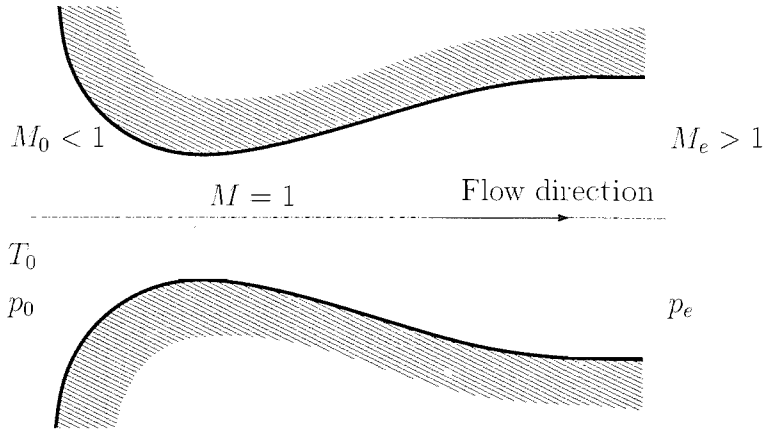

Figure 1.1: A converging-diverging nozzle. Subscripts $_0$ and $_e$ represent combustion chamber conditions and nozzle exit plane conditions respectively.

relation

$$\frac{p_0}{p} = \left(\frac{T_0}{T}\right)^{\gamma/(\gamma-1)} \tag{1.5}$$

holds, where, for a perfect gas, p is pressure, T temperature and γ the ratio of specific heats. The temperature ratio is

$$\frac{T_0}{T} = 1 + \frac{\gamma - 1}{2} M^2, \tag{1.6}$$

where M is the Mach number. Substituting (1.6) into (1.5) gives

$$\frac{p_0}{p} = \left(1 + \frac{\gamma - 1}{2} M^2\right)^{\gamma/(\gamma-1)}. \tag{1.7}$$

The exhaust velocity $v_e = M_e a_e$, where a_e is the sound speed: $a_e = \sqrt{\gamma R T_e}$. So, if we solve (1.7) for M_e we obtain

$$M_e^2 = \frac{2}{\gamma - 1}\left[\left(\frac{p_0}{p_e}\right)^{(\gamma-1)/\gamma} - 1\right] \tag{1.8}$$

for the Mach number. In a similar way, the exit plane temperature is given by re-arranging (1.5) so that

$$T_e = T_0 \left(\frac{p_e}{p_0}\right)^{(\gamma-1)/\gamma} \tag{1.9}$$

Multiplying the above by the sound speed gives us the exit plane velocity:

$$v_{ex}^2 = M_e^2 a_e^2 \tag{1.10}$$

$$v_{ex}^2 = \frac{2}{\gamma - 1}\left[\left(\frac{p_0}{p_e}\right)^{(\gamma-1)/\gamma} - 1\right]\gamma R T_0 \left(\frac{p_e}{p_0}\right)^{(\gamma-1)/\gamma} \tag{1.11}$$

$$= \frac{2\gamma R T_0}{\gamma - 1}\left[1 - \left(\frac{p_e}{p_0}\right)^{(\gamma-1)/\gamma}\right]. \tag{1.12}$$

This result is particularly revealing. Optimisation of v_{ex} by way of maximising the pressure ratio is somewhat futile because v_{ex} will converge to some finite ideal value as $p_0/p_e \to \infty$ (i.e. $p_e/p_0 \to 0$). The utility of very high pressure ratios is not sufficient to overcome the structural cost: it is sufficient for $p_e/p_0 < 0.01$. The result is that for a given propellant, the only variable that significantly affects v_{ex} is the chamber temperature T_0. v_{ex} is fundamentally limited by the maximum temperature that the chamber walls can endure. Such temperature limits are around 3500 K for some Titanium-Nickel alloys, perhaps higher for ceramics.

As an example, a liquid oxygen-liquid hydrogen rocket has $\gamma = 1.22$ and $R = 519.6$ for the chemically reacting gas in the chamber. The pressure ratio is around 2500 – 25 atm chamber pressure and 10^{-2} atm exit pressure [2]. The theoretical upper limit of v_{ex} is

$$\begin{aligned} M_e &= \sqrt{\frac{2(2500^{0.180} - 1)}{1.22 - 1}} = 5.31 \\ a_e &= \sqrt{(1.22)(519.6)(3500)(0.0004^{0.180})} = 735.65 \text{ m s}^{-1} \\ v_{ex} &= (5.31)(735.65) = 3905 \text{ m s}^{-1} \\ \Rightarrow \mathrm{I}_{sp} &= 3905/9.81 \approx 398 \end{aligned}$$

Where equation (1.11) was used instead of (1.12) so the magnitudes of the Mach number and sound speed can be shown. Clearly then, chemical rockets have fundamental limits associated with them. The question is, how do electric rockets overcome this? Since there is a wide range of EP systems, specific explanations are best left to following section. It suffices to say that in each case, energy is added to the the exhaust by other methods instead of, or in addition to, the nozzle expansion.

1.2 Electric propulsion

There are two good review papers that attempt to cover the field of electric propulsion: Fearn [32] and, more recently, Martinez-Sanchez and Pollard [58]. The most logical mechanism for classification of the many devices is by the primary mechanism

by which the exhaust velocity is raised. This classification system is not flawless since there is inevitably much overlap. Indeed, all gas propellant electric rockets derive some thrust via conventional gas dynamic expansion of the propellant, albeit a very small amount in some cases. Classification by (electro-) thermal, magnetic and electrostatic propulsion; this ordering is broadly representative of increasing complexity and performance, although there are several exceptions. An overview of the most successful devices is given in table 1.2. The table, adapted and abbreviated from Martinez-Sanchez and Pollard [58] is shown to give an impression of the comparative performance and requirements of current operational devices.

In the interests of brevity and relevance, detailed discussions of thermal and magnetic EP devices are omitted. Since the work in this study relates to the hollow cathode, it is only worth considering devices that make most use of it. The **ion thruster** is such a device, and can be considered a flagship device within electric propulsion. Brewer [16] is the author a book on the ion thruster, and much research has been conducted since. The success of the NASA Deep Space 1 mission that used an ion thruster for primary propulsion is further testimony to this [97]. In an gridded ion thruster, a plasma is formed behind an electrode grid, as the name suggests. The grid serves accelerate ion beam-lets. Since positive ions are expelled, the thruster becomes negatively charged. This means that the plume must be neutralised. This is true of both gridded and grid-less ion thrusters (hall effect thrusters). A hollow cathode is used to expel electrons corresponding to the ion current of the plume, so that the spacecraft does not become electrically charged. The method by which the working plasma is generated in a gridded ion thruster is used to further classify the device. The working plasma can be formed using a DC discharge (Kaufmann type), with the current provided by a hollow cathode, by RF excitation of the propellant (the German Radio-frequency Ion Thruster - RIT) or by tuned Electron Cyclotron Resonance (ECR). Figure 1.2 shows a diagram of a Kaufmann type ion thruster. The solenoid (or fixed magnet) plays an important role: the magnetic field in the chamber serves to increase the ionising electron residence time because the electrons follow helical paths along magnetic field lines. The increase in electron path results in a corresponding increase in ionisation probability. From a practical point of view, it is generally accepted that gridded ion thrusters require a more complex collection of auxiliary equipment compared to grid-less (hall effect) thrusters. Hence, currently Hall effect ion thrusters are considered the more robust general purpose EP device. Nonetheless, ion thrusters are still competitive for station keeping applications and offer higher I_{sp} generally than HETs (see table 1.2).

Name	Resistojet	Arcjet	HET	Ion	PPT	FEEP	self-field MPD
Main type	Thermal	Thermal	Magnetic	Electrostatic	magnetic	Electrostatic	Magnetic
Sub type		(magnetic)					(Thermal)
Propellant	NH_3	NH_3	Xe	Xe	Teflon	Cs	-
Power (W)	500	500-20 K	300-6000	200-4000	1-200	10^{-5}-1	200-4000 K
I_{sp} (s)	350	500-800	1600	2800	1000	6000	2000-5000
η	80%	27-36%	50%	65%	7%	80%	30%
Peak Voltage (V)	28	200	300	900	1000-2000	6000	100
Mass (kg/KW)	1-2	0.7	2-3	3-6	120	-	-
Lifetime (h)	500	> 1000	> 7000	10000	10^7 pulses	-	-
status	Operational	Operational	Operational	Operational	Operational	Development	Lab
Mission	NSSK	NSSK	NSSK, mid ΔV	NSSK, large ΔV	precision orbit correction	precision orbit correction	large ΔV

Table 1.2: Key features and performance of current electric propulsion devices. Adapted and condensed from Martinez-Sanchez and Pollard [58]

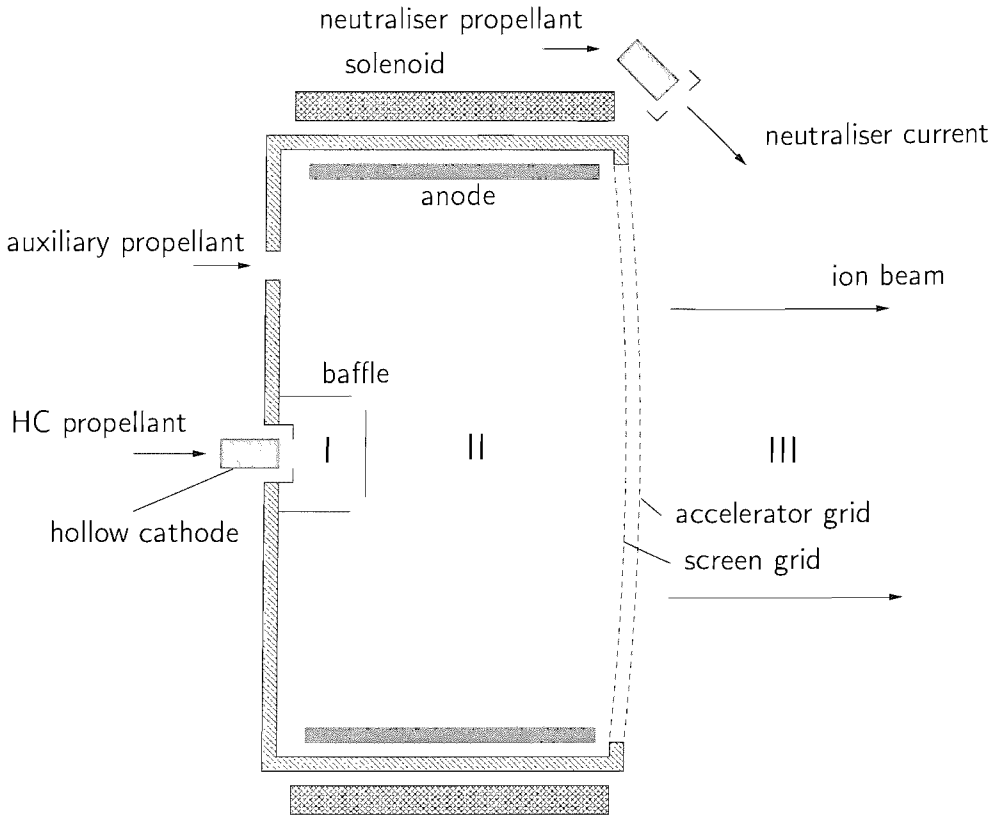

Figure 1.2: Illustration of an ion thruster. Regions denoted I,II and III refer to the coupling plasma, main discharge plasma and ion beam plasma respectively.

1.2.1 Micro-propulsion

Kaufmann type ion thrusters suffer a particular problem when miniaturisation is considered. Consider the ionisation chamber of a DC discharge gridded ion thruster. Here, the efficiency is closely related to obtaining the maximum ionisation probability for each electron passing from the cathode to the anode (see fig 1.2). A magnetic field is applied to the main discharge region so that the electron follow a longer path (since in the presence of a magnetic field, charged particle trajectories are helical [40]). If the chamber size is reduced, the same ionisation efficiency can only be retained with a corresponding increase in magnetic field strength.

Research has been carried out by Wirz et al. [108] on a 3.3 cm chamber diameter thruster. Even for such a small discharge chamber, the group managed to produce reasonable performance through modification of the configuration: the accepted wisdom regarding chamber configuration, particularly magnet configuration, does not hold for these smaller chambers. Wirz et al. [108] concluded that more extensive testing with electromagnets is required to optimise performance. It was also concluded that along with a small discharge chamber, there was a requirement for a small current source, either field emitter or hollow cathode (for simplicity, the research of Wirz et al.

[108] used a planar cathode). Hutchinson and Gabriel [46] provided an analysis of a micro ion thruster, and also built and tested a micro hollow cathode. Initial results are promising, although the complexities of experimental testing of hollow cathodes remain.

Similarly, Sandonato's group have completed experimental and numerical analysis of a variety of configurations of a 5 cm chamber diameter thruster [90 91]. The group developed numerical and analytical models of the discharge chamber in order to study the effects of miniaturisation on propellant utilisation efficiency. Again, by the use of unconventional magnetic confinement configurations, they feel that the 5 cm chamber ion thruster would provide good theoretical performance.

As the requirement for even smaller thrusters presents itself, some have suggested that due to exotic behaviour of the hollow cathode under some conditions (it can produce very high energy ions) that the main chamber be dispensed with completely. This leaves the hollow cathode thruster, which is still only in the design stage. Gessini et al. [38] is, at the time of writing, attempting to measure the thrust (and hence predict specific impulse) for hollow cathodes operating without any auxiliary discharge chamber and accelerator grids.

For even lower thrust applications, the Colloid thruster array concept – currently receiving much interest – remains the primary micro-thrust candidate. Of course, at this stage of the development of micro thrust devices, several other possibilities exist, most notably MEMS arrays of resistojets [47] and perhaps the MEMS ion thruster concept of Yashko et al. [110].

1.3 Objectives of this research

The fundamental problem with understanding hollow cathode physics is that taking experimental measurements within the cathode is fraught with difficulty; these problems will be discussed in detail in the literature review (following chapter). Although there are several reasonable hypotheses, there is no generally accepted description of the structure of the internal plasma. In addition, there have been almost no experiments relating to the neutral gas flow, and the comparative effect neutral gas dynamics may have on the discharge characteristics. In light of these statements, the objectives of the work presented here follow. Using a numerical description of gas and plasma mechanics:

- Investigate and characterise the neutral gas flow in the hollow cathode:
  - Condition of the flow,
  - Scale of the forces present,
  - effect of varying cathode geometry/configuration.
- Try to understand the structure of the internal plasma during discharge:
  - Fundamental structures (sheaths, etc),
  - Scaling with mass flux and current,
  - effect of varying cathode geometry/configuration.

As long as the standard methods for validation and comparison of the numerical model with experimental data are followed, it is expected that the conclusions that can be drawn regarding the gas and plasma structure can help clarify current understanding of hollow cathode physics.

1.4 Thesis structure

The following chapter contains a review of previous hollow cathode research, including the more important experimental studies, analytical models and numerical models. This is followed by a discussion of the problems associated with computational fluid dynamic (CFD) analysis of very low pressure plasmas. Since the cathode is fundamentally simply a current source, a section covering electron emission is provided. Finally, in order to quantify the collisional processes present in the bulk discharge, a mean free path analysis concludes the chapter.

Chapter 3 contains a presentation of the numerical method used in this research. This begins with a description of the fundamentals of particle simulation CFD, and is followed by the methods by which inter-particle collisions are implemented. Finally, the particle-in-cell (PIC) plasma model is presented. The chapter is concluded with some notes on the numerical stability of the model and adjustments that can be made to ensure stability.

Any study that involves the use of a simulation to further the understanding of a device is of little value unless the simulation code is validated. Chapter 4 is devoted to the presentation of various numerical studies of both individual components and the overall validity of the code. This includes comparing the output to some standard, well-known cases. Of course, this represents only a part of the verification of the model, comparison to hollow cathode data is contained in the discussion chapter.

Results produced for this work are divided into two chapters: 5 and 6. This division occurs because chapter 5 contains results obtained in the first phase of the research that characterise the neutral gas flow in the hollow cathode. Some of the work presented in chapter 5 is contained in references [24] and [25]:

- Francis T. A. Crawford and S. B. Gabriel. "Modelling small hollow cathode discharges for ion micro-thrusters." *AIAA paper 2002-2101*, May 2002.

- Francis T. A. Crawford and S. B. Gabriel. "Microfluidic model of a micro hollow cathode for small ion thrusters (**invited**)." *AIAA paper 2003-3580*, June 2003.

Chapter 6 contains results that describe the full discharge. Reference [26] contains some early results similar to those in chapter 5:

- Francis T. A. Crawford and S. B. Gabriel. "Numerical simulation of the hollow cathode plasma using a PIC-DSMC code." In *Proceedings of the 28th International Electric Propulsion Conference, Toulouse, France*, March 2003.

The seventh chapter contains a discussion and analysis of the results in the previous two chapters, including comparison to experimental cathode data. Conclusions and recommendations for further work are as usual found in the final chapter.

Chapter 2

Background and Literature review

This chapter contains a review of some of the more important previous hollow cathode research (including experimental results, and modelling). Next, approaches to computational fluid dynamic analysis of very low density gases and plasmas are reviewed. Since the hollow cathode is fundamentally an efficient electron emitter, there is a section devoted to an analysis of electron emission at the particle energies expected within the hollow cathode plasma. Finally, an analysis and comparison of the mean free paths for the many collision types present in the hollow cathode plasma is presented. This allows a first estimate as to which types of collision will dominate the cathode plasma transport, and thus the collision models that are needed in a realistic hollow cathode simulation.

2.1 The Hollow Cathode

The hollow cathode is a special type of cathode used in several different types of electric propulsion rocket. For example, two HCs are used in DC discharge ion thrusters (one discharge and one neutraliser), while hollow cathode neutralisers are used for any thruster that produces an ion beam, Hall Effect Thrusters (HET), for instance. In addition to the obvious importance of the hollow cathode in terms of convention electric propulsion, there is also much interest in the use of hollow cathodes — or micro hollow cathodes — for micro-propulsion applications. An illustration of the HC is shown in figure 2.1, the diagram is a cut-through of the device. It is essentially an open ended cylinder with a constricted end (the 'tip' or 'orifice'). Neutral gas is injected from the left. During operation, a plasma is present in the enclosed region (sometimes referred to as the 'insert region'). The diagram shows an example of a 'keeper' anode, this is a familiar arrangement, but by no means a necessary component.

The 'insert' or 'dispenser' is made of a low work function material to enhance electron emission. In early cathodes, no insert was used, and the interior was simply

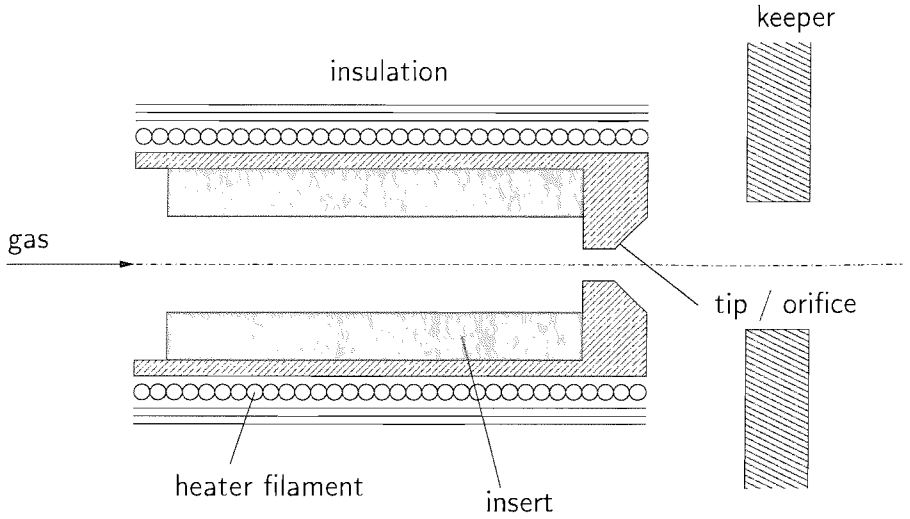

Figure 2.1: The hollow cathode. This a slice through the cathode: it is cylindrical.

coated with Barium. Modern cathodes make use an insert made of porous tungsten impregnated with a 4:1:1 mixture of BaO, CaO and Al_2O_3. This lowers the work function, allowing a very large current density to pass from the cathode to the plasma located adjacent to the insert. The current flow within the cathode is then axial, through the tip and outward towards either an anode (or keeper/anode pair), or in some cases, another plasma. Depending on whether or not there is a free space between cathode and keeper through which plasma can escape, the configuration is referred to as 'open keeper' or 'closed keeper'.

Some of the components in the diagram are relevant only during initialisation of the discharge, but are shown for completeness. In particular, the heater is essential for startup, but switched off during operation: ion bombardment from the plasma heats the cathode. The keeper itself is of primary use only during startup; it is possible remove the keeper completely if required.

There are two categories downstream configurations to consider: a coupling plasma and an anode. In the first case, for instance during operation of an ion thruster, the current from the cathode flows to another plasma. In the second case, an anode of some type is placed downstream. For experimental investigation of the hollow cathode, the anode configuration is generally used, although there is also research conducted using the cathode in the 'in-thruster' configuration. In the case of most experiments discussed later, the cathode is in the diode (simple cathode-anode, no external coupling plasma) configuration. This setup has traditionally been used for cathode experimentation, although doubts have recently been raised over the validity of this approach due to the proposed interference of the external discharge structure with the discharge characteristics (see Rudwan [88]).

The hollow cathode often exhibits two different operational modes, shown in figure

2.2. At low discharge current and mass flow rate, the 'plume' mode exists. Typical of this mode is a noisy, high voltage discharge with a characteristic bright plasma plume downstream of the tip. In contrast, 'spot' mode occurs at higher currents and mass flow rates. This mode is associated with low anode voltage and an intense spot adjacent to the orifice. The presence and explanation of these modes is still under investigation; there is a long list of theories and contradicting experimental evidence (see Edwards [31] for further discussion). It has been variously proposed that the modes result from

- location and size of the keeper - Mandell and Katz [57]
- structure of the tip plasma sheath - Csiky [29]
- extent of emitting area within cathode - Philip [79], also Wells and Harrison [106]
- location and configuration of anode - Rudwan [88]

The final item (the most recent) attempts to fully explain the phenomena using a more unified theory; in fact Rudwan [88] suggests that the transition characteristics and presence of the modes is entirely down to the external configuration of the anode, rather than cathode parameters. Rudwan is the only author to critically compare his theory to that of the others. One factor that contributes to the complexity of spot/plume transition is the presence of an hysteresis effect. Figure 2.2 demonstrates the spot and plume regions, along with an impression of the transition hysteresis.

The following subsections cover some elements of research carried out into hollow cathode physics. Hollow cathodes are mentioned by many authors since they are used as an auxiliary device for both gridded and Hall effect ion thrusters: the most rigorously researched devices within the field of electric propulsion.

2.1.1 Experimental Cathode Research

Measurements of the plasma properties within and in the plume of the hollow cathode have been carried out by several authors for many years. Since the design has evolved only to a small extent, many of the early work (late 1970s) remains valid for the trends it shows, see in particular that of Siegfried and Wilbur [94, 95, 96]. Unfortunately, the transition from mercury propellant to inert gas propellant represents something of a discontinuity in data available on the hollow cathode plasma. Today, inert gases are exclusively used for Ion thrusters, so further design optimisation of the mercury device is not required. Table 2.1 shows some of the characteristic results obtained by various experimenters from 1969 to the present day.

	Propellant	Cathode	$\dot{m}$ (Aeq)	I_d (A)	T_e (eV)	n_e (m^{-3})			V_p (V)
Internal Plasma									
Fearn and Phillip (1973)	Hg	UK-10 (T5)	0.014–0.12	1.5	0.1-0.6	10^{16}	–	10^{18}	4.5–7.5
Siegfried (1978)	Hg	SERT-II	0.1	2.0	0.6	1.5×10^{16}	–	4×10^{20}	4.0–8.0
				6.0	0.4–0.6	10^{17}	–	3×10^{20}	4.0–6.0
Ahmed Rudwan (2002)	Xe, Ar, Kr	UK-25 (T6)	0.1–4	1–5	0.8	1×10^{20}	–	6×10^{22}	–
External Plasma									
Csiky (1969)	Hg	SERT-II	0.05	0.3	1.5-2.0	10^{16}	–	10^{17}	11
				2.0	0.5	10^{17}	–	10^{18}	11
Siegfried (1978)	Hg	SERT-II	0.1	2.0	2–5	5×10^{16}	–	3×10^{18}	12–28
				6.0	0.6	3×10^{17}	–	5×10^{18}	11–12
Singfield (1990)	Xe	UK-10 (T5)	0.175–0.475	1.1	0.4–2.2		no data		8–13
Friedly (1992)	Xe	High I_d	0.37	20	2.0			1.5×10^{19}	12
				40	2.9			5×10^{19}	13–15
				60	3.8			6×10^{19}	15–17
Monterde (1997)	Xe	UK-25 (T6)	1.67–3.11	5–10	0.3–0.4	1×10^{17}–	5×10^{17}	$\approx 10^{20}$	12
Edwards (1997)	Xe, Ar, Kr	UK-25 (T6)	0.3–1.2	2–15	1–5	5×10^{17}	–	3×10^{18}	16–30
Crofton (2003)	Xe	UK-25 (T6)	0.2–2.4	4–8	1.15		no data		10

Table 2.1: Experimental data describing plasma properties of the hollow cathode. This table was originally complied by Edwards [31] and has been adapted and extended to include among others the results of the original author. Cathodes with names such as UK-25 refer to the those that accompany RAE/DERA/Qinetiq gridded ion thrusters where the number refers to the main discharge chamber diameter in centimetres. SERT refers to the US 'Space Electric Rocket Test' program that evolved to provide the NASA DS-1 ion thruster.

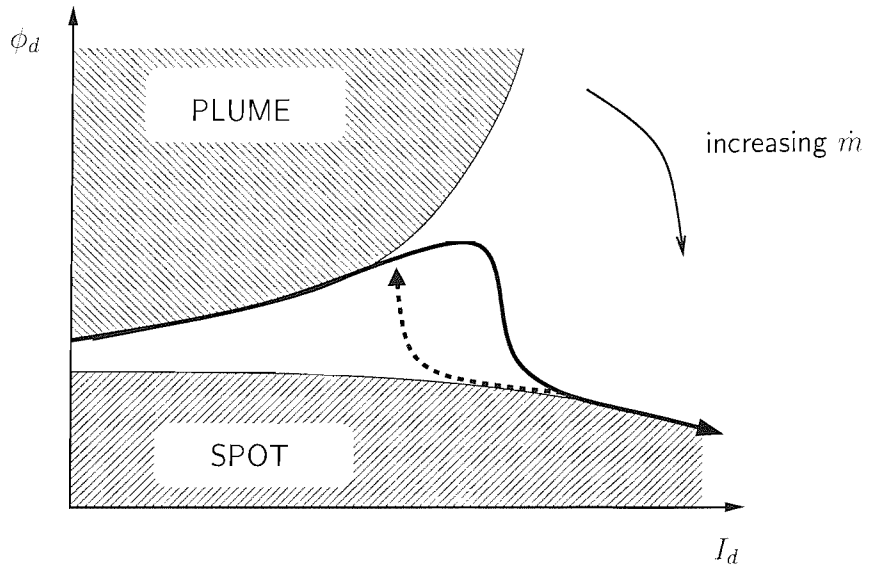

Figure 2.2: Hollow cathode operating modes. ϕ_D is discharge voltage (i.e. the anode voltage when operating in diode configuration), I_d is the discharge current.

It is clear from table 2.1 that most data collected on the hollow cathode plasma relates to the external plasma (the 'coupling' or 'discharge' plasma). The lack of data on the internal plasma is not due to a lack of interest in internal conditions: it is because experimental investigation of the hollow cathode internal plasma is fundamentally very difficult; indeed that is a major motivation for theoretical/numerical research. The use of Langmuir probes within the cathode itself is dismissed by some authors because the presence of a probe would be expected to significantly alter the internal plasma in any case. It is difficult to imagine a way of inserting Langmuir probes right into the heart of the internal discharge without significantly disrupting the discharge itself. Such disruption occurs because the interior is very small and there is a limit to the minimum size that Langmuir probes can be manufactured to without them being destroyed by the plasma. As can be seen from the table, however, there are two instances of internal hollow cathode measurement using probes: Fearn and Philip [34] and Siegfried and Wilbur [96]. Monterde et al. attempted such measurements, but they remain unpublished; the data in the table is for external (coupling plasma) probe measurements, see ref. [63].

Fearn and Philip [34], in a continuation of the initial research of Philip [79] attempted to take measurements of plasma conditions in the hollow cathode interior. A cylindrical Langmuir probe was inserted into the cathode. Fearn notes in [34] that the electron density values (shown in table 2.1) were too low due to the interaction of the probe with the discharge plasma. The probe also indicated plasma potentials in the

range 4.5 to 7.5 V. It should be noted that the cathodes used in the Fearn and Philip experiments were both (a) running with Mercury propellant and (b) much smaller than the current cathodes (tip diameter of the order 0.3 mm compared to around 1.2 mm for the T6 cathode).

Siegfried and Wilbur [96], using a cathode similar to Fearn and Philip (small, Mercury propellant) took measurements using a Langmuir probe, approaching the tip of the cathode from both the upstream and downstream directions. Unfortunately, the discharge tended to be extinguished by the probe when the probe came closer than 4–5 mm from the tip. This indicates the degree of interference such measurements encounter; as can be seen in the table, the values for electron density within the cathode vary by a great extent.

The final set of data relating to internal conditions is that of Rudwan. This data was obtained without an intruding probe via spectroscopic analysis of the plasma. The design of the cathode itself tends to complicate spectroscopic techniques, but it is still possible to examine the cathode along its axis, and this was precisely what was done by by Rudwan. Unfortunately, examination of the plasma in this manner results in an integration of the spectroscopic data along the axis: axial gradients, that must exist, are not resolved. Rudwan records a variation in density of nearly four orders of magnitude for a 5 A discharge as mass flow rate is varied from 1.0 to 3.0 mg s^{-1}, and admits that the process of integrating the spectroscopic measurements axially introduces significant error. Current research by Pottinger and Gabriel [80] is aimed at overcoming the difficulties of spectral analysis by using a modified cathode. In these experiments, the 'instrumented cathode' first suggested by Rudwan is introduced. This modified cathode has had holes drilled in a radial direction through the cathode insert section and tip. This means that spectroscopic probes can examine the plasma through a range of radial holes that extend axially along the cathode. The advantage is clear: axial gradients can now be resolved. Unfortunately, these results are very recent and are yet to be published.

What can be concluded is that experimental investigation of the interior of the hollow cathode is wrought with difficulties and uncertainties. Although some of the new non-intrusive techniques such as the instrumented cathode may yield more reliable data, these methods are still in their infancy.

There are almost no experimental data relating to the gas dynamics of the hollow cathode. The two exceptions are Gessini et al., who measured thrust, and Fearn and Patterson, who attached a pressure transducer in the propellant supply pipe of their cathode. The experiments of Gessini et al. are ongoing at the time of writing, although some preliminary discussions and results are contained in [37] and [38]. Fearn and Patterson [33] found a linear relationship between propellant supply pipe pressure and

mass flow rate in the case of cold gas flow (no discharge) — this would be expected. They also reported significant divergences from the linear depending on the condition of the discharge.

2.1.2 Analytical

For the purposes of this discussion, the distinction between analytical and numerical modelling of the hollow cathode refers to the complexity of the solution; although the models that are discussed in this section are often solved numerically, the solution method is trivial. This is not true of the CFD methods discussed in the following section.

Analytical models of the hollow cathode consist largely of phenomenological 1-D models that assume the length of the cathode is much longer than its diameter. A summary of analytical cathode models is given in table 2.2. The original cathode model was proposed by Siegfried and Wilbur [95] in response to their excellent experimental results (recall from table 2.1 that the Siegfried and Wilbur [94] results are considered some of the highest quality). The model is based on a current balance that describes the emission scaling. The research of Katz et al. [50, 51, most recently] has been concerned with the development of an analytical quasi 1D model of the hollow cathode. The quality of the research of the Katz group is very highly regarded, and the degree of correlation between the 1-D cathode model and experimental data is very good. The experimental data against which the model was compared was that of the neutraliser cathode for the NSTAR (NASA Deep Space 1) gridded ion thruster. Of primary concern was the accurate prediction and improvement of cathode component lifetime. Interestingly, the recent work of Katz et al. [51] suggests that resonant charge exchange (CEX) collisions are essential to understanding the behaviour of the hollow cathode; this conjecture is considered in this thesis. Some of the theoretical ideas contained in the Katz model were also proposed by Malik and Fearn [56]. Recently, a new analytical model has been developed by Rossetti et al. [86]. While this model lacks the maturity of that of Katz *et al.*, it has been demonstrated that it provides a better description of the thermal environment (the Katz *et al.* method currently assumes uniform temperature within the cathode).

2.1.3 Numerical

In contrast to the analytical approach, numerical models tend to make fewer assumptions by specifying a more fundamental description of the cathode. In doing so, the method of numerical solution becomes much more complex. Detail regarding the various methods is contained in the following chapter, but a broad overview is provided

Author(s)	ref.	description
Siegfried and Wilbur (1983)	[95]	Current balance; Fixed length emission region; internal ionisation. Sheath approximation to obtain emission current.
Katz et al. (2002–)	[51]	1-D incompressible continuity, momentum, energy; sheath approximation; predicts CEX collisions essential.
Rossetti et al. (2002–)	[86]	Thermal balance; similar to Katz et al. [51], but model based primarily on thermal gradient in cathode.

Table 2.2: Analytical cathode models. All are 1D models that assume some emitting plasma region and model the three species flow using 1D approximations.

by table 2.3. By comparison to table 2.1 it is clear that the body of experimental research far outweighs that of numerical research. Given the rather dire assessment given previously of the problems associated with experimentally investigating hollow cathode internals, this may be surprising; It turns out that there are some significant problems faced by the numerical experimenter as well.

The research of Murray et al. [65, 66] and Jugroot and Harvey [48, 49] remain the only attempts at numerical modelling of the internal hollow cathode plasma (barring that contained in this document, of course). Murray et al. [65] built a Navier-Stokes solver and succeeded in modelling the internal neutral gas flow. The research may have yielded some very useful research, but was abandoned during the implementation of the plasma code (some initial results can be found in Murray et al. [66]). The work of Jugroot and Harvey [49] was aimed more at the ion thruster discharge chamber, so the cathode research was limited to a basic neutral gas only flow model using a commercial CFD code. Results were similar to Murray et al. [65]. The problem that both encountered was that while a Navier-Stokes model is usually applicable in the insert region of the cathode, it is not in the near vacuum downstream region. These issues of local and global rarefaction of the fluid flow are discussed in detail in the following section.

There has been somewhat more research carried out regarding the plume region of the cathode. As was mentioned, Jugroot and Harvey [49] modelled the discharge chamber of the UK-25 ion thruster. More relevant to the hollow cathode is the collaboration(s) of Crofton and Boyd. Here, experimental research of Crofton has been compared to numerical modelling of the downstream plume region of the hollow cathode. The most recent available data [28] shows good comparison between experimental and numerical in some respects, although there are still some major discrepancies.

Many authors using particle methods to model Hall effect thruster and ion engine plumes resort to very primitive cathode models that are good enough for a neutraliser

	Code type	Model	Geometry
Internal			
Murray (1997)	Navier-Stokes, 3-species plasma	Neutral	UK-10 (T5)
Jugroot (2000)	Navier-Stokes	Neutral	UK-25 (T6)
External			
Jugroot (2000)	PIC-DSMC (particle electron)	Plasma	UK-25 (T6)
Boyd (2003)	PIC-DSMC (fluid electron)	Plasma	n/a (keeper only)

Table 2.3: Numerical modelling of the hollow cathode. The MHD plasma simulations of Murray were never completed, but initial results can be found in Murray et al. [65].

current source but of little relevance here. An example of the implementation of these basic cathodes can be found in Szabo [100].

2.2 Approaches to CFD in electric propulsion

Many electric propulsion thrusters are characterised by being very low flow rate (an exception would be a large MPD, for instance). The ion thruster is just such a device. Since the thruster is operated only in a vacuum, a rather interesting fluid environment ensues. Fundamentally, the problem is that the flow is expected to start off at very low pressure and expand to become nearly collisionless on the downstream boundary. Flows under these conditions can behave in a very different manner to 'conventional' (atmospheric pressure) fluid flows. The degree of rarefaction of a gas is measured with a non-dimensional characteristic known as the Knudsen number, defined as

$$\mathrm{Kn} = \frac{\lambda}{L}, \tag{2.1}$$

where λ is the mean free path between collisions and L is the characteristic dimension of the problem. It turns out that although this definition gives a good impression of the *overall* Knudsen number, it is possible to have regions of high rarefaction. In this case, L can be defined as the scale length of the macroscopic flow gradients [9]

$$L = \frac{\rho}{d\rho/dx}. \tag{2.2}$$

It is necessary to ask the question: what does Knudsen number mean in terms of the selection of a given method to solve a problem? First, it is necessary to classify the methods available. Table 2.4 summarises the methods that are to be considered. Three models are shown. All are fundamentally solutions to the Boltzmann (or Boltzmann-Vlasov equations in the case of a plasma), but differing assumptions are made. In the case of the Navier-Stokes equations, only first order terms in Kn are

Name	plasma	validity	Notes
Navier-Stokes [35]	MHD	$0 < \mathrm{Kn} < 0.1$	The standard modern fluid description. Solution (comparatively) easy and fast.
Burnett [1]	MHD	$0 < \mathrm{Kn} < 0.3$	Higher order accurate version of Navier-Stokes. Increase in complexity.
DSMC [9]	PIC	$0 < \mathrm{Kn} < \infty$	Direct solution to Boltzmann equation, efficient only for rarefied flows.

Table 2.4: Fluid models.

included. The Burnett equations include first and second order Kn terms, while the direct simulation Monte Carlo (DSMC) method takes a completely different approach and makes no assumption about the rarefaction of the flow.

The fundamental difference between a method like the DSMC and the Navier-Stokes equations is that the former describes the gas using a molecular model, while the latter assumes the gas to be a continuum. It is best to understand the limits using a diagram such as that shown in figure 2.3 [9]. The diagram simply gives an overview of the validity of traditional fluid models (Navier-Stokes, Euler) compared to the Boltzmann equation. It is necessary to point out that solution of the Boltzmann equation for non-trivial cases — those in more that one dimension with — is nearly impossible on current computers; if it were possible this entire discussion would be moot as a direct Boltzmann solver would simply be used. For more discussion on the reasons why solving the Boltzmann equation either analytically or numerically, see Bird [9, §3.2]. In a continuum model, a Maxwellian velocity distribution is assumed at all points in the flow: the intermolecular spacing is small and the collision rate is high. In a Molecular model, the velocity distribution of the gas molecules can vary by an arbitrary amount from a Maxwellian. Figure 2.3 can be further explained by the fact that Navier-Stokes equations represent an expansion of the Boltzmann equation to first order Kn. This was demonstrated by the Chapman-Enskog theory [21]. Not shown in the figure is the second order accurate in Kn model: the Burnett equations [1]. This model extends upon the validity of the Navier-Stokes equations somewhat. But only, of course, until third order Kn terms become significant. The Burnett equations are more complex – and more troublesome – to solve that the Navier-Stokes equations. This increase in complexity of higher order continuum Kn methods leads to a catch-22 situation: the Boltzmann equation is prohibitively expensive to solve numerically, but so are increasingly high order Kn Chapman-Enskog expansions. Either approach is very difficult even on todays computers. An efficient model that is independent of Kn is what is required. In fact, this model need not be particularly

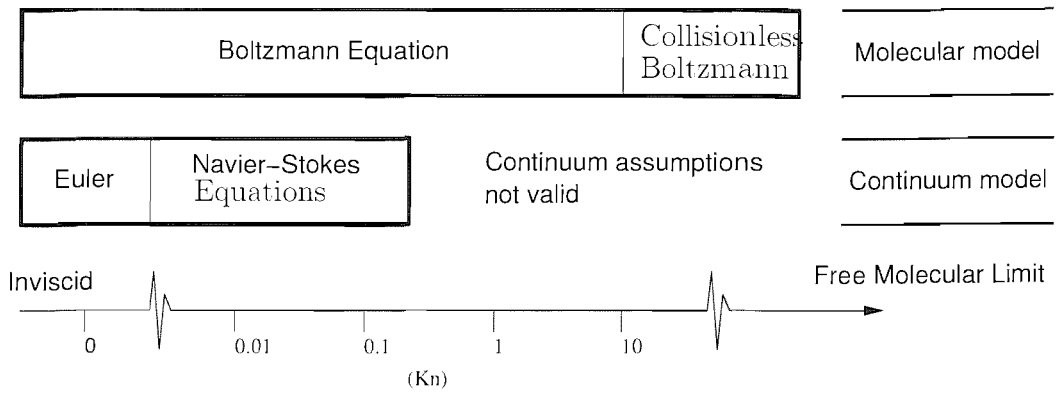

Figure 2.3: Rarefaction limits of various fluid descriptions.

efficient at Kn < 0.1 because the Navier-Stokes or Burnett equations are both efficient and accurate in that region. Of course this model is the aforementioned Direct Simulation Monte Carlo method of Bird [9]. There also exists an extension to the method that accounts for charged particle dynamics (i.e. the presence of a plasma). It is for the reasons outlined above that the vast majority of CFD modelling of electric propulsion components tends to use molecular simulation models.

2.3 Electron Emission

When used in an ion thruster in either discharge chamber or neutraliser role, the hollow cathode is still basically a cathode. The fundamental electron emission processes are now introduced and discussed. There are two classes of emission to consider: space-charge limited and emission limited. The former occurs when the current density is limited by the presence of excess electrons near the emitting surface. In the context of hollow cathodes, we are usually considering the latter: emission-limited current, although it is uncertain as to whether the hollow cathode is emission limited for all surfaces. Most authors choose to simply consider the ubiquitous Richardson-Dushmann model with Schottky correction [31 88] for field-enhanced emission (FEE). The fact that the predicted current is often underestimated by this model is attributed to factors such as secondary electron emission and filament formation and subsequent emission [33]. In fact, it is worth critical examination of the accuracy of the Richardson-Dushmann model in terms of a more comprehensive description; this is covered in the following section. The subsequent section discusses the phenomena of secondary electron emission in the context of hollow cathode plasma conditions.

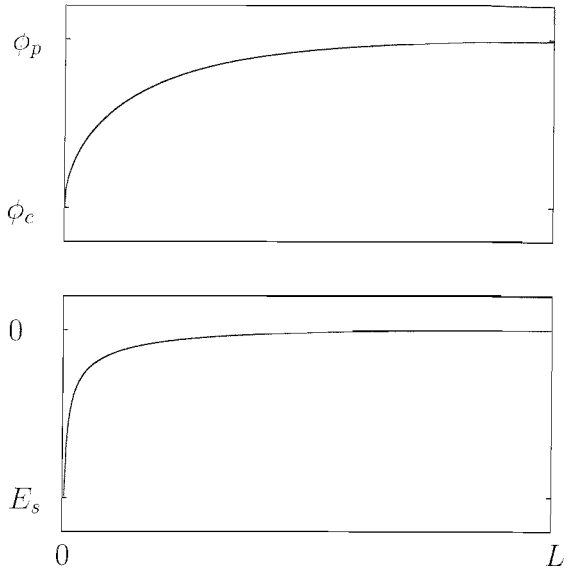

Figure 2.4: Thermo-Field electron emission. The cathode is at ϕ_c and is located at $x = 0$. The plasma is quasi-neutral for $x > L$ at at potential of ϕ_p. The sheath has shape $\sqrt{1 - e^{-x}}$ and the electric field is shown in the lower plot for this sheath via $E = -d\phi/dx$. The field strength that enhances electron emission is then E_s.

2.3.1 Thermo-Field

Consider a plasma adjacent to an emitting surface. The space charge present within the sheath can in some cases enhance the emission current from the surface. Under these conditions, the current density from a cathode is limited by the maximum rate at which it can expel electrons: a function of temperature, surface field strength and the work function of the material. There are two models that are typically used to predict electron emission as a function of the thermal and electric field environment. The commonly used Richardson-Dushmann equation is in essence simply a model of thermionic emission that can be corrected for the presence of an electric field via the Schottky correction. The model of Murhpy and Good [64] is a more accurate — albeit more abstract — description of the phenomena. Figure 2.4 illustrates the situation. A sheath of potential $\phi_p - \phi_c$ forms adjacent to the cathode. The emission is enhanced by the fact that there is an electric field applied at the surface due to the sheath. The emission enhancing field strength is simply $-d\phi/dx$ evaluated at the cathode surface.

Murphy-Good Model

Fundamentally, electron emission from a metal lattice is described in terms of a (potential) barrier at the surface that an electron must overcome to leave the surface. The general form of the Murphy-Good description follows this description. Two values are needed: the probability of a given electron of energy W surpassing the barrier, and

the number of electrons incident on the barrier in some time period. In this way, the emission current density can be related to the surface temperature T_s, work function ϕ and electric field strength E_s by [64]

$$j_{T-F}(E_s, T_s, \phi) = e \int_{-W_a}^{\infty} D(E_s. W)\ N(W, T_s, \phi)\ dW, \tag{2.3}$$

where W is the total energy of an electron incident on the potential barrier at the conductor surface and $-W_a$ is the effective constant potential of the electron within the metal. $D(E_s, W)$ is the probability that an electron of energy W penetrates the barrier and $N(W, T_s, \phi)$ is the so-called 'supply function': the number of Fermi-Dirac electrons within energy range dW incident on the barrier per unit time and surface area. While this model provides a useful description of electron emission from first principles, constructing and integrating expressions of $D(E_s, W)$ and $W(W, T_s, \phi)$ is prohibitively complex. Further details concerning the numerical evaluation of D and N over a range of conditions can be found in Murhpy and Good [64]. It turns out that for a limited range of T_s and E_s, simplifications can be made and an approximate model constructed.

Richardson-Dushmann Model

The Richardson-Dushmann equation is actually the limit of eq. (2.3) as $E_s \to 0$. The result is that the purely thermionic emission current density emitted from a material of work function ϕ is given by

$$j_T(T, \phi) = \frac{4\pi m_e e (kT)^2}{h^3} \bar{t} \exp\left(\frac{-e\phi}{kT}\right), \tag{2.4}$$

where T is the temperature of the material and $\bar{t}$ is the average transmission coefficient, which is in turn a function of the reflection coefficient, r. $\bar{t}$ is assumed to be unity in the basic form of the Richardson equation [62]. The constant

$$A_0 = \frac{4\pi m_e e k^2}{h^3}, \tag{2.5}$$

that constitutes the first few terms of eq. (2.4) is referred to as Richardson's constant (1.2×10^6 A m^{-2}). For a constant non-unity value of $\bar{t}$, it is often assumed that $\bar{t}$ is still unity and that A_0 is modified, although for the vast majority of metals, A_0 does not deviate from the theoretical value, see Shelton [93] for a discussion of the experimental evaluation and accuracy of A_0. The Schottky correction is introduced to adjust the work function of the material due to the presence of an electric field.

The correction is

$$\Delta\phi(E_s) = \sqrt{\frac{eE_s}{4\pi\epsilon_0}}, \tag{2.6}$$

where E_s is the electric field strength at the surface of the emitting material. This value is subtracted from the original work function to give the *effective work function*: $\phi_e(E_s) = \phi - \Delta\phi$. The final description of the current density becomes

$$j_{T-F}(T_s, \phi, E_s) = A_0 T_s^2 \exp\left(\frac{-e\phi_e(E_s)}{kT}\right), \tag{2.7}$$

As can be seen, a large E_s will tend to reduce ϕ and so increase the net current density. Indeed, due to the location of the effective work function (within the exponentional), small values of $\Delta\phi$ can result in much increased j.

Coulombe and Meunier [23] give an excellent comparison of the M-G and R-D models. The authors conclude that since the R-D equation is fundamentally and adjusted thermionic emission model, it is at its most valid under circumstances where there is significant thermal emission. Significant inaccuracies in the case of cold cathodes in an almost pure field-emission mode were also noted: in all cases the R-D equation underestimates the emitted current. It should be noted that there is a description of cold field emission known as the Fowler-Nordheim equation [22]. This expression is not relevant here because it is applicable only in very strong field cases. For the purposes of this research, the cathode is always 'hot'; i.e. some significant fraction of the current density is due to thermal effects. This indicates that the discrepancy between the true (M-G) and approximate (R-D) descriptions will be small. The other factor is that the R-D model is far more tractable for the purposes of practical numerical implementation [23].

Energy of emitted electrons

The final point to consider regarding *direct* electron emission is the energy distribution of the electrons that are successfully emitted. We know that electrons far inside the emitting solid are governed by Fermi-Dirac statistics, the question is, how does this relate to the emitted energy distribution? Mondinos [62] gives a review of the theory of thermionic emission and notes that while there is a large body of theoretical and experimental work relating to the energy distributions of electrons emitted in the pure field (Fowler-Nordheim) emission regime, there is little relating to thermal emission (with the exception of some brief data given by Shelton [93] and some theoretical work of Terpigor'yev [103]).

Mondinos does, however, give a derivation of the emitted electron energy distribution based on the premise that only the tail of the Fermi-Dirac distribution contains

electrons whose energy exceed the work function of the material (this is in fact the basis of the Richardson-Schottky equation in any case). Before presenting the expression, the reflection coefficient, r, needs to be defined. r, strictly $r(E, \mathbf{k}_{\|})$, is the probability that an electron travelling toward the surface of a metal from within is *reflected* back, where E is the electron energy and $\mathbf{k}_{\|}$ is the electron wave-factor parallel to the surface. Recall from the discussion of the Richardson-Schottky expression that the transmission coefficient is usually assumed to be unity, and that this was as a result of setting r to unity. In this case, something similar is done: first we assume r is a function of the internal electron energy only, $r(E)$, and hence that $r = 1$ for $E > E_F + \phi_w$ and $r = 0$ otherwise (where E_F is the Fermi level). In other words, if the electron energy exceeds that of the barrier then the electron is emitted, otherwise it is reflected back into the metal. According to this assumption, it can be shown that the energy distribution of emitted electrons is [62]

$$f(E) = \begin{cases} \dfrac{m_e}{2\pi^2\hbar^3} \exp\left(-\dfrac{\phi_w - \sqrt{e^3 E_s}}{kT_s}\right) \epsilon \exp\left(-\dfrac{\epsilon}{kT_s}\right), & \text{for } \epsilon > 0, \\ 0, & \text{for } \epsilon < 0 \end{cases} \tag{2.8}$$

where E_s and T_s are as before the surface field and temperature and ϵ is defined

$$\epsilon = E - V_{max}. \tag{2.9}$$

E is the electron energy and $V_{max} = (E_F + \phi_w)$ is the total potential barrier that the electron must pass, thus $\epsilon > 0$ must be true for an emitted electron. What this means it that if energies, E, are selected from the Fermi-Dirac energy distribution, then it is possible to find ϵ and hence obtain the post-emission energy of the electron. The $\epsilon > 0$ case of equation (2.8) now provides the energy of an emission capable electron after it has penetrated the potential barrier. If it is assumed that the surface of the material is a uniform crystalline layer, then the electron will have a single velocity component perpendicular to the surface. In the case of a 'rough' surface, the velocity vector may include a deflection angle away from the surface face perpendicular. It is worth noting that the emission energy as predicted by equation 2.8 is a function of the surface temperature, the work function, and the local surface field.

The Fermi energy of Tungsten can be deduced using information such as the base electron configuration and atomic radius to be $E_F = 5.94$ eV (see Harrison [41]). Using this, it is possible to construct an example to demonstrate the shape and comparative magnitude of emitted electrons using eq. (2.8). Figure 2.5 demonstrates this. In the figure, only electrons that have positive ϵ are emitted. The full barrier (including Fermi energy and work function) is shown. It is worth noting that the shape of the distribution is fundamentally part of the tail of the Fermi-Dirac distribution, and

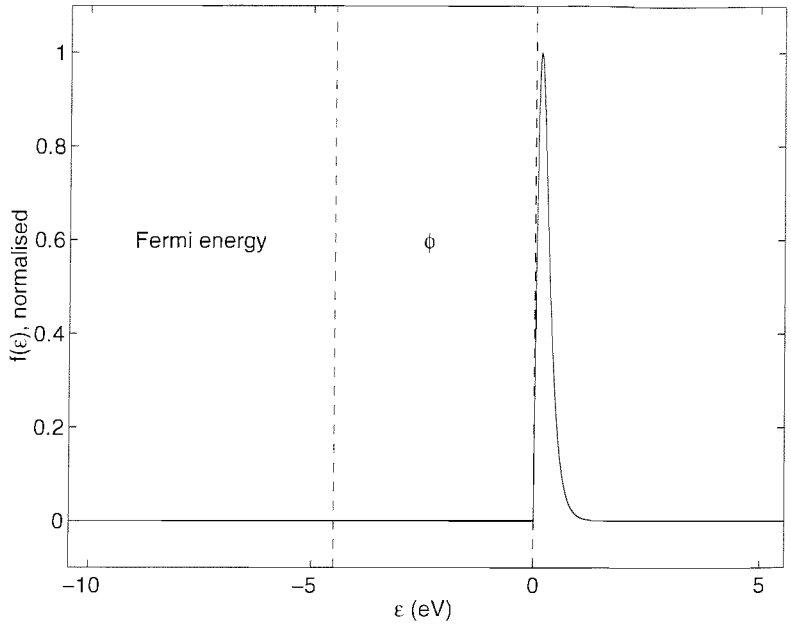

Figure 2.5: Electron energy distribution for electrons that are successfully emitted. Plotted using equation (2.8) using Tungsten as the material ($\phi = 4.5$ eV) at 1500 K. There is assumed to be no external applied field ($E_s = 0$), so this example is representative of pure thermionic emission. Note: the energy distribution is essentially the tail of a Fermi-Dirac distribution, and additionally that the bulk of the emitted electrons have energies a few times higher that the thermal energy of the emitting surface.

the the energies of the majority of emitted electrons tend to be a few times the thermal energy of the emitter (1500 K, equivalent to 0.13 eV in this example). This is consistent with both the few experimental observations that exist, see Hirsch [44] for example, and the alternative theoretical analysis of Terpigor'yev [103].

A final point with respect to figure 2.5 is that the plot is intentionally constructed on a scale that gives a qualitative impression of the comparative magnitudes of the Fermi, work function and emission energies. Unfortunately, at this scale it is difficult to make out the energy of electrons for which ϵ is only slightly above V_{max}. For such electrons, as eq. 2.8 shows, as $\epsilon \rightarrow 0$, so $f(E) \rightarrow 0$ as well. In practical terms, this means that there will be a vanishingly small number of electrons emitted with very low energies. This means that the energy distribution of emitted electrons is neither a pure Fermi-Dirac shape, nor a Maxwell-Boltzmann shape.

2.3.2 Secondary electron

Electron emission can occur as a result of particles impacting the emitting surface. This secondary electron emission (SEE) takes the form of either emission due to heavy particle bombardment (usually ions) or that due to electron impact. A comprehensive

review that covers both electron and ion impact is given by Schou [92]. The degree of SEE is generally measured in terms of the yield δ_{se}, which is the number of emitted electrons per external impact. In stable discharges (i.e. electron and/or ion impact cases) the value of δ_{se} will have an effect on the absolute current emitted from a conducting surface. For transient cases, a significant δ_{se} can lead to violent cascade effects: the breakdown characteristics of plasma discharges will change significantly for small variation in δ_{se}. The value of SEE δ_{se} is a function of both emitter material and impact energy, although to a first approximation, for similar materials, δ_{se} can be considered to be related to impact energy only.

When an energetic ion (KeV range) strikes a surface, a cascading recoil in the material atoms occurs. During this cascade, some bound electrons may become excited and disassociate. This is the underlying process that leads to electron emission [92]. Chalise et al. [19] recently conducted experiments with an ion gun to measure secondary electron emission. The setup used Helium ions and a stainless steel cold cathode. The authors state that secondary electron emission due to kinetic ion impact does not occur below ion velocities of around 10^7 cm s^{-1}; equivalent to about 200 eV. This correlates with authors such as Carter [17] and Rudwan [88] who, while discussing the possibility of SEE contributing to hollow cathode current, estimate that δ_{se} will be much less than 0.02 for ion impact energies in the 10 eV range (for Carbon).

In contrast to the single (atomic cascade) mechanism for ion-induced SEE, there are two ways in which electron emission occurs under electron impact. The direct form of emission occurs when a high energy electron ionises a surface atom. The second form is more complex and occurs when a combination of electron scattering, conduction and excitation occurs due to electron impact [92]. In terms of practical data, Ordonez and Peterkin Jr. [76] give a review of current theory surrounding SEE for a variety of materials. As is the case for ion impact, SEE yields only become significant at high impact energies. For typical metals, it is shown that $\delta_{se} < 1$ for electron impact energies below 100 eV. For the energy ranges encountered in the hollow cathode, up to 25 eV depending on cathode-keeper voltage, δ_{se} would be much less than 0.1. It is likely, however, that secondary electron emission due to electron bombardment may be significant in the hollow cathode. This is because the insert coating, that contains Barium Oxide, is known to exhibit very high δ_{se}, exceeding 9 in some cases at 500 V [109].

2.3.3 ion wall neutralisation

There is a final method by which current may be drawn from a conductor to a plasma, albeit indirectly. When an ion strikes a conducting wall, there is a possibility that

it may be neutralised. There are two mechanisms by which this can occur. First, the ion may be neutralised on contact with a metal atom if an Auger electron is present at the time of ion impact. In metals, at ion impact energies in the 1-100 eV range, the probability of this state occuring is typically around 0.5, although the exact chemistry of the surface will have a significant effect on the rate of Auger electron ion neutralisation. The second method by which neutralisation can take place is via ion tunneling. If the incident ion has sufficient energy to overcome the potential barrier at the metal surface, then direct chemcial reaction can occur between the ion and metal atoms. Again, this method clearly depends heavily on the reactivity of the ion and metal atoms, although in this case there is a minimum energy below which the neutralisation cannot take place. It is generally reasonable approximate the cut-off energy as being the same as the work function of the metal under consideration.

In a similar way to the above discussion regarding secondary electron emission, the complexity of solid atomic structure and complex chemistry is replaced by a coefficient representing the yield or neutralisation probability, and this value can be determined experimentally. The neutralisation probability, δ_n, is then a coefficient with value in the range $[0, 1]$. For insulators, or materials that will not tend to for chemical bonds. δ_n can be effectively zero. In the case of metals, δ_n could be high, although is typically around 0.5.

Unfortunately, there is very little research available in the literature that covers the value of δ_n. This is probably because in many cases the value is unimportant compared to other forms of electron emission and ion loss. Another reason for the apparent lack of data for this value is the difficulty of making experimental measurements to quantify it. Barbeau and Jolly [4] state that it is possible to take measurements about neutrals generated via ion neutralisation. However, in their experiments, Barbeau and Jolly are only able to take measurements of returned neutrals in excited states: this means that their data takes only high energy ion collisions into account. Further, in the case of the sheath expected in a hollow cathode, electron energy gain in the sheath may cause collision ionisation of free neutrals, so that despite the neutralisation rate, there is also an ion source within the sheath. Considering the sheath and surface together, the overall neutralisation yield may be small. Other authors of numerical simulation codes for low pressure plasma have expressed dismay at the absence of such data; Nanbu et al. [72] for instance, in simulations of etch rate in semiconductor processing plasmas tried cases for $\delta_n = 0.1$ and $\delta_n = 0.5$ and found significant differences in the macroscopic results of the simulation. Clearly, the way to proceed in this case is to be aware that the effect of varying δ_n must be characterised fully and the results of the characterisation taken into account when results are presented.

This concludes the discussion of electron emission mechanisms. The following

section proceeds to analyse and compare the collision processes that occur within the bulk plasma of the hollow cathode.

2.4 Mean free path analysis

Particles in a plasma undergo a wide variety of collisions. In order to produce a realistic model of the hollow cathode plasma, the dominant collision types need to be included. By extrapolating mean free paths based on collision cross sections, an assessment of the more important collision types can be made.

The reaction rate between two species can be written

$$R_{12} = n_1 n_2 v_{12} \sigma_{12}, \tag{2.10}$$

where, for collision between species 1 and 2, n is density, v_{12} is relative velocity and σ is collision cross section. This leads to another definition of collision rate, namely that the interaction frequency of species 1 with species 2 is

$$\nu_{12} = \frac{R_{12}}{n_1} = n_2 v_{12} \sigma_{12}. \tag{2.11}$$

From this, we can define approximate expressions that are useful for estimating mean free path. In particular, depending on the comparative magnitudes of v_1 and v_2, there will be three cases:

1. $v_1 \gg v_2 \Rightarrow v_{12} \approx v_1$

$$\lambda = \frac{v_1}{\nu_{12}} \approx \frac{v_1}{n_2 v_1 \sigma_{12}} = \frac{1}{\sigma_{12} n_2}, \tag{2.12}$$

2. $v_1 \approx v_2$

$$\lambda \approx \frac{1}{\sigma_{12} n_2}, \tag{2.13}$$

3. $v_2 \gg v_1 \Rightarrow v_{12} \approx v_2$

$$\lambda = \frac{v_1}{\nu_{12}} \approx \frac{v_1}{n_2 v_2 \sigma_{12}}. \tag{2.14}$$

Before the mean free path analysis of the various collision types can begin, some reference parameters need to be set. It is acceptable to assume that the most collisional region of the entire hollow cathode flow will be the emitting plasma region in the insert. Of course, data on this region is sparse; it is precisely the objective of this research to learn more about this region. For the purposes of this analysis, approximate experimental estimates suffice. Re-examining table 2.1 we can see that the upper estimates for plasma density are of the order 10^{20} m^{-3} (Siegfried and Wilbur [95] or Monterde et al. [63]). From the same table of experimental data, it is clear that

the electron temperature should be 1 eV. We choose to set the ion temperature to 1 eV also, and the neutral temperature equal to the wall temperature, set to 1700 K (0.15 eV). This results in a Debye length of $\lambda_D \approx 1$ μm. But what of the neutral density? Again, it is one of the objectives of this research to predict the neutral density. Since estimates of ionisation fraction based on experimental data vary so widely, there is little consensus regarding neutral density. If it were assumed that the ionisation fraction is 1, then $n_i = n_n$ would be implied. It is clear, however, from analyses of particle fluxes at hollow cathode mass flow rate values, that the density is likely to be perhaps 10^{22} m^{-3}. In any case, an upper estimate will tend of overestimate the importance of collisions involving neutrals, which is preferable to ignoring potentially relevant collision processes. In the following discussion, it should be noted that while these conditions are for the bulk plasma, the energies and densities will be different within the sheath. In particular. the electron energy will be higher than 1 eV and that, by definition, $n_i \neq n_e$.

The final point to make is to specify how we rate a mean free path as denoting collisions that are significant. If, on average, a particle undergoes one collision of a certain type while moving from one side of the insert section of the cathode to the other, then it has $\lambda = 2$ mm. Thus, if λ is significantly bigger than 2 mm, then the collision is insignificant, if it is much less than 2 mm it is significant.

2.4.1 Electron-Neutral

There are three types of collision to consider when examining electron-neutral collisions: elastic momentum transfer, inelastic momentum transfer (excitation) and ionisation. These can be seen in figure 2.6. In the figure, only the excitation level with the highest σ is shown. It is a standard trick to include on the dominant few excitation levels in cases such as this (in fact, there is insufficient data available on the lower levels for Xenon).

Elastic momentum transfer

For elastic collisions between electrons and neutrals, case one applies because $v_e \gg v_n$. The expression is

$$\lambda_{en}^{elastic} \approx \frac{1}{\sigma_{en} n_n}, \tag{2.15}$$

where the maximum elastic $\sigma_{en}^{elastic}$ at a temperature of of 1 eV (from fig. 2.6) is $\approx 10^{-18}$ m^2 and $n_n \approx 10^{22}$ m^{-3}. In the figure, very low energy electrons have much higher cross-sections; we assume that there are few of these, even in the thermal

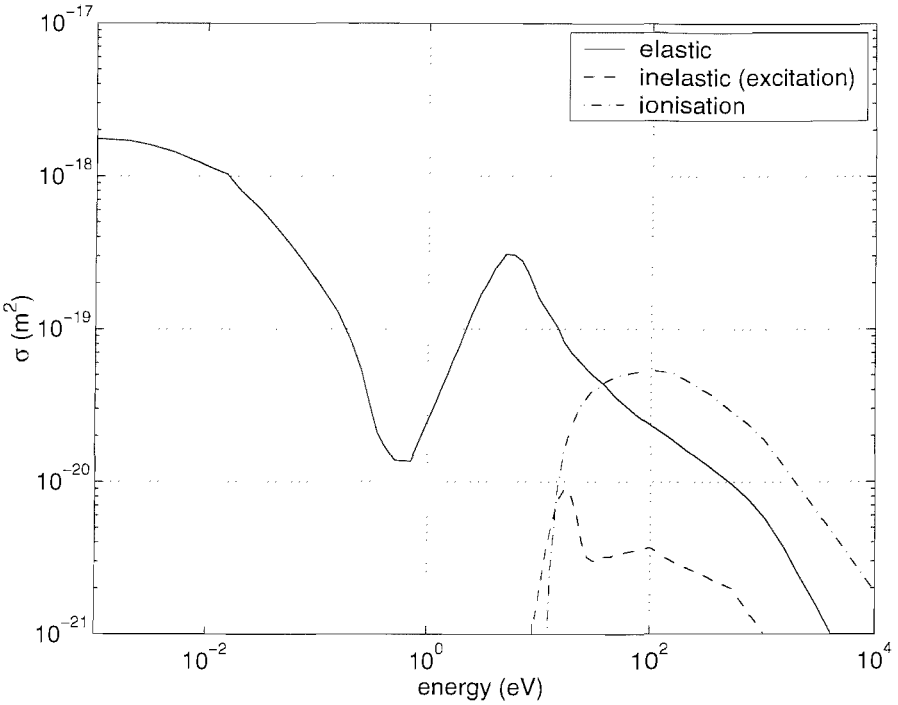

Figure 2.6: e-Xe collision cross sections vs. electron energy. The data is from Puech and Mizzi [82]

plasma. The path length is

$$\lambda_{en}^{elastic} \approx \frac{1}{10^{-18} \cdot 10^{22}} \approx 0.1 \text{ mm}. \tag{2.16}$$

Clearly, many collisions may take place within the volume of the maximum plasma density region. This means we expect momentum transfer between the electrons and neutrals to be significant, and such collisions should be included in the model.

Inelastic momentum transfer

Examining figure 2.6 again, it is clear that the excitation collision level shown has a cross section an order of magnitude lower than even the ionisation collision. The same expression applies as before, except that the maximum $\sigma_{en}^{excitation}$ is $\approx 10^{-20}$ m^2, so that

$$\lambda_{en}^{excitation} \approx \frac{1}{10^{-20} \cdot 10^{22}} \approx 10 \text{ mm}. \tag{2.17}$$

As expected, this is much larger, and if this collision were particularly difficult to evaluate, it could be omitted. It is however wise to include excitation for two reasons. First, many authors suggest that ionisation within the cathode emitting region occurs is a multilevel phenomena where the neutrals are first excited before being ionised. This suggests that this collision type would be more important within the sheath and

so should be included. Secondly, as will be seen in the discussion of the implementation of such collisions, if *elastic* momentum transfer collisions are evaluated (previous subsection), then it becomes trivial to do so for inelastic collisions.

Ionisation

Since we plan to develop a model that is self consistent in the generation of a plasma, we must have some provision for ionisation. This means that ionisation collisions have to be included, and we would expect the mean free path to be small enough else the analysis is incorrect. Figure 2.6 shows that the mean free path is an order of magnitude larger than that for the dominant excitation:

$$\lambda_{en}^{ionisation} \approx \frac{1}{10^{-19} \cdot 10^{22}} \approx 1 \text{ mm}. \tag{2.18}$$

The analysis verifies this: ionisation should be included. Indeed, this verifies the assumption that the neutral density needs to be as high as 10^{22} within the cathode for internal ionisation to take place at all (although there is some debate as to the amount of ionisation that takes place in the hollow cathode interior). Clearly, the density may be much lower externally, but the distance over which collisions need to occur is much larger (the cathode-anode separation distance).

2.4.2 Charged particle

There are four different types of collisions that involve exclusively charged participants. In this section the various types of Coulomb collision are considered, before analysing the significance of recombination collisions. In order to evaluate Coulomb collisions, a couple of auxiliary constants need to be calculated. The 90° scattering impact parameter, b_0 is given by

$$b_0 = \frac{e^2}{4\pi\epsilon_0 m v^2}, \tag{2.19}$$

for single charged species (i.e. assuming $Z = 1$). We will also need the value of the ratio of the maximum to 90° impact parameters:

$$\Lambda \equiv \frac{b_{max}}{b_0} \approx \frac{\lambda_D}{b_0}, \tag{2.20}$$

so that the Coulomb logarithm, $\ln \Lambda$, can be calculated. The above approximation to Λ is given by Goldston and Rutherford [40] and is valid where $n\lambda_D^3 \gg 1$; i.e. a plasma. For reference parameters, $n\lambda_D^3 \approx 100$. In a departure from the method used for other collision types, it makes more sense for Coulomb collisions to consider the

mean collision rate, $\langle\nu\rangle$, instead of collision cross section. This is because Coulomb collisions 'always occur' by their nature. It is still possible to derive an effective mean free path so that the significance can be compared to other collision types.

Electron-ion Coulomb

From the definition of mean free path, $\lambda_{ei} = v_e/\nu_{ei}$. An expression for the mean electron-ion collision frequency is given by [40]

$$\langle\nu_{ei}\rangle = \frac{2^{\frac{1}{2}} n_i e^4 \ln\Lambda}{12\pi^{\frac{3}{2}}\epsilon_0^2 m_e^{\frac{1}{2}}(eT_e)^{\frac{3}{2}}}, \tag{2.21}$$

where T_e is in units of eV. The 90° scattering impact parameter for this case is

$$b_0 = \frac{e^2}{4\pi\epsilon_0 m v^2} = \frac{e^2}{4\cdot\pi\cdot\epsilon_0\cdot 9.1\times10^{-31}\cdot 600000^2} \approx 7.0\times10^{-10}\text{ m}. \tag{2.22}$$

Further, the Coulomb logarithm is then $\ln\Lambda \approx \ln(\lambda_D/b_0) = 7.26$. The collision rate is then

$$\langle\nu_{ei}\rangle = \frac{2^{\frac{1}{2}}\cdot10^{20}\cdot e^4\cdot 7.26}{12\cdot\pi^{\frac{3}{2}}\cdot\epsilon_0^2\cdot\sqrt{9.1\times10^{-31}}e^{\frac{3}{2}}} \approx 2.1\times10^9, \tag{2.23}$$

so that the mean free path becomes

$$\lambda_{ei}^{coulomb} = \frac{v_e}{\nu_{ei}} = \frac{600000}{2.1\times10^9} = 0.3\text{ mm} \tag{2.24}$$

As expected, this is then significant; electron-ion Coulomb collisions should be included. If electron-ion collisions are significant and the plasma is both quasi-neutral and has similar electron and ion temperatures, then we should expect electron-electron and ion-ion collisions to be significant. An assessment of this follows.

Electron-electron Coulomb

The same method used above to develop the electron-ion Coulomb collision mean free path could be applied here, except that now an approximation can be introduced to simplify the analysis [40]:

$$\langle\nu_{ee}\rangle \approx \frac{n_e e^4 \ln\Lambda}{\epsilon_0^2 m_e^{\frac{1}{2}}(eT_e)^{\frac{3}{2}}} \approx \frac{\langle\nu_{ei}\rangle}{n_i/n_e}, \tag{2.25}$$

so that

$$\langle\nu_{ee}\rangle \approx \frac{2.1\times10^9}{10^{20}/10^{20}} = 2.1\times10^9. \tag{2.26}$$

So the Coulomb collision frequency $\nu_{ee} \approx \nu_{ei}$. The mean free path is the same: $\lambda_{ee}^{coulomb} = v_e/\nu_{ee} \approx 0.3$ mm; if electron-ion Coulomb collisions are included, then so should electron-electron collisions. This is not necessarily generally true; it is true in this case because the plasma is both quasi-neutral, and has single charge electrons and single charge ions (or at least $n_i^{2+} \ll n_i$).

Ion-ion Coulomb

The ion-ion collision frequency for single charged ions is given by [40]

$$\langle \nu_{ii} \rangle = \frac{n_i e^4 \ln \Lambda}{12\pi^{\frac{3}{2}} \epsilon_0^2 m_i^{\frac{1}{2}} (eT_i)^{\frac{3}{2}}}. \tag{2.27}$$

If this is compared to equation (2.21), it is clear that since

$$\langle \nu_{ei} \rangle m_e^{\frac{1}{2}} (eT_e)^{\frac{3}{2}} / \sqrt{2} = \langle \nu_{ii} \rangle m_i^{\frac{1}{2}} (eT_i)^{\frac{3}{2}}, \tag{2.28}$$

then the ion-ion rate in terms of the electron-electron rate is

$$\langle \nu_{ii} \rangle = \frac{\langle \nu_{ei} \rangle}{\sqrt{2}} \sqrt{\frac{m_e}{m_i}} \left(\frac{T_e}{T_i} \right)^{\frac{3}{2}} = \frac{2.1 \times 10^9}{\sqrt{2}} \cdot \sqrt{\frac{9.1 \times 10^{-31}}{218 \times 10^{-28}}} = 3 \times 10^6. \tag{2.29}$$

This means that the mean free path is $\lambda_{ii}^{Coulomb} = v_i/\nu_{ii} \approx 0.4$ mm. This is around a third larger than the electron-electron mean free path, but is still significant as predicted.

Recombination

There are two types of recombination collision to consider: three body and radiative. It is reasonable to suggest that if we are considering a plasma in which some types two-body collisions occurring so infrequently that they can be ignored, then the reaction $e + e + A^+ \Rightarrow e + A$ is likely to be very insignificant. Such an assumption is confirmed by Goldston and Rutherford [40], who state that in the few-eV range, the three body recombination rate will never exceed that of radiative recombination for plasmas of density lower than 10^{22} m^{-3}. Further, *multistage* three body recombination ($e + A \Rightarrow A^-$, then $A^+ + A^- \Rightarrow A$) will not be considered because the attachment coefficient that specifies the number of electron-neutral collisions that must take place to create one negative ion is infinite for the inert gases: only inert gas propellants are used in this work. Radiative recombination is likely to be far more likely as the collision relies on only one electron and one ion being in close proximity: $e + A^+ \Rightarrow A + p$, where p is a photon. Regardless of the intricacies of the collisions, it turns out that there is an alternative argument based on experimental observations that can help predict

the comparative significance of recombinations. Following the logic of Szabo [100], the relative importance of recombination can be assessed in terms of a comparison of the recombination rate with the ion and neutral fluxes. Mitchner and Kruger [61] state that for a bulk recombination rate given by

$$\frac{dn_e}{dt} = -\alpha n_e n_i, \tag{2.30}$$

the constant α is

$$\alpha = 1.09 \times 10^{-20} n_e T^{-\frac{9}{2}}, \tag{2.31}$$

although there is some debate over the accuracy of this expression for very high temperatures [100]. For the above parameters ($T_e = 1$ eV and $n_i = n_e = 10^{20}$ m^{-3}),

$$\alpha = 1.09 \times 10^{-20} \cdot 10^{20} \left(\frac{e}{k}\right)^{-\frac{9}{2}} = 5.60 \times 10^{-19}, \tag{2.32}$$

so

$$\frac{dn_e}{dt} = -5.60 \times 10^{-19} \cdot 10^{20} \cdot 10^{20} = -5.6 \times 10^{21} \text{ m}^{-3} \text{ s}^{-1}. \tag{2.33}$$

Consider the maximum plasma density region within the cathode: a cylinder of radius 1 mm and length 4 mm, for example. The volume of this shape is 12.6×10^{-9} m^3, so the recombination rate within the cathode is approximately -7.1×10^{13} s^{-1}. For Xenon, this is equivalent to a neutral mass flux of 15.4×10^{-6} mg s^{-1}. This is about five orders of magnitude smaller than typical cathode propellant flow rates. Further, even if the ionisation fraction is as small as 10^{-3}, typical ion fluxes within the plasma will still be several orders of magnitude larger than the recombination mass flux. Based on this logic, recombination collisions need not necessarily be included as part of a cathode model. It turns out, however, that a simple radiative recombination method based on the rate quoted above is trivial to implement in the context of the other collision types that are necessary. For this reason, radiative recombination is included.

2.4.3 Neutral-Neutral

To evaluate the mean free path for neutral-neutral collisions, the cross section can be approximated to πd^2, where d is the molecular diameter. Of course, more elaborate expressions for cross section exist, but this suffices for now. Xenon has a molecular diameter of 5.74×10^{-10} m, so $\sigma_{nn} \approx 10^{-18}$ m^2. Using case two this means that

$$\lambda_{nn} \approx \frac{1}{10^{-18} \cdot 10^{22}} \approx 0.1 \text{ mm}, \tag{2.34}$$

so it appears that neutral-neutral collisions should be included. It should be noted, however, that depending on the mass flux and orifice dimensions, neutral-neutral collisions could become very insignificant (leading to nearly collisionless flow).

A discussion of dissociation phenomena in high temperature neutral-neutral collisions is moot because (a) we consider only mono-atomic insert gas species, so that vibrational dissociation of non-trivial molecules does not apply and (b) the energies of the neutral particles will never be high enough to cause ionisation; only electron-neutral collisions involve such energies, since electrons may gain energy in the anode-cathode fall and sheaths. The second point raised could be debated on the grounds that ions may gain high energies and, via charge exchange, transfer such energies to neutrals. Even taking this into account, it remains unreasonable to suggest that based on experimental estimates of ion energies and charge exchange rates, many neutrals will gain energy equivalent to their own ionisation energy. This issue of ion-neutral collisions is now addressed.

2.4.4 Ion-Neutral

There are two types of ion-neutral collision to analyse: momentum transfer and charge exchange. The former is similar to neutral-neutral momentum transfer considered above, with the exception that fact that the two species may not be in thermal equilibrium with one another. The latter involves the reaction

$$A_{slow} + A^+_{fast} \Leftrightarrow A_{fast} + A^+_{slow}, \tag{2.35}$$

and an alternative method must be used for the cross section.

Momentum transfer

The significance of neutral-neutral collisions does not imply significance of ion-neutral collisions, although, of course, ignoring one implies ignoring the other; for example, even if the ions do not perturb the neutrals to a great extent, if the neutrals have a large impact on ion dynamics then the collision must be included. The asymmetry in the significance of collisions can stem from different densities ($n_n \gg n_i$ in this case) or possibly that the two species are not in thermal equilibrium ($T_i \neq T_n$). In light of this, the analysis for both $T_i = T_n$ and $T_i \neq T_n$ is considered. As for the neutral-neutral case, the collision cross section can be approximated based on molecular diameter. For Xenon, $\sigma_{ni} = \sigma_{in} = \pi d^2 \approx 10^{-18}$ m^2. For the case of no thermal equilibrium, at 1700 K, the neutral $v_{th} = \sqrt{2kT/m} \approx 500$ m s^{-1} and at 1

eV, the ion $v_{th} = \sqrt{2kT/m} \approx 1200$ m s^{-1}. This means $v_i \approx 3v_n$, so

$$\lambda_{ni} = \frac{v_1}{\nu_{12}} = \frac{v_n}{\sigma_{12} v_i n_i} \approx \frac{1}{3}\frac{1}{\sigma_{12} n_i} \approx \frac{1}{3}\frac{1}{10^{-18} \cdot 10^{20}} \approx 3.3 \text{ mm}, \tag{2.36}$$

and

$$\lambda_{in} = \frac{v_1}{\nu_{12}} = \frac{v_i}{\sigma_{12} v_n n_n} \approx \frac{3}{\sigma_{12} n_n} \approx \frac{3}{10^{-18} \cdot 10^{22}} \approx 0.3 \text{ mm}. \tag{2.37}$$

Clearly, then, in the presence of higher energy ions, neutral-ion collisions become more frequent. Given that even faster ions may be found in sheath regions, it would be wise to include neutral-ion momentum transfer; it should be noted that ion-neutral collisions are a border-line case that depend heavily on the ionisation fraction n_i/n_n. The result of this analysis is interesting in the context of the interior cathode dynamics. Due to the high ratio n_n/n_i, it seems that the ion species may tend to be forced toward the neutral temperature and bulk velocity within a short distance inside the cathode (a millimetre, perhaps), while the comparatively low ion density means that ion heating of the neutral gas may not be very significant. Of course, it is questions such as these that are intended to be examined in the results section later. At this stage all that is required is to say that if question of ion-neutral dynamics is important (as it may be), then ion-neutral collisions need to be included. Although we have concluded that ion-neutral collisions are relevant already, it should be remembered that there is still charge exchange to consider.

Charge exchange

The classical expression describing the variation of resonant charge exchange (CEX) collision cross section for is given by Rapp and Francis [83] as

$$\sigma_{cex} = (k_1 \ln(v_i) + k_2)^2 \times 10^{-20} \text{ m}^2, \tag{2.38}$$

where v_i is the ion velocity and for Xenon the constants are: $k_1 = -0.8821$ Å^2 and $k_2 = 15.1262$ Å^2. In fact, there is a more recent model based on new experimental results given by Miller et al. [60]. For the purposes of this discussion we use the Rapp and Francis model since the cross sections given by the two methods do not differ considerably for the parameters specified. The Miller et al. model is discussed further later. For CEX collisions, $v_i \gg v_n$, by definition, so case 1 is used:

$$\lambda_{in}^{cex} \approx \frac{1}{\sigma_{cex} n_n} = \frac{1}{(-0.8821 \cdot \ln(500) + 15.1262)^2 \times 10^{-20} \cdot 10^{22}} \tag{2.39}$$

$$\approx \frac{1}{7.9 \times 10^{-19} \cdot 10^{22}} \approx 0.13 \text{ mm}, \tag{2.40}$$

Collision type	Model	ref.	σ (m^2)	λ (mm)	Included
§2.4.1 **Electron-Neutral**					
e-A (elastic)	MCC	[68]	10^{-18}	0.1	$\star$
e-A (inelastic)	MCC	[68]	10^{-20}	10	$\star$
e-A (ionisation)	MCC	[68]	10^{-19}	1	$\star$
§2.4.2 **Coulomb**					
e-e	MCC	[40]	n/a	0.3	$\star$
e-A^+	MCC	[40]	n/a	0.3	$\star$
A^+-A^+	MCC	[40]	n/a	0.4	$\star$
§2.4.2 **Recombination**					
e-A^+ (radiative)	empirical	[101]	-	-	$\star$
e-e-A^+ (3 body)	"	[101]	-	-	
§2.4.3 **Neutral-Neutral**					
A-A	DSMC	[9]	10^{-18}	0.1	$\star$
A-A (dissociation)	-	[9]	n/a	-	
§2.4.4 **Neutral-Ion**					
A-A^+	MCC	[70]	10^{-18}	0.3	$\star$
A-A^+ CEX	MCC	[70]	10^{-19}	0.13	$\star$

Table 2.5: Collision types, cross sections and mean free paths of the various collisions that can occur in the maximum density region of the hollow cathode discharge. The right-hand column denotes those collision types that are included in the numerical model; the adjacent column shows the preceding section number that contains the discussion of that collision type.

where the ion thermal velocity at 1700 K is used. Interestingly, the CEX m.f.p. will increase as the ion temperature increases. This implies that charge exchange collisions are likely to be significant in determining the interaction between neutral gas and plasma. In fact, in sheath regions where the ions density is significantly higher, it has been proposed that such collisions are integral to understanding the cathode [51].

Summary

Table 2.5 gives an overview of the collision types included in the discussion contained in previous sections. The table gives the characteristic mean free path for the cathode interior plasma and accordingly whether they need to be included in an accurate cathode model.

- All forms of electron-neutral collisions are significant, although only the excitation level(s) with the largest σ need be used. It seems that electron-neutral

elastic momentum transfer may be one of the most significant forms of collision present.

- It was shown that Coulomb collisions are relevant. The table omits information on Coulomb collision cross sections because these have little meaning: the Coulomb 'collision' is a continuous phenomena rather than a discrete binary collision. The effective mean free over which Coulomb collisions perturb the plasma was determined from the mean collision rate.

- It is possible to rule out the inclusion of recombination collisions, because the recombination mass flux is several orders of magnitude smaller than the neutral and ion fluxes.

- Neutral-neutral collisions would not be deemed important if we adhere to the assumption that the ionisation fraction is 1 and the plasma density is 10^{20}. In fact, the neutral density is likely to be two to three orders of magnitude larger, so such collisions are relevant.

- Ion-neutral momentum transfer collisions are likely to be significant only in regions of abnormally high ion energy, but are included because they are trivial to compute once neutral-neutral collisions are calculated.

- Resonant charge exchange collisions are found to be significant particularly at comparatively lower ion energies. Some authors consider resonant CEX essential to hollow cathode plasma dynamics: it would be wise to assess this assertion.

Chapter 3

Numerical Method

3.1 Direct particle simulation

The approach that shall be taken to presenting the numerical methods used in this research is one where the underlying collisionless 'billiard ball' kinetic model is first described, with the plasma and collisional elements discussed second. There is a certain degree of replication present between the work of Birdsall and Langdon [12] (collisionless plasma), Bird [9] (rarefied gas dynamics) and Hockney and Eastwood [45] (general theory of sparse many-particle systems). Underlying both the Monte Carlo models for collision selection and the particle-in-cell (PIC) model for electrical (plasma induced) forces is a basic, collisionless, charge neutral direct gas model. Presentation of this trivial collisionless charge neutral model first is useful because it removes any chance of repetition and sets up a common terminology that the following chapters that cover Monte Carlo and PIC methods can use. The greatest level of clarity is obtained by extracting the simple underlying model and then later applying the complexities of the MCC and PIC methods. Indeed, this separation of the common material from the specifics of either method is how the code is structured; in essence it is a collisionless charge neutral gas model. Subroutines for Monte Carlo collisions and plasma effects can simply be 'plugged in' to enhance the physical description as required.

Accordingly, this section covers collisionless un-charged particle kinetics, including particle transport, wall interactions, sampling and numerical methods used to replicate equilibrium conditions in a set of particle velocities and positions.

3.1.1 Particle model

The collisionless charge neutral model of particle dynamics is really painfully trivial. It is simply the case of points moving around in a domain, the only complexity is

that which involves the interaction with surfaces and inlet/exit boundaries. Particle information that needs to be stored includes position vector $\mathbf{x}$, velocity vector $\mathbf{v}$, and species s. For more complex models, data such as internal energy may also be stored. Using this information, a set of particles may be moved around in a domain.

For a particle at $\mathbf{x}_t$, we move the particle for a given time period Δt at velocity $\mathbf{v}$ to find the new position $\mathbf{x}_{(t+\Delta t)}$:

$$\mathbf{x}_{t+\Delta t} = \mathbf{x}_t + \mathbf{v}\Delta t \tag{3.1}$$

The exact value of Δt is irrelevant in this section; clear limits will be described according to the stability of the collisional or plasma parameters described later. Motion in Cartesian space is trivial,

$$\mathbf{x}_{t+\Delta t} = \begin{pmatrix} x_t + u\Delta t \\ y_t + v\Delta t \\ z_t + w\Delta t \end{pmatrix}, \tag{3.2}$$

movement in cylindrical coordinates with rotational symmetry is not.

Particle motion in axially symmetric cylindrical coordinates

Birdsall and Langdon [12, §14.12] give information regarding position advance in cylindrical co-ordinates, but not for axi-symmetry. Bird [9] provides an entire chapter on the subject, including validation cases for his method. The exact procedure is now presented.

A particle initially at r_1, z_1 and, by axi-symmetry $\theta = 0$, moves with with velocity u_1, v_1, w_1. Motion in the $\hat{\mathbf{z}}$ direction is as before:

$$z_{t+\Delta t} = z_t + u_1\Delta t. \tag{3.3}$$

The problem is associated with motion in the r–θ plane: the particle has an initial circumferential velocity component, w_1 that will take the particle out of the plane of symmetry.

Figure 3.1 shows a schematic of the movement of a particle from r_1 to r. The movement out of plane by a distance $w_1\Delta t$ has the effect of increasing the new radial position r. This is not a problem: evaluation of the new radial position r can be seen on the diagram. The problem lies in the fact that the new velocity vector now lies out of the plane of symmetry.

This means that before the move is complete, v_1, w_1 must be rotated through θ so that they again lie in the r–z plane. This is accomplished using the standard method

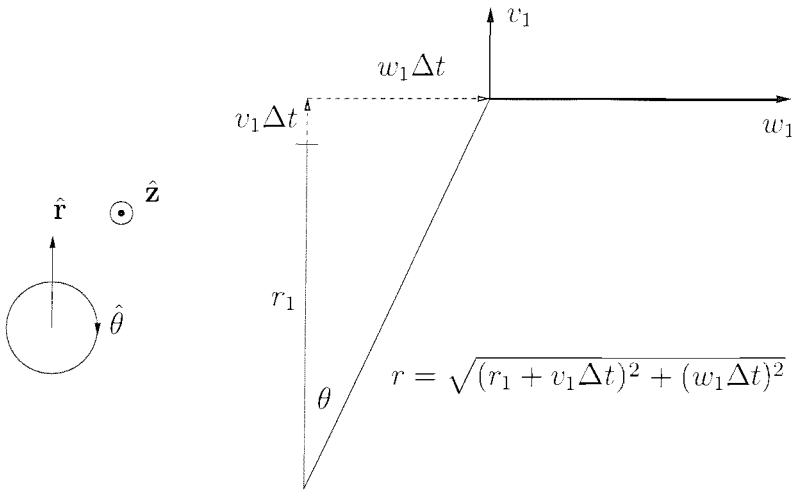

Figure 3.1: Movement of a particle in the r–θ plane.

for rotation of a vector through angle ϕ [107, §14.5]:

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \cos\phi & \sin\phi \\ -\sin\phi & \cos\phi \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \tag{3.4}$$

Hence, the new velocity components become

$$v = v_1 \cos\theta + w_1 \sin\theta, \tag{3.5}$$

$$w = -v_1 \sin\theta + w_1 \cos\theta. \tag{3.6}$$

This form is not particularly useful in the context of a numerical simulation: it is best to remove all explicit evaluations of trigonometric functions. Examining figure 3.1 again, notice that

$$\sin\theta = w_1\Delta t/r, \tag{3.7}$$

$$\cos\theta = r_1 + v_1\Delta t/r. \tag{3.8}$$

The final expression for the rotation of the velocity components back into the plane of symmetry becomes

$$\mathbf{v} = \begin{pmatrix} u_1 \\ (v_1(r_1 + v_1\Delta t) + w_1^2\Delta t)/r \\ (-v_1 w_1\Delta t + w_1(r_1 + v_1\Delta t))/r \end{pmatrix} \tag{3.9}$$

where, from fig. 3.1

$$r = \sqrt{(r_1 + v_1\Delta t)^2 + (w_1\Delta t)^2}. \tag{3.10}$$

Notice in this final equation that the condition $r \geq 0$ is always maintained regardless of the magnitude or direction of v_1 and w_1. Movement of particles in cylindrical coordinates with rotational symmetry can be entirely prescribed using eq. (3.9). Results presented in this document do not extend to coordinate systems beyond axially symmetric cylindrical polar.

Surface interaction

Surface interaction is an important element of particle simulations as they provide boundary conditions to the simulation. The properties and behaviour of various boundary interactions are now presented. For the purposes of the ensuing discussion, the terms 'wall' and 'surface' are used interchangeably. The term 'boundary' is a broader expression not necessarily imply a solid surface/wall. Within the context of wall collisions, all surfaces conserve mass: for any collision the mass influx must be the same as the mass out-flux. The only exception to this is that in a plasma model, walls that absorb a net current also by definition absorb a net mass flux.

The thermal accommodation coefficient, a_c is a measure of the extent by which a particle assumes the energy of the surface that it collides with. Specifically,

$$a_c = \frac{q_i - q_r}{q_i - q_w}, \tag{3.11}$$

where q_i and q_r are the incident and reflected energy fluxes and q_w represents the energy flux that would be carried from the wall if the wall and gas were in thermal equilibrium. The TMAC is a function of the wall material. If the surface is irregular on a microscopic level, a_c is closer to unity. If the surface is very smooth, the a_c becomes smaller. Typically, $a_c \approx 0.9$. In particle simulations, implementation $a_c = 0$ is simple: the velocity and position relative to the particle direction are simply reversed. In the case of fully diffuse reflection, the particle is given a randomly sampled Maxwellian velocity based on the wall temperature (for isothermal walls) or based on the particle's energy (for adiabatic walls). If a_c is between 0 and 1, then the reflection is diffuse with probability a_c; this approach is valid for many particles and is a method used by many authors [9 101].

Surfaces where $a_c = 0$ are very useful in particle simulations. These surfaces involve zero energy and mass flux across them. This means that the gas conditions on one side of such a surface are as though there is an identical gas occupying the other side. These surfaces are referred to as specularly reflecting surfaces and are powerful because they can be used to represent planes of symmetry. Figure 3.2 provides a qualitative view of the difference between diffuse and specular surfaces. In a plasma simulation, if a specularly reflecting surface is coupled with a Neumann boundary

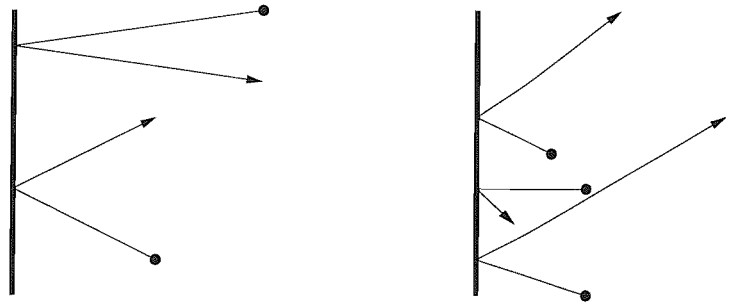

(a) Specular, $a_c = 0$ (b) Diffuse, $a_c \approx 1$

Figure 3.2: Illustration of wall collisions at different thermal accommodation coefficient.

for the field equation(s), then a plane of symmetry can be reproduced. Further, in a plasma context, $a_c = 0$ boundaries are literally zero current density boundaries because the reflected velocity is simply the incident velocity in reverse direction.

Sampling

The procedures for displaying results from particle simulations are fairly trivial. At the most basic level, a scatter plot of particle positions can be produced, although this serves little purpose other than a qualitative impression of density. Generally, macroscopic quantities are extracted using a mesh. Particle positions are interpolated onto the mesh and data extracted. There is usually no need to define a mesh explicitly for particle sampling when one is available due to the PIC or MCC algorithms. It is possible to sample many macroscopic quantities by storing only five pieces of data per cell per species:

- Number of real particles, n,equivalent to F_N, the super-particle ratio).
- Velocity components, u_z, u_r, u_c.
- Square of velocity magnitude, $V^2 = u_z^2 + u_r^2 + u_c^2$.

For each particle in a cell the above data is stored. When enough samples have been taken, typical macroscopic quantities can be calculated. For a sample of N *computational* particles of species a, these are shown in table 3.1. Further data can be extrapolated from these fundamental values. Procedures for a gas mixture are not relevant in this document, but are not significantly more complex; see Bird [9, section 1.5].

One form of information that it is possible to extract from particle simulations is velocity probability density functions (or electron energy distribution functions).

Property	Expression	
Density	$n_a = n/NV_C$	m^{-3}
Mass density	$\rho_a = nm/NV_C$	$\mathrm{kg\ m}^{-3}$
Kinetic energy	$E_a = \frac{1}{2}mV^2/N$	J
Stream velocities	u_i/N	$\mathrm{m\ s}^{-1}$
Temperature	$V_0^2 = (u_z/N)^2 + (u_r/N)^2 + (u_c/N)^2$	
	$T_a = m(V^2 - V_0^2)/3k$	K

Table 3.1: Extracting macroscopic fluid properties from particle data. m is molecular mass of species a. Further properties can be extrapolated from these using gas laws, etc.

When a traditional Navier-Stokes CFD code is used, velocity PDFs are assumed to be a perfect Maxwellian, by definition. Indeed, it can be argued that if PDFs that are perfectly Maxwellian are extracted from a particle simulation, then the user should switch to a Navier-Stokes code! PDFs, be they of particle velocities or electron energies, are extracted in a similar way to that described above. A volume of interest is selected, and particles are sampled from that volume. The difference is that particles are delineated by their energy into data bins rather than by position into spatial cells. It should finally be noted that if sufficient particles are used in a simulation, the PDFs can be extracted directly from the particle data at the post-processing stage: the regions of interest need not be specified at run-time.

Finally, it is useful to record data concerning particle impact on surfaces. This allows both shear and normal data to be used for the purposes of estimating heat flux and surface forces.

Particle interaction algorithms

While Birdsall and Langdon [12] provide cover plasma particle modelling (PIC), and Bird [9] provides collisional particle simulation (DSMC), it is Hockney and Eastwood [45] who give an overview of particle simulation from a 'neutral' perspective. This view allows for a more concise definition of the computational problems and solutions that exist when particle interaction is to be considered. The two fundamental algorithms in their most simplistic forms are now presented:

Particle-particle collision: Consider calculating collisions between a collection of n particles contained in a domain as previously described. If we know the collision cross section of the particles, then it is possible to use billiard ball mathematics to determine collisions over a period of time. First, we would evaluate the trajectories of all the particles. Second, for each trajectory, we would test for intersection with every other. Any intersecting trajectories constitute collisions. This algorithm is fundamentally

inefficient for large n. The number of operations per time-step can be expressed as the sum of the trajectory calculation phase (n operations) and the intersection calculation $\frac{1}{2}(n-1)n$, a total of $n + \frac{1}{2}n(n-1)$ operations; such an algorithm scales approximately as $O(n + n(n-1))$.

Self consistent particle field interaction: In the case of collisionless charged particles, a similar method naturally presents itself: based on basic electrostatic theory, the total force on each particle is the sum of the forces exerted by all others. So on each time-step, the force exerted by each particle is calculated by evaluating the force from all the others. It is clear to see that in this case, such calculations scale as $O(n^2)$ for n particles.

Since, for reasons such as sampling and numerical stability it is best to use many particles (large n), something more efficient than the above $O(n^2)$ 'particle-particle' (PP) methods discussed. The methods discussed in the following sections are called 'particle-mesh' (PM), because the n^2 relationship is replaced with more tractable computational costs (typically something like $n \log n$) by utilising some more other auxiliary mesh based algorithm. Individual descriptions of the methods (DSMC, MCC, PIC) are contained in the following sections. Hockney and Eastwood [45] do note the possibility of using hybrid 'particle-particle-particle-mesh' (P^3M) methods for greater accuracy, although this has not been extensively researched.

3.2 Particle distributions

The fundamentals of the direct particle approach to the solution of the Boltzmann-Vlasov equation have now been covered. Essentially, we are modelling a collection of many real particles using a few simulated particles and relying on the statistical nature of the solution being the same. In future sections, the collision term modelling and plasma-dynamics will be covered, but it is first necessary to consider the method by which simulated particles can be generated based on macroscopic 'external' definitions such as temperature. Prior to presenting methods for generating the statistical representation of the plasma, it is first necessary to discuss random numbers.

Pseudo random number generator

Computers have trouble generating truly random numbers. Instead, they are able to produce sequences that are truly random, but that are identical in that the same sequence can be repeatedly generated. This is a problem in cryptography, where truly random numbers are required, but no problem at all for numerical simulations: we

are not concerned over the repeatability characteristics of a sequence, but rather the randomness of the sequence itself. In light of this, a high quality *pseudo* random number generator (PRNG) is a necessity for the statistical validity in numerical models that use Monte Carlo techniques. The PRNG needs to provide a very long sequence of random numbers that are uniformly distributed and uncorrelated at low computational cost. A popular free PRNG, used for all the results presented in this document, is the 'Mersenne Twister' of Matsumoto and Nishimura [59]. This generator produces a sequence that is random until the 2^{19937}th number. Once initialised, when the generator is called, it returns a number which is equivalent to the occurrence of R in the following discussions. R is a uniformly distributed random number in the range $[0, 1]$. Clearly, it is important that we have confidence in such a generator; it would be unwise to use one without checking for correctness. Some tests for the generator of Matsumoto and Nishimura [59] are contained later in the validation section (§4.1).

Random Maxwellian velocity

A random sample from a Maxwellian velocity distribution of a given temperature can be constructed based on the statistical analysis in either [9, App. C] or [12, §16.4]. For a temperature T, the most probable thermal velocity is

$$v_{mp} = \sqrt{\frac{2kT}{m}}, \tag{3.12}$$

where m is the particle mass. The magnitude of the 2-D Maxwellian velocity $\mathbf{v}$ for this temperature can be found via

$$|\mathbf{v}| = v = v_{mp}\sqrt{-\ln(R)} \tag{3.13}$$

where R is as usual a uniformly distributed random number of range $[0, 1]$. From v, it is then possible to extract any two components of velocity that are normal to each other using the following:

$$\begin{aligned} \theta &= 2\pi R \\ v_1 &= v \sin\theta \\ v_2 &= v \cos\theta \end{aligned} \tag{3.14}$$

where R is another (fresh) random number. Since particle velocities are three dimensional vectors, this method is first applied twice to obtain v_{1-4}, then the particle given $v_x = v_1, v_y = v_2, v_z = v_3$, while v_4 is carried over so that the above method is applied only once for the next particle. In this way, only one R is required per particle velocity component (so in general three calls to the random number generator are made to construct one Maxwellian particle).

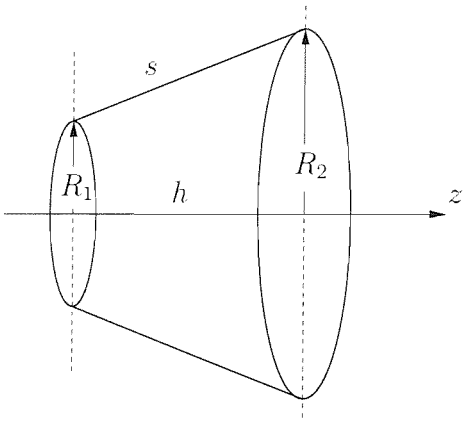

Figure 3.3: A conic frustum.

Random particles on objects in cylindrical co-ordinates

Generation of particle positions randomly on the surface of a cylinder that extends from z_0 to z_1 at radius r_c is trivial:

$$z = z_0 + hR \tag{3.15}$$

$$r = r_c \tag{3.16}$$

$$\theta = 2\pi R \tag{3.17}$$

$h = z_1 - z_0$ is the length of the cylinder. Recall that R is a uniform random number on $[0, 1]$ To generate a uniform covering of random positions on the surface of a disk, the result of the derivation presented in Bird [9, Appendix C] is used:

$$r = \sqrt{a^2 + (b^2 - a^2)R} \tag{3.18}$$

where a and b represent the radii of two circles making up the (hollow) disk, $b > a$. z is constant and if needed, θ is as eq. 3.15. For a non-hollow disk simply use $a = 0$.

Figure 3.3 shows a conic frustum. For this shape, standard results for s, the total area and the description of the shape $r(z)$ are

$$s = \sqrt{(R_1 - R_2)^2 + h^2}, \tag{3.19}$$

$$A = \pi(R_1 + R_2)s, \tag{3.20}$$

$$r(z) = R_1 + (R_2 - R_1)\frac{z}{h}. \tag{3.21}$$

Shapes such as disks, cones and cylinders can be described as 'flat'. Although counter-intuitive, this description refers to the idea that they can be constructed from a flat plane (a cylinder is a rolled up rectangle or square) [54]. Hence, to obtain our required

uniform random point, we 'flatten' the conic frustum so that it is trivial to select the z point using eq (3.18). r follows using eq. (3.19). Hence, a uniformly distributed random point on a conic frustum is simply:

$$z(z_0, z_1) = \sqrt{z_0^2 + (z_1^2 - z_0^2)R} \tag{3.22}$$

$$r(z) = R_1 + (R_2 - R_1)\frac{z}{h} \tag{3.23}$$

$$\theta = 2\pi R \tag{3.24}$$

where $h = z_1 - z_0 > 0$. These techniques are used when generating uniformly injected particles from a given surface.

3.3 Direct Simulation Monte Carlo

Nanbu [68] chooses to classify inter-particle collision algorithms into three classes according to the nature of the particles involved: electrons, e, neutrals atoms A and ionised atoms A^+;

1. Short range collisions between unlike particles e.g. e-A, A-A^+ CEX,
2. Short range collisions between like particles e.g. A-A, A-A^+ momentum transfer,
3. Coulomb collisions (between charged particles) e.g. e-e, e-A^+ etc...

Neutral gas flow consists entirely of type 2 collisions. Further, the particular algorithm associated with these collisions—Bird's DSMC—is perhaps the most elegant and efficient of the various collision methods.

The basic ideas behind DSMC were first introduced by Bird [6]. At this stage, the fundamental molecular models were simply those inherited from classical kinetic gas theory. Only with the introduction of more realistic molecular models such as the Variable Hard Sphere that accounted for variable viscosity with temperature did the model begin to show near perfect accuracy. A combination of Bird's inability to prove that the DSMC provided a solution to the Boltzmann equation, coupled with insufficient computing power, held back the method until the early 1990s. Doubts concerning the theoretical validity of the model were cast away when Wagner [104] demonstrated that in the limit of zero time-step and mesh spacing, the DSMC formulation represents an exact solution to the Boltzmann equation. Since this time, the method has been embraced by many as the standard tool for rarefied gas dynamics. The number of references presented previously relating to thruster plume modelling gives only a small idea of the range of application of the technique. As has been mentioned, the DSMC has recently been successfully applied to micro gas dynamics, a comparatively recent but rapidly expanding field. There are two comprehensive review papers covering the method: Bird [7] and more recently Oran et al. [75]. The definitive work on the method remains the book *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Bird [9]. A good review of the various experimental validations of the DSMC is given by Harvey and Gallis [42].

An introduction to the core numerics of the DSMC is now presented. This consists of two primary parts: the method for sampling collision events and the method for determining post-collision properties based on the molecular description. The first part introduces the 'no time counter' (NTC) sampling method of Bird [9]. During the early stages of the evolution of the method (1970s), a slightly different method was used: the 'time counter'. There advantages of the NTC are such that the TC is only of historical interest; in particular, the NTC uses a constant time-step and, as will

be seen in the following section, allows for the prediction of the number of colliding particles in a cell explicitly. The NTC is also inherently vectorisable. The second part of the introduction to the DSMC covers the molecular models that are used describe collision cross section and derive post collision properties.

3.3.1 DSMC collision sampling procedure

Consider the simulation space described in the early part of this chapter: a collection of super-particles contained within the simulation volume, moved periodically at some interval Δt, interacting with boundaries. This is the basic collisionless simulation. The simulation domain is first subdivided into cells. Consider a cell with volume V_C; contained within are N super-particles. Each super-particle represents F_N real particles. Hence, there are $F_N N$ real particles contained in V_C. The average *number density* in the cell is $n = F_N N / V_C$.

Consider a test particles moving with velocity $\mathbf{g}$ relative to another. For the purposes of the discussions regarding collisions in the following sections, $\mathbf{g} = \mathbf{v}_p - \mathbf{v}_q$ is the relative velocity between particles p and q. g is as usual simply $|\mathbf{g}|$. The collision probability between p and q during time Δt is equal to the volume swept out by their total cross section divided by the cell volume:

$$P = F_N \sigma_T g \frac{\Delta t}{V_C}. \tag{3.25}$$

Now, in theory the full set of collisions could be evaluated by selecting all $N(N-1)/2$ potential collision pairs in the cell and calculating the collisions with probability P each time. Since P is usually very small, this method is inefficient and the best approach is only a fraction of the pairs are tested and the probability of each collision increased by dividing eq. (3.25) by the fraction. The fraction is given by

$$P_{max} = F_N (\sigma_T g)_{max} \frac{\Delta t}{V_C}. \tag{3.26}$$

The number of pair selections is obtained by multiplying this by $N^2/2$. In fact, for a steady flow, since N is varying, it is best to multiply by $\bar{N}N$, where $\bar{N}$ is a time averaged value of cell population. The method is then that

$$\frac{1}{2} N \bar{N} F_N (\sigma_T g)_{max} \frac{\Delta t}{V_C} \tag{3.27}$$

prospective collision pairs are selected in the time-step, and collisions are evaluated with probability

$$\frac{\sigma_T g}{(\sigma_T g)_{max}}. \tag{3.28}$$

It is worth noting that the extent to which $\bar{N}F_N$ (which is equivalent to $\langle n \rangle$ in the cell) is averaged depends on whether the flow is steady or not. In the validation case for the DSMC given later, an unstable flow is used, so no time averaging is used. In general, if the gas consists of species of similar mass and energy, the above procedure holds. In the case of a single species plasma, the ion mass is effectively identical to the neutral mass, but the ion energy may be significantly different to the neutral energy. In this case, it is worth using the DSMC procedure for gas mixtures, where,

$$\frac{1}{2} N_p \bar{N}_q F_N \{(\sigma_T g)_{max}\}_{pq} \frac{\Delta t}{V_C} \tag{3.29}$$

collisions are selected between species p and q and collision probability evaluated as

$$\frac{\sigma_T g}{\{(\sigma_T g)_{max}\}_{pq}}. \tag{3.30}$$

The collision models that describe the evaluation of σ_T and the post-collision properties of a pair of molecules now follow.

3.3.2 Collision Models

Once a collision has been selected, there are various methods to resolve collision mechanics. The the most basic is the hard sphere model, which was superseded by the Variable Hard Sphere model of Bird [8], and refined further as the Variable Soft Sphere model of Koura and Matsumoto [53]. All three are now presented and compared.

Hard sphere model

The hard sphere model is a basic description of molecular interaction defined so that the inter-molecular force during collision becomes effective at distance

$$r = \frac{1}{2}(d_1 + d_2) = d_{12} \tag{3.31}$$

where d_i are molecular diameters. r is physically the apse line between the centres of the colliding spheres shown in figure 3.4. χ is the deflection angle due to the collision, it is this angle that allows us to calculate the post-collision velocities. In the figure, the distance b is

$$b = d_{12} \sin \theta_A = d_{12} \cos(\chi/2) \tag{3.32}$$

and

$$\left| \frac{db}{d\chi} \right| = \frac{1}{2} d_{12} \sin(\chi/2). \tag{3.33}$$

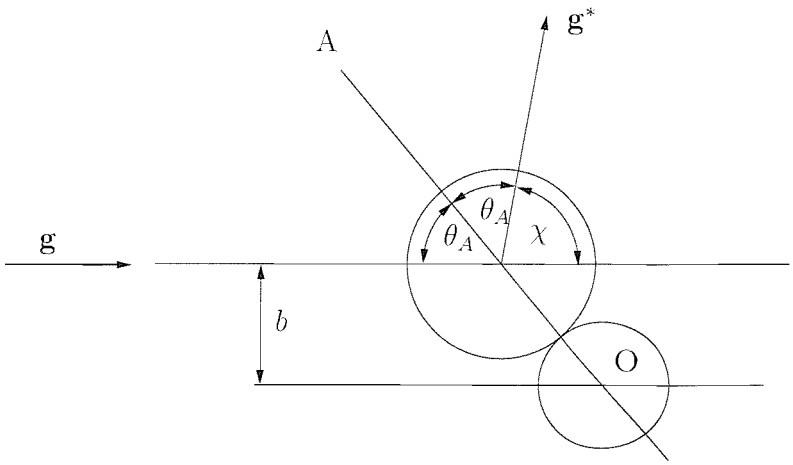

Figure 3.4: Collision geometry of hard sphere molecules (from Bird [9]). $\mathbf{g}$ is the pre-collision relative velocity (notice that the co-ordinate system is aligned with the frame of reference due to $\mathbf{g}$). χ is the resulting deflection angle and $\mathbf{g}^*$ is the post-collision velocity.

Now, from basic theory of binary elastic collisions, the collision cross section, σ, is defined as

$$\sigma = \frac{b}{\sin\chi}\left|\frac{db}{d\chi}\right| \tag{3.34}$$

[9, eq. 2.13], so that

$$\sigma = \frac{b}{\sin\chi}\frac{1}{2}d_{12}\sin(\chi/2) = \frac{1}{4}d_{12}^2. \tag{3.35}$$

Hence, the cross section is independent of χ: the scattering is isotropic in the centre of mass frame of reference (as can be inferred by examination of fig. 3.4). The total cross section is defined as

$$\sigma_T \equiv \int_0^{4\pi} \sigma d\Omega, \tag{3.36}$$

where $d\Omega$ is the unit solid angle around vector $\mathbf{g}^*$. The total cross section for the hard sphere collision model is then

$$\sigma_T = \frac{1}{4}d_{12}^2 \int_0^{4\pi} d\Omega = \pi d_{12}^2. \tag{3.37}$$

Variable Hard Sphere (VHS) model

Although the hard sphere model is very simple, it is unrealistic in two respects. First, the scattering is real gas collisions is not necessarily isotropic. This does not represent a major problem, as for simple molecules changes in the scattering law are not significant. Of greater importance is the fact that for a real gas, the effective cross section varies with relative velocity and relative translational energy. In light of these

Gas	Degrees of freedom, ξ	mass, m	viscosity coefficient, μ	viscosity index, ω	diameter, d
		10^{-27} kg	10^{-5} N s m^{-2}		10^{-10} m
Helium	3	6.65	1.865	0.66	2.33
Neon	3	33.5	2.975	0.66	2.77
Argon	3	66.3	2.117	0.81	4.17
Krypton	3	139.1	2.328	0.8	4.76
Xenon	3	218.	2.107	0.85	5.74

Table 3.2: VHS parameters for the inert gases. Diameters are referenced to a temperature of 273 K, properties are given for standard conditions (101,325 Pa and 0°C). Degrees of freedom is 3 in all cases since all the elements are insert monatomic species: the ratio of the specific heats is then $\gamma = (\xi + 2)/\xi = (3 + 2)/3 = 1\frac{2}{3}$. Radon, the most massive of the insert gas species, is omitted for two reasons: (a) there is no gas dynamic data available and (b) it is not available in sufficient quantities to be a viable propellant for the hollow cathode.

objections, the variable hard sphere (VHS) model was introduced by Bird [8]. The VHS model allows d to vary with g, using an inverse power law

$$d = d_{ref} \left(\frac{g_{r.ref}}{g} \right)^{\nu} \tag{3.38}$$

where subscript $_{ref}$ refers to reference values. Via the Chapman-Enskog viscosity and basic kinetic theory, it can be shown that the molecular diameter can be related to the translational energy by

$$d = \left[\frac{(15/8)(m/\pi)^{\frac{1}{2}}(kT_{ref})^{\omega}}{\Gamma(9/2 - \omega)\mu_{ref} E_t^{\omega - \frac{1}{2}}} \right]^{\frac{1}{2}} \tag{3.39}$$

Where ω is the viscosity index and E_t is the translational energy of the molecule. μ_{ref} and T_{ref} are tabulated for many gases. In particular, see table 3.2 for values for the inert gases considered in this document. The collision deflection angle is given in the same way as the hard sphere case

$$\chi = 2\cos^{-1}(b/d), \tag{3.40}$$

except that now d is a function of g so the scattering is no longer isotropic.

Variable Soft Sphere (VSS) model

The VSS was introduced by Koura and Matsumoto [53] to address the fact that the VHS still inherits the deflection angle of the hard sphere model. It is identical to the

VHS in all other respects. The deflection angle is given as

$$\chi = 2\cos^{-1}[(b/d)^{1/\alpha}] \tag{3.41}$$

where α is the VSS coefficient that can be related directly to the Schmidt number by [9, eq. (3.74,3.76,A2)]

$$\mathrm{Sc} = \frac{2+\alpha}{(3/5)(7-2\omega)\alpha}, \tag{3.42}$$

where ω is the viscosity index. The VSS is less efficient than the VHS and produces only marginally better results under most circumstances. Gimelshein et al. [39], compared the various collision models for the case of a plasma reactor, and conclude that in the case of a plasma the exact reproduction of polyatomic rarefied gas flow does not effect the overall plasma behaviour: plasma dynamic effects dominated. Further, the gases that are considered in this work are exclusively mono-atomic inert species (Kr, Ar. and Xe; perhaps a mixture two) i.e. they have only translational degrees of freedom. This means that for our purposes, the VSS gains almost nothing, at some additional cost (in mathematical operations per collision).

3.3.3 Producing a free stream flow in the DSMC

Previously, it was noted that the code written for this research was designed so that inflow conditions for gases are set as a mass flow rate so that the results can be easily compared to experimental data (this is because mass flux can be measured in a laboratory accurately and easily). This method is valid for internal fluid dynamics, but the DSMC is frequently used for external flows, as in the case of this example. While we will be considering internal flows exclusively in this work, there is still a case where 'free stream' conditions need to be generated: the downstream boundary. For the purposes of experimental comparison, 'vacuum' chambers tend to have an ambient particle density that could be significant. This is particularly true for experimental data taken from the hollow cathode running within the ion thruster assembly as the discharge chamber acts to contain a plasma at a density much higher than the ambient vacuum chamber density. The method for determining the particle flux through boundaries given free stream conditions is now briefly presented.

Starting with an expression for the equilibrium Maxwellian velocity distribution function, written in the notation of Bird [9] as

$$f(v) = \frac{\beta^3}{\pi^{\frac{3}{2}}} \exp(-\beta^2 v^2), \tag{3.43}$$

where Bird tends to use the abbreviation for $\sqrt{m/2kT}$ terms by defining

$$\beta = \sqrt{\frac{m}{2kT}}, \tag{3.44}$$

as must be the case by comparison of eq. (3.43) to any standard text on statistical mechanics. Second, the *most probable* thermal velocity, the maximum of $g(v)$, turns out to be simply the inverse of β:

$$v_{mp} \equiv v_{th} = \sqrt{\frac{2kT}{m}} \equiv \frac{1}{\beta}. \tag{3.45}$$

By examining fluxal properties in a gas, it is possible to write an expression for the *inward number flux* to a surface element due to a gas of density n as

$$N_i = \frac{n}{2\beta\sqrt{\pi}} \left\{ \exp(-s^2 \cos^2\theta) + \sqrt{\pi} s \cos\theta [1 + \operatorname{erf}(s\cos\theta)] \right\}. \tag{3.46}$$

This expression is the result of a lengthy integration that is omitted for brevity, see Bird [9, §4.2] for further detail. θ is the angle between the element normal vector and the stream velocity vector and s is the molecular speed ratio:

$$s = u_\infty \beta = \frac{u_\infty}{v_{mp}}. \tag{3.47}$$

N_i is simply the number of particles that pass through an element of unit area per unit time. This expression is used to determine the particle flux from uniform free stream conditions. Equation (3.46) is of course much easier to implement when boundaries are orthogonal to $\mathbf{u}_\infty$. For a boundary perpendicular to the stream velocity, $\theta = 0$, so

$$N_i = \frac{n}{2\beta\sqrt{\pi}} \left\{ \exp(-s^2) + \sqrt{\pi} s [1 + \operatorname{erf}(s)] \right\}, \tag{3.48}$$

and for a boundary parallel to the stream velocity, $\theta = \pi/2$, (or a stationary gas so $s = 0$),

$$N_i = \frac{n}{2\beta\sqrt{\pi}}. \tag{3.49}$$

This concludes the discussion on how a free stream flow is affected in a DSMC simulation.

3.4 Monte Carlo Collisions (MCC)

This section covers both electron-atom collisions and Coulomb collisions. The theory behind the collision selection and post-collision particle properties is covered by the

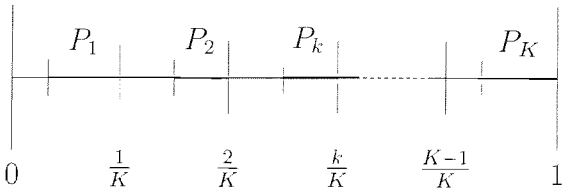

Figure 3.5: Geometric representation of Nanbu's technique for the sampling of an electron-atom collision event.

excellent review paper of Nanbu [68], which incidentally also includes a good but brief discussion of the DSMC method.

3.4.1 e-A collisions: Event selection

Due to the complex nature of electron-molecule collision cross section variation with energy, it is necessary to use a database of experimental data rather than a numerical model such as the VHS. This greatly increases confidence in electron-ion and electron-neutral collisions compared to approximating the electron super particles as a DSMC species. The MCC system follows the selection methodology found in Nanbu [68]. The model allows for an arbitrary number of electron collision events (momentum transfer, atomic excitation, ionisation) to be used. Nanbu's modification allows that both the event, and the collision probability of that event, be selected using a single random number. The constant time-step method is presented here: this technique is analogous to the DSMC NTC technique. The probability of an electron collision with a background gas in time Δt is $P_T = \nu \Delta t$, where $\nu = n_g v \sigma_T$ is the collision rate. n_g is the background gas number density (which we already have from our DSMC background flow) and v is the velocity of the electron under consideration. Moreover, if there are K possible collisional events (elastic, inelastic (excitation) or ionisation), then

$$P_T = \sum_{k=1}^{K} P_k, \tag{3.50}$$

where P_k, the probability of the kth event occurring is given by

$$P_k = n_g v \sigma_k \Delta t. \tag{3.51}$$

Nanbu [68] presented a method by which both the event and collision probability can be selected from a single random number, increasing efficiency. Equation 3.50 is rewritten

$$1 = P_T + (1 - P_T) = \sum_{k=1}^{K} \left[P_k + \left(\frac{1}{K} - P_k \right) \right], \tag{3.52}$$

as represented in figure 3.5.

Sample a random number A between 0 and 1 (in figure 3.5). If A lands in one of the right hand parts, collision event k occurs. If A is located in a left hand part of any interval, no collision event takes place. The integer part of $1 + AK$ determines the interval, k, in which A falls. Now, only P_k need be evaluated to determine whether the collision occurs. The cross section of event k occurring at velocity v is sampled from a database of experimental data (an example is given later in the results section for electron-Xenon collision cross section).

This technique allows the modelling of ionisation, recombination and excitation as well elastic momentum transfer from the electron cloud to the gas. Knowledge of atomic excitation raises the prospect of extracting data based on photon emission from super particles in the model for comparison with experimental data.

3.4.2 e-A collisions: Post collision properties

For electron with velocity $\mathbf{v}$ and molecule with velocity $\mathbf{V}$, the relative pre-collision velocity is $\mathbf{g} = \mathbf{v} - \mathbf{V}$. Now, since in general $v \gg V$, the magnitude of the relative velocity g is effectively just v. The differential collision cross section is a function of the deflection angle χ and velocity $g \approx v$: $\sigma(v, \chi)$. The integral cross section $\sigma_k(v)$ is defined as [68]

$$\sigma_k(v) = \int_0^{2\pi} d\psi \int_0^{\pi} \sigma(v, \chi) \sin\chi \; d\chi, \tag{3.53}$$

where $\sin\chi \; d\chi d\psi$ is the solid angle into which the post-collision velocity of the electron is directed, χ is the polar angle and ψ the azimuthal angle. Post-collision properties are denoted with a prime: electron post-collision velocity is $\mathbf{v}'$. The probability that $\mathbf{v}'$ is in the solid angle $\sin\chi \; d\chi d\psi$ is

$$\sigma \sin\chi \; d\chi d\psi / \sigma_k, \tag{3.54}$$

the product of $(2\pi)^{-1} d\psi$ and $2\pi(\sigma/\sigma_k)\sin\chi \; d\chi$. This means that the coefficient $(2\pi)^{-1}$ of $d\psi$ is the probability density function of ψ. A random sample of ψ is obtained as simply

$$2\pi R. \tag{3.55}$$

Similarly, the coefficient $2\pi(\sigma/\sigma_k)\sin\chi$ of $d\chi$ is the probability density function of the deflection angle χ. In this case, a random sample of χ can be obtained from

$$\frac{2\pi}{\sigma_k(v)} \int_0^{\chi} \sigma(v, \chi') \sin\chi' \; d\chi' = R. \tag{3.56}$$

Hence, χ is a function of v and R. Unfortunately there is insufficient data on *differential* cross sections available for most collisions. The only option left in the absence

of such data is to assume that $\sigma(v, \chi)$ does not depend on χ. Extensive data exists for the collision cross section in the $\sigma(v)$ form, this is discussed in detail later in this section. This means that (3.53) reduces to

$$\sigma_k = 4\pi\sigma. \tag{3.57}$$

Further, this means that the post-collision velocity $\mathbf{v}'$ is in a random direction: such random scattering is *isotropic*. Such an assumption has been shown not to be the source of significant error [72 84]. Hence, assuming isotropic scattering and using fundamental conservation laws, it is possible to derive expressions from which post-collision properties extracted.

Elastic and Inelastic momentum Transfer

The masses of the electron and atom are m and M respectively. For a collision that is either elastic (post-collision energy equals pre-collision energy) or some form of exciting interaction that involves energy E, we have conservation of momentum and energy as

$$m\mathbf{v}' + M\mathbf{V}' = m\mathbf{v} + M\mathbf{V}, \tag{3.58}$$

$$\frac{1}{2}\mu(\mathbf{v}' - \mathbf{V}')^2 + E = \frac{1}{2}\mu(\mathbf{v} - \mathbf{V})^2 \tag{3.59}$$

respectively, where $E = 0$ for elastic momentum transfer and μ is the *reduced mass* in the same sense as referred to in the DSMC method; $\mu = mM/(m + M)$. The energy equation, (3.59), becomes

$$\mathbf{v}' - \mathbf{V}' = \left[(\mathbf{v} - \mathbf{V})^2 - \frac{2E}{\mu}\right]^{\frac{1}{2}} \mathbf{R} \tag{3.60}$$

where $\mathbf{R}$ is a vector in a uniformly random direction with $|\mathbf{R}| = 1$. Put another another way, $\mathbf{R}$ describes a random point on the surface of a sphere of unit radius that has its centre at the origin. This results in isotropic scattering of the relative pre-collision velocities. The momentum equation (3.58) is rearranged to

$$\mathbf{v}' = \mathbf{v} + \frac{M}{m}\mathbf{V} - \frac{M}{m}\mathbf{V}', \tag{3.61}$$

which, substituted into (3.60) gives

$$\mathbf{v} + \frac{M}{m}\mathbf{V} - \frac{M}{m}\mathbf{V}' - \mathbf{V}' = \left[(\mathbf{v} - \mathbf{V})^2 - \frac{2E}{\mu}\right]^{\frac{1}{2}} \mathbf{R} \tag{3.62}$$

from which $\mathbf{V}'$ can be found. $\mathbf{v}'$ follows directly from eq. (3.61). These final two equations are all that remain to be evaluated in order to determine post-collision properties for elastic momentum transfer and excitation collisions. Computation assuming three-dimensional velocity vectors is not too problematic, even for large numbers of collisions. The pre-collision relative velocity that is calculated by necessity during the event selection phase, thus the square root in the RHS of (3.60) need only be evaluated once because $g^2 = (\mathbf{v} - \mathbf{V})^2$. Indeed, in MCC code, the constant factor f is defined so that

$$f = \left[(\mathbf{v} - \mathbf{V})^2 - \frac{2E}{\mu}\right]^{\frac{1}{2}}. \tag{3.63}$$

Ionisation

For an e-A ionising collision, we are considering the reaction

$$e(\mathbf{v}) + A(\mathbf{V}) \;\Rightarrow\; e(\mathbf{v}') + e_p(\mathbf{v}_p') + A^+(\mathbf{V}'), \tag{3.64}$$

where the particles are written in terms of their velocities, primes denote post-collision velocity and e_p is electron freed from the atom during the ionisation. This collision can only occur if the electron energy based on $\mathbf{v}$ exceeds the ionisation energy of the ion. This is necessarily the case because the collision cross section $\sigma(E)$ evaluates to zero if $E(|\mathbf{v}|)$ is too small. In the same way as above, it is possible to write equations for conservation of momentum and energy. The centre of mass velocity, $\mathbf{W}$, is defined as

$$\mathbf{W} = \frac{m\mathbf{v} + M\mathbf{V}}{m + M}. \tag{3.65}$$

If this velocity is subtracted from all velocities, then the pre-collision momentum is zero and the momentum equation is written

$$m\tilde{\mathbf{v}}' + m\tilde{\mathbf{v}}_p' + (M - m)\tilde{\mathbf{V}}' = 0, \tag{3.66}$$

where the $\tilde{}$ refers to velocities that have had $\mathbf{W}$ subtracted. Conservation of energy gives

$$\frac{1}{2}m\tilde{v}'^2 + \frac{1}{2}m\tilde{v}_p'^2 + \frac{1}{2}(M - m)\tilde{V}'^2 + E = \frac{1}{2}\mu(\mathbf{v} - \mathbf{V})^2 \tag{3.67}$$

where E is the ionisation energy. For the cases of $m \ll M$, the momentum equation (3.66) can be re-written

$$\tilde{\mathbf{V}}' = -\frac{m}{M}(\tilde{\mathbf{v}}' + \tilde{\mathbf{v}}_p') \tag{3.68}$$

If the above equation is substituted in equation (3.67) it is clear that the third term on the LHS is of negligible magnitude compared to the first two: it would contain an

additional factor of Mm. Hence, re-writing (3.67) gives

$$\frac{1}{2}m\tilde{v}'^2 + \frac{1}{2}m\tilde{v}_p'^2 = \frac{1}{2}\mu(\mathbf{v} - \mathbf{V})^2 - E(\equiv \Delta E) \tag{3.69}$$

where ΔE is the excess energy after ionisation. Generally, there is little published data regarding the division of energy between e and e_p, so it is divided randomly so that

$$\frac{1}{2}m\tilde{v}'^2 = A\Delta E \tag{3.70}$$

$$\frac{1}{2}m\tilde{v}_p'^2 = (1 - A)\Delta E \tag{3.71}$$

where A is simply a random number $A = R$. If it is assumed that $\tilde{\mathbf{v}}'$ and $\tilde{\mathbf{v}}_p'$ are isotropic, then it is possible to generate the post-collision velocities as

$$\tilde{\mathbf{v}}' = \tilde{v}'\mathbf{R}_1 \tag{3.72a}$$

$$\tilde{\mathbf{v}}_p' = \tilde{v}_p'\mathbf{R}_2 \tag{3.72b}$$

where $\mathbf{R}_{1,2}$ are a pair of independently generated uniform random unit vectors. Once $\tilde{\mathbf{v}}'$ and $\tilde{\mathbf{v}}_p'$ are obtained, $\tilde{\mathbf{V}}'$ follows from (3.68). Velocities are converted $\tilde{\mathbf{v}} \rightarrow \mathbf{v}$ by adding $\mathbf{W}$.

Some minor assumptions were made during the derivation of the methods for determining post-collision properties. In light of this, it is worth noting that results generated using the above procedures were compared to the experimental data of Nakamura and Kurachi [67] by Nanbu and Kondo [71] and showed excellent agreement. See also Nanbu [68] for further analysis and Phelps [78] for the collision cross section data for Xenon.

3.4.3 Ion-neutral Collisions

The most basic approach suggested by Nanbu [68] to the evaluation of ion neutral collisions is to treat them as though they are neutral-neutral collisions, but take into account the additional cross section for charge exchange. This means that in addition to the elastic scattering

$$A^+(\mathbf{v}_A) + B(\mathbf{v}_B) \Rightarrow A^+(\mathbf{v}_A') + B(\mathbf{v}_B'), \tag{3.73}$$

there is charge exchange (CEX), where

$$A^+(\mathbf{v}_A) + B(\mathbf{v}_B) \Rightarrow A(\mathbf{v}_A') + B^+(\mathbf{v}_B'). \tag{3.74}$$

It is reasonable to assume the the background gas is at rest and in equilibrium *when compared to the ion gas.* This assumption is valid in the case where the ions have a significantly higher temperature and higher stream velocity component compared to the background neutrals. Based on this quasi-stationary background neutral gas assumption, the probability that an ion A^+ collides with a neutral in time interval Δt is

$$P = n_B g \sigma_T(g) \Delta t \tag{3.75}$$

where n_B is the background neutral density, $g = |\mathbf{v}_A - \mathbf{v}_B|$ is the relative velocity and σ_T is the *total* collision cross section, made up of $\sigma_{n-i}^{elastic}$ and σ_{n-i}^{CEX}. For the elastic cross section, it is reasonable to use πd_{AB}^2 (the hard sphere model). The charge exchange cross section is a function of relative velocity and two models fitted to experimental data exist: the original Rapp and Francis [83] and the recent Miller et al. [60]. An expression for the former was given during the mean free path analysis in section 2.4.4. The latter takes the form

$$\sigma^{CEX} = A - B \log(E) \tag{3.76}$$

where, for Xenon, $A = 87.3$ Å, $B = 13.6$ Å and E is the ion energy $\frac{1}{2} m_A \mathbf{v}_A^2$. Although similar, the Miller et al. expression was fitted to a wider range of more accurate and recent data, so it seems wise to use this one.

For an elastic collision, the scattering is assumed to be isotropic, so,

$$\begin{aligned} \mathbf{v}'_A &= \frac{1}{m_A + m_B}(m_A \mathbf{v}_A + m_B \mathbf{v}_B + m_B g \mathbf{R}) \\ \mathbf{v}'_B &= \frac{1}{m_A + m_B}(m_A \mathbf{v}_A + m_B \mathbf{v}_B - m_A g \mathbf{R}) \end{aligned} \tag{3.77}$$

where as usual $\mathbf{R}$ is a uniformly distributed random unit vector. In the case of charge exchange, the simplest method – shown to be reasonably accurate [68] – it to use an identity switch. Using this method, $\mathbf{v}'_A = \mathbf{v}_A$ and $\mathbf{v}'_B = \mathbf{v}_B$, but charge is transferred from A to B. This is actually equivalent to changing the relative velocity from $\mathbf{v}_A - \mathbf{v}_B$ to $-(\mathbf{v}_A - \mathbf{v}_B)$ and hence, the deflection angle of the collision is π. Nanbu states that a reasonable approximation is that half the collisions be elastic and half be CEX. While this may be true under conditions where $\sigma^{CEX} \approx \sigma^{elastic}$, in this work the type of collision is determined based on the comparative magnitudes of the cross sections. The probability that a given collision is CEX is given by

$$P = \frac{\sigma^{CEX}}{\sigma_T}, \tag{3.78}$$

and clearly the probability that it is elastic is then $(1 - P)$. It is felt that this is

a better method that the half and half approximation because there are collision energies that upon substitution into eq. (3.76) yield σ^{CEX} significantly larger than $\sigma^{elastic}$. Incidentally, for Xenon under these conditions, the CEX cross section never falls to significantly *less* than the elastic cross section: CEX collision tend to dominate.

3.4.4 Coulomb Collisions

The Coulomb collision, also known as Rutherford scattering, is very different in nature compared to all of the collision types considered until this point. Indeed, the term collision is misleading because the interaction is not a binary event but a many body 'collision' that is continuous. It is nonetheless common to model Coulomb collisions in plasma simulations using a Monte Carlo methodology: the charged particles undergo a series of small angle binary collisions equivalent to the many-body force at a distance Coulomb collision. This means that the method uses exactly the same ideas behind the DSMC and e-A MCC collisions with regard to the fixed volume cell and fixed time-step over which collisions manifest themselves. Due to the continuous collision nature of the Coulomb collision, the method is somewhat easier to construct because no collision selection is required: all particles 'collide' during each time-step. This means that unlike with the DSMC and MCC electron-neutral collisions, the only section of the algorithm that needs to be presented is the determination of post-collision properties. The method is summarised by Nanbu [68], but a more complete review of the underlying theory is given in [69]. Verification of the method by examination of relaxation rates in plasmas is given by Nanbu and Kondo [71].

For collisions between particles of species α and β (the notation of Nanbu [68]), we construct two arrays, one containing the αs and the other the βs. If the cell populations are not equal, $N_\alpha \neq N_\beta$, or the particle weightings are different, then the Δt that follows can be adjusted so that all the particles collide to the correct extent (see discussion on p.987 in Nanbu [68]). Once the two arrays are constructed, we have N pairs of particles that can now experience the small angle binary collision that will make up part of the ongoing Coulomb collision. The impact parameter

$$A_{\alpha\beta} = \frac{n}{4\pi}\left(\frac{q_\alpha q_\beta}{\epsilon_0 \mu}\right)^2 \ln \Lambda_{\alpha\beta}, \tag{3.79}$$

is defined, where n is the *total* number density, $n = n_\alpha + n_\beta$, q_α the charge on particle α, μ is again the reduced mass $m_\alpha m_\beta/(m_\alpha + m_\beta)$ and $\ln \Lambda_{\alpha\beta}$ is the Coulomb logarithm pertaining to the particles in question. $\Lambda_{\alpha\beta}$ is given by

$$\Lambda_{\alpha\beta} = \frac{4\pi\epsilon_0 \mu \langle g_{\alpha\beta}^2 \rangle \lambda_D}{|q_\alpha q_\beta|}, \tag{3.80}$$

where λ_D is the electron Debye length based on the local temperature of α and β and $\langle g_{\alpha\beta}^2 \rangle$ is the mean square relative velocity ($\mathbf{g} = \mathbf{v}_\alpha - \mathbf{v}_\beta$). $\langle g_{\alpha\beta}^2 \rangle$ is the mean square of all the particles in the cell, so $A_{\alpha\beta}$ is only calculated once per cell not once per collision pair. Once we have $A_{\alpha\beta}$, it is possible to obtain the value of τ, a parameter that is used to select the scattering angle. For the ith pair of particles, τ is

$$\tau = \frac{A_{\alpha\beta}\Delta t}{|\mathbf{v}_{\alpha i} - \mathbf{v}_{\beta i}|^3}, \tag{3.81}$$

which allows us to evaluate $\cos\chi$ using

$$\cos\chi = \frac{1}{A(\tau)} \ln\left(e^{-A(\tau)} + 2R\sinh A(\tau)\right), \tag{3.82}$$

where R is as usual a uniformly distributed random number between 0 and 1 and $A(\tau)$ is such that

$$\coth A - A^{-1} = e^{-\tau}. \tag{3.83}$$

In general, to sufficient accuracy for a particle simulation, the Newton-Rhapson method [107] can be used to tabulate $A(\tau)$ between 0.01 and 3, while for $\tau < 0.01$, $A \approx 1/\tau$ and $\tau > 3$, $A \approx 3e^{-\tau}$. A verification of the tabulation used for this research is contained in the validation section; it matches that given by Nanbu [69]. Once the scattering angle, χ, is found, the post collision velocities are

$$\mathbf{v}'_\alpha = v_\alpha - \frac{m_\beta}{m_\alpha + m_\beta}[\mathbf{g}(1 - \cos\chi) + \mathbf{h}\sin\chi] \tag{3.84}$$

$$\mathbf{v}'_\beta = v_\beta + \frac{m_\alpha}{m_\alpha + m_\beta}[\mathbf{g}(1 - \cos\chi) + \mathbf{h}\sin\chi] \tag{3.85}$$

where $\mathbf{h}$ is given as

$$\mathbf{h} = \begin{pmatrix} g_\perp \cos\psi \\ -(g_x g_y \cos\psi + g g_z \sin\psi)/g_\perp \\ -(g_x g_z \cos\psi - g g_y \sin\psi)/g_\perp \end{pmatrix}, \tag{3.86}$$

where $g_\perp = \sqrt{g_y^2 + g_z^2}$. The angle ψ is the azimuthal angle of the collisional plane, $\psi = 2\pi R$. For the case of unequally weighted particles α and β, the less weighted particle always receives a new velocity, while the heavier one does with probability W_{light}/W_{heavy}. The method is can be arbitrarily extended to interaction between any pair of charged species, including $\alpha = \beta$. This means that all three of e-e, e-A^+ and A^+-A^+ are covered by the above. Nanbu verified this binary collision approximation to Coulomb collision method against a number of test cases. As would be expected, the inclusion of Coulomb collisions has a heating effect on the plasma. For more

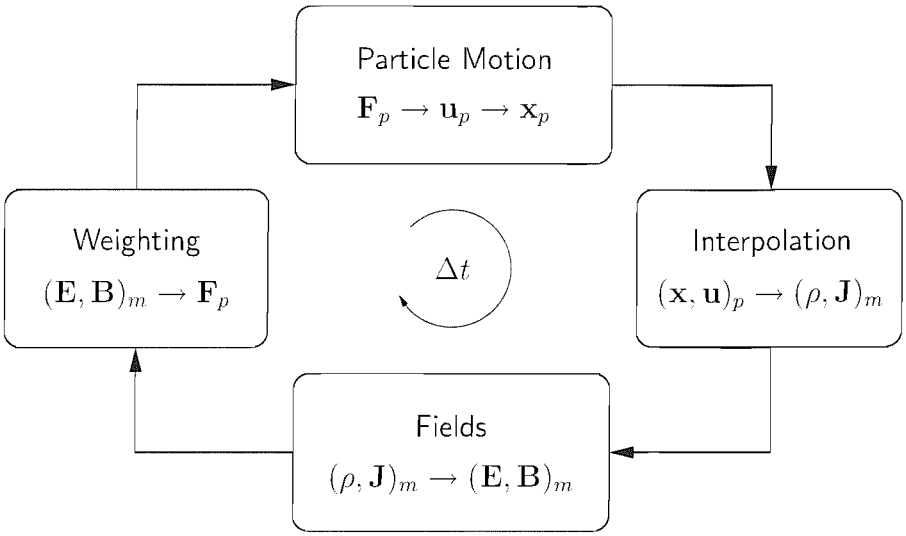

Figure 3.6: The Particle-Mesh algorithm applied to plasma physics. Subscripts $_p$ and $_m$ refer to particle and mesh information respectively. Notice that the 'Particle Motion' stage takes place entirely in 'particle space', while the field solution stage is entirely in 'mesh space'. Interpolation and weighting transfer information between mesh and particle.

details on the validation of the method see Nanbu [68].

3.5 Particle-in-Cell (PIC) Plasma model

The particle in cell plasma model is somewhat older than the DSMC. Since the 1950s, PIC computational plasma physics has been studied. It is fair to say that compared to the Monte Carlo collision particle models, the PIC model has been more widely accepted for a longer period of time. This is largely due to the fact that the underlying proof of the method has been available for longer: the relationship between the PIC model and the Vlasov equation is clearer than, for instance, that between the DSMC and the collision term of the Boltzmann equation (recall that Wagner [104] only provided mathematical proof that the DSMC is a direct solution of the Boltzmann equation as recently as 1990). The best historical review of the method is given by the most important contributor to it: Birdsall [10]. The primary text describing the method is *Plasma Physics via Computer Simulation*, Birdsall and Langdon [12].

While the Monte Carlo collision methods are referred to as particle-mesh methods because they employ meshes and rapid selection techniques to calculate particle interaction while avoiding the $O(N^2)$ problem, the PIC method could be described as a being a more pure particle-mesh algorithm, in the terminology of Hockney and Eastwood [45]. Figure 3.6 illustrates the core algorithm.

Forces on the mesh are electric and magnetic fields. These are weighted to particles

via the Lorentz force:

$$\mathbf{F} = \mathbf{F}_{\text{electric}} + \mathbf{F}_{\text{magnetic}}, \tag{3.87}$$

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}). \tag{3.88}$$

The application of such forces to the direct particle model covered in §3.1 is done using the leap-frog method. Although higher order trajectory integrations would be more accurate, data storage becomes excessive and as such all useful PIC codes known to the author have used the low-storage leap-frog trajectory integration. For each particle, then, the equations that describe motion in the presence of a force become

$$m\frac{d\mathbf{v}}{dt} = \mathbf{F}, \tag{3.89}$$

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}. \tag{3.90}$$

This covers the 'Weighting' and 'Particle Motion' stages in fig. 3.6. The leap-frog integrator has error, but since all PIC codes use it, there is extensive analysis of the accuracy of the method. Such analysis is covered later once the whole algorithm has been described.

Chronologically, the next stage in fig. 3.6 is 'Interpolation'. The simplest form of interpolation is the nearest grid point (NGP) method. The name explains the nature of this 'zeroth order' interpolation method. It has been found, however, that the stability and accuracy of PIC simulations are increased by using smoother approximations to particle charge. In particular, the first order Cloud-in-Cell (CIC) method that interpolates charge to two or more cells is often used since it is a good compromise between computational expense and charge/current density smoothing. Details regarding interpolation in cylindrical co-ordinates are given later in this section.

Finally, some form of force field equations are solve on a mesh to provide the values for $\mathbf{E}$ and $\mathbf{B}$. Typically the emphasis is on speed in the field equation solutions. Many PIC codes compromise boundary flexibility so that spectral methods can be used, although on modern computers iterative linear algebra solvers are the norm.

The following sections cover the the solution of the field equations, weighting of forces to particles, and interpolation of charge and current density, with particular emphasis on the co-ordinate system and problem we plan to solve. The field equation section is presented first, because only with knowledge of the field structure is it possible to write expressions of particle acceleration.

3.5.1 Field Equations

The electric and magnetic fields present in the plasma are governed by Maxwell's equations, expressed here in MKS units:

$$\nabla \cdot \mathbf{B} = 0. \tag{3.91}$$

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0}, \quad \text{(Poisson's equation)} \tag{3.92}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad \text{(Faraday's law)} \tag{3.93}$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{j} + \frac{1}{c^2}\frac{\partial \mathbf{E}}{\partial t}. \quad \text{(Ampère's law)} \tag{3.94}$$

where $\mathbf{E}$ and $\mathbf{B}$ are the electric and magnetic fields, $\rho = e(n_e - n_i)$ the charge density, $\mathbf{j}$ the current density, ϵ_0 the permittivity of free space, c the speed of light and μ_0 is the permeability of free space. Poisson's equation (3.92) can be expressed in terms of the electric potential using $\mathbf{E} = -\nabla\phi$, so that

$$\nabla^2 \phi = -\frac{\rho}{\epsilon_0} \tag{3.95}$$

The current density can be expressed in terms of the movement of charge so that for a plasma with ion density n_i and electron n_e (where n_i is not necessarily equal n_e),

$$\mathbf{j} = e(n_i \mathbf{v}_i - n_e \mathbf{v}_e) \tag{3.96}$$

where $\mathbf{v}_e$ and $\mathbf{v}_i$ are the drift velocities of the electron and ion species. In the PIC algorithm, the problem is that of obtaining the fields $\mathbf{E}$ and $\mathbf{B}$ given knowledge of current density and, if we wish to model transient fields, the rate of change of the fields. There are a number of simplifying assumptions that are often made to reduce the complexity of the problem.

The most broad assumption that is commonly used in PIC simulation is that the magnetic field is constant; i.e. there is no significant current density and no significantly fast changes in the electric field. This is the 'electrostatic' PIC model. Note that particles in the simulation may still be affected by a magnetic field applied by some boundary conditions, it is simply that the field does not change at all during the simulation. The vast majority of previous PIC simulations of HET thrusters and charged thruster plumes make this assumption. Clearly, for steady-state expansion of plasmas at low current density, such assumptions are acceptable.

Without introducing the full complexity and computational expense of a transient electromagnetic model, it is possible to solve for induced magnetic fields in a steady state condition. In this case, the final term in both eqs. (3.93) and (3.94) are omitted,

leaving a steady state form of Ampère's law:

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{j}. \tag{3.97}$$

This assumption is often valid for DC discharges (or very low frequency AC discharges). Using this, plus Poisson's equation, allows for a steady state model that includes induced magnetic effects. Further discussion of the relative importance of magnetic effects follow the next section, in which the structure of the fields in our chosen coordinate system is discussed.

The standard boundary condition in the solution of 3.95 are either a Dirichlet (ϕ =constant) on a solid wall, for instance, or a Neumann condition (ϕ' =constant). No other conditions are used for the solution of the Poisson problem in this work. Dirichlet conditions are applied to the cathode surfaces and anode surfaces. Neumann conditions are applied to upstream, downstream and radial extent boundaries where applicable. The axis $r = 0$ is dealt with separately and described in the following subsection.

Field equations in Axially symmetric cylindrical coordinates

In cylindrical coordinates, Curl **A** is

$$\left(\frac{\partial A_z}{r\partial\theta} - \frac{\partial A_\theta}{\partial z}\right)\hat{\mathbf{r}} + \left(\frac{\partial A_r}{\partial z} - \frac{\partial A_z}{\partial r}\right)\hat{\theta} + \left(\frac{\partial (rA_\theta)}{r\partial r} - \frac{\partial A_r}{r\partial\theta}\right)\hat{\mathbf{z}} \tag{3.98}$$

and the Laplacian of ψ is

$$\nabla^2\psi = \frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial\psi}{\partial r}\right) + \frac{1}{r^2}\frac{\partial^2\psi}{\partial\theta^2} + \frac{\partial^2\psi}{\partial z^2}. \tag{3.99}$$

(See appendix C, Kreyszig [55]). Since we are considering an axially symmetric problem, scalar quantities are simply functions of r and z: $\rho(r, z)$ and $\phi(r, z)$. There is no variation azimuthally; in the above expression, eq.(3.99), the second term vanishes:

$$\nabla^2\psi = \frac{\partial\psi}{r\partial r} + \frac{\partial^2\psi}{\partial r^2} + \frac{\partial^2\psi}{\partial z^2}. \tag{3.100}$$

The first term in eq. (3.99) has been expanded via the product rule. The Laplacian in Poisson's equation takes on this form; the result is a second order linear elliptic PDE which poses no particular problems to solve numerically. The solution method is discussed in the appropriate section later.

For the steady state induced magnetic case, things are not so simple. Essentially, we would like to express equation 3.97 so that given **j** we can find **B**. At first, this

may not seem a problem: we can obtain $j_r\hat{\mathbf{r}}$ and $j_z\hat{\mathbf{z}}$ from the particle side of the simulation, and due to rotational symmetry $j_\theta\hat{\theta} = 0$. So we can solve for $\mathbf{B}$. In fact, things seem easier when the curl of $\mathbf{B}$ is taken:

$$-\frac{\partial B_\theta}{\partial z}\hat{\mathbf{r}} + \left(\frac{\partial B_r}{\partial z} - \frac{\partial B_z}{\partial r}\right)\hat{\theta} + \frac{\partial(rB_\theta)}{r\partial r}\hat{\mathbf{z}}. \tag{3.101}$$

Here, all $\partial/\partial\theta$ terms are zero (see appendix C for differential operators in cylindrical coordinates). We may omit the $\hat{\theta}$ term due to rotational symmetry. The problem is that the equation

$$-\frac{\partial B_\theta}{\partial z}\hat{\mathbf{r}} + \frac{\partial(rB_\theta)}{r\partial r}\hat{\mathbf{z}} = j_r\hat{\mathbf{r}} + j_z\hat{\mathbf{z}} \tag{3.102}$$

is of no use. The system of PDEs in $\hat{\mathbf{r}}$ and $\hat{\mathbf{z}}$ for the solution of B_θ is inconsistent. It turns out that there is a further transformation that will yield an acceptable PDE for the solution of B_θ. The original steady state Ampère law is that

$$\nabla \times \mathbf{B} = \mu_0\mathbf{j}. \tag{3.103}$$

Now, it is possible to simply take the curl of both sides to give

$$\nabla \times \nabla \times \mathbf{B} = \mu_0\nabla \times \mathbf{j}. \tag{3.104}$$

using the vector identity

$$\nabla \times (\nabla \times \mathbf{A}) = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2\mathbf{A} \tag{3.105}$$

and since due to Maxwell's equations the magnetic field must be divergence free ($\nabla \cdot \mathbf{B} = 0$) the problem becomes

$$-\nabla^2\mathbf{B} = \mu_0\nabla \times \mathbf{j} \tag{3.106}$$

This result, expanded into cylindrical polars with the $\hat{\mathbf{r}}$ and $\hat{\mathbf{z}}$ terms removed, gives

$$\left(\nabla^2 B_\theta - \frac{B_\theta}{r^2}\right)\hat{\theta} = \mu_0\left(\frac{\partial j_r}{\partial z} - \frac{\partial j_z}{\partial r}\right)\hat{\theta}. \tag{3.107}$$

Again, B_r and B_z terms disappear because they are constant in the axi-symmetry case. Equation (3.107) is second-order linear and numerical solution is of the same level of complexity as the Poisson equation. As a check of the mathematics, it is worth considering the physical reality of (3.107). Consider an infinite current carrying wire of constant radius r (for very small r) lying on the z axis in (r, θ, z) space. The nature of the problem implies $\partial j_z/\partial z = 0$. This means that the magnetic field induced by

the wire is present, according to eq. (3.107), will be azimuthal. Further, the field is induced by the change in current density radially moving off the surface of the wire. This is entirely consistent with basic electromagnetic physics.

The above analysis has been deliberately presented in detail because the steady-state induced magnetic method is not often used in the literature. It is rather strange that such a formulation of the PIC model is not used more often. The calculation of the induced magnetic field is not significantly more troublesome than the electric field, and the additional physics would surely enhance most electrostatic PIC codes. A method similar to the steady-state induced magnetic one is mentioned although not extensively presented by the primary authors on PIC simulations [12]. Birdsall and Langdon's book is on pure plasma physics using direct numerical simulation; not practical plasma engineering. It is reasonable to conclude that the steady state induced magnetic method is of little interest to pure plasma physicists who require either a full electrostatic model or a full electromagnetic model. However, for plasma device engineering, the steady state form becomes attractive: the highly restrictive conditions of a transient electromagnetic model are removed. The nature of the stability criteria for PIC simulations is discussed once a description of each element of the algorithm is complete.

3.5.2 Weighting (particle acceleration)

Particles in a plasma experience a force that is the result of both magnetic and electrostatic element. Specifically the Lorentz force is

$$\mathbf{F} = q\left(\mathbf{E} + \mathbf{v} \times \mathbf{B}\right), \tag{3.108}$$

where q is the charge on the particle. This means that the (non-relativistic) acceleration that the particle experiences due to the field is simply

$$\mathbf{a} = \frac{q}{m}\mathbf{E} + \frac{q}{m}(\mathbf{v} \times \mathbf{B}). \tag{3.109}$$

In order to accelerate particles due to both electric and magnetic fields, Birdsall and Langdon [12] proposed a time centred scheme so that the particles experience a half acceleration due to electric field, then a rotation due to magnetic field followed by the remaining half acceleration. This is deemed acceptable because the $q(\mathbf{v} \times \mathbf{B})$ term is

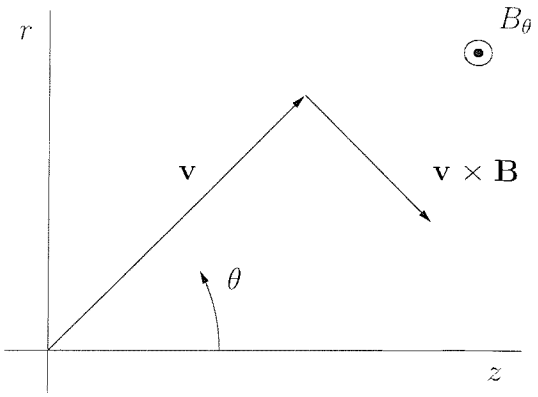

Figure 3.7: Illustration of the $\mathbf{v} \times \mathbf{B}$ rotation of velocity in cylindrical coordinates due to azimuthal magnetic field (actually the case $B_\theta > 0$). The method of Boris [13] states that $\Delta\theta = -\omega_c \Delta t$, i.e. $\dot{\theta} = -\omega_c$.

a *rotation* of $\mathbf{v}$, not a change in magnitude. Thus, the method is that

$$\mathbf{v}_{t+\Delta t/2} = \mathbf{v}_t + \frac{q}{m}\mathbf{E}\frac{\Delta t}{2}, \tag{3.110}$$

$$\mathbf{v}' = \mathbf{v}_{t+\Delta t/2} + \frac{q}{m}(\mathbf{v}_{t+\Delta t/2} \times \mathbf{B}), \tag{3.111}$$

$$\mathbf{v}_{t+\Delta t} = \mathbf{v}' + \frac{q}{m}\mathbf{E}\frac{\Delta t}{2}. \tag{3.112}$$

In the axially symmetric co-ordinate system used for the hollow cathode model the half acceleration is

$$\mathbf{v}_{t+\Delta t/2} = \begin{pmatrix} v_r + (q/m)E_r \frac{1}{2}\Delta t \\ v_\theta \\ v_z + (q/m)E_z \frac{1}{2}\Delta t \end{pmatrix}, \tag{3.113}$$

v_θ is unchanged because $E_\theta = 0$ by virtue of axi-symmetry. The $\mathbf{v} \times \mathbf{B}$ force can be considered a rotation of $\mathbf{v}$. Specifically, the method of Boris [13] (see fig. 3.7) is then that

$$\begin{pmatrix} v_z' \\ v_r' \end{pmatrix} = \begin{pmatrix} \cos\omega_c\Delta t & \sin\omega_c\Delta t \\ -\sin\omega_c\Delta t & \cos\omega_c\Delta t \end{pmatrix} \begin{pmatrix} v_z(t+\Delta t/2) \\ v_r(t+\Delta t/2) \end{pmatrix} \tag{3.114}$$

where ω_c is defined as a *local* cyclotron frequency

$$\omega_c = \frac{q}{m}B_\theta(z,r), \tag{3.115}$$

and in a similar way to the previous equation, since we have only azimuthal magnetic field in an axi-symmetric problem, v_θ is unaffected. The acceleration is completed by an additional half acceleration the same as eq. (3.113). The above method is shortened to the case of a single $q\mathbf{E}/m$ acceleration if magnetic fields are ignored.

Were there to be any source of azimuthal 'swirl' (net v_θ) in the problem, then the

above would not be valid. Of course, it is possible that instabilities in the plasma could generate azimuthal currents and hence axial and radial magnetic fields, this is of course the classical 'Flute' or 'sausage' instability that occurs in z-pinch plasmas (see [40, §19.7]). Modelling the Flute instability correctly requires that the problem become fully three dimensional and computationally intractable at the expected densities[1]. In any case, the objective of this research is to attempt to observe the steady state internal plasma of the hollow cathode in the first place, investigation of any potential three dimensional instabilities must surely wait. The assumption of steady state $j_\theta = 0$ and correspondingly induced $B_z = B_r = 0$ is not unreasonable for an axially symmetric device. According to the PIC algorithm (fig 3.6), the natural progression of the discussion at this point would be to examine particle motion. This has already been covered in the description of the direct particle model, §3.1.1 on p. 40, so the discussion now moves directly to interpolation of particle charge and current density onto a mesh.

3.5.3 Interpolation

The interpolation phase is simply the phase involving the construction of ρ and perhaps also $\mathbf{j}$ from the particles in the simulation. There are various methods of increasing complexity that can be used to produce the interpolated description of the particle distribution. Generally, the more complex the interpolation method, the smoother the resulting source functions and hence the more stable the simulation tends to be. For a weighting shape factor S_g, where subscript $_g$ refers to grid quantities and subscript $_p$ refers to particles, for charge conservation during the interpolation from particle to mesh, we must have

$$\sum_g S_g(\mathbf{x}_p) = 1 \tag{3.116}$$

for all $\mathbf{x}_p$ [89]. From this, the mesh-based charge density in the simulation domain becomes

$$\rho_g \equiv \frac{1}{V_g} \int S_g(\mathbf{x}_p) d\mathbf{x}_p. \tag{3.117}$$

Two methods that each use different shape factors are are now presented, higher order (more complex shape factor) methods are only of academic interest because while they are theoretically more stable, the computational cost is prohibitive.

The first method is the so-called 'nearest grid point' (NGP) scheme. In this scheme, a particle's weight is simply added to the mesh point that it is nearest to. This is illustrated in figure 3.8. The NGP scheme has a particle shape factor $S(x)$

[1] Of course, within the main chamber of the ion thruster, the *applied* magnetic can be considered constant and hence axial and radial magnetic flux are used. Jugroot and Harvey [49] did this for the ion thruster, and the fixed magnetic field method has often been used in HET simulations.

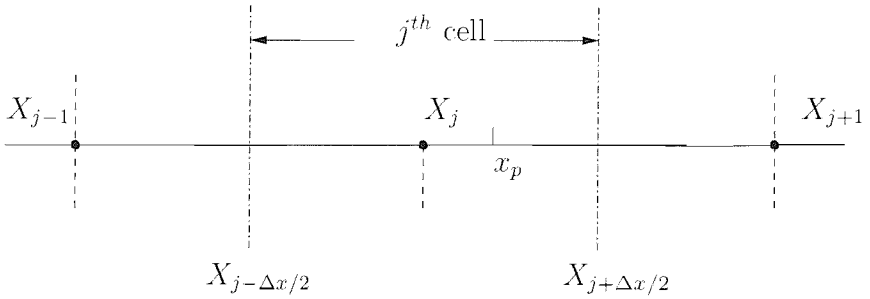

(a) Zeroth order 'nearest grid point' particle weighting. The j^{th} grid point is located at X_j and a sample particle is positioned at x_p. Cell spacing is Δx.

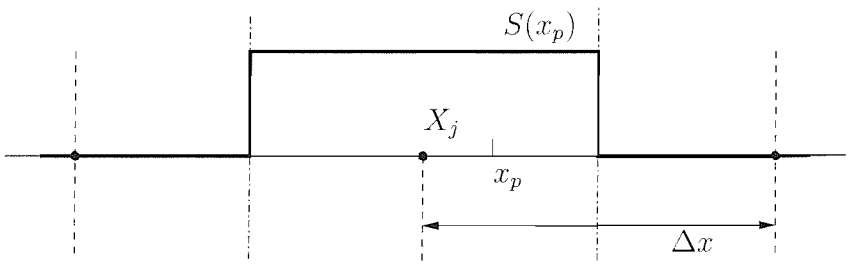

(b) Shape factor seen by mesh as particle x_p moves within cell X_j.

Figure 3.8: Nearest grid point (NGP) PIC interpolation. Grid points are located at X_j with spacing Δx. The density profile generated by a particle located at x_p is considered. See Figure 2-6a of Birdsall and Langdon [12]

that is Δx wide and square in nature (this method is sometimes called the square cloud method). The advantage of the NGP interpolation/weighting is that it is very fast computationally. This is offset by the fact that due to the simplicity of the density description, many particles are required in order that smooth density profiles are produced.

First order weighting, the so-called 'Cloud-in-cell' (CIC) method provides a far more realistic density description without too much prohibitive extra computation. Higher order weighting is possible by use of quadratic and cubic splines that describe particle shape, and this further reduces noise, but at greatly increased computation. This is why CIC interpolation is the choice of most PIC algorithms: it is the ideal trade-off. Figure 3.9 illustrates the method. The time-averaged shape factor $S(x_p)$ perceived by the grid as the particle moves about in $X_{j-1} < x_p < X_{j+1}$ is now a much better approximation to the 'cloud' of particles that the single particle represents: the shape factor is now $2\Delta x$ wide and triangular. Cell j receives

$$q_j = q_p \left[\frac{\Delta x - (x_p - X_j)}{\Delta x} \right] = q_c \frac{X_{j+1} - x_p}{\Delta x}, \tag{3.118}$$

while cell $j+1$ receives

$$q_{j+1} = q_c \frac{x_p - X_j}{\Delta x}. \tag{3.119}$$

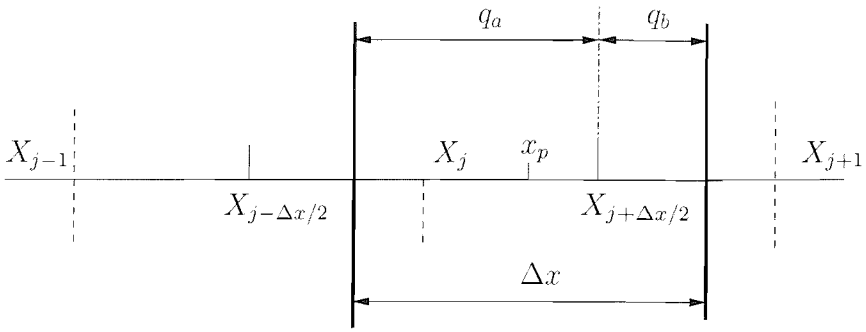

(a) 1st order 'cloud-in-cloud' particle weighting. The j^{th} grid point is located at X_j and a sample particle at x_p. The particle represents a uniformly charged cloud that contributes q_a charge to cell j and q_b charge to cell $j+1$.

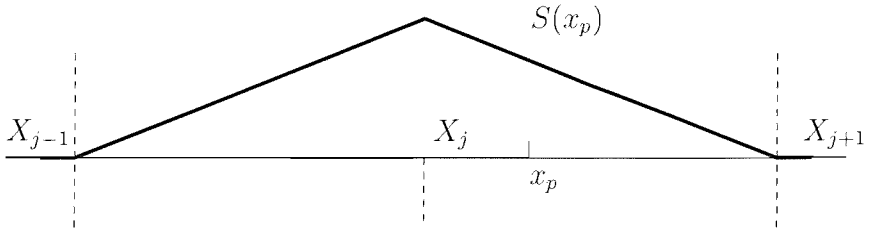

(b) Shape factor produced as a particle *moves through* CIC weighted cells.

Figure 3.9: Cloud-in-Cloud (CIC) PIC interpolation. In the same way as the illustration of the NGP scheme (fig. 3.8), mesh points are located at X_j and a test particle is located at x_p. Cell spacing is Δx. For further explanation, see figure 2-6b of Birdsall and Langdon [12] and associated description.

There are some additional reasons for using higher order weighting schemes for describing the particle distribution on a mesh. In general, such higher order schemes tend to reduce some of the non-physical effects that can appear due to noise. Such phenomena are discussed in the following section.

Weighting in cylindrical co-ordinates

For cylindrical co-ordinates, no modification needs to be made when using NGP interpolation/weighting. In the case of CIC weighting, however, the radial element of the cylindrical co-ordinate system needs to be taken into account. The shape factors given by Ruyten [89] are used. Ruyten's method has been tested and used for Hall thruster simulations, see for example Koo et al. [52]. In the axial direction, CIC weighting is the same as above. For a particle located at r, radially, the following are used:

$$S_j(r) = \frac{(r_{j+1} - r)(2r_{j+1} + 3r_j - r)}{2(r_{j+1}^2 - r_j^2)} \tag{3.120}$$

$$S_{j+1}(r) = \frac{(r - r_j)(3r_{j+1} + 2r_j - r)}{2(r_{j+1}^2 - r_j^2)} \tag{3.121}$$

Note that these expressions are charge conserving for an arbitrary arrangement of monotonically increasing r_j with j in the sense of eq. (3.116): it can be shown that

$$S_j(r) + S_{j+1}(r) = 1. \tag{3.122}$$

As must be the case for a charge conservative weighting.

3.5.4 Stability

It is possible to derive expressions that constrain the accuracy of the particle motion in a PIC model based on the relationship between the particle oscillation frequency and the time-step over which motion is descretised. When the leapfrog method is applied to a simple harmonic operator, it can be shown that there is no amplitude error in the particle motion for $\omega\Delta t < 2$ [12]; $\omega\Delta t > 2$ is inherently unstable in terms of the finite difference representation of particle motion. This means that the first absolute limit on the stability of PIC simulations is

$$\omega\Delta t < 2. \tag{3.123}$$

In fact, for accurate description of particle orbits in strong magnetic fields, and for dense sheath regions, it is usually wise to keep $\omega\Delta t < 1$. The type of frequency ω represents depends on the nature of the simulation. For the case of a sparse plasma in the presence of a strong external magnetic field, the cyclotron frequency will be larger than the plasma frequency, so the stability conditions is $\omega_c\Delta t < 1$ to resolve electron orbits. In an alternative scenario, $\omega_p > \omega_c$, the stability criterion becomes $\omega_{pe}\Delta t < 1$. Put simply, the condition is that $\omega\Delta t < 2$ where $\omega = \max(\omega_c, \omega_p)$.

In terms of the grid, it is more difficult to define such an exact limit on stability. In general, for charge-in-cell (CIC) interpolation, it is sufficient for a cell spacing of Δx that $\lambda_D < \Delta x$. It should also be noted that the phenomenon of 'numerical heating' introduced by rounding errors and other numerical artifacts is of some concern for certain PIC plasma simulation cases. This is often not a problem for electric propulsion device plasmas for several reasons. In particular, numerical heating is at its worst when periodic (specularly bounded) plasma simulations are performed. Under these conditions, numerical errors are retained and grow because particle velocities are never 'reset' to reference values (for example by exiting and re-entering the simulation, or by a thermally accommodating wall collision). In our case, particle collisions with walls are frequent. In the case of the hollow cathode (and several other EP devices), a DC discharge exists that involve injection and extraction of particles. None the less, this effect should still be considered if results show abnormal temperatures.

3.5.5 Finite difference discretisation

Standard methods are used for the discretisation of the various field equations via the finite difference approximation. These are now briefly presented, with additional time spent covering the nuances present when cylindrical coordinates are used. Using the notation shown in fig 3.10, the forward difference scheme (FDS), backward difference scheme (BDS) and central difference scheme (CDS) approximations to the first derivative are

$$\left(\frac{\partial\phi}{\partial x}\right)_i \approx \frac{\phi_{i+1}-\phi_i}{x_{i+1}-x_i}, \tag{3.124}$$

$$\left(\frac{\partial\phi}{\partial x}\right)_i \approx \frac{\phi_i-\phi_{i-1}}{x_i-x_{i-1}}, \tag{3.125}$$

and

$$\left(\frac{\partial\phi}{\partial x}\right)_i \approx \frac{\phi_{i+1}-\phi_{i-1}}{x_{i+1}-x_{i-1}}. \tag{3.126}$$

These are presented without derivation, see Ferziger and Perić [35], for greater detail. On boundaries, higher order schemes may be used. This is particularly relevant to this case with reference to the Neumann condition (ϕ' =constant) of rotational symmetry. A second order scheme derived from a parabolic fit to the boundary point and two inner points again given by Ferziger and Perić [35] is that

$$\left(\frac{\partial\phi}{\partial x}\right)_1 \approx \frac{-\phi_3(x_2-x_1)^2+\phi_2(x_3-x_1)^2-\phi_1\left((x_3-x_1)^2-(x_2-x_1)^2\right)}{(x_2-x_1)(x_3-x_1)(x_3-x_2)} \tag{3.127}$$

This notation can become cumbersome when using grid positions rather than assuming a constant grid spacing Δx. Since we are not considering constant mesh spacing in the final problem, notation of the type shown in eq. (3.127) is used in the code; full notation for variable mesh spacing is given.

Finally, a second-order accurate approximation to the second derivative is [35]

$$\left(\frac{\partial^2\phi}{\partial x^2}\right)_i \approx \frac{\phi_{i+1}(x_i-x_{i-1})-\phi_i(x_{i+1}-x_{i-1})+\phi_{i-1}(x_{i+1}-x_i)}{\frac{1}{2}(x_{i+1}-x_{i-1})(x_{i+1}-x_i)(x_i-x_i-1)} \tag{3.128}$$

The notion of a computational molecule (see figure 3.10) is used to ease the understanding of the notation. The point P is that where the discretisation is taking place, positions N,E,S and W give the surrounding points corresponding to $a_{i,j+1}$. $a_{i+1,j}$, $a_{i,j-1}$ and $a_{i-1,j}$ respectively.

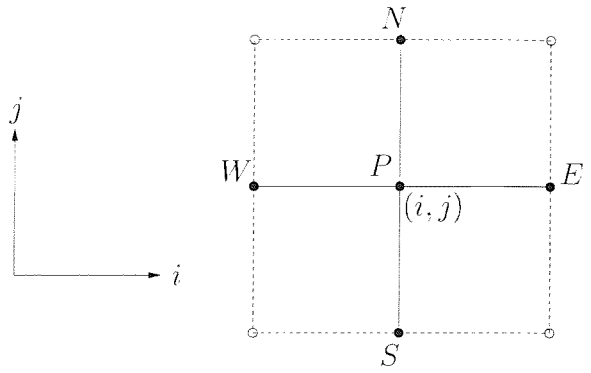

Figure 3.10: Cartesian grid computational molecule.

The Poisson Equation

Using this notation, we derive the discretisation of the cylindrical polar Poisson equation with rotational symmetry. The original PDE is

$$\frac{\partial^2\phi}{\partial z^2} + \frac{\partial^2\phi}{\partial r^2} + \frac{\partial\phi}{r\partial r} = \frac{\rho}{\epsilon_0}, \tag{3.129}$$

Using (3.128) for the second derivatives and the CDS – (3.126) – for the third term, we arrive at

$$\begin{aligned}\frac{\phi_N(x_P - x_S) - \phi_P(x_N - x_S) + \phi_S(x_N - x_P)}{\frac{1}{2}(x_N - x_S)(x_N - x_P)(x_P - x_S)} + \\ \frac{\phi_E(x_P - x_W) - \phi_P(x_E - x_W) + \phi_W(x_E - x_P)}{\frac{1}{2}(x_E - x_W)(x_E - x_P)(x_P - x_W)} + \\ \frac{\phi_N - \phi_S}{r(x_N - x_S)} = \frac{\rho_P}{\epsilon_0}.\end{aligned} \tag{3.130}$$

Defining

$$\alpha = \frac{1}{\frac{1}{2}(x_N - x_S)(x_N - x_P)(x_P - x_S)}, \tag{3.131}$$

$$\beta = \frac{1}{\frac{1}{2}(x_E - x_W)(x_E - x_P)(x_P - x_W)}, \tag{3.132}$$

$$\gamma = \frac{1}{r(x_N - x_S)}, \tag{3.133}$$

and

$$\begin{aligned} A &= (x_P - x_S)\alpha, \\ B &= (x_N - x_S)\alpha, \\ C &= (x_N - x_P)\alpha, \\ D &= (x_P - x_W)\beta, \\ E &= (x_E - x_W)\beta, \\ F &= (x_E - x_P)\beta, \\ G &= \gamma, \end{aligned}$$

we obtain

$$A\phi_N - B\phi_P + C\phi_S + D\phi_E - E\phi_P + F\phi_W + G\phi_N - G\phi_S = \frac{\rho_P}{\epsilon_0}, \tag{3.134}$$

which reduces to the final form

$$(A+G)\phi_N + (C-G)\phi_S + D\phi_E + F\phi_W - (E+B)\phi_P = \frac{\rho_P}{\epsilon_0}. \tag{3.135}$$

There remains a question regarding how the axis $r = 0$ is dealt with. Here, it would appear that the original equation (3.129) contains a discontinuity as $r \to 0$ because the third term is

$$\frac{\partial \phi}{r \partial r}. \tag{3.136}$$

L'Hôpital's rule [107] is that

$$\lim_{x \to 0} \frac{f(x)}{g(x)} = \lim_{x \to 0} \frac{f'(x)}{g'(x)}. \tag{3.137}$$

For the expression in eq. (3.136)

$$\lim_{r \to 0} \frac{\partial \phi / \partial r}{r} = \lim_{r \to 0} \frac{\partial^2 \phi}{\partial r^2} \tag{3.138}$$

This means that for grid points where $r = 0$, we solve

$$\frac{\partial^2 \phi}{\partial z^2} + 2\frac{\partial^2 \phi}{\partial r^2} = \frac{\rho}{\epsilon_0}, \tag{3.139}$$

which, for $\phi_N = \phi_S$ due to axial symmetry, descretises to

$$2(A+C)\phi_N + D\phi_E + F\phi_W - (2B+E)\phi_P = \frac{\rho_P}{\epsilon_0}. \tag{3.140}$$

The finite difference schemes presented above are certainly the most basic; there are many more exotic, higher order schemes available. In the context of PIC simulations, however, Birdsall and Langdon [12, p. 327] note that when examining the error it is most important to reduce the maximum error in the whole system $\mathbf{x} \rightarrow \rho \rightarrow \phi \rightarrow \mathbf{E} \rightarrow \mathbf{F}$ within acceptable tolerances. Investigations by Birdsall and Fuss [11] using a nine-point FD approximation showed that the improvements in the reproduction of the final plasma physics were negligible: other parts of the algorithm introduce comparatively greater error.

The Steady-state Ampère equation

The equation from §3.5.1 ($-\nabla^2\mathbf{B} = \mu_0 \nabla \times \mathbf{j}$) that allows us to find the steady state (no time derivative) azimuthal induced magnetic field based on the axial and radial current density expressed in cylindrical coordinates is

$$\nabla^2 B_\theta - \frac{B_\theta}{r^2} = \mu_0 \left(\frac{\partial j_r}{\partial z} - \frac{\partial j_z}{\partial r} \right), \tag{3.141}$$

where we solve for B_θ. Compared to the Poisson equation examined previously, the magnetic field equation is very similar: there is a source term (actually just the curl of the current density) and a Laplacian of the subject, B_θ. The difference is that there is an additional term: $(-B_\theta/r^2)$. This is important because on $r = 0$, B_θ must be zero. It is certainly true that the on-axis azimuthal magnetic field must be constant since due to Maxwell's equations, the magnetic field must be divergence free ($\nabla \cdot \mathbf{B} = 0$). The upshot of this extra term is largely that the on-axis boundary condition is trivial: $B_\theta = 0$. The source term, the curl of the current density, is found on the mesh using a central difference approximation to $\partial j_z/\partial r$ and $\partial j_r/\partial z$, where both are assumed to be zero on the $r = 0$ axis due to azimuthal symmetry. The equation for B_θ descretised as before (abbreviating B_θ to just B) is:

$$\begin{aligned} &\frac{B_N(x_P - x_S) - B_P(x_N - x_S) + B_S(x_N - x_P)}{\frac{1}{2}(x_N - x_S)(x_N - x_P)(x_P - x_S)} + \\ &\qquad \frac{B_E(x_P - x_W) - B_P(x_E - x_W) + B_W(x_E - x_P)}{\frac{1}{2}(x_E - x_W)(x_E - x_P)(x_P - x_W)} + \\ &\qquad\qquad \frac{B_N - B_S}{r(x_N - x_S)} - \frac{B_P}{r^2} = \mu_0 \nabla \times \mathbf{j}. \end{aligned} \tag{3.142}$$

On free space boundaries (upstream, downstream and radial extent), a Neumann condition (ϕ' =constant) is applied to the magnetic field. On solid surfaces we assume the magnetic field to be zero — a Diriclet condition *phi* =constant. This is obviously incorrect in the case of current carrying or even simply ferrous materials, but it is

assumed that the plasma induced magnetic field will be significantly stronger than that generated by the cathode body (for example), and that none of the solid objects in the simulation are highly magnetically susceptible, so that magnetisation is not significant. Clearly however, in future work, a code to predict magnetic conditions due to the cathode, anode, chamber etc.. would be useful.

3.5.6 Adjustments that enable PIC stability

This section covers the adjustments that are made to the simulation in order that the PIC stability criteria (e.g. $\omega_{pe}\Delta t < 2$) are fulfilled. The reason of applying these standard modifications is to accelerate the solution: without them the mesh size and time-step would be spectacularly unfeasible. First, modifications to allow for a given mesh size are described, followed by an assessment and solution to the problem of large m_i/m_e.

Mesh stability

Recall from section §3.5.4 that in order that the plasma physics is accurately reproduced using the PIC methodology $\Delta x < \lambda_D$. In the case of plasma device engineering this can frequently cause problems because the nominal λ_D dictates meshes that are completely impractical given current computational capabilities. Consider the Debye length of the plasma in the maximum density region of the hollow cathode (T_e in eV):

$$\lambda_D = \sqrt{\frac{\epsilon_0 T_e}{n_e e}} = \sqrt{\frac{8.85 \times 10^{-12} \cdot 1}{10^{20} \cdot 1.6 \times 10^{-19}}} = 0.74\ \mu\text{m}. \tag{3.143}$$

This implies a minimum of 1000 cells radially between the central axis and the cathode insert. For a 10 mm long section of insert, this results in a requirement for approximately 10^7 cells: computationally intractable. The question that inevitably arises is: how can there exist such a large volume of existing electric propulsion device research that uses this method, if it is apparently impossible? This answer is that there is a modification that can be made to the fundamental physics that serves to stretch the length scale of the plasma so that λ_D becomes similar to Δx for a reasonable mesh. The method is simple: the permittivity, as used in the solution of the Poisson equation for electric potential, is increased by some factor a^2 so that the artificial Debye length is increased by a. Thus, the approach is to set a based on the discrepancy between the real Debye length and the practical mesh. Of course, the effect of this modification must be carefully considered when analysing the results, but the experience of many past EP device simulations shows that it is valid. Despite such an apparently major modification to the physics of the problem, it is important to emphasise that

the method is a standard practice. For example, every single piece of published PIC-simulation research in the field of electric propulsion uses at least this modification to yield the problem tractable. Primary examples include Szabo [100], Gatsonis and Yin [36], Wang et al. [105], Boyd et al. [15], Sullivan et al. [99] and Celik et al. [18], all of whom also generally used the time-step acceleration modification presented in the following section.

Time-step stability

As with the above section, it can easily be shown that some modifications are required to ensure stability for the plasma that is expected. Consider some reference parameters of interior hollow cathode conditions: for a discharge current in the region of 5 A, the electron temperature will be 0.83 eV at an estimated density of approximately 10^{20} m^{-3}. These data are arbitrarily selected to be that of Monterde et al. [63], see table 2.1 on p. 14. Under these conditions, the plasma frequency is

$$\omega_{pe} = \sqrt{\frac{n_e e^2}{\epsilon_0 m_e}} = \sqrt{\frac{10^{20} \cdot (1.6 \times 10^{-16})^2}{8.85 \times 10^{-12} \cdot 9.1 \times 10^{-31}}} = 5.64 \times 10^{11} \text{ Hz}, \tag{3.144}$$

which implies that the period is of the order 10^{-12} s. The PIC time-step needs to be less than the plasma period. The requirement for such a small time-step makes PIC simulations of plasmas with densities similar to those found in electric propulsion devices impossible for any volume of interest. This is further exacerbated by the fact that a Xenon ion is about 240,000 times heavier than an electron. This again raises the question: how can there be such a large body of successful research in EP that uses the PIC method, if it is computationally so inefficient? There are two answers.

First, some EP simulation codes use fluid electron models with particle neutral/ion simulations. This overcomes the electron mass problem at the cost of assuming that the electron gas is a continuum (i.e. Maxwellian energy distribution) and that quasi-neutrality is maintained for all $\mathbf{x}$. Clearly, we are interested in both non-neutral regions and non-Maxwellian electron velocity distribution functions, so this approach is useless.

The second solution to the mass problem is quite simply to alter the mass. Bizarre as this may sound, it is a very common trick used in all varieties of PIC plasma physics. Increasing the electron mass not only allows for stability of the PIC method, but also reduces the difference in simulation residence time between electrons and heavy particles. Clearly, such modification will have a major impact on the results, but if the correct precautions are taken, meaningful results are still produced. An example extensive EP research using a mass-modified PIC code is the simulations at

MIT (see Szabo Jr et al. [101], Sullivan et al. [99], Celik et al. [18]). Results have frequently compared very well with experimental data. The MIT code operates by speeding up the neutrals and ions (decreasing mass) so that they have simulation residence times similar to the electrons. An example of research in EP that has retained neutral/ion mass and increased electron mass is the work of Okawa and Takegahara [74] who ran PIC simulations of ion thruster grid optics with a wide variety of electron masses. Since their computational domain was quite small, they were able to run the simulation with the real electron mass (although this took weeks) and with modified masses. Almost no difference was found between the 'heavy electron' case and the case for real electron mass. Clearly then, if the correct compensations and modifications are made, then results using smaller mass ratios can be perfectly valid. The compensations used in this research follow.

Modifications that compensate for larger electron mass

At a fixed kinetic energy, as the electron mass is increased, so the velocity must fall: by changing the mass, we alter the absolute velocity for a given energy. Changes in absolute velocity manifest themselves in changes to the fluxes at boundaries. It is possible to evaluate the change in flux and derive a method to compensate [101]. Consider the expression in §3.3.3 for the number flux though an element of unit area in an equilibrium gas,

$$N_e = \frac{n}{2\beta\sqrt{\pi}}. \tag{3.145}$$

Using $\beta_e = \sqrt{m_e/2kT_e}$, and the factor f by which the mass is changed, this can be rewritten

$$N_e = \frac{n}{2\sqrt{\pi}}\sqrt{\frac{2kT}{fm_e}}, \tag{3.146}$$

$$N_e\sqrt{f} = \frac{n}{2\sqrt{\pi}}\sqrt{\frac{2kT}{m_e}}. \tag{3.147}$$

Clearly then, fluxes vary according to $\sqrt{f}$, so when electron currents (*not* ion currents) are measured, the results need to be multiplied by $\sqrt{f}$. Similarly, when electrons are injected into the simulation via the field enhanced emission routine, emitted current needs to be reduced by multiplying by $\sqrt{1/f}$.

It is also worth mentioning that care has been taken in the implementation of the MCC routine so that the electron mass issue does not change the collision dynamics. In particular, recall that the collision probability for an individual event, k, is given by

$$P_k = nv\sigma_k(E)\Delta t, \tag{3.148}$$

where n is the neutral density, $v = |\mathbf{v}_e|$, E the electron energy and Δt the time-step (equation (3.51) in §3.4.1). The question is: should v_r, the real electron velocity (based on real mass) or v_s the simulated velocity be used? Since $v_r = \sqrt{f}v_s$, this expression does not need altering at all. The energy, E will be the same because although we increase m_e, the velocity is correspondingly lower: $E = E_s = E_r$. In the above expression, if we were to use the $v = v_r = \sqrt{2E/m_e}$, this would be incorrect unless multiplied by $\sqrt{1/f}$. Put another way, the collision probability is based on v in terms of the electron flux relative to the neutral background; if we were to calculate v_r from the energy based on the real electron mass, then the probability would be too high by a factor $\sqrt{f}$. Thus, we can either evaluate v_r from E and compensate by $\sqrt{1/f}$ *or* we can simply use v_s directly, since $\sqrt{1/f} \cdot \sqrt{2E/m_e} = v_s$ in any case.

When evaluating post-collision properties, it is necessary to scale the electron velocities so that they are correct with respect to m_e by multiplying by $\sqrt{f}$. If this were not done, then relative velocity magnitudes would be incorrect. For example, as $fm_e \rightarrow m_i$, $v_e \rightarrow v_i$, so $v_r \rightarrow 0$ on average. In this case a large number of ionising collisions would be rejected on the grounds of adverse neutral energy. Once electron-neutral collisions are complete, the electron(s) velocity is scaled back down by $\sqrt{f}$.

Chapter 4

Validation

The following subsections cover various validation studies that have been conducted on the components of the code to verify their correctness. A distinction is made between validation and experimental comparison: this section covers validation against known results, some of which are known experimentally; the results section contains comparisons with real experimental hollow cathode data.

4.1 Random number generator

The random number generator is clearly a critical component of any Monte-Carlo based numerical code: it is essential to check that it is accurate. A program to test the PRNG can be found in `tests/prng/testran`. There are two quick tests that are made. First, we check that the generator does produce a uniform result, i.e. $\langle R \rangle = 0.5$ over a large number of samples. For the reference sample size of 10^8 calls to the generator, $\langle R \rangle = 0.49998$. If the sequence seed is changed, the result is consistently 0.5 ± 10^{-4}. This verifies that the generated sequence is uniform about $[0, 1]$.

Secondly, it is important that successive numbers are not correlated. Sequential correlation can be a major source of error in all Monte Carlo based numerical models. In this case we are lucky since the problem is fundamentally three dimensional (three velocity components). Hence, it is necessary to check for non-correlation over triplets of random variates. In fact, for multi-dimensional problems solved using the Monte Carlo method, it is a requirement that the RNG be non-correlated is at least as many dimension. It was this requirement that led to the development of the Mersenne twister PRNG used here. The MT is actually equidistributed (non-correlated) for problems in up to 623 dimensions, so it should suffice for this work! Szabo [100] provides a method for checking for any correlation in triplets — directly related to 3-velocity component particle codes. The method follows: three random numbers are

generated successively and stored as

$$\zeta_1 = \sqrt{2}\cos(2\pi R_1), \tag{4.1}$$

$$\zeta_2 = \sqrt{2}\cos(2\pi R_2), \tag{4.2}$$

$$\zeta_3 = \sqrt{2}\cos(2\pi R_3). \tag{4.3}$$

Over a large number of generated numbers, mean values for the following should hold:

$$\langle\zeta_1\rangle = \langle\zeta_2\rangle = \langle\zeta_3\rangle = 0, \tag{4.4}$$

$$\langle\zeta_1^2\rangle = \langle\zeta_2^2\rangle = \langle\zeta_3^2\rangle = 1, \tag{4.5}$$

$$\langle\zeta_1\zeta_2\rangle = \langle\zeta_1\zeta_3\rangle = \langle\zeta_2\zeta_3\rangle = 0. \tag{4.6}$$

For the reference sample size of 10^8, the test program produces:

$$\langle\zeta_1\rangle = -0.00012063, \quad \langle\zeta_2\rangle = -0.00006481, \quad \langle\zeta_3\rangle = 0.00003310; \tag{4.7}$$

$$\langle\zeta_1^2\rangle = 0.99998223, \quad \langle\zeta_2^2\rangle = 0.99989672, \quad \langle\zeta_3^2\rangle = 1.00001385; \tag{4.8}$$

$$\langle\zeta_1\zeta_2\rangle = 0.00000001, \quad \langle\zeta_1\zeta_3\rangle = 0.00000001, \quad \langle\zeta_2\zeta_3\rangle = 0.00000001. \tag{4.9}$$

This verifies that the random number generator of Matsumoto and Nishimura [59] provides sequences of pseudo random numbers that are both uniformly distributed between 0 and 1, and exhibit no correlation. The results all fall within 10^{-4} of the target, which for a sample size of 10^8 is expected. Having shown that the PRNG is working correctly, the numerical methods that rely upon it can be tested.

A final note relating to random number generators is that of period. The period is simply the number of numbers that can be output before the quality of the stream begins to degrade (or repeat, become non-uniform, etc...). This is a non-issue for the PRNG used here (the MT). The MT can be mathematically proven to have a period of $2^{19937} - 1$. For perspective, this is enough for a theoretical stream producing a billion numbers per second for the current lifetime of the universe: there is no computer today that can come remotely close to exhausting the MT period.

τ	A	τ	A
0.01	100.50	0.3	3.845
0.02	50.50	0.4	2.923
0.03	33.84	0.5	2.448
0.04	25.50	0.6	2.067
0.05	20.50	0.7	1.779
0.06	17.17	0.8	1.530
0.07	14.79	0.9	1.363
0.08	13.01	1.0	1.207
0.09	11.62	2.0	0.4105
0.10	9.60	3.0	0.1496
0.20	5.28	4.0	0.05495

Table 4.1: Tabulation of $A(\tau)$ between 0.01 and 4 generated by the test function within the code. This matches the data of Nanbu [69]

4.2 Coulomb collision coefficients

Recall that to find the scattering angle of the binary Coulomb collision, we need the value $A(\tau)$ where A is given by $\coth A - A^{-1} = e^{-\tau}$. The approximations

$$A(\tau) \approx \begin{cases} 1/\tau, & 0 < \tau < 0.01 \\ \text{tabulated}, & 0.01 \leq \tau \leq 4 \\ 3e^{-\tau}, & 4 < \tau < \infty \end{cases} \tag{4.10}$$

have been shown to be accurate enough for particle simulations [71], but we need to verify that our tabulation is the same as that provided in [69]. For the Newton-Rhapson iteration, we use

$$f(A) = \coth A - A^{-1} - e^{-\tau} \tag{4.11}$$

$$f'(A) = 1 - \coth^2 A + A^{-2} \tag{4.12}$$

and a starting value $A_0 = 10^{-4}$ for all values of τ; the convergence seems very stable for the range $0.01 \leq \tau \leq 4$. Table 4.1 shows data generated using the Newton-Rhapson solver, the results match those printed in [69] to within numerical floating-point accuracy. The test function is contained in `mcc.c`, the source for the Monte Carlo collision code. The output is simply table 4.1. When comparing the above to the original tabulation of Nanbu, in [71], note that the notation change between 1997 and 2000, so that in the original paper, τ is equivalent to s (the notation used in this document follows more recent version contained in the review paper, reference [68]).

4.3 DSMC and direct particle transport

There are, of course, many papers covering methods of validation of DSMC codes. An excellent review of the attempts to verify the DSMC against experimental data is given by Harvey and Gallis [42]. They conclude that for neutral flows, and flows involving simple (plasma) ionising reactions, the DSMC provides very high quality description of rarefied gas dynamics; their only concerns are over the implementation of complex chemical reactions. Thankfully this does not apply to this research. There are several factors to consider when choosing a test case against which to verify the implementation of the DSMC presented in this thesis. Since the code developed for this research is exclusively used in a 2D axially symmetric case, it is worth using an example that is axi-symmetric. While verifying the DSMC routines, such an axi-symmetric test case has the additional benefit of verifying the correctness and accuracy of the particle move code that is non-trivial in the 2D axi-symmetric co-ordinate system (recall §3.1.1). Bird [9, §15.3–15.5] gives three examples of axially symmetric flow: hypersonic flow past a flat nosed cylinder, a Taylor-Couette flow, and the impact of a supersonic jet on a flat plate. The second of these examples, [9, §15.4], is a good validation case because it has been studied by several others, and is an example of an unsteady flow. Compared to the flow in a hollow cathode (steady, low vorticity), reproducing an example as complex as the Taylor-Couette flow would seem over-ambitious. This is not the case: the reasoning is that if the code validates against a complex flow, then it should be correct for the comparatively simple hollow cathode flow. The results against which the code developed for this research are compared are those provided by Bird [9] and Stefanov and Cercignani [98]. In summary, the reasons for choosing this case are:

- It is an axially symmetric flow: this tests the axially symmetric direct particle transport, mesh generation and sampling procedures;
- Bird's code (that provides the results for comparison) is the most tested and validated of any DSMC code ever written: it is the quintessential 'text-book' example;
- The results have been previously verified by both Reichelmann and Nanbu [85] and Stefanov and Cercignani [98] and found to be correct;
- The instabilities present in such flows are a standard case in fluid dynamics, and the nature of the onset of instability is well known via the Taylor number.

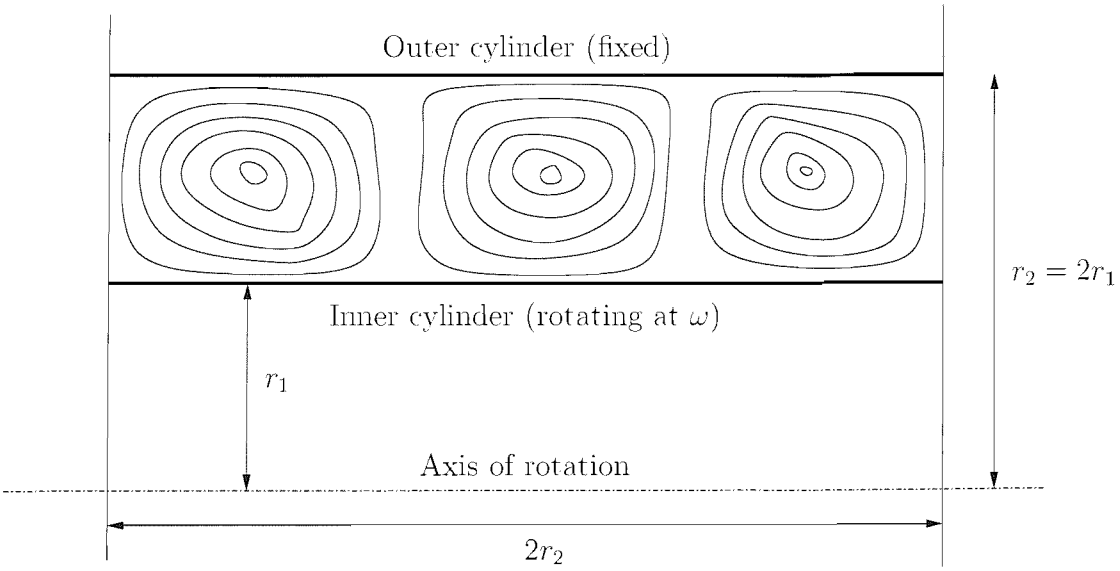

Figure 4.1: The configuration of the Taylor-Couette flow. The concentric shapes represent annular vorticies; this illustrates a stable three-vortex configuration.

4.3.1 Supersonic Taylor-Couette flow

The problem involves two cylinders of radii r_1 and r_2, where $r_2 = 2r_1$. An initially stationary uniform gas is placed between the cylinders at $t = 0$, and the inner cylinder is instantaneously provided with a constant angular velocity, ω. The end conditions are symmetric: as though the problem is mirrored on either side. End conditions are not periodic. The Taylor number for such a flow is given by

$$\mathrm{Ta} = \frac{4\rho^2\omega^2 r_1^4}{\mu^2\{1-(r_1/r_2)^2\}^2}, \tag{4.13}$$

where ρ is the gas density (kg m^{-3}), ω is the angular velocity and μ is the viscosity. It is know that for Ta > 33.100, the flow is unstable and vorticies will form. For practical purposes, the flow must have a finite length axially, Bird chooses this to be the diameter of the outer cylinder, giving the flow an aspect ratio of 4. Figure 4.1 illustrates the setup.

The specific configuration details used by Bird [9] now follow. The density is set so that the mean free path is $(r_2 - r_1)/50$. For $r_1 = 0.1$ m, $\lambda = (0.2 - 0.1)/50 = 0.02$ m. This is chosen so that the Knudsen number based on r_1 is 0.02. It is possible to evaluate the density required to yield such conditions from λ,

$$n = \frac{1}{\sqrt{2}\lambda\sigma_T} = \frac{1}{\sqrt{2}\lambda\pi d^2}, \tag{4.14}$$

given $\sigma_T = \pi d^2$ for a hard-sphere gas (see [9] and [98]). The reference density is then

$$n = \left[\sqrt{2}\frac{0.2 - 0.1}{50} \cdot \pi \cdot \left(4.17 \times 10^{-10}\right)^2\right]^{-1} = 6.4719 \times 10^{20} \text{ m}^{-3}, \tag{4.15}$$

where the gas is Argon (the molecular diameter is taken from table 3.2 in §3.3.2). Bird also states that the circumferential velocity of the inner cylinder, v_c, is to be set to three times the most probable molecular velocity, v_{mp}. The factor by which the cylinder circumferential velocity varies compared to the thermal velocity is the speed ratio, S (in fact, Stefanov and Cercignani [98] give results for a range of S from 1 to 15). The question is, what temperature should be used? Bird quotes a value for (Ta) of 521,600 for the results of Stefanov and Cercignani [98]. Given the density specified above, it is possible to reverse equation (4.13) to give a value for angular velocity of $\omega = 13362.6$ rad s^{-1}. This means that $v_c = 1336.26$ m s^{-1} and the temperature can be derived from the definition of v_{mp}:

$$v_c = S v_{mp} = S\sqrt{\frac{2kT}{m}}, \tag{4.16}$$

rearranged to give

$$\begin{aligned} T &= \frac{m}{2k}\left(\frac{v_c}{S}\right)^2 \\ &= \frac{66.3 \times 10^{-27}}{2 \cdot 1.38 \times 10^{-23}}\left(\frac{1336.26}{3}\right)^2 \\ &= 476.362 \text{ K}. \end{aligned} \tag{4.17}$$

So, for the chosen dimensions of the flow, it is necessary to set the reference temperature to 476.362 K so that the inner cylinder angular velocity results in a Taylor number the same as that used by Stefanov and Cercignani [98] and subsequently Bird [9].

The configuration files needed for repeating this test can be found in the directory `tests/Taylor-Couette/` from the source directory. Running the validation test is easy; Bird limited the examples contained in [9] to fit within 8Mb of RAM and run in under 24 hours on a 66Mhz 486. It was possible to reproduce the example using the code developed for this research running on a 1.6 Ghz P4 in under 20 minutes. Bird [9] limited the presented results to 20 revolutions of the inner cylinder at $S = 3$, which further limited his ability to judge whether the three vortex structure is stable. Such limits are not present today, and the stability of the vortex structure is discussed as an aside in the following section.

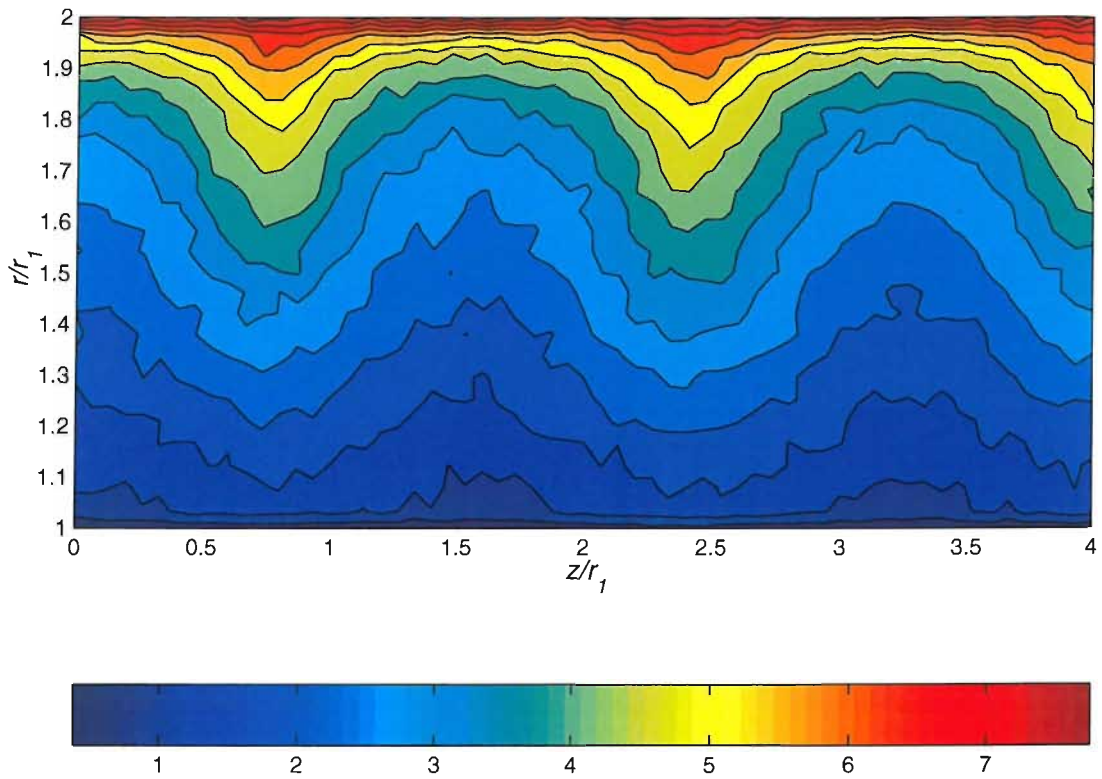

Figure 4.2: Contours of density ratio in a Taylor-Couette flow. Ta = 521637 and Kn = 0.02. This snapshot of the flow field is taken after 30 rotations of the inner cylinder (located at r=1). The aspect ratio of the plot is 2:1, the same as is used in [98] and [9].

Results

The results of Stefanov and Cercignani [98], confirmed by Bird [9] and Reichelmann and Nanbu [85] suggest that for Kn = 0.02 and Ta ≈520,000, there may be a stable structure of five vorticies after at least 20 complete revolutions. Figures 4.2, 4.3 and 4.4 show density ratio, temperature ratio and axial-radial velocity vectors respectively. The temperature and density ratio plots can be compared to those in [9, §15.4] or [98, fig 3], the quiver plot can be compared to those shown in [98]. Notice that the plots are shown at an aspect ratio of 2:1, this is why the vorticies do not appear circular. Plots are shown in this way because they are displayed in the same way in both [9] and [98]. All of the results compare very well: vorticies of the same magnitude and direction are formed in the same positions as shown by both Stefanov and Cercignani [98] and Bird [9]. Stefanov and Cercignani [98] investigated a range of speed ratios for stability. They list the case shown in the figures ($S = 3$, Kn = 0.02) as being unstable in the sense that Taylor vorticies will form, but stable in so far as the structure of the

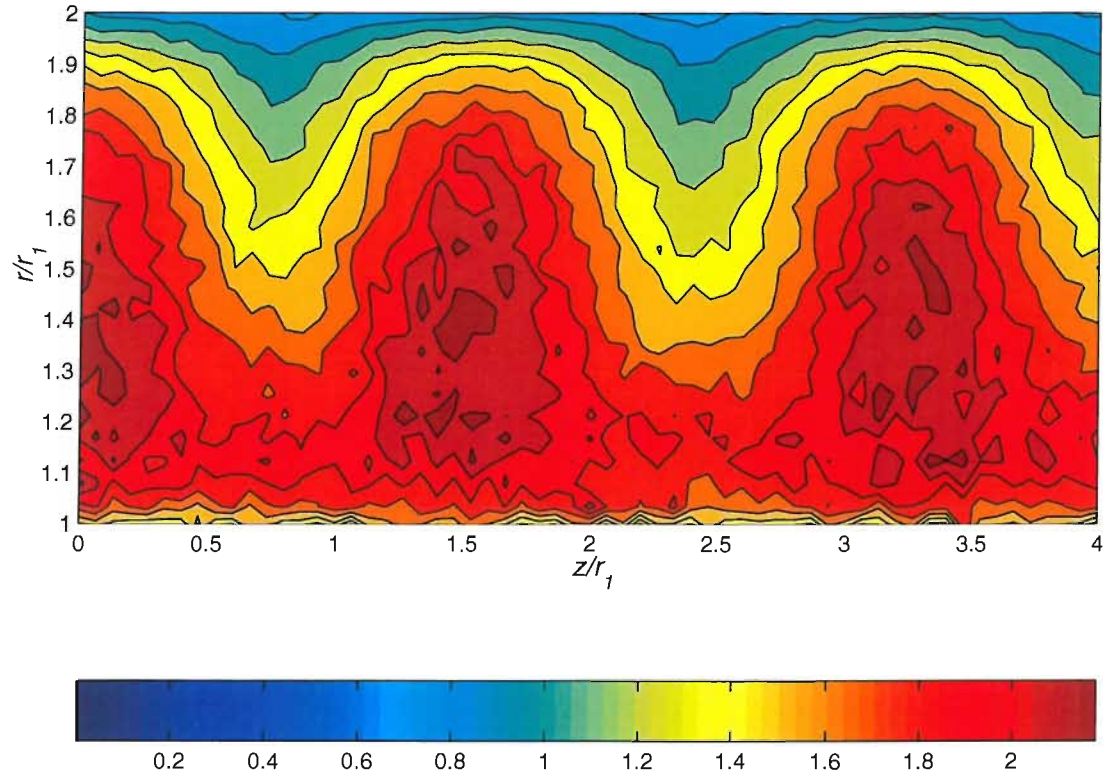

Figure 4.3: Contours of temperature ratio in a Taylor-Couette flow. Ta = 521637 and Kn = 0.02. This snapshot of the flow field is taken after 30 rotations of the inner cylinder (located at r=1).

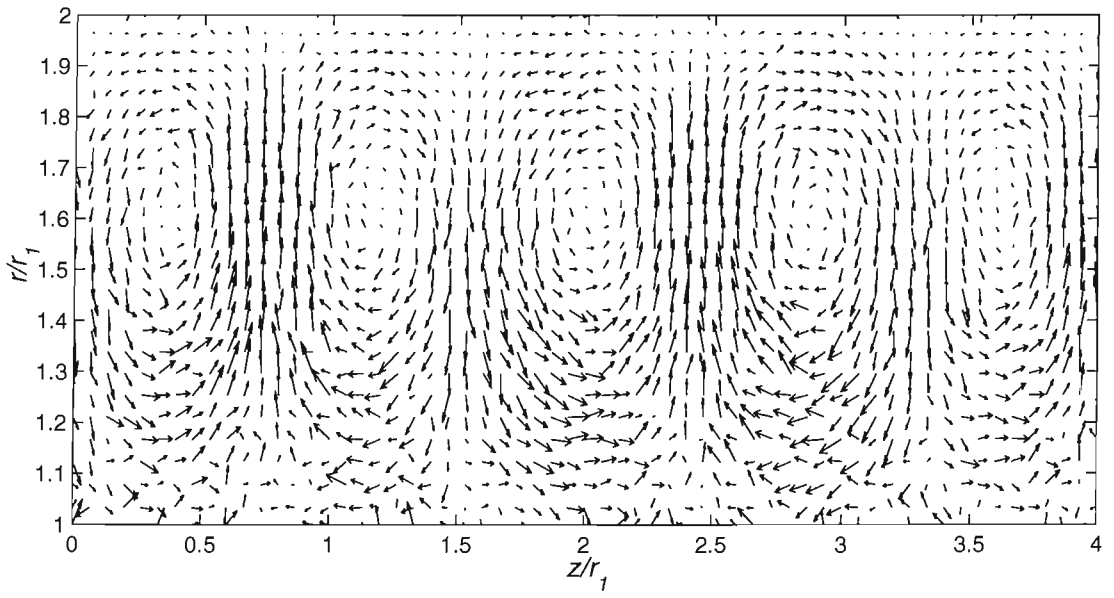

Figure 4.4: Axial-radial velocity vectors due to vorticies in a Taylor-Couette flow. This snapshot of the flow field is taken after 30 rotations of the inner cylinder (located at r=1).

vorticies tends to be regular and stable[1].

The excellent agreement of results for a flow problem that is significantly more complex than that we intend to investigate in future sections provides a good level of confidence in the results that are produced. The vast majority of the elements that make up the code have been tested. In summary:

- the code for providing a uniform startup particle distribution is correct (both the method discussed in §3.3.3 and the acceptance-rejection Maxwellian sample generating code);
- the 2D axially symmetric particle mover described in §3.1.1 is working correctly. This is, for obvious reasons, of great importance;
- the DSMC routines must be working correctly. As has been alluded to, this validation case asks the code to reproduce complex unsteady annular vortex structures: any minor errors present in the numerics of the DSMC code and/or direct particle code would most likely show up under such stress. It is logical to suggest that if this flow is reproduces flawlessly then the code should work well on the 'simpler' hollow cathode flow.
- the method for constructing and comparing analytical and numerical velocity distribution functions appears to be numerically correct.
- the sampling procedures and data processing code is working correctly.

This means that we have verified more than two thirds of the code. The remaining section is the PIC algorithm. This is divided into a validation of the iterative solver, and a verification of the particle→mesh algorithm.

4.4 Field solver and numerical discretisation

This section covers both a verification of the accuracy of the iterative field solver developed for use in the PIC model, and an analysis of the quality of the discretisation of Poisson's equation. It is useful to directly validate the field solver against an example similar to that with which it is used. In particular, it is important to use the cylindrical polar Laplace equation because this will produce an asymmetric $\mathbf{A}$ matrix in the linear algebra problem. Chao et al. [20] provide an example for the validation

[1]Stefanov and Cercignani [98] found that for $S = 1$ and $S = 2$ (but not $S = 1.5$!), no Taylor vorticies are formed. At all other values of S from 3 to 15, the flow is unstable, in some cases exhibiting transition to molecular chaos ($S = 12$ and 15). In some cases ($S = 2$ at Kn = 0.005 and $S = 1.5$ at Kn = 0.02), no clear vortex structure is observed. Results such as these were confirmed for air and compared very well with experimental results (see Reichelmann and Nanbu [85]).

of a Poisson solver in cylindrical co-ordinates but provide no verification or derivation of the analytical solution; this is now presented. The analytical result is evaluated and compared to the numerical one to assess the accuracy of both the solver and the discretisation scheme.

Analytical Solution to Poisson equation

The particulars of the validation problem follow. We find a solution to Poisson's equation on a rectangular (r, z) domain of length L and radial extent R. Neumann boundary conditions ($\partial\phi/\partial z = 0$) are applied on $z = 0$ and $z = L$. On the z axis, the axi-symmetry condition is applied: another Neumann condition so that $\partial\phi/\partial r|_{r=0} = 0$. Finally, we choose to set the far wall potential so that

$$\phi(r, z)|_{r=R} = V_0 \cos\left(\frac{2\pi z}{L}\right). \tag{4.18}$$

i.e. a complete cosine wave scaled to some reference potential V_0.

Beginning with the cylindrical polar form of the Poisson equation with rotational symmetry (see appendix C),

$$\frac{\partial^2\phi}{\partial z^2} + \frac{\partial^2\phi}{\partial r^2} + \frac{\partial\phi}{r\partial r} = \frac{\rho}{\epsilon_0}, \tag{4.19}$$

where we solve for $\phi(r, z)$. The charge density is zero: $\rho(r, z) = 0$, so this becomes the Laplace equation. The PDE is a second order linear and elliptic. It is separable, so that a general form of the solution is

$$\phi(r, z) = f(z)g(r), \tag{4.20}$$

where

$$f''(z) - c^2 f(z) = 0 \tag{4.21}$$

and

$$g''(r) + \frac{1}{r}g'(r) - c^2 g(r) = 0. \tag{4.22}$$

Both of these ODEs have solutions, the first (eq. 4.21)

$$f(z) = \alpha_1 \sin(cz) + \alpha_2 \cos(cz). \tag{4.23}$$

The second equation is in the form of the Bessel equation so the solution can be expressed in terms of Bessel functions;

$$g(r) = \gamma_1 I_0(cr) + \gamma_2 K_0(cr), \tag{4.24}$$

where I_ν and K_ν are νth order modified Bessel functions of the first and second kinds respectively. Combining these with the original general form gives the initial general solution

$$\phi(r,z) = (\alpha_1 \sin(cz) + \alpha_2 \cos(cz)) (\gamma_1 I_0(cr) + \gamma_2 K_0(cr)) . \tag{4.25}$$

By applying the conditions mentioned above, it is possible to deduce the values of the unknown co-efficients. First, taking the derivative of eq. (4.25) with respect to r and evaluating at $r = 0$ gives

$$\left.\frac{\partial\phi(r,z)}{\partial r}\right|_{r=0} = 0 = (\alpha_1 \sin(cz) + \alpha_2 \cos(cz)) (-\gamma_1 c I_1(cr) - \gamma_2 c K_1(cr)) . \tag{4.26}$$

Since $K_1(cr) \to -\infty$ as $r \to 0$, and $I_1(0) = 0$ we can say that for a physically realistic solution, $\gamma_2 = 0$. The solution becomes

$$\phi(r,z) = \mu_1 \sin(cz) I_0(cr) + \mu_2 \cos(cz) I_0(cr). \tag{4.27}$$

where the coefficients have been condensed using $\mu_1 = \gamma_1\alpha_1$ and $\mu_2 = \gamma_1\alpha_2$. We can now apply (4.18) on $z = 0$, say, so that

$$\phi(r,z)|_{r=R,z=0} = V_0 = \mu_2 I(cR) \tag{4.28}$$

which implies

$$\mu_2 = \frac{V_0}{I_0(cR)} \tag{4.29}$$

By taking the derivative w.r.t. z on $z = 0$ we obtain

$$\left.\frac{\partial\phi(r,z)}{\partial z}\right|_{z=0} = 0 = \mu_1 c I_0(cr). \tag{4.30}$$

so $\mu_1 = 0$. The current state of the solution is

$$\phi(r,z) = \frac{V_0}{I_0(cR)} \cos(cz) I_0(cr). \tag{4.31}$$

Applying eq. (4.18) gives

$$\phi(r,z)|_{r=R} = V_0 \cos\left(\frac{2\pi z}{L}\right) = \frac{V_0}{I_0(cR)} \cos(cz) I_0(cR), \tag{4.32}$$

which means that $c = 2\pi/L$. The final solution, which matches that given by Chao et al. [20] is then

$$\phi(r,z) = \frac{V_0}{I_0(2\pi R/L)} \cos\left(\frac{2\pi z}{L}\right) I_0\left(\frac{2\pi r}{L}\right) . \tag{4.33}$$

Poisson's equation is discretised as explained in section 3.5.5 and solved using the bi-conjugate gradient, the conjugate gradient and Jacobi solvers described previously in this section.

Comparison with numerical solution

The analytical equation is evaluated numerically at the same grid points that the numerical solution uses. Numerical values for the zeroth order modified Bessel function of the first kind, I_0 are obtained using the approximation of Press et al. [81], p.237. Although there is no significant comment on the accuracy of the approximation, the authors do state that the polynomial used to approximate I_0 is accurate to double precision. This is good enough for us because the natural discretisation error that will be observed in the following plots is several orders of magnitude greater than the computational floating point accuracy (if it were the case that the computer accuracy was worse than the quantifiable finite difference (FD) truncation error then we should give up any hope of a real solution immediately!). As with many modules of the source code, the field solver source contains a tester program that will reproduce the results that now follow.

All plots here use $L = 0.08$ m and $R = 0.04$ m. The wall potential V_0 is set to 1.0. The numerical solution is carried out on a uniform 64 by 64 grid to an absolute residual tolerance of $\omega_r = 10^{-6}$. ω_r is defined as the sum of the total change in the value of the solution between one timestep and the next. Hence, bearing in mind that there are grid points on $z = 0$ and $z = L$ so that there are 63 'cells', $\Delta z = L/63 = 1.267$ mm and $\Delta r = R/63 = 0.635$ mm. For this grid, the FD truncation error should be orders of magnitude greater than the solver residual error ω_r. Figure 4.5a shows the numerical evaluation of (4.33); figure 4.5b shows the equivalent numerical solution. The similarity indicates that both the solver and discretisation scheme show excellent agreement with the analytical solution.

Since it is possible to evaluate the magnitude of the FD truncation error, if we subtract fig. 4.5a from 4.5b, and find that the result matches the FD error, then it is possible to conclude that the BiCGStab solver is correctly written. Figure 4.6 shows the error between analytical and numerical solution. On the radial limit, the difference reduces to zero. This must be the case because the evaluation of eq. (4.18) is identical for both cases. Away from the wall, there is comparatively little radial error. The can be explained due to the nature of high order derivatives of modified Bessel functions: the radial variation in the original solution (figure 4.5a or b) is due to the $I_0(cr)$ term in eq (4.33). The discretisation is much more accurate for the radial term than the axial term.

Axially, the error clearly follows the shape of the boundary condition. At first, this

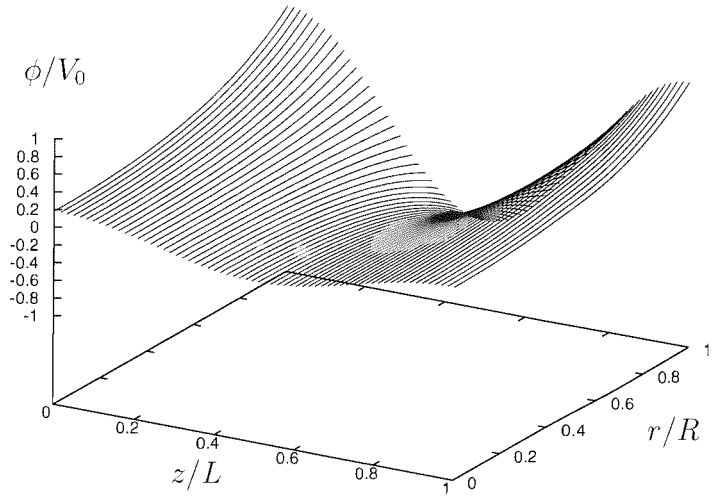

(a) Analytical

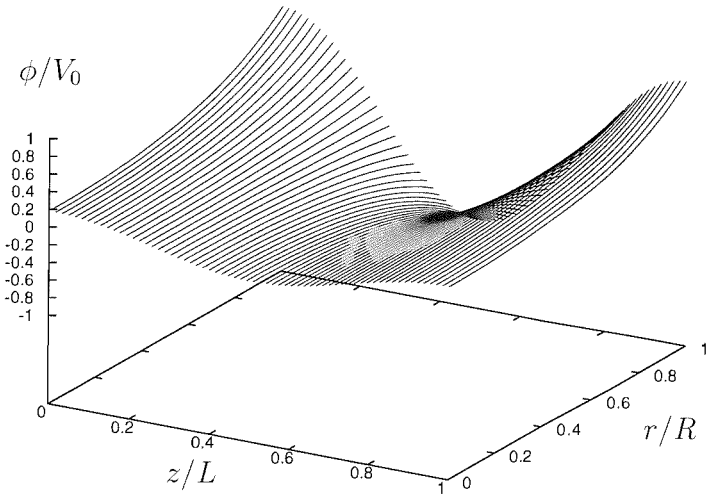

(b) Numerical

Figure 4.5: Comparison of analytical and numerical solution of cylindrical Laplace equation. The wall potential on $r = R$ is $V_0 \cos(2\pi z/L)$

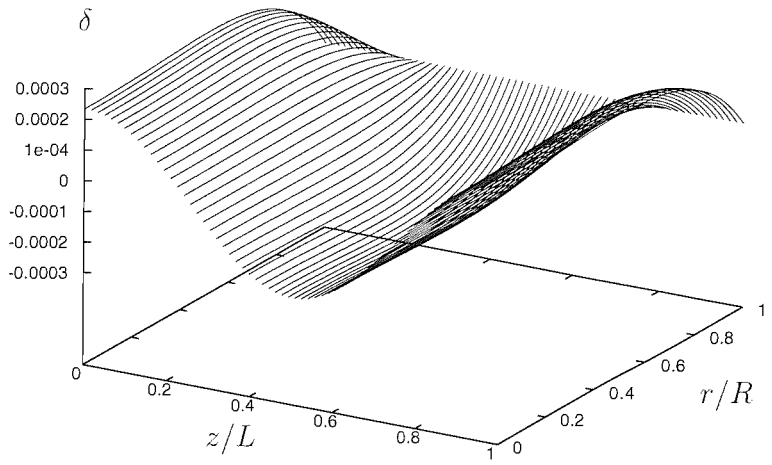

Figure 4.6: Finite difference truncation error between analytical and numerical solution.

may seem strange, because it might be expected that the error would be greatest where the derivative of the solution was greatest. More precisely, it could be expected that the leading $\partial^3\phi/\partial z^3$ error would dominate and the error would be a sine wave between the axial limits. In fact, since a uniform grid is used, the third order truncation error disappears so that the the leading truncation error term is fourth order; the shape of the error is consistent with what would be expected from FD truncation [35, §3.4]. When a finer mesh is used, smaller error is observed.

Comparison of numerical solvers

Figures 4.7a and 4.7b are included to further demonstrate the accuracy of the field solvers. By way of self-validation, two reference solvers are included in the code, against which the best (fastest) can be compared. In fig. 4.7, the same problem has been computed by all three iterative solvers. The Jacobi solver represents an elementary stationary iterative linear algebra solver of trivial complexity. The BiCG solver is an example of a non-stationary method for asymmetric matrices. Since the BiCG algorithm can exhibit highly erratic convergence behaviour, the stabilised variant is used as the final solver. For more discussion, see Barrett et al. [5] or Ferziger and Perić [35]. The algorithms for these methods are shown in appendix B. As can be seen, the difference between the obtained solutions is very small: of the order of 10^{-11}. This is much smaller that the convergence residual of 10^{-7}. It is possible to conclude that either (a) all three solvers are implemented correctly or (b) all three

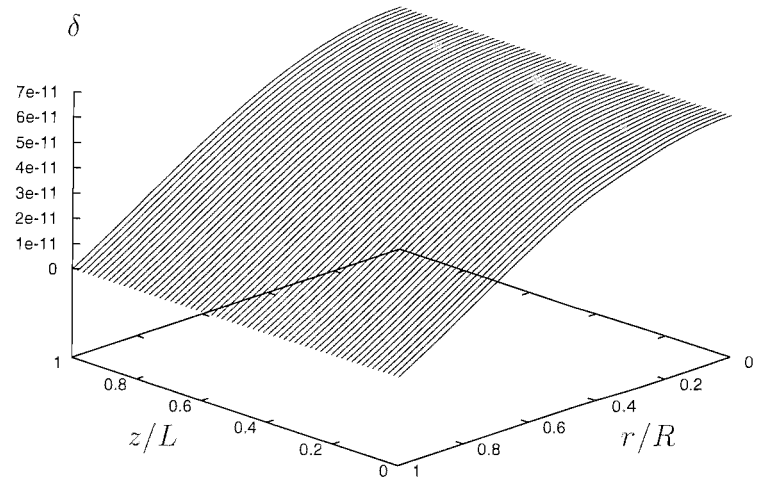

(a) Jacobi vs. BiCGStab

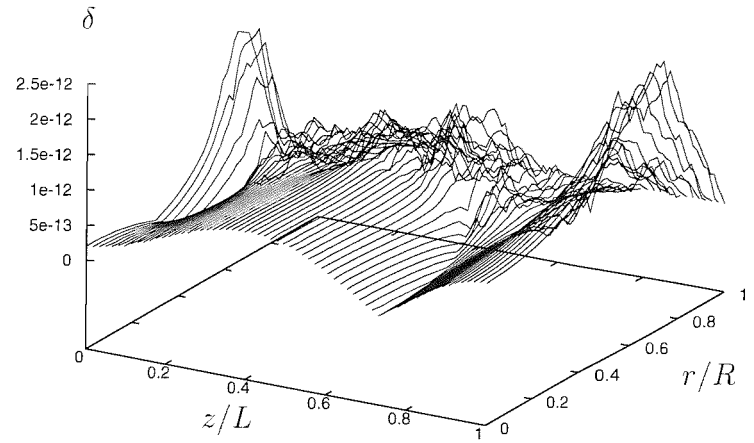

(b) BiCG vs. BiCGStab

Figure 4.7: Comparison of various iterative solvers. Shown are the reference Jacobi and BiCG solvers, and the final BiCGStab solver. Note reversed r axis in the top figure.

are implemented incorrectly so that they introduce very similar error. Given the previous discussion regarding the origin and magnitude of the FD truncation error in figure 4.6, coupled with the fact that the Jacobi and CG algorithms are very different, it is reasonable to assume (a) is correct.

4.5 PIC validation

Thus far, we have shown tests that have verified a large fraction of the code: a test for the particle transport and neutral-neutral collisions, small tests of various components and an examination of the validity of the numerical solver and PDE discretisation. The parts that remain unverified are few, but they are important. The charged particle weighting and acceleration are unchecked, as is the method for injection of current into the simulation. A classic case that presents itself that can check for all of these is a simple discharge between a pair of flat plate electrodes in a vacuum (no neutrals). A case similar to this has been studied using a PIC code by Asano et al. [3], see also Takamura et al. [102]. In such a configuration, depending on the nature of emission occurring (that in turn depends on the cathode temperature), the results can be compared to known experimental/analytical results such as the thermionic emission equation or the Child-Langmuir law based space-charge limited emission equation. By running the code at a variety of temperatures, we expect to see a transition from pure thermionic emission (following the Richardson-Dushmann equation) to space charge limited emission. For the results to match the above equation in the space-charge emission regime, the PIC code needs to correctly form a sheath adjacent to the cathode that will limit the current density. This also verifies the assumption that electrons incident on the cathode should be absorbed.

4.5.1 Flat plate discharge: Configuration and theory

The configuration is given in figure 4.8. Two disc shaped flat plates of radius L are placed *in vacuo* with separation L. A potential of $\phi_c = 0$ V is applied to the cathode and $\phi_L = 50$ V applied to the anode. The other boundaries are specified as specularly reflecting: i.e. as non-intrusive as possible while maintaining a zero mass flux out of the simulation. Under this configuration, we are examining current flow as a function of cathode temperature. Hence, the independent variable that is altered is simply the temperature of the plate at z=0. The temperature of the other plate is irrelevant to the problem for two reasons. First, being an anode, it is assumed that the plate emits zero current. Second, electrons in particle simulations are generally assumed to be completely absorbed upon contact with conducting walls. If the discharge voltage were a few orders of magnitude higher, then secondary electron emission could occur

and the anode temperature would become important. For ≈ 50 eV electron impact, SEE can be safely ignored. As the cathode temperature is increased from zero, the current density will begin to increase according to the thermionic emission equation:

$$j = A_0 T_c^2 \exp\left(\frac{-e\phi_w}{kT_c}\right) \tag{4.34}$$

given in §2.3.1. Here, T_c is the cathode temperature and the cathode material is assumed to be a metal with a regular surface lattice, so $A_0 = 1.2 \times 10^6$ A m^2 (the theoretical maximum Richardson constant). Recall that the work function, ϕ_w, can be reduced by the Schottky correction in the case of a favourable (enhancing) electric field; this will not be the case here as this is simply a discharge in a vacuum: no positive charge. The electric field strength produced by the 50 V anode-cathode potential drop is insufficient to produce field emission because the plate separation is large (1 cm). Once the electron density in the free space between cathode and anode is such that $\lambda_D < L$, a very sparse electro-negative plasma can be said to exist. This follows the commonly held definition of a plasma, that $L \gg \lambda_D$ [12]. As T_c and thus j increases, so n_e will increase and λ_D will continue to fall. Once the density is high enough, a *space charge limiting* sheath will form that will inhibit electron emission. The current density will then scale according to the well known Child-Langmuir law as

$$j = -\frac{4\epsilon_0}{9L^2}\left(\frac{2e}{m_e}\right)^{\frac{1}{2}} \phi_L^{3/2}, \tag{4.35}$$

where ϕ_L is the potential applied to the anode plate and L is the plate separation [40]. For the scaling to be correct, the PIC code must produce a sheath of the correct magnitude to shield the current so that it will not vary with temperature. There will be a temperature for which the thermionic current is equal to the maximum space charge limited current. This intersection point is worth investigating.

In summary, we expect the current density to scale with T_c using eq. (4.34) for low temperature. At high temperature, we expect the current density should remain constant as the cathode temperature is increased (according to eq. (4.35) with constant $\phi_L = 50$ V and $L = 1$ cm).

4.5.2 Thermionic to space charge limited emission

The configuration file for the code can be found in the test subdirectory. Different data points are run by changing the cathode temperature. In addition to the test itself, each simulation data point was run with two sample electron masses: real and heavily modified. The setup is fairly standard: for the real electron mass examples, a time-step of 10^{-11} s was used. A data set using an artificial electron mass was

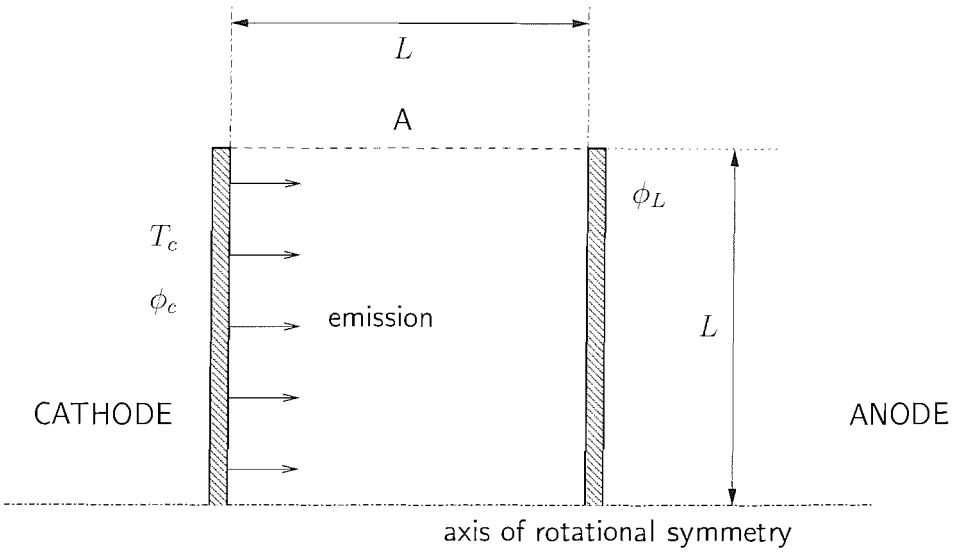

Figure 4.8: PIC validation case: configuration. Boundary A is a Neumann condition for the solution of Poisson's equation, and is set to reflect particles at $a_c = 0$ so that there is no radial current at B.

also run to verify the correctness of the compensations outlined in §3.5.6. In this case, the modification to mass is 2.4×10^5, which makes the electron have mass approximately the same as a Xenon atom. The compensation to electron fluxes is then $\sqrt{240000} = 489.89$. The mass modification allows for a time-step of 10^{-8} s to be used. 6000 computational electrons were modelled in both cases. The PIC mesh was 64 by 64 cells, uniformly distributed radially and scaled using a geometric progression axially so that the space-charge limiting sheath can be accurately resolved. Data points were run for each electron mass between 1250 K and 1450 K in 50 K intervals. An additional simulation was run at a temperature of 1325 K since this is where the theory predicts that the transition from a thermionic (no electron plasma) to a space charge limited discharge should take place. The initial condition is simply $n_e = 0$ everywhere.

Figure 4.9 shows a plot of current density as a function of temperature. The solid line is the Richardson-Dushmann equation evaluated at $\phi_w = 3$ eV as a function of temperature. The dashed line is the maximum current density expected due to space-charge limited emission at an inter-electrode spacing of 1 cm and voltage drop of 50 V. The data points are obtained directly from the simulation by setting the cathode temperature, circular points for real electron mass, crosses for artificial electron mass.

The first thing to notice is that changing the electron mass appears to have no discernible effect on the results. This verifies that the compensations discussed in §3.5.6 are working correctly and that changing m_e does not invalidate macroscopic results such as current.

In terms of the location of the results, at low, thermionic-saturated tempera-

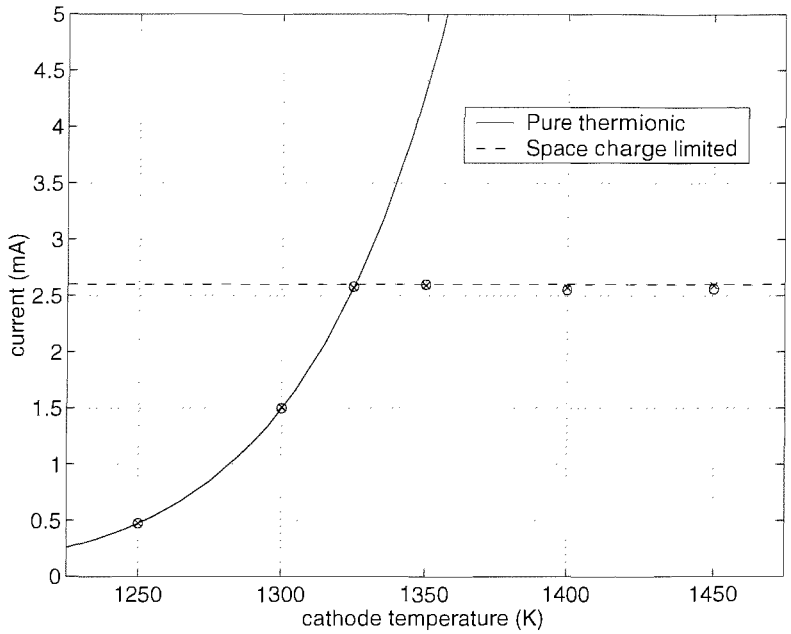

Figure 4.9: Thermionic and Space charge emission. The anode potential is 50 V and the separation is 2 mm. The plate material has a work function of $\phi_w = 3$ eV. Circular data points are for the discharge simulated using the *real* electron mass ($m_e = 9.1 \times 10^{-31}$ kg) and the crosses represent simulation runs at artificial electron mass of $240000 m_e$, or approximately one Xenon atom.

tures that lead to a discharge with a Debye length greater than the plate separation, the discharge current matches the Richardson-Dushmann model exactly. This is hardly surprising, given that the code contains the R-D model explicitly so that thermionic emission can be modelled; the most this shows is that there are no bugs in the thermionic emission model.

At temperatures higher than around 1325 K, the model again matches well known analytical observations exactly: the PIC code reproduces the space-charge limiting effect perfectly so that the current density matches that predicted by the Child-Langmuir expression. This means that we can conclude that the model is correct for the main emission regimes. To verify that the transition occurs at the correct temperature, a data point was generated at 1325 K. This point correctly lies on the intersection of the thermionic and field-saturated limits.

The final question to ask is: exactly what is happening to the electric potential profile to induce this limitation? Figure 4.10 shows plots of potential as a function of axial position near the cathode (which is located at $z = 0$). Two temperatures are shown: 1300 K and 1400 K, these are deliberately chosen as thermo and space-charge limited cases respectively. The retarding sheath can clearly be seen where the dashed line drops below 0 V momentarily: this is what is causing the limitation. It is also worth noticing that in the space charge limiting case, the increased current (1.5 mA

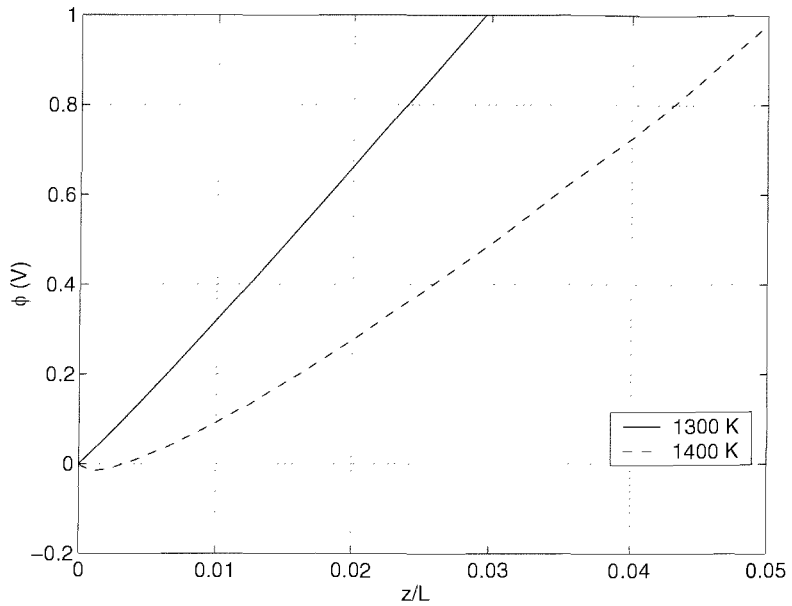

Figure 4.10: Electric potential on axis for a flat plate discharge at two cathode temperatures.

at 1300 K compared to 2.55 mA at 1400 K) means that the potential in the bulk space between the electrodes is depressed somewhat.

The results demonstrate the following:

- routines for electron emission appears to work well in thermo-emission and space-charge limited emission.
- This implies that chosen electron particle boundary condition is correct (i.e. absorb electrons at cathode surface) .
- Compensations for variable mass ratio appear to work: the heavy electron data points are indistinguishable from the real electron points.
- The PIC particle-mesh algorithm appears to work well: formation of the space-charge limiting sheath works correctly. In particular, this means the the axi-symmetric CIC charge density routines are working well.
- Methods for recovering simulation currents based on particle fluxes work correctly (also that the mass ratio compensations in this area are correct).

It is also worth noting that the Coulomb electron-electron collisions were used for these results. It would be wrong to claim that this implies that the Coulomb collision routines are correct, because under these conditions it would not be expected for Coulomb collisions to be significant, but clearly they are not drastically incorrect. This completes the verification of the PIC model for vacuum emission.

Conclusions

This concludes the discussions of verification of the code components. It has been shown that the DSMC performs well when asked to reproduce a complex time-dependant flow that is expected to be much more difficult to model accurately compared to the simple hollow cathode expansion. The PIC model and associated routines has been shown to reproduce the key emission regime exactly according to analytical expressions. Verification of the main components of the code by comparison to standard results is now complete, and the following chapter begins the presentation of the analysis of the hollow cathode.

Chapter 5

Results: Neutral propellant flow

The first set of results presented cover the characterisation of the hollow cathode neutral flow. This section covers comparison of collision models, verification in terms of grid and time-step and comparison to experimental data. Of course, the difficulty is that there is very little known about the neutral flow environment, and the number of experimenters who have gas data (usually a pressure measurement) is very small. The first thing to do is to estimate the condition of the flow so that the time-step and mesh configuration can be predicted. Second, the time-step and mesh need to be verified by varying them. Once all of the verification and comparison is complete, sets of exploratory data are given. This covers the effect of varying:

- mass flow
- cathode tip internal radius
- cathode temperature
- gas

For the purposes of initial verification, we use a reference set of parameters that forms a central point for each of the four variables in the above list. Table 5.1 shows the reference parameters for the neutral flow. Figure 5.1 show the large-scale and small scale geometry of the cathode.

It is possible to estimate the condition of the flow, and the computational parameters such as grid size and time-step, based on the basic input parameters such as mass flow rate and cathode temperature. Assuming the conditions upstream of the simulation are those of a uniform reservoir of gas (the approach of Murray et al. [65]), we can use the derivation of number flux through an element given in §3.3.3 to estimate the upstream density based on the (known) mass flux and temperature. For a stationary reservoir, $u_\infty = 0$, so the molecular speed ratio, s is zero. Hence, from

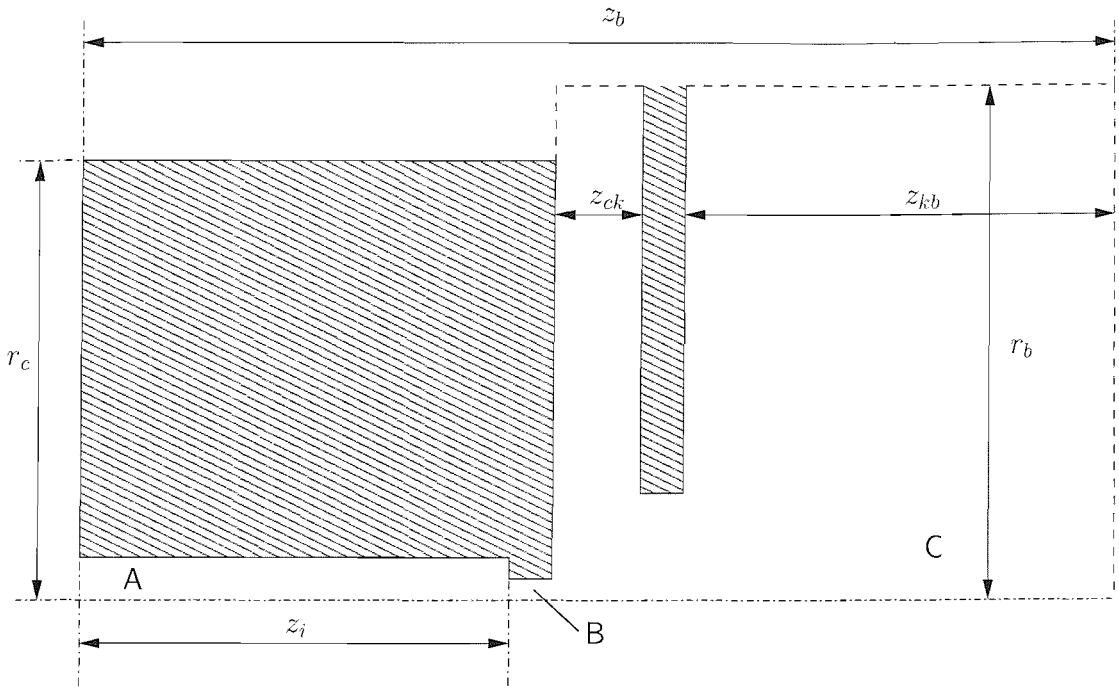

(a) Boundary geometry

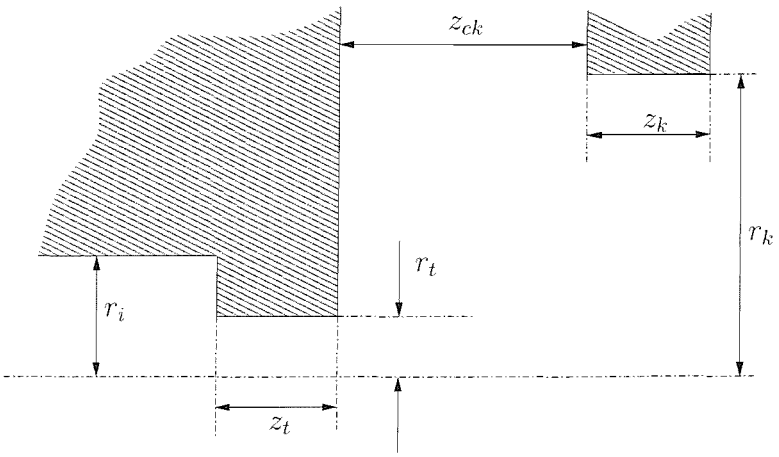

(b) Tip and keeper geometry

Figure 5.1: Cathode geometry. The scale for both sub-figures is as for the reference case: $z_t = z_k = r_i = 1$ mm, $z_{ck} = 2$ mm, $r_t = 0.5$ mm, $r_k = 2.5$ mm. The dimensions are essentially that of the T6 cathode. See table 5.1 for further reference dimensions. The points A, B and C are reference sample points where velocity distribution functions are recorded. Taking the bottom left of the upper figure to be at (0,0) mm, A is at (0.5, 0.5) mm, B at (10.5,0.5) mm and C at (15,1) mm.

Parameter	Value	Distance	mm
gas	Xenon	total axial length, z_b	24.00
mass flux, $\dot{m}$ (mg s^{-1})	1.0	total radial length, r_b	12.00
cathode T_c (K)	1200	cathode-keeper separation, z_{ck}	2.00
keeper T_k (K)	1200	keeper-boundary separation, z_{kb}	10.00
DSMC σ	VHS	cathode outer radius, r_c	9.45
		tip radius, r_t	0.50
		tip thickness, z_t	1.00
		insert radius, r_i	1.00
		insert length, z_i	10.00
		keeper thickness, z_k	1.00
		keeper inner radius, r_k	2.50

Table 5.1: Neutral flow: reference case parameters. Lengths are in mm.

eq. (3.46), the number flux is

$$N_i = \frac{n}{2\beta\sqrt{\pi}}. \tag{5.1}$$

N_i, the number flux *per unit area* can be written in terms of mass flow rate (in kg s^{-1}) as simply $\dot{m}/mA$, so the density can be written in terms of $\dot{m}$ as

$$n = 2\beta\left(\frac{\dot{m}}{mA}\right)\sqrt{\pi}, \tag{5.2}$$

where, for the reference case,

$$\beta = \sqrt{\frac{m}{2kT}} = \sqrt{\frac{218\times10^{-27}}{2\cdot1.38\times10^{-23}\cdot1200}} = 2.565\times10^{-3}. \tag{5.3}$$

Given and ambient mass flux of the same order of magnitude to that specified in table 5.1 ($\dot{m} = 1$ mg s^{-1}) and $A = \pi(10^{-3})^2$, so the predicted upstream density will be

$$n = 2\cdot2.565\times10^{-3}\cdot\left(\frac{1\times10^{-6}}{218\times10^{-27}\cdot\pi\cdot10^{-6}}\right)\sqrt{\pi} = 1.328\times10^{22}\ \text{m}^{-3}. \tag{5.4}$$

Of course, this is a very approximate estimate, as it does not take into account the constriction of the cathode tip. Nonetheless, a density of 10^{22} m^{-3} seems approximately right based on experimental experience [30]. The DSMC time-step needs to be of the same order of magnitude (or preferably less than) the mean collision period. Based on a density of 10^{22} m^{-3} and the reference parameters in table 5.1, the collision rate is

$$\nu = nv\sigma = 10^{22}\cdot390\cdot10^{-18} = 3.9\times10^{6}, \tag{5.5}$$

where we have used $v = v_{th} = \sqrt{2kT/m}$ and the hard sphere $\sigma = \pi d^2 \approx 10^{-18}$ for Xenon. The mean time between collisions is then

$$\tau = \nu^{-1} \approx 2.5 \times 10^{-7} \text{ s.} \tag{5.6}$$

This means the DSMC time-step needs to be around 10^{-7} s. For an initial estimate of mesh spacing (that in turn dictates particle population) recall from the mean free path analysis in §2.4 that for collisions between species of similar velocities, $\lambda \approx (n\sigma)^{-1}$ so the mesh spacing needs to be less than

$$\lambda \approx \left(10^{-18} \cdot 10^{22}\right)^{-1} \approx 0.1 \text{ mm.} \tag{5.7}$$

This indicates that for a density of 10^{22} m^{-3}, only around 10 cells are required to correctly resolve collisions radially between the axis and insert section. While this may be the case for $n = 10^{22}$ m^{-3}, it is clear that for an order of magnitude increase in density (that we expect for small tip radius cases) λ is an order of magnitude smaller and so 100 cells may be required. The other factor that will effect mesh spacing is collision cross section; the following section contains results covering the effect of using the various collision models discussed in §3.3.2.

The final factor that it is worth estimating before examining the results is the approximate time (and hence number of time-steps) it will take for a stable flow to form. A conservative estimate of the mean axial drift velocity is 50 m s^{-1} (this will become clear when the results are examined). This means that the gas travels $50\Delta t$ m per time-step or around $\frac{1}{10}$ of a cell. Of course, we expect this; in reality, the mean free path and mean collision time are inexorably linked in such a way that if one is valid, then the other will probably also be. Returning to the original point, the transit time of a particle is the total time it takes to pass from the upstream boundary to the downstream boundary. This time, assuming $z_b = 24$ mm is then about $\frac{1}{2}$ ms. For a 'jet' flow such as this, at least two transit times must pass before we can expect stability. Hence the code must simulate at least 2 ms of real time.

Boundary Conditions

The upstream boundary condition consists of an injection of neutrals as per the numerical method described in §3.3.3. The upstream temperature is set to experimental conditions; i.e. room temperature 300K. Particles that diffuse in the upstream are replaced during the following timestep with an equivalent particle at the 300K inlet temperature. This maintains the mass flow rate at a constant value. In the results presented here, particles that diffuse past the downstream boundary are simply deleted from the simulation. This is equivalent to a hard vacuum. The code also

provides for a downstream "vacuum chamber" gas. Where this is used (primarily in the plasma simulation results) it is noted, although this conditions simply serves to replicate vacuum chamber conditions recorded by experimenters for a particular test.

5.1 Numerical Validity

This section covers numerical studies of the effect of finite time-step and mesh spacing. These numerical analyses are carried out at the reference conditions explained previously.

5.1.1 Time-step

An expression was developed in the introduction to these results for the approximate time-step based on a density of 10^{22} m^{-3} and temperature of 1200 K. The result was a time-step of $\approx 2.5 \times 10^{-7}$ s. To verify this, it is necessary to run the simulation with time-steps both above and below Δt_{ref}, so that an indication of the degree of time-step independence can be obtained.

It is useful at this point to quickly discuss timings. The symbol T_c refers to the total *computational* time taken to run a given case, for example 20 minutes, 12 hours or however long. In contrast, T_p is the real, physical time in seconds that is computed *within the simulation.* This will typically be very small, generally on the order of milliseconds. For a given case, the two are related because to simulate from $T_p = 0 \rightarrow 1$ ms, say, may take $T_c = 15$ minutes, for example, so if we estimate that the simulation comes into equilibrium after 4 ms, then we should allocate approximately one hour CPU time. Of course, if, for example, the collision rate increases (thus increasing computation per time-step), then computation may take longer than one hour before $T_p = 4$ ms. Some confusion of terminology can arise in discussions involving these quantities, which is why the above definitions have been given. To re-iterate: physical, real or simulation time is the time that has passed within the simulation, while computational time is the time you have to wait between starting the simulation and viewing the results.

Having defined T_p and T_c, it is now possible to clearly explain a particular problem that arises in the case of numerical studies of time-step. The code is written so that the usual method of operation is to simply set the compute time T_c and allow the simulation to continue ad nauseam during that period: T_p is allowed to run on indefinitely. Even if the code reaches equilibrium quickly, additional sampling time during equilibrium is still beneficial (due to the statistical nature of the model the results will be smoother). There are however some cases in which it is useful to run the code for a fixed T_p, and this is one such case. For this numerical study — the verification of

time-step independence — it is imperative that all other parameters remain the same. This means that we must run for the same time period, and start sampling for a set period at the same T_c, as the flow structure may be exhibiting some time-dependant fluctuations that would invalidate the analysis of time-step independence.

The extent by which the simulation has reached equilibrium can be seen in the time history plots shown in figure 5.2. 'Time history' plots such as these are the main mechanism by which flow stability is judged. Notice also that there is some statistical ripple; this is to be expected from a numerical model that describes the flow physics in a fundamentally particulate manner. The upper plot shows that the mass flux out of the simulation is equal to the mass flux in (set to 1 mg s^{-1} for this reference case). The lower plot in figure 5.2 gives an impression of the degree of stability in the collision calculations: the DSMC net collision frequency has stabilised at just over 8×10^{21} s^{-1} neutral-neutral collisions in the simulation volume. Note that this value has not physical meaning: it is a computational measure of the number of neutral-neutral collisions that are being calculated by the DSMC routines per second physical time (multiplied by the weighting of the colliding particles). In light of the previous discussion regarding physical time over which the simulation is run, it was chosen that for this test, the code would run for 3 ms before sampling would begin and continue for an additional 1 ms. This seems reasonable by examination of fig. 5.2 since the flow seems stable after only 2 ms. The value of 3 ms can further be seen to be in agreement with the previous analytical estimate for particle residence (or transit) time.

In order to evaluate differences in flow conditions with change in time-step, three different sample points were chosen. Point A is located 1mm in from the upstream boundary and $\frac{1}{2}$ mm radially from the axis. Point B is located at the tip of the cathode, $\frac{1}{10}$ mm from the axis and sample point C is placed in the expansion plume, 10 mm downstream of the tip and 1 mm from the centreline axis radially. These can be seen in figure 5.1(a). Values at these points are now presented and discussed for a range of time-steps including the reference value (0.2 μs), a run at 0.1 μs and runs at 0.4, 0.6, 0.8 at 1 μs. This selection of values covers a data point at a time-step smaller that the reference case, as well as a range that covers an order of magnitude. This should give a good indication of the magnitude and onset of and non-physical descretisation effects.

Figure 5.3 shows plots of axial stream velocity at the three points as a function of time-step. It is clear that there is very little variation in velocity as the time-step is changed. With extended analysis, it may be that the variation could be shown to be less than the statistical scatter present in the model. It seems that the maximum velocity in the plume does fall somewhat as the time-step is increased, this is possibly numerical. At our reference time-step, however, there is very little difference in the

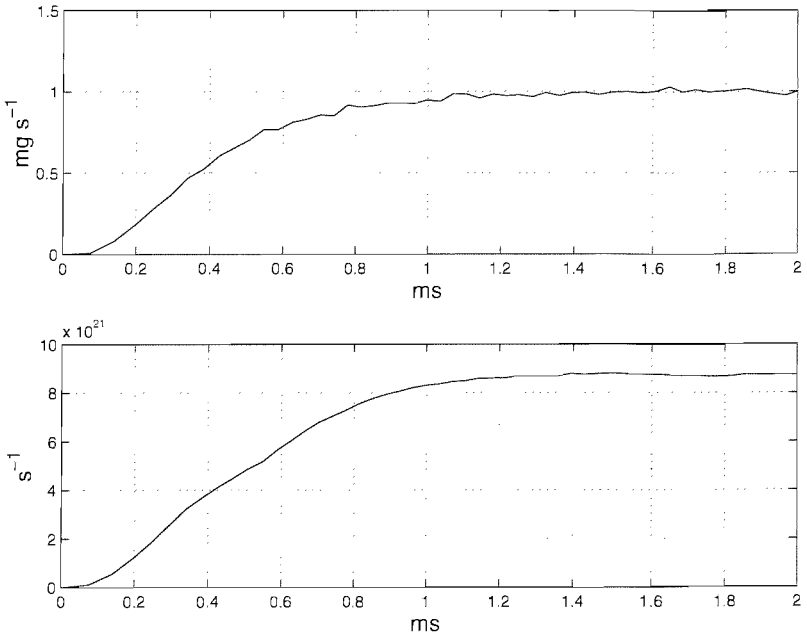

Figure 5.2: Neutral gas time-step resolution test: $\dot{m}$ and ν_{DSMC}. These time history plots are taken from the reference $\Delta t = 0.2$ μs case. ν_{DSMC} is a computational value that indicates the number of collisions calculated per second simulation time, *not* the real collision rate in particles per second.

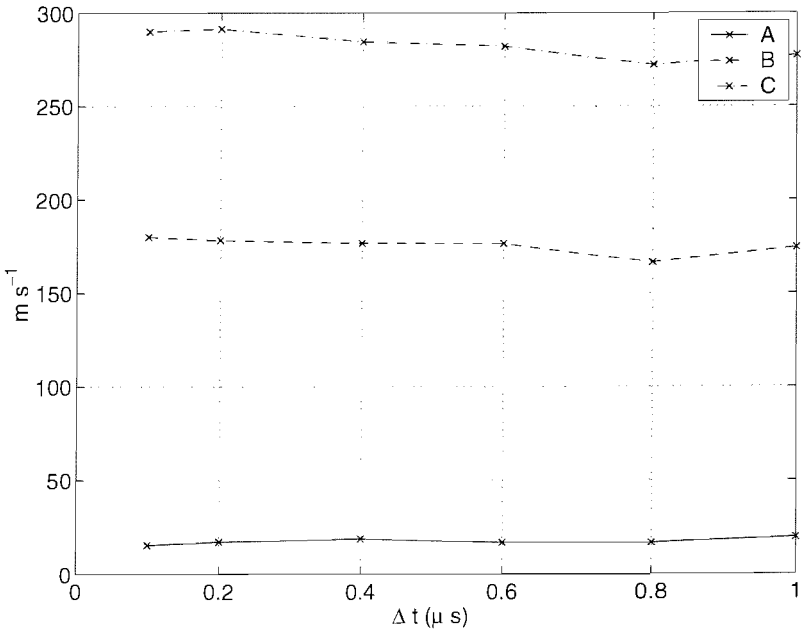

Figure 5.3: Neutral gas time-step independence test: axial stream velocity. Simulation results at three sample points for $\Delta t = 0.1$, 0.2, 0.4, 0.6, 0.8 and 1 μs.

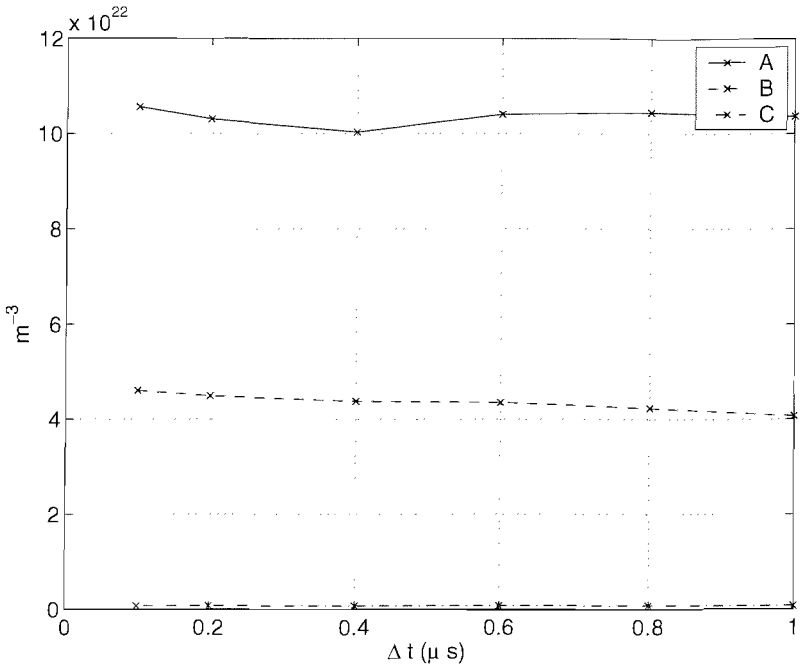

Figure 5.4: Neutral gas time-step independence test: density. Simulation results at three sample points for $\Delta t = 0.1$, 0.2, 0.4, 0.6, 0.8 and 1 μs.

result compared to a run with $\Delta t/2$. This gives a good initial verification that the chosen time-step is valid, but it is worth checking some additional results.

Figure 5.4 shows results for density at the three sample points as a function of run Δt. Again, there is very little variation between runs. At point A (near the upstream boundary, hence the high density) the variation is less than 5%. What is interesting is that the values for density vary very little between the 0.2, 0.6, 0.8 and 1 μs cases, but the variation is a couple of percent for the 0.1 and 0.4 μs cases. At position B however, it does appear that there is a continued trend toward higher density at smaller timesteps. This means that we cannot be fully sure that timestep indenpendence is reached. Regardless of this, it is clear that the level of numerical dependence on time-step looks to be at most the of the same order as the statistical scatter, equivalent to $\pm 0.3 \times 10^{22}$ m^{-3} in this case.

The results show that the DSMC code seems to produce results well outside the recommended time-step based on the predicted collision rate. This is to be expected; regardless of the recommended τ, Monte Carlo collision models often exhibit excellent levels of time-step independence well outside the theoretical range [9]. These results show that *for these conditions* a time-step of $\Delta t = 0.2$ μs is acceptable. Everything else remaining the same, an increase of collision rate could result for increasing mass flux or shrinking tip diameter (both of which will be done). At the reference time-step, the results presented in this section indicate that there is some room for increased

collision rate before numerical problems are observed.

5.1.2 Grid

Next, it is necessary to check whether the DSMC grid has an impact on the development of the flow. Since the research DSMC was completed first, a mesh suitable for MCC collision evaluation was first developed. The PIC mesh has different requirements, so *two* meshes are used for full plasma runs. In contrast to the traditional discussions of meshes in CFD, where the geometric *structure* of the mesh can have a significant effect on the stability and accuracy of the result, MCC meshes are very simple. MCC meshes are analogous to a finite volume descretisation in traditional CFD, but provided the simulation volume is descretised sufficiently for flow gradients to be observed, factors such as the connectivity of the mesh are irrelevant [9]. It is fair to say that the absence of instability and inaccuracy originating from the mesh is a major attraction of direct particle methods when compared to continuum CFD.

Much of the discussion in the previous section on time-step independence holds for this case, in particular the issue of holding T_c the same. The difficulty that is encountered here that is not encountered with the previous results is that as the mesh density is increased, so the number of particles should be increased. The recommendation of Bird [9] is that the number of particles be calculated from the mesh based on the rule that there should be at least 2.5 computational particles per cell. Bird show that this number is high enough to remove computational effects from the DSMC calculations. What this means is that the statistical sampling scatter will *change* as the grid is refined, so rather than encountering an approximately similar level of scatter between runs, it will increase as the grid is made finer.

Since most of the discussion is complete, results are now presented for three grids, shown in figure 5.5. All of the grids use a DSMC sub-cell arrangement of 2 axial and 2 radial sub-cells per main cell. This means that the reference grid, fig. 5.5(a), has 5600 cells, but Monte Carlo collision cell volumes generally occur between particles in the same sub-cell, so effectively the simulation volume is divided into 22400 regions. Two grids with fewer cells are chosen: one with 3464 cells and one with 1550 cells. The configuration files used to input these mesh descriptions can be found on the data disc in `data/neutral/res-grid/` under `normal`, `sparse` and `very-sparse`.

The same run time was used as per the time-step study, and for all three meshes, the flow was steady well before the sampling time was reached. Results of axial stream velocity and density are shown in figures 5.6 and 5.7 respectively.

Overall, the results show that there is no serious variation in macroscopic flow data as the grid is changed. For the densest sample point (the one near the upstream boundary, point A), there is very little variation in either stream velocity or density.

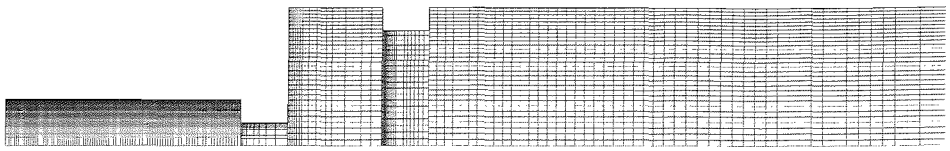

(a) Reference grid (22,400 sub-cells)

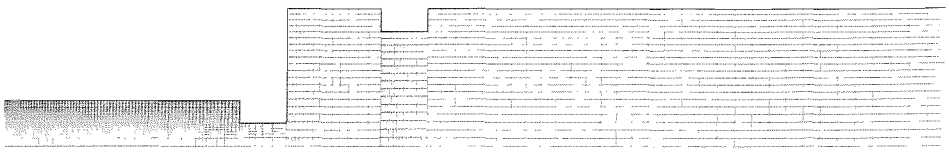

(b) Sparse grid (13,856 sub-cells)

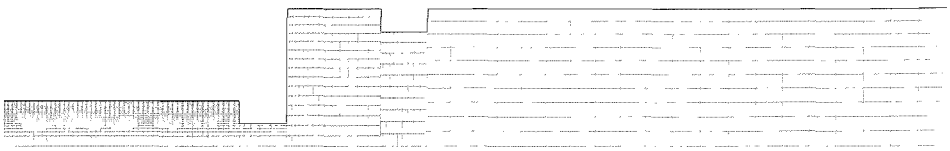

(c) Very sparse grid (6,200 sub-cells)

Figure 5.5: Three levels of grid refinement used to evaluate grid independence in the MCC collision routines.

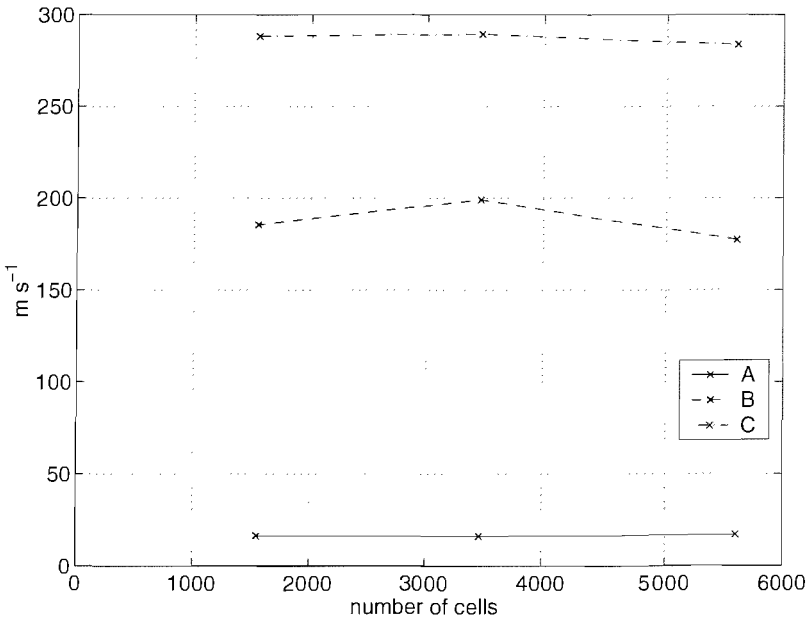

Figure 5.6: Neutral gas mesh refinement test: axial stream velocity. Simulation results at the three sample points for three levels of mesh density.

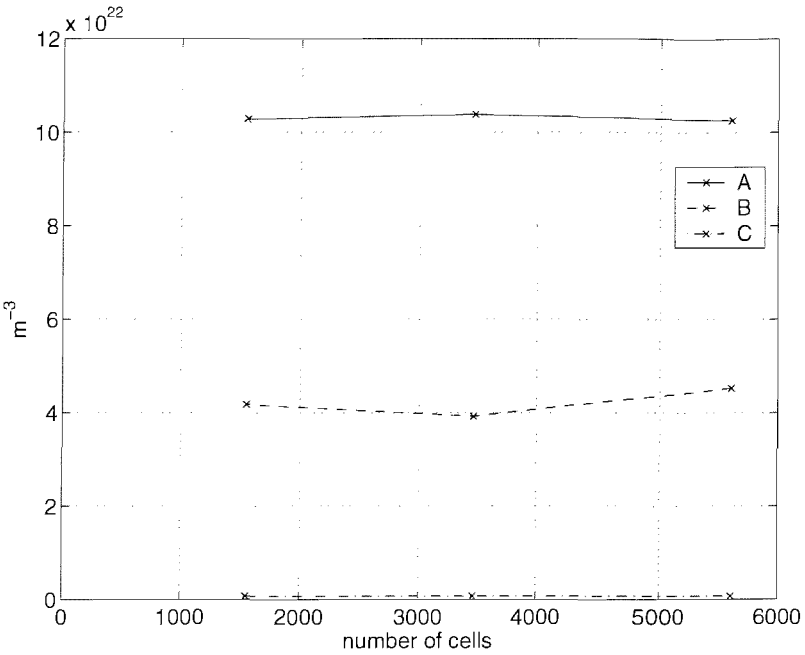

Figure 5.7: Neutral gas mesh refinement: axial stream velocity. Simulation results at the three sample points three levels of mesh density.

In both plots, there does appear to be a slight anomaly in the intermediate grid, this is manifested as a slightly lower density and higher velocity for the sample point located in the tip. It is reasonable to attribute this to a temporary instability in the flow that is picked up particularly well by this mesh. It is also true to say that since the variation between data points is typically much less that 5%, and there there is no discernible pattern to the variation — excluding the aforementioned exception — the variation could be statistical scatter (for this case/position this was estimated at around $\pm 0.3 \times 10^{22}$). Of course, the comments contained in the previous section regarding the statistical nature of particle simulations hold. In general, it is possible to make a similar conclusion to that made for time-step: any numerical effect of the mesh on the results is of a magnitude at least as small as the statistical scatter, under the parameters we are using. A final note on the use of non-orthogonal elements in the mesh follows.

A note on the chamfered tip cathode

It should be noted that the T6 cathode is generally used with some level of chamfer drilled into the tip (from a downstream direction, see the diagram of a hollow cathode in the introductory section, page 12). It was decided that the increase in mesh complexity was not worth the additional accuracy of including the chamfer. This decision is based on a number of factors. The first point to be made, albeit a somewhat

weak one is that traditionally, such meshes have been used. For instance, the work of Murray et al. made the same decision (to omit the chamfer), based on mesh complexity and the comparative unimportance of the small geometric difference on the bulk cathode physics; some discussion is given in ref. [65]. The fact that it was proposed by someone else is not sufficient reason to accept the omission of the chamfer. The question must be asked: what is it there for? In the progress of research conducted with Rudwan this question was address. There are some results provided by the code presented in this document contained in the thesis of Rudwan that address the structure of the electric field due to the absence or otherwise of the chamfer. Given that a primary object of the work of Rudwan was to investigate the breakdown (initiation) of the cathode, we determined that chamfering allows for a significantly higher initial electric field to penetrate the cathode for the purposes of the breakdown current extraction. This in turn allows for the discharge starting at much lower voltages. The conjecture that the chamfer is primarily an aid to breakdown was backed up at an early stage in this research when a neutral gas run was completed using a chamfered surface (at 45°) but a regular grid. It was found that the axial pressure varied by less than 10% between the chamfered and non-chamfered cases. As a footnote, in terms of comparison with experimental data, it should be remembered that the T5 cathode and earlier SERT cathodes were always unchamfered in any case (so that the early Fearn and Philip and Siegfried and Wilbur work used unchamfered cathodes); chamfering the tip was introduced in the T6 cathode.

5.2 DSMC cross section model

Having concluded the discussion of mesh and time-step, there is one final element that it is interesting to examine before a detailed analysis of the flow structure. Recall that there are two different models of the collision cross section for neutral-neutral collisions (section 3.3.2). In this section, results are briefly presented for the reference case for runs with the two different models, these are also compared to a collisionless flow. To speed up compute time, a shortened cathode insert was used: 5mm rather than 10mm. This is the only deviation from the reference case stated previously. In summary, the following are compared:

$$\begin{array}{lll} 1. \text{ Collisionless} & \sigma = 0 & \\ 2. \text{ Hard Sphere} & \sigma = \pi d^2 & \\ 3. \text{ Variable Hard Sphere} & \sigma = \dfrac{\pi d}{\Gamma(5/2-\omega)} \left(\dfrac{2kT_{ref}}{\mu v^2} \right)^{\omega - \frac{1}{2}} & \end{array} \tag{5.8}$$

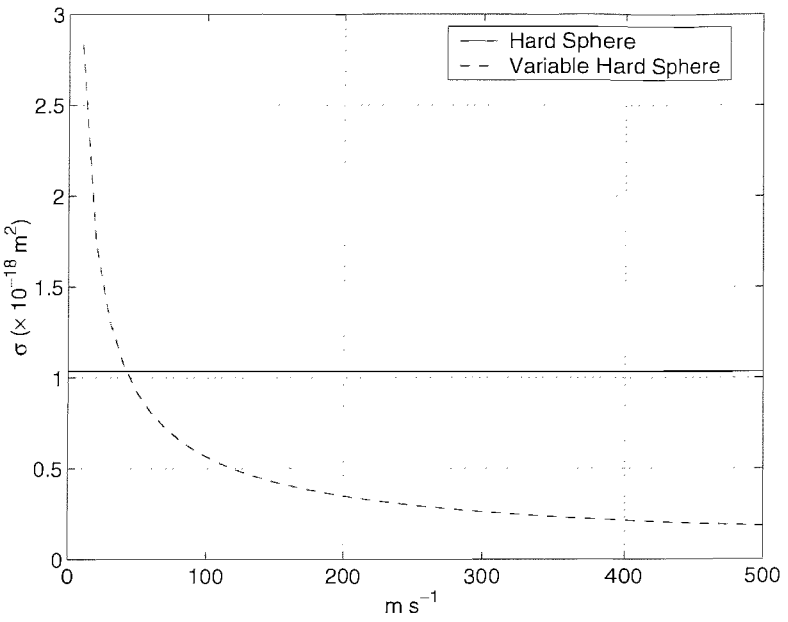

Figure 5.8: Comparison of collision cross section σ for HS and VHS DSMC collision models as a function of collision velocity.

where the final expression is that given by Bird [9] for the variable hard sphere collision cross section. T_{ref} is given for d_{ref} [9, app. A], while g is the relative collision velocity between two particles, $\mu = m_1 m_2/(m_1 + m_2)$ is, as usual, the reduced mass, ω is the viscosity index and Γ the gamma function. Refer to table 3.2 on p. 54 for values of ω, T_{ref} and d_{ref} for the inert gases. The shape of the variation in σ for case three can be seen in figure 5.8, where cases two and three are compared. With the VHS, we expect far higher collision probabilities in cold gases, while for large g (hot gas) the collision probability will be less than that predicted by the constant HS approximation.

Figures 5.9 and 5.10 show axial centreline plots of density and axial stream velocity respectively for the three cases. Considering the plot of figure 5.9 first, the density profile consists of an initial expansion that persists for 1mm from the upstream boundary (the injected gas is at room temperature, while the cathode is several hundred degrees warmer), followed by a gradual drop until the inner lip of the cathode orifice is reached. Until the inner face of the tip, all three models give similar results. Immediately prior to the expansion there is a slight anomaly in the collisionless version, and once the flow is clear of the outer face of the tip (6 mm) there is a major divergence between the collisionless and collisional models. Downstream boundary density is around 7×10^{20} m^{-3} for both the collisional models and 1.5×10^{21} m^{-3} for the collisionless case. These results tell us two things: first, it seems that it would be incorrect to assume that the flow is collisionless, as the expansion appears incorrect under such assumption (of course, we knew this already, but here it is confirmed).

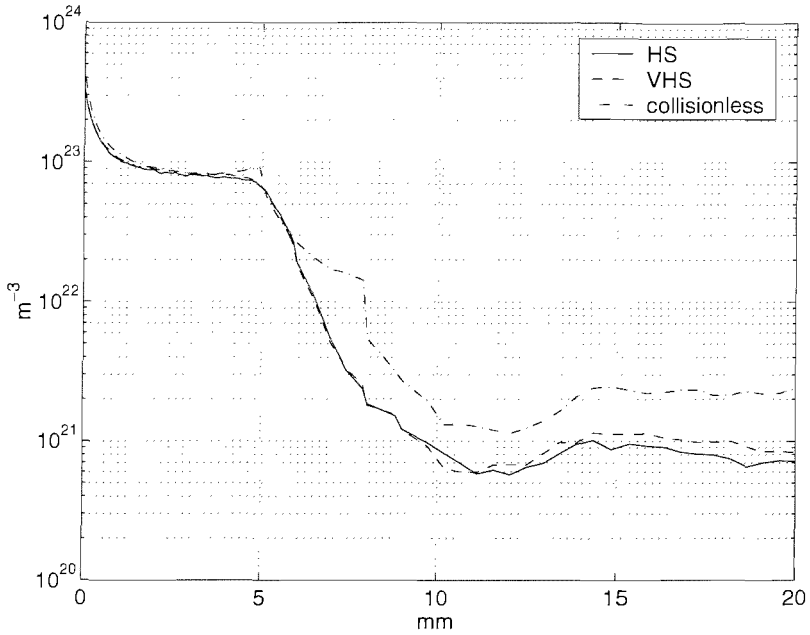

Figure 5.9: Axial centreline plot of density for the three collision models.

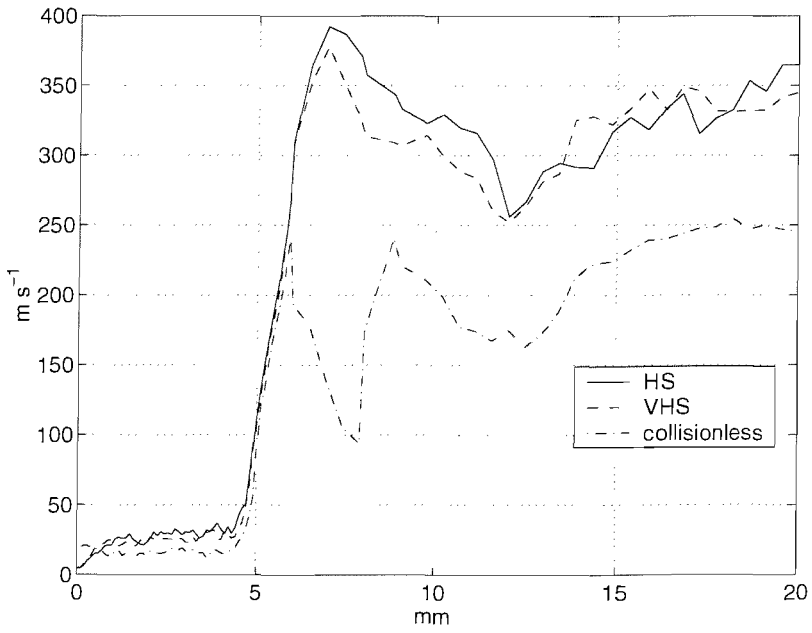

Figure 5.10: Axial centreline plot of axial stream velocity for the three collision models.

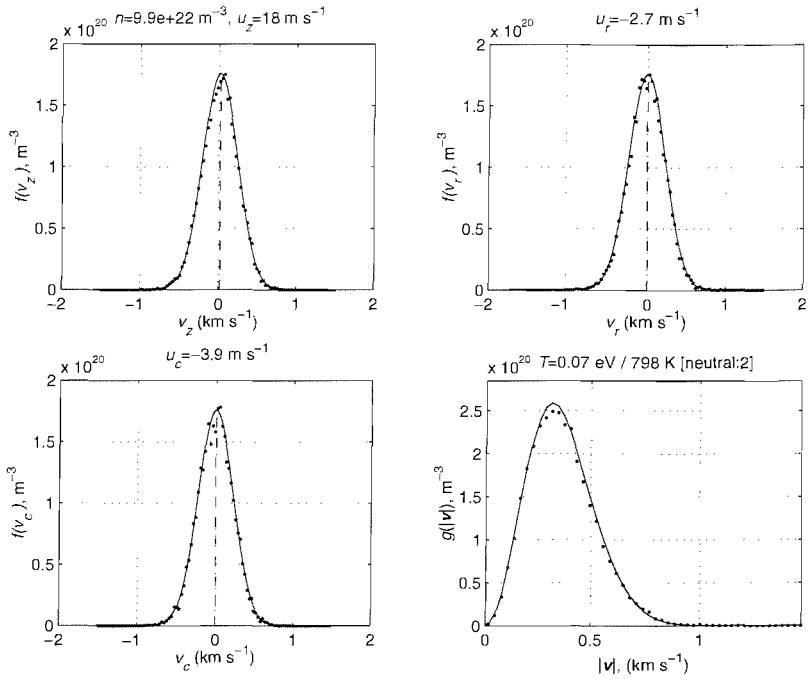

Figure 5.11: Velocity distribution plots of particles simulated using the hard sphere collision model near the upstream boundary.

Secondly, there is very little difference between the two collision models the VHS and HS, except that the VHS model takes somewhat longer to run.

The plot of axial stream velocity (fig. 5.10) appears to back up these sentiments: the expansion is incomplete in the collisionless case compared to the collisional ones. The peak velocity for both collision models is around 380 m s^{-1}, located 2 mm downstream of the tip. In fact, the logarithmic y-axis of the density plot was hiding the fact that the initial expansion near the upstream boundary ($z = 0$ to 0.5 mm) also differs from what is predicted by either of the collisional models.

Finally, in order to fully appreciate the difference in the flows as collisions are switched on, and to provide a little more evidence that the collision routines appear to be operating correctly, velocity distribution functions can be examined for the hard sphere and collision models. For sample point A (near the upstream boundary) these are shown in figures 5.11 and 5.12. Since this is this first time such plots are used in the analysis, a little explanation is required. These plots consist of four diagrams: the upper two are axial and radial velocity distribution functions (v_z and v_r) the lower two are circumferential velocity and velocity magnitude functions (v_c and $|v|$). The *solid lines* are plotted as per theoretical Maxwell-Boltzmann statistics [40], while the *data points* are real samples from the simulation. The Maxwellian plots are based on the sampled macroscopic temperature from the simulation and shifted according to the macroscopic sampled stream velocities. The *dashed line* is a line to show where the

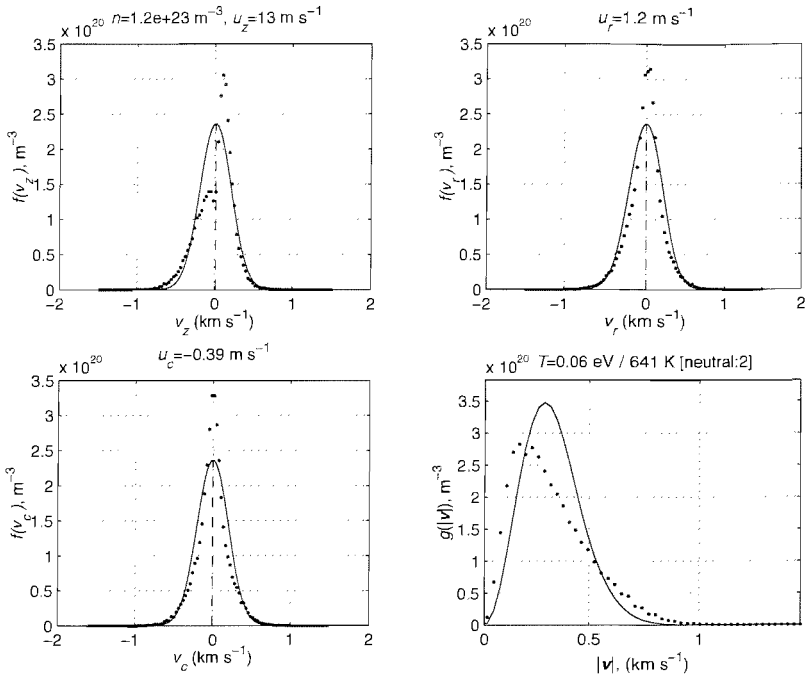

Figure 5.12: Velocity distribution plots of particles simulated using the collisionless model model near the upstream boundary.

stream (mean) velocity lies. Figure 5.11 is an excellent example to start with, because this is an example of a nearly stationary gas (the stream velocities are shown above the individual plots) in thermal equilibrium. More importantly, it is the collision routines that are causing the gas to tend to have a Maxwellian distribution in velocity: this is what would be expected for Xenon at this temperature. In this way, the fact the the data points lie on the theoretical lines is good because we would not expect a gas under these conditions to behave otherwise. Indeed, it could be interpreted as though the plots say that it would be fair to use a continuum representation of the gas here. As as the sample point is moved downstream, and the rarefaction increases, so the points will not tend to correlate with the lines and it would be certainly not be fair to use a continuum representation — this has all been covered in the section on the validity of the continuum assumption. What is interesting now is to compare with the collisionless plots in figure 5.12. Here, clear divergences from the Maxwell-Boltzmann distribution can be seen, particularly in the axial velocity. We have a two temperature effect due to the difference in temperature of injected gas and gas that has obtained the temperature of the insert wall. Such non-Maxwellian profiles are not uncommon in plasmas, particularly within non-neutral (beam-creating) regions. In this case, however, it is clear that collisions are very important and must be included.

This concludes the analysis of the model in terms of numerical accuracy and the collision model. This chapter now continues to consider the structure of the flow in

detail for the reference case, before evaluating the effects of varying some of the flow parameters and cathode geometry. The reference case and all other results that follow use the VHS model for neutral-neutral collisions.

5.3 Reference case: flow structure

This section provides a detailed analysis of the structure of the flow under reference case conditions. Following sections proceed to alter various parameters to quantify the effect they have on the flow character.

Since the flow is axially symmetric, the clearest picture is usually obtained by plotting values axially. Once these have been analysed, further plots including 2-D contour plots are also provided. Figure 5.13 shows axial centreline plots of density, temperature, pressure and Mach number. The numerical noise that can be seen near the downstream end of the plots (particularly in the plot of Mach number) is a sampling artefact: the data are sampled on a uniformly spaced PIC mesh, so due to the variation in number density, there are correspondingly fewer computational particles in the low density plume. Recall that this noise has no impact on Monte Carlo collision sampling since MCC calculations are based on a different mesh, see for example figure 5.5.

It can be seen that density falls monotonically throughout the length of the flow. There is an initial expansion from reservoir upstream conditions due predominantly to the fact that freshly entering gas on the upstream boundary has a temperature of 300 K while the cathode has temperature 1200 K. Within the insert section, the temperature is broadly constant, although the density falls steadily, as does the pressure. For a continuously expanding, accelerating flow in the positive z direction, there must be a negative pressure gradient to drive the expansion. This can clear be seen in the third from top figure and the final plot which shows axial stream velocity.

Once the tip is reached (located at $z = 10$ mm) there follows a powerful expansion. The density falls about two orders of magnitude within the space of 3 mm. In addition, there is a rapid cooling of the gas, and the scalar pressure drops three orders of magnitude. Such an expansion is of course associated with the flow becoming supersonic. The sonic point occurs precisely on the outer lip of the tip ($z = 11$ mm). This is the minimum area point of the flow, and so as expected, the hollow cathode is operating as a choked convergent-divergent nozzle. The expansion of the plume is not as steady and uniform as might be expected. This is entirely due to the presence of the keeper protruding into the plume. The keeper, whose reference temperature is set to 600 K, heats the plume by around 100 K (from approximately 350 K to 450 K). This directly affects the local sound speed and hence the Mach number is seen to

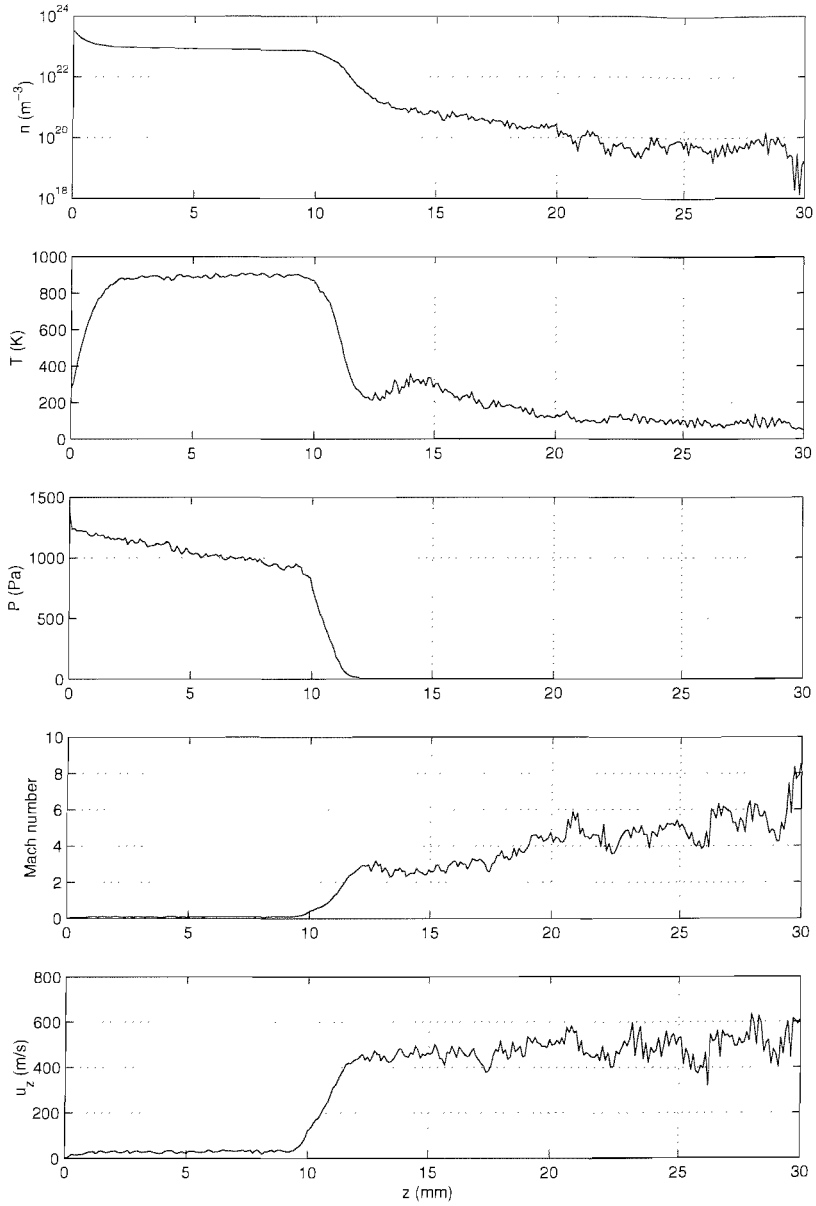

Figure 5.13: Axial centreline plots of number density, temperature, pressure and Mach number for the neutral gas flow in the reference case hollow cathode.

dip from just above 2.0 to a minimum of 1.5. Once the plume clears the keeper, the expansion continues. Recall that no downstream condition is applied: particles are simply extracted as though a hard vacuum is present. In this case, the density has dropped to around 2×10^{19} m^{-3} and the pressure to around 25 mPa. The vacuum chamber used for hollow cathode experiments at the University of Southampton (see any of Rudwan [88], Pottinger and Gabriel [80] or Gessini et al. [38]) with no propellant flow pumps down to about 10^{-7} mbar, equivalent to 0.01 mPa. When operational — when the cathode is supplied with 1 mg s^{-1} of propellant — the chamber pressure rises by a couple of orders of magnitude. This means that it is unlikely that the flow

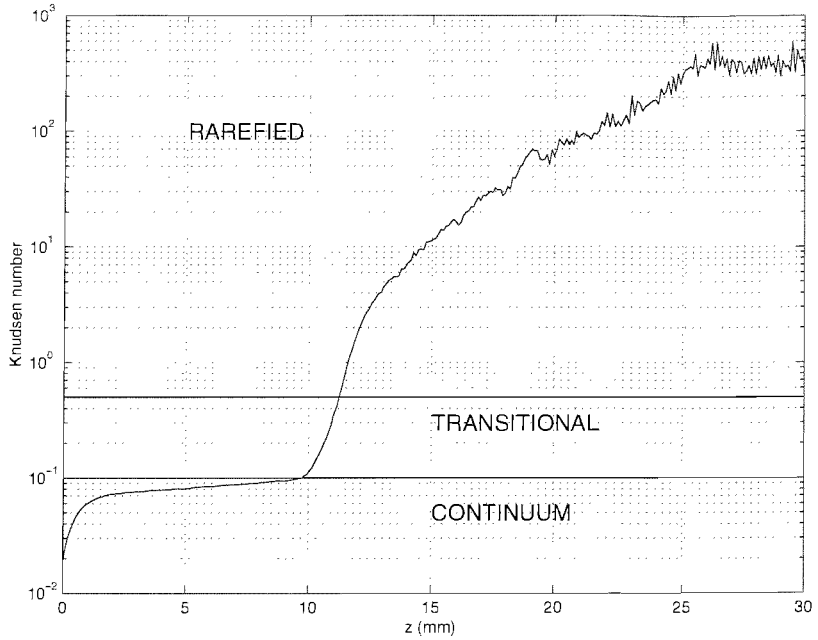

Figure 5.14: Axial centreline plot of Knudsen number for the hollow cathode under reference case conditions. The characteristic length is assumed to be equal to the tip radius: 700 μm.

described by the figure would be altered significantly by the application of a small downstream pressure condition.

The general flow structure is the same as that presented by Murray et al. [66]; i.e. a shallow drop in pressure between upstream boundary and tip, followed by a rapid expansion (with the sonic point located at the outer face of the tip). The work of Murray represents the best continuum results for cathode interior neutral flow. The peak Mach number reached is lower than that predicted by Murray et al. [66]. This can be explained for two reasons. First, the mass flow rate in the Murray et al. [66] reference case is 3.38 mg/s (more than 3 times larger than this case). Second, the Murray et al. [66] work included no keeper. Murray predicts that the model is likely to be accurate because at the flow rates he considers, rarefaction is unlikely to affect his results.

An important question to ask is: how rarefied is the flow? One of the justifications for the choice of particle method over continuum method was that the flow could be significantly rarefied in the cathode expansion. Figure 5.14 shows an axial centreline plot of Knudsen number based on a characteristic length of 0.7 mm (the tip radius). The figure shows that the flow within the cathode is continuum (Kn $<$ 0.1), while the expansion is transitional: tricks like slip flow wall conditions or the use of the Burnett equations would be reasonable here. In the far plume, the flow becomes very rarefied and approaches free-molecular. There is some debate over exact terminology;

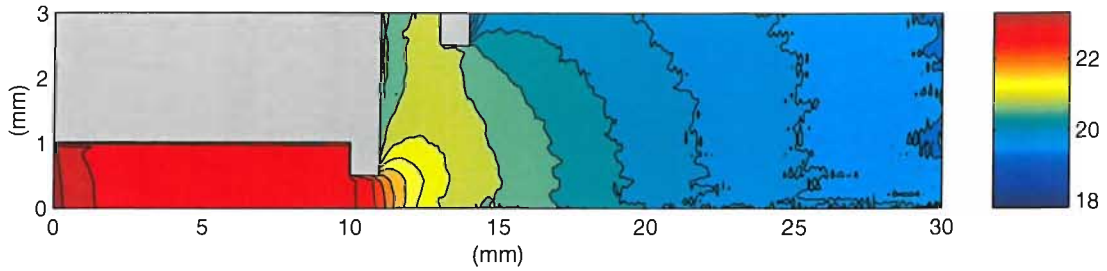

Figure 5.15: Contours of number density for the reference case.

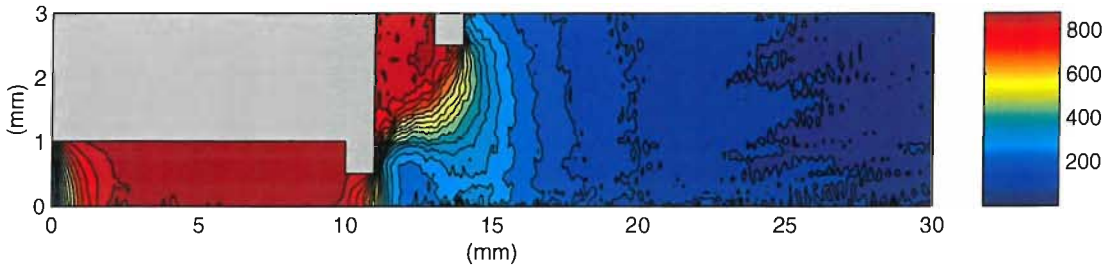

Figure 5.16: Contours of temperature (K) for the reference case.

some prefer to refer to the rarefied regime as $0.1 < \mathrm{Kn} < 10$, with the flow being collisionless for $\mathrm{Kn} > 10$. This is of no real concern: the justification for using a particle model was that the hollow cathode is a transitional device in the context of the rarefaction of the propellant and thus, a model that is independent of Kn is the wise choice. As was mentioned in the early discussions on the matter, the most computationally efficient approach would be to use a Navier-Stokes solver for $z < 10$ mm and a DSMC code for $z > 10$ mm (in this case). Unfortunately, the question of what to do at $z = 10$ on the interface between the models is one that is not fully answered yet[14]: the objectives of this research are based on further understanding of the hollow cathode; not on further developing hybrid continuum-particle methods for the sake of increased computational efficiency.

Axial centreline plots of flow data are the most informative, and easy to interpret. An overall picture of the flow structure can be provided using 2-D contour plotting. This can be seen in figures 5.15, 5.16, 5.17 and 5.18 that show density, temperature, axial stream velocity and scalar pressure respectively.

5.4 Mass Flow rate

It is regrettable that there is no consistency in the unit used to measure mass flow rate in hollow cathode studies. The two units that are used are SI base units, kg s^{-1} for example, and equivalent current (Aeq). The reason for the two measures stems from the fact that it is is often useful, when trying to make sense of the complex current-

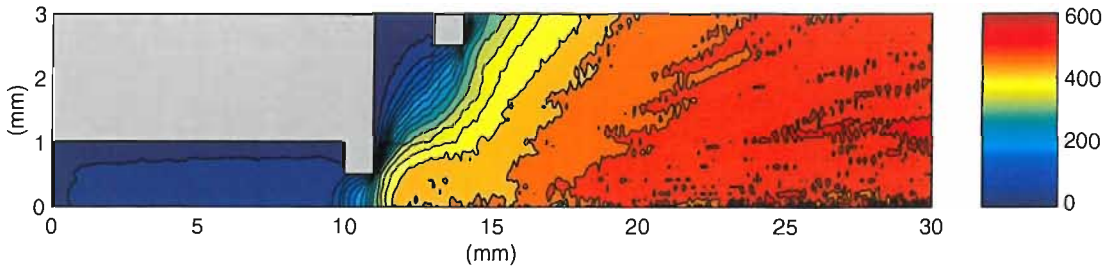

Figure 5.17: Contours of axial stream velocity for the reference case.

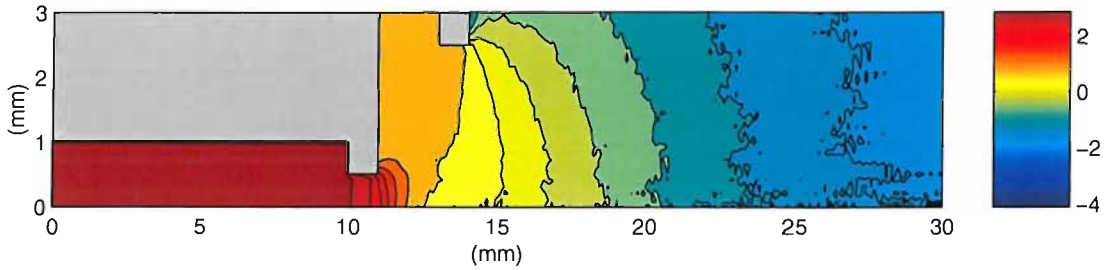

Figure 5.18: Contours of scalar pressure (Pa, log scale) for the reference case.

voltage-mass flux relationship in hollow cathode physics, to express neutral flow in terms of a current, so that current balances may be constructed for all plasma species. Thus, in theoretical studies of HC physics it is usual to find neutral influx quoted in equivalent Amperes. In terms of macroscopic performance (thrust, efficiency, etc...), mass flux is more frequently found quoted in SI units. In order that the reader may easily use values presented here in comparison with either measure of mass flux, the independent variable is presented in mg s^{-1} on the lower x axis and in Aeq on the upper x axis; this is of course simply a conversion factor e/m.

Values are chosen to cover the entire range of values given by Rudwan [88], from 0.5 mg s^{-1} to 4.0 mg s^{-1} (for Xenon) in steps of 0.5 mg s^{-1}. This is equivalent to Xe mass fluxes of 0.367 Aeq to 2.94 Aeq in 0.367 Aeq steps. There is one practical note worth making if the reader wishes to run sets of results while changing the mass flux. To a much greater extent than any of the other parameters examined in this chapter,

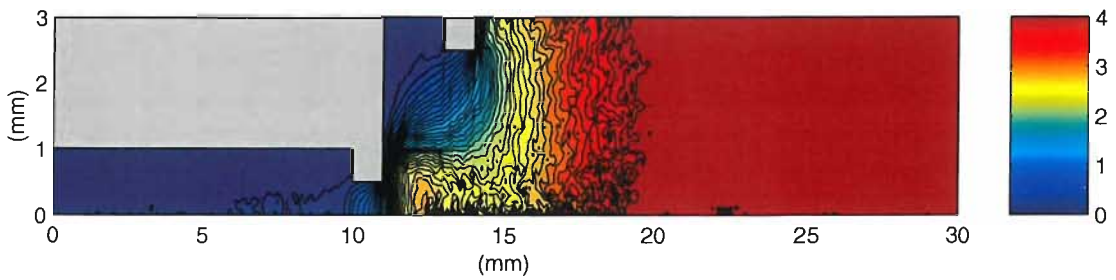

Figure 5.19: Contours of Mach number for the reference case.

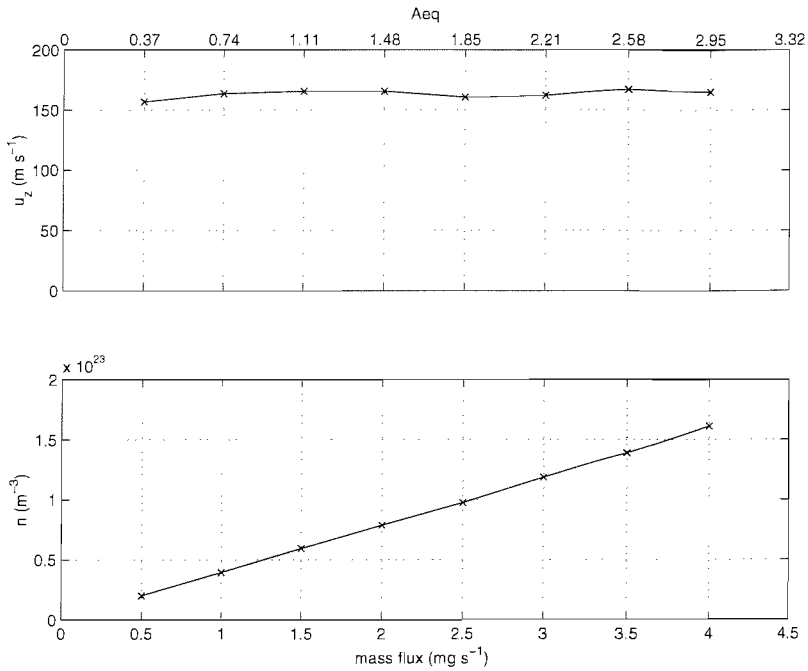

Figure 5.20: velocity and density plotted as functions of mass flow rate

increases in mass flux induce large increases in interior density and hence collision rate. The computation speed of Monte Carlo collision based codes is inevitably linked very strongly to the collision rate. In this case, for instance, the 4.0 mg s^{-1} case takes more than six times longer to run compared to the 0.5 mg s^{-1}. This example is a stark reminder of how quickly Monte Carlo based techniques become unusable as the continuum flow regime is approached.

Figure 5.20 shows data from the sample point near the tip. The two plots show axial stream velocity and number density respectively. The primary trend that is expected is that the density in the cathode interior will increase with mass flow rate. This is shown in the lower plot. In fact, the relationship between interior density and mass flux is very clear. It seems that density is proportional to mass flux. Given the very clear linear relationship, a simple expression can be constructed for the variation of mass flow rate with density. Based on doubling the mass flow rate from 2 to 4 mg s^{-1} yielding an increase in density from 7.87×10^{22} m^{-3} to 1.61×10^{23} m^{-3},

$$n = 4.12 \times 10^{28} \dot{m} - 3.6 \times 10^{21} \tag{5.9}$$

for $\dot{m}$ specified in kg s^{-1}. Of course, this expression is specific to the density at the tip of the T6 cathode with a 0.7 mm tip radius, so it is in no way generally applicable. What is shows is that in the absence of a discharge, the interior density scales in proportion to the mass flux. This is an important result that has been tacitly assumed for the hollow cathode geometry but never explicitly demonstrated.

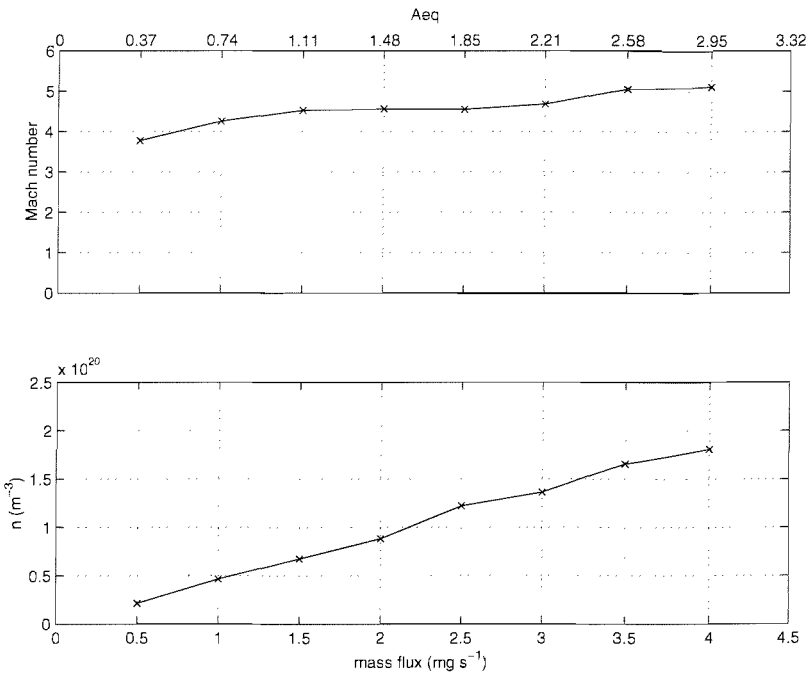

Figure 5.21: Mach number and density plotted as functions of mass flow rate

Figure 5.21 provides plots of Mach number and density for the sample point within the plume. Here, the same trend is shown for the density-mass flux relationship. The Mach number is seen to rise slightly. The absolute axial velocities rise as mass flux increases, but in addition, due to the higher density upstream of the expansion, plume temperature is actually cooler for higher mass flux. Since the Mach number is seen to rise slightly, this means that the rising stream velocity is marginally more strongly related to increasing mass flow rate compared to the fall in local sound speed due to the cooler plume. In summary: higher mass flow rates seem to result in an increasingly cooler plume at a higher absolute velocity and Mach number.

A note was made near the start of this section about the increase in compute time, particularly for the very high mass fluxes (greater than 3 mg s^{-1}). It was stated that as the mass flux and hence density increases, so the flow regime rapidly becomes more collisional: it approaches continuum flow. To what extent does this occur? Figure 5.22 shows Knudsen number at the three sample points as a function of mass flow rate. The Knudsen number is calculated based on the tip radius ($L = 0.7$ mm) from macroscopic flow data using

$$\lambda = \frac{RT}{\sqrt{2}\pi d^2 N_A p} \tag{5.10}$$

A critical continuum validity line is included in the plot, below which it would be expected that significant differences could be found between a Navier-Stokes and DSMC (or Boltzmann) solution. It can be seen that the Knudsen number falls monotonically

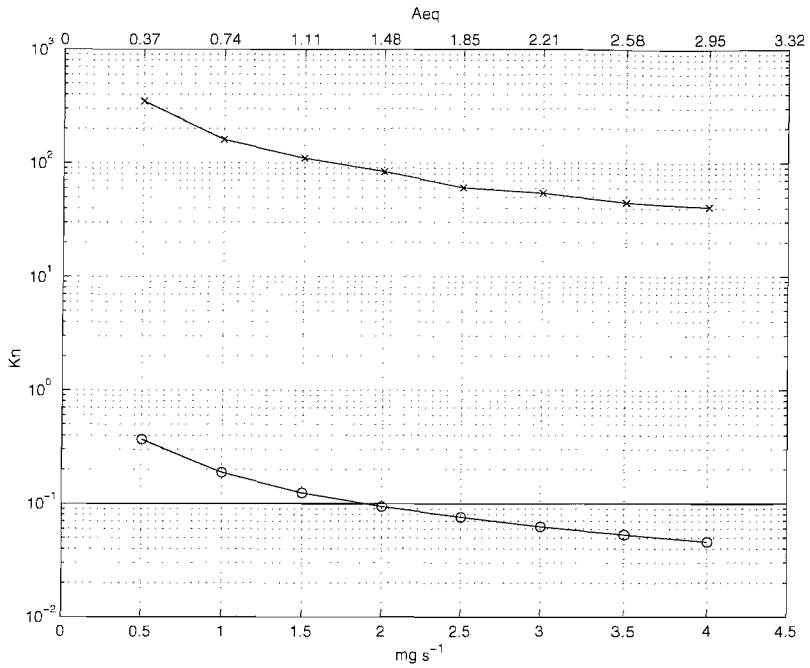

Figure 5.22: Knudsen number at the tip (o) and in the plume (×) of the cathode as a function of Xenon mass flow rate.

as mass flux rises. In the plume, the Knudsen number is always far outside the limit for validity of the continuum description. In the cathode interior, the flow falls within the Kn= 0.1 limit for mass flow rates larger than 2 mg s^{-1}. This is important as it shows that if a very high mass flow rate cathode (with a $\dot{m}$ range of 5 to 10 mg s^{-1}) was used, it would be safe to use a continuum fluid model to describe the flow. Indeed, from examining the figure, it seems as though the use of the Burnett equations [1] would probably be reasonable down to $\dot{m}$ = 1 mg s^{-1}. This is however not the full truth. The precise definition of Knudsen number is really that it is the ratio of mean free path to the *length scale of the smallest flow gradient.* In fact, the primary expansion at the tip of the cathode flow extends to less than the cathode radius in this case. In other words, the Knudsen number based on the correct resolution of flow gradients is probably higher than that shown in the fig. 5.22. The main point that it is possible to be precise about in this case is that the flow is *transitional* in the cathode and *rarefied* in the plume.

5.5 Tip radius

In experimental studies of hollow cathode physics, the tip is often drilled out to progressively larger diameters. In the context of the present analysis of neutral gas flow in the cathode, altering the tip diameter is likely to have a large effect on the flow

conditions. This is because convergent-divergent expansion is known from elementary theory to depend heavily on the minimum or throat area. Recall that in the introductory chapter of this thesis, §1.1.1, the specific impulse limit of a chemical rocket was derived, and some basic quasi-1D isentropic relations were given. Using these relations, it can further be shown that the area ratio of two parts of a variable area duct is a function of Mach number:

$$\left(\frac{A}{A^*}\right)^2 = \frac{1}{M^2}\left[\frac{2}{\gamma+1}\left(1+\frac{\gamma-1}{2}M^2\right)\right]^{(\gamma+1)/(\gamma-1)} \tag{5.11}$$

where A^* refers to throat (minimum) area and A the area at some other point. γ is the ratio of the specific heats. Complete derivation of the above is given by Anderson [2, §5.4]. It is clear to see that $M = f(A/A^*)$, so if the cathode tip radius is changed, so $A^* = \pi r^2$ will change and the nature of the flow could deviate significantly from reference case conditions.

The range of radii usually considered experimentally generally spans a minimum of approximately 0.3 mm to a maximum operational radius of 0.8 mm. The cathode tends to become very inefficient and difficult to operate as $r_t \rightarrow r_i$. This perhaps to be expected: it is assumed that as the tip is drilled away completely there is no geometric confinement of the gas or plasma adjacent to the insert and the neutral and plasma density that can be sustained drops significantly. The results presented imminently demonstrate this explicitly for the first time. In light of these factors, simulations are performed as per the reference case but varying r_t from 0.3 to 1 mm in 0.1 mm steps. Of course, for the purposes of analysis the absolute value of the tip radius is irrelevant: plotting results against the ratio of tip radius to cathode internal radius is far more meaningful. Accordingly, the minimum, 0.3 mm, corresponds an *area ratio* of 0.09.

It seems logical to start by examining the data sampled at the tip itself to quantify the effect of varying the tip radius. Figure 5.23 shows plots of axial stream velocity and density sampled at the sample point in the tip for the full range of tip radii (or area ratios). There are some clear trends. As the radius increases, the velocity that is obtained within the tip increases monotonically, with a plateau at just above 200 m s^{-1} for ratios greater than 0.5. The lower plot shows that the density decreases monotonically, and continues without a plateau for area ratios greater than 0.5. This plot (of number density) shows that the tip radius has a very significant effect on interior density. The density varies by more than an order of magnitude as the tip radius is increased from 300 to 1000 μm.

There is however a question that is raised by the configuration of these results. Since the sample point is located in the same position for all cases, the plots shown

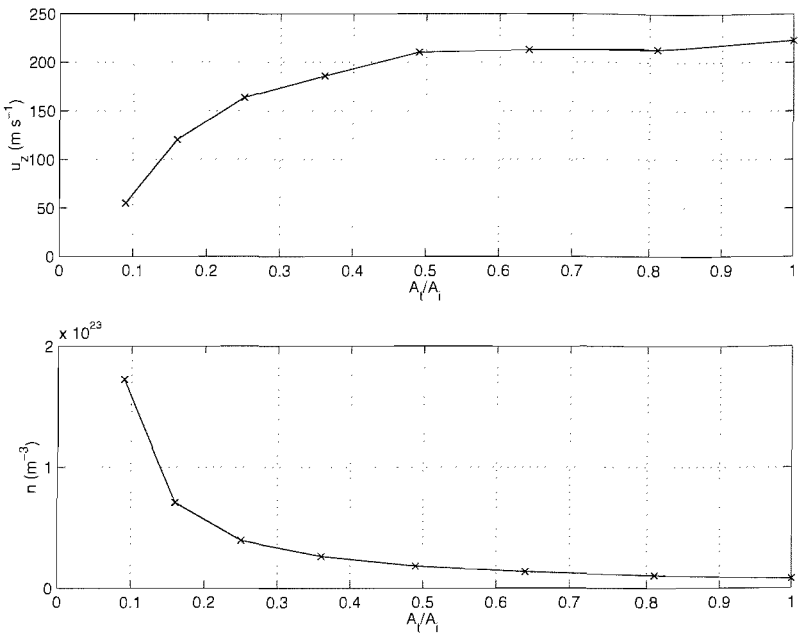

Figure 5.23: Axial stream velocity and density at the tip as a function of tip to cathode area ratio.

may not be telling the whole story. This is because if the tip radius is changed while leaving the sample point in the same place, if there is a radial profile to the flow data, then the sample point is recording both changes in the radial profile as well as the effect of the wider tip. It is very difficult to separate these effects. If the motion of the profile were known through further analysis, then the sample point could be moved accordingly so that it always occupies the same place within the profile. While this would seem the logical way to proceed, such analysis of the change in flow profile would be complex and there is a more efficient way to gain an accurate view of the effect of varying tip radius. Rather than go to the trouble of trying to normalise the effect of the moving profile as the radius is altered, it is instead best to examine data from the other sample points that are far removed from the tip: both upstream and downstream. These sample points should give an indication of the effects of varying tip radius without incurring error due to changes in the flow structure near the tip. This approach would be flawed if the nature of the flow was changing catastrophically as the tip was opened; it is not: the flow remains a stable and choked expansion for all values of r_t.

What should be expected to happen to the flow upstream and downstream of the tip as the radius is changed? It is reasonable to propose that we will see significant changes in the upstream (subsonic) region: an increase in axial flow velocity and a decrease in density as the throat area ratio rises. Downstream of the expansion, it may be that the flow is fairly similar regardless of area ratio.

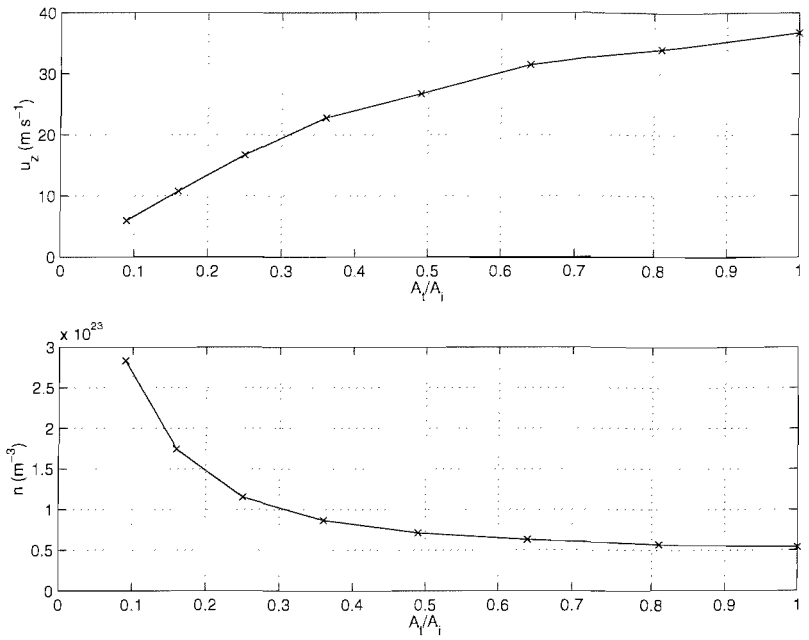

Figure 5.24: Axial stream velocity and density near the upstream boundary plotted as a function of tip to cathode area ratio.

Figure 5.24 shows conditions near the upstream boundary. Although trends are similar to those shown in fig. 5.23, there are some important differences. Axial velocity still increases monotonically, but the plateau region is now absent. This is strong evidence to suggest that the radial flow profile was interfering with the results shown in figure 5.23. The situation is mirrored in the lower plot (that of density). Here, the trend is the same, but the data point at $A_t/A_i = 0.09$ (the 300 μm case) lies more accurately on an exponential extrapolation from the other points (not shown here: this can be seen on a plot with log scale). These plots verify that there is some interference from the varying radial profile. It is worth considering this when tip radius variation is consider in the plasma discharge results.

It is worth examining plots of conditions downstream of the expansion to answer the question regarding the effect of cathode tip radius on exterior dynamics. Figure 5.25 shows velocity and density as per the previous two figures. The difference is very marked: the stream velocity does not vary significantly as the area ratio is changed. This was what was predicted: far downstream in the supersonic plume, there is little difference between a tip expansion at $A_t/A_i = 0.1$ and $A_t/A_i = 1$ since in both cases there is a stable, choked flow of gas at 1 mg s^{-1}.

The final consideration is to plot Mach number as a function of area ratio so that the code may be compared to the analytical relation given in equation (5.11). Figure 5.26 shows Mach number at the three sample points as a function of area ratio. The upper plot is for the plume sample point, while the lower plot shows both the

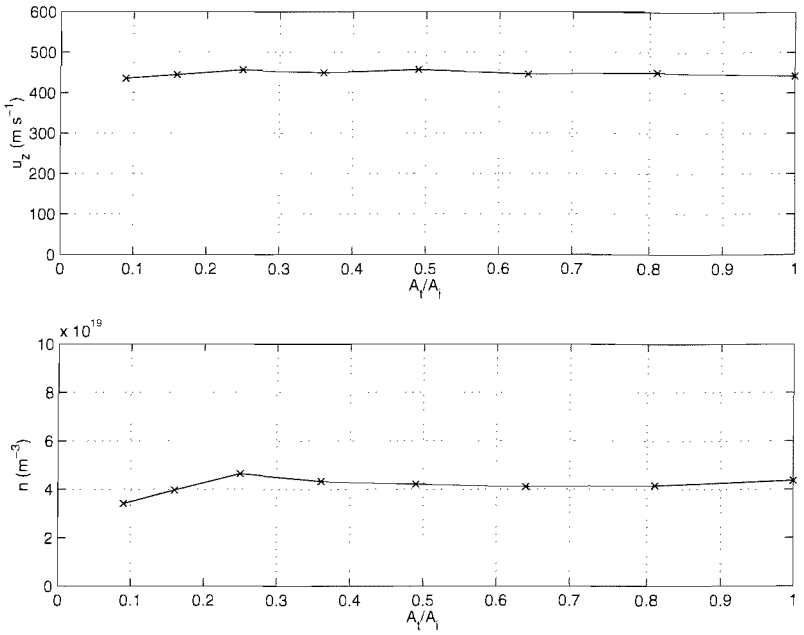

Figure 5.25: Axial stream velocity and number density sampled downstream of the expansion plotted as a function of tip to cathode area ratio.

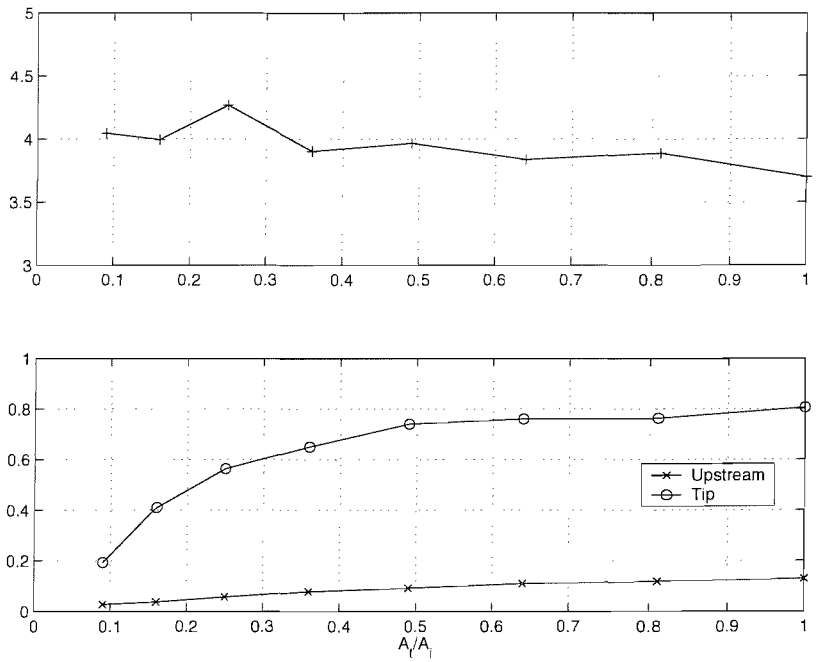

Figure 5.26: Mach number as a function of expansion area ratio at three different sample points. Upper plot shows plume Mach number (position C), lower plot shows interior and tip points (positions A and B). See fig. 5.1.

upstream and tip Mach number values. The trend here is broadly similar to that presented previously: there is very little variation in Mach number in the plume. In fact there is a slight *fall* in Mach number as the area ratio is increased. Since it has already been shown that there is almost no variation in axial stream velocity in the plume for the difference area ratios, the slight drop in Mach number implies that the gas must be hotter for the wider area ratios ($M = v/a$ and v remains approximately constant, so since $a = \sqrt{\gamma RT}$, falling M with A_t/A_i implies rising T). In other words, the plume is warmer for wider tip cases. This is possibly due to the fact that the expansion becomes weaker as A_t/A_i increases, so the post-expansion temperature does not drop so much.

In the upstream region, there is an increase in Mach number with rising area ratio. Recall that for these sample points, the stream velocity increases with area ratio; since the trend is very similar to the previous plots, it is reasonable to first conclude that a, and thus T is not changing significantly as the area is changed.

5.6 Gas

Another parameter that is clearly a candidate for examining is the propellant. Xenon is generally used as the propellant for ion thrusters because the thrust per ion generated at a given accelerator grid voltage is greater (since the atomic mass of Xenon is greater than both Argon and Krypton). Efficiency of ion-bombardment ion propulsion is strongly related to the ion creation cost in the main chamber [16]. Xenon is, however, the most expensive gas commercially, and either Krypton or Argon are often used in experimental studies to reduce cost. Although all three propellants are similar in that they are all inert gases, cathode performance will necessarily vary significantly at the same injected mass flux. This is because the equivalent current of the lighter propellants is greater for a given $\dot{m}$.

All three propellant gases are run under reference case conditions. Since all three are inert gases, the ratio of the specific heats is unchanged at $\gamma = 5/3$. The only differences between the three gases are then the mass and collision cross section (diameter and viscosity-temperature index). The data are contained in table 3.2, on page 54.

Figure 5.27 shows axial centreline plots of key gas conditions for the three gases. The results are as as would be predicted. Since the mass flow rate is fixed for the three gases, it would be expected that the gas with the lowest atomic mass would exhibit the highest number density within the cathode. The top plot in figure 5.27 shows this to be true: the highest number density is Argon (atomic mass 18) followed by Krypton (36) and then Xenon (54).

There is almost no discernible difference in the gas temperature between the three

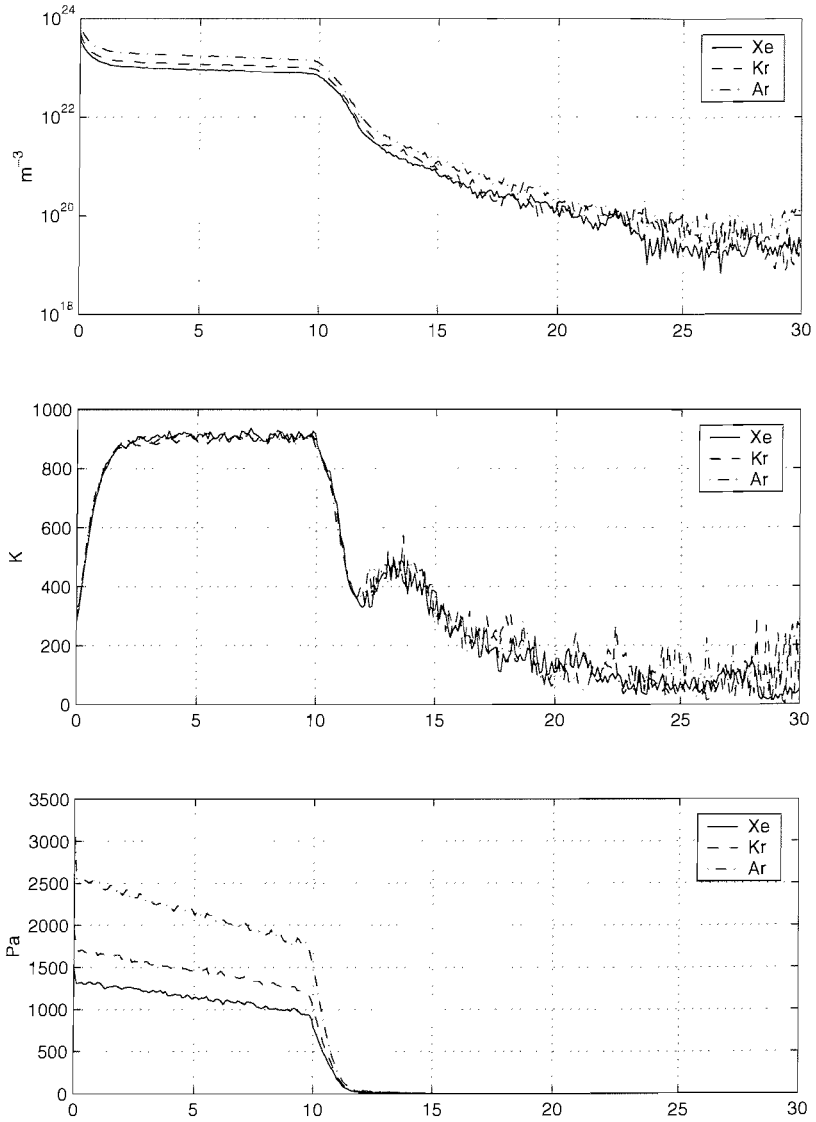

Figure 5.27: Comparison of gas conditions for three different propellants at reference case configuration. The upper plot shows density, the middle plot temperature and the lower plot scalar pressure. All data are axial centerline values.

cases. This is because the model makes no provision for an external thermal simulation. Recall that walls are assumed to be isothermal, and when set to full accommodation (with respect to wall-particle collisions), gas molecules will always be returned from an encounter with a surface with temperature equal to the surface temperature. If the code were modified so that the heat flux to the gas was regulated (rather that the absolute temperature being held constant), then there could be differences in temperature between the gases. Although the code is unable to use external thermal modelling, the heat fluxes through the walls are recorded, so that the magnitude of this difference can be observed.

Figure 5.28 demonstrates this. By default, heat flux is recorded for each surface

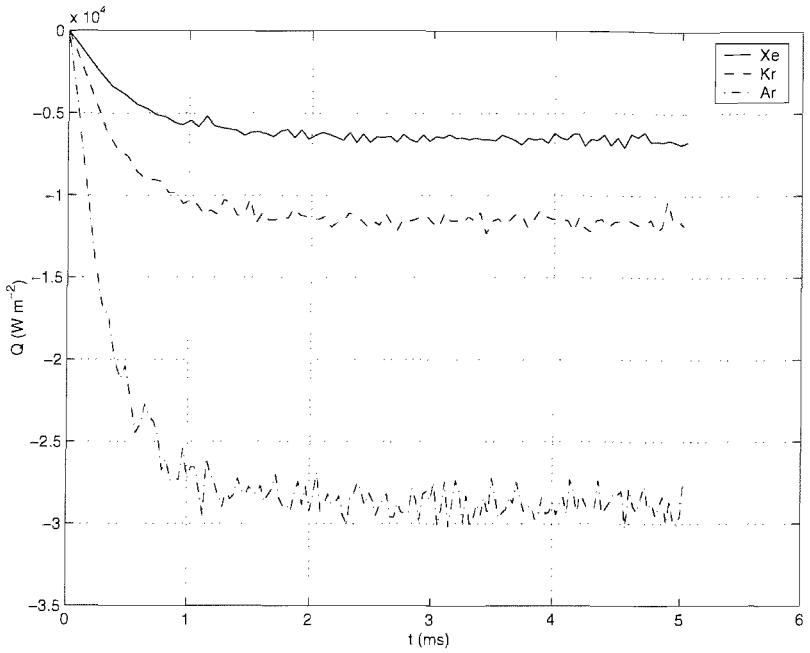

Figure 5.28: Time histories of heat flux between the cathode insert and working gas for three different propellant at a fixed mass flow rate and wall temperature.

present in the simulation: this is done via a simple normal energy balance that is calculated upon every particle-wall interaction. The figure shows the heat flux for the cylindrical insert surface. The plot is similar to the time-histories examined previously. In fact, a record of heat flux could be considered as valid a candidate for judging convergence as exhaust mass flux or collision rate. The convergence can be seen in this plot: for the reference conditions the simulation is run for 4 ms, at which time it is assumed the flow is steady. Data is then sampled for a further 1 ms. A negative surface flux denotes a gas that is cooler than the surface (the wall is heating the gas), a positive flux indicates a surface that is cooler that the gas, so the wall is taking heat from the gas. The plot shows that the flux *from the insert to the gas* in the three cases is stable at approximately 7.0, 11.5 and 29.0 kW m^2 for Xenon, Krypton and Argon respectively. Do these values seem realistic? A quick calculation may be performed to estimate the heat flux required to heat Argon based on the specific heat. The area of the insert is $2\pi r_i z_i = 62.8 \times 10^{-6}$ m^2 for a 10 mm long insert section. This means that for the Argon case, the power input to the gas is 1.82 W. How does this compare to the minimum predicted heat flux based on the specific heat and volume of gas? The density is approximately 3×10^{23} m^{-3}, so the total mass of the gas resident in the insert section is 6.1×10^{-10} kg (based on a volume of 31×10^{-9} m^3 and the mass of an Argon atom: 66.3×10^{-27} kg). The temperature change between the injection temperature (300 K) and free stream insert temperature (approx. 900 K) is then approximately 600 K. Given the specific heat of Argon at constant volume

is 0.3122 kJ kg^{-1} K^{-1}, the minimum energy required to heat the mass of gas present is 0.12 mJ. Of course, the gas is resident within the cathode for only a fraction of a second before being completely replaced: the stream velocity is of the order of 30 m s^{-1}. This implies a residence time of the order 10^{-4} s, so that the 1.82 W power input is equivalent to 0.18 mJ energy flow to the gas. This is very similar to the minimum predicted by elementary physics of specific heats above (0.12 mJ), so the heat fluxes quoted seem reasonable. In particular, the energy flow to the gas exceeds the minimum required to heat the gas by 600 K. This is a good verification of the correctness of the fundamental thermodynamics reproduced by the code.

The difference between the values for Q confirms the suggestion that the heat flux would be different between the gases, and further, this trend makes sense. Since the number of Ar atoms is greatest, this means that the Argon flow at 1 mg s^{-2} requires a greater heat flux to reach a given temperature. In practical terms, this means that the lighter gas at a fixed mass flow rate (and hence higher number density) is more effective at cooling the cathode.

The final point to consider is the third plot in figure 5.27: pressure. This is calculated using the ideal gas law $p = nkT$ (n and T are sampled directly from raw particle data). At the given mass flow rate, and, due to the isothermal walls, fixed temperature, the pressure is progressively greater for the lighter propellants. It is clear to see that this must be the case not simply by definition through evaluation of the gas law, but in a fundamental way since the heat flux is much greater for the lighter propellant.

A final note that is worth including is to state that there is no provision for heat transfer via thermal radiation. This is examined at greater length with respect to these and other results in the discussion chapter.

5.7 Temperature

In light of the findings relating to the use of isothermal walls and heat flux, it is necessary to examine the relationship between wall temperature and free stream flow temperature. For this, the reference case flow is computed at a range of temperatures typical of cathode operation. A range of 900 K to 1500 K in 100 K steps should give a broad impression of the relationship between cathode body temperature and flow structure. This range covers both the initialisation (breakdown) temperature regime and the main discharge regime. Specifically, Rudwan [88], in experiments aimed at determining mininum breakdown temperatures, was working in the range 900 to 1100 K. Typical cathode operating temperatures are estimated to be in the 1300 to 1500 K, although precise experimental measurements are sparse. Finally, a data point is

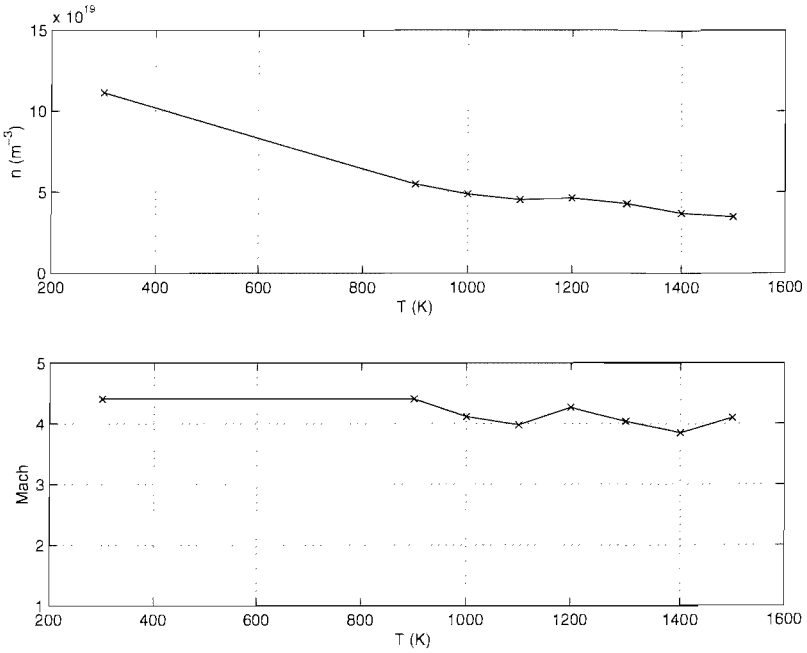

Figure 5.29: Density and Mach number in the cathode plume plotted as a function of cathode temperature.

computed for a cold case, a room temperature 300 K. This additional cold case is included because recently Gessini et al. [38] has been conducting experiments to measure the thrust produced by the hollow cathode. To calibrate the thrust measurement experiment, a cold gas example is run.

From a practical standpoint, computation time was found to increase as the temperature is raised, although the change is nothing as drastic as for the variation in mass flux. Compute time is increased by approximately 35% for the simulation at 1500 K compared to the 300 K case.

Figure 5.29 shows plots of density and Mach number in the cathode plume as temperature is changed. Plume density shows a steady decrease as cathode temperature is increased, while Mach number remains broadly constant, although there is a slight drop in Mach number at the higher temperatures. Figure 5.30 shows pressure and temperature in the plume. By comparing the trend in temperature to that in Mach number, it is clear that as the cathode temperature is raised, so the plume temperature increases, but due to the increased interior pressure, the expansion is stronger yielding a higher absolute plume velocity. For comparison, the plume velocity at the sample point for a cathode at 300 K is 240 m s^{-1}, while the plume velocity for the 1500 K case is more than twice that value, at just above 500 m s^{-1}. The variability toward the high values of temperature is almost certainly a result of statistical scatter: noise. Data in this region is best interpreted as though a trend line were used.

The relationship describes above is of course classic nozzle behaviour. The cathode

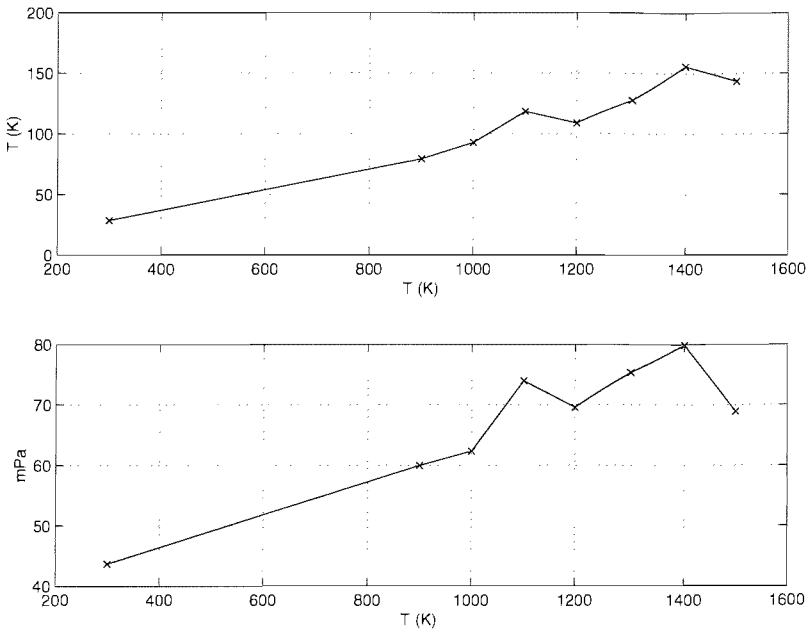

Figure 5.30: Scalar pressure and temperature in the cathode plume plotted as a function of cathode temperature.

temperature can be considered the 'combustion chamber', while the tip is effectively a convergent-divergence nozzle. As was shown in the introductory part of this thesis the exhaust velocity of a con-di nozzle is effectively just related to chamber temperature if the pressure ratio is close to zero. Recall equation 1.12 from the introduction:

$$v_{ex} = \sqrt{\frac{2\gamma RT_0}{\gamma - 1}\left[1 - \left(\frac{p_e}{p_0}\right)^{(\gamma-1)/\gamma}\right]}. \tag{5.12}$$

It is possible to numerically evaluate this expression and compare it to the recorded v_{ex}. T_0 is set to the temperature at the sample point upstream of the tip (i.e. the 'chamber' temperature). p_e and p_0 are simply scalar pressures evaluated using $p = nkT$ for density and temperature in the plume and chamber (subscripts $_e$ and $_0$ respectively). Since there is a vacuum downstream condition (particles are simply extracted from the simulation), the exit pressure (see fig. 5.30) is of the order of tens of mPa. As has been seen, for the reference case, chamber pressures are approximately a kPa. This means that the pressure ratio p_e/p_0 is of the order 10^{-4}, and hence the exhaust velocity should not be very sensitive to pressure, rather it will be nearly totally dependent on chamber temperature. Figure 5.31 shows both direct velocity data, and equation (5.12).

The first thing to notice is that both data show a very similar trend. The direct numerical result underestimates the exhaust velocity compared to the analytical relation

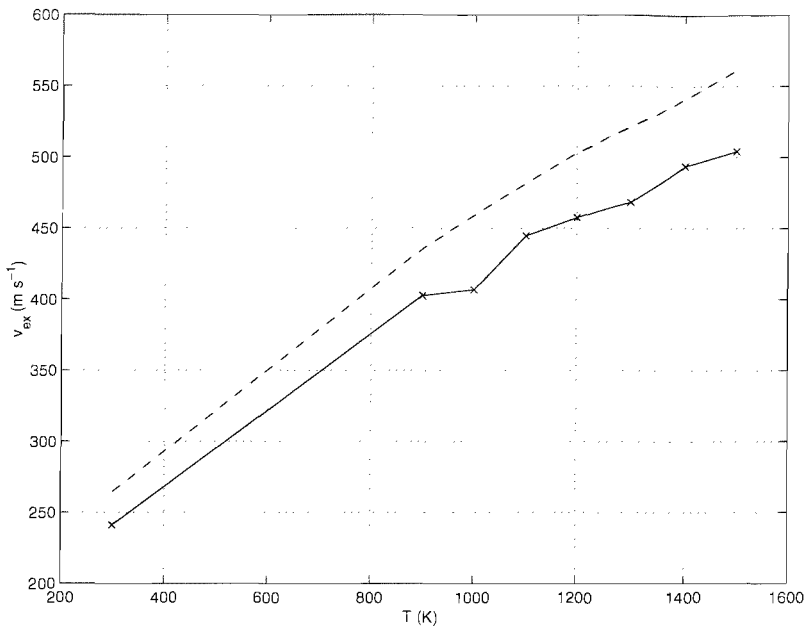

Figure 5.31: Exhaust velocity as a function of chamber temperature. The dashed line is numerically evaluated based on chamber temperature, chamber pressure and plume pressure, while the solid line is directly sampled velocity data from the numerical simulation.

by about 10%. This is of course to be expected: equation 5.12 takes no account of the presence of the keeper interfering with the plume. In fact, it should be expected that the value produced numerically underestimate v_{ex} compared to the analytical result. There is also a slight divergence between the analytical and computational results at higher temperatures. This can be further explained due to keeper interference, or more specifically, the fact that in the simulations, the cathode surface temperature is varied, but the keeper temperature is held constant. Little is known about the keeper temperature during discharge, not least because the thermal environment it is placed within is quite complex. The keeper is never seen to glow [88] during normal operation (unlike the cathode tip) and the keeper is necessarily thermally insulated from the cathode itself, since by definition it must be electrically insulated. In any case, with respect to the divergence in v_{ex} with increasing T_c, the presence of a constant temperature 'cool' keeper is an adequate explanation. In terms of the accuracy of the code, the reproduction of the expected nozzle expansion trend like this provides good evidence that the code is valid.

Chapter 6

Results: Plasma discharge

The results in this chapter can be considered a continuation of those from the previous since here, the neutral gas flow is retained, while plasma modelling is included. The code is the same as that used in the previous section, except that the plasma elements are now included. This means that the level of confidence built up in the code through validation in the neutral flow chapter applies here, although of course the new parts that handle plasma mechanics need to be carefully considered.

The aim of the results presented here is then to investigate the hollow cathode interior plasma, and experiment with changes to parameters and geometry.

Since this chapter is a continuation, this introduction can be very brief, as there is no need to cover the detail of geometry and reference case conditions; these are the same. The chapter begins by examining the structure of the discharge produced. This is followed by an examination of the current-voltage behaviour of the configuration, and then a comparison to experimental data. In order to assess the validity of the numerical model, the results of some numerical studies of computational parameters are then presented. Finally, the effect of varying the anode position and tip radius are examined.

6.1 Additional configuration

There are several additional configuration parameters that now need to be set due to the introduction of the plasma. It is necessary to choose the modification factor to the electron mass (see §3.5.6), and choose a smaller time-step compared to the neutral case. Of course, any possible effects of using these computational parameters are examined first, before presentation of the main results. In the plasma model, an arbitrary number of anodes can be included in the geometry if desired; these are cross-referenced to the surfaces and mesh segments. For each additional anode, the 'current demand' and capacitance must be set, along with initial voltage. The current demand

of an anode, in the computational context, can be imagined as though it is the current selected on a power supply attached to the anode. Choosing these values was done by trial and error, until values that were stable were found. In addition to the criterion that parameters provide numerical stability, a secondary motivation in the choice of such values is that they tend to induce a rapid convergence to equilibrium. In terms of electron emission, the surface temperature is simply used as it was in the previous chapter, although the work function of the surface now needs to be set. Common estimates are that $\phi_{eff} \approx 2$, so a value of 2.0 eV is used.

Next, additional surface boundary conditions need to be set, including particle impact values such as recombination fraction and secondary electron emission fraction yield. The value of the ion recombination fraction was discussed in §2.3.3 and it was concluded that strongly emitting surfaces (such as the insert) are liable to have a value between 0.5 and 1, while metal surfaces with comparatively high work functions (such as the cathode casing) may be considered as nearly insulating, so that the value is likely to be close to 0. Based on the reasoning in §2.3.3, then the standard configuration is for the insert surface to have $\delta_i = 0.5$, while surfaces making up the cathode casing have $\delta_i = 0$. As for secondary electron emission, it was concluded in §2.3.3 that secondary electron emission due to ion bombardment can be assumed to be zero, due to the comparatively low energies encountered, while based again on electron energy arguments from §2.3.2, the SEE electron yield from the cathode surface will not tend to exceed 1. Thus, by default, $\delta_{see,i} = 0$ and $\delta_{see,e} = 1$ for the cathode, while all other surfaces are assumed not to yield secondary electrons.

Next, the PIC mesh needs to be defined. In the neutral gas flow chapter, the PIC mesh was used for recording 2-D data because the PIC code includes the 'smooth' CIC interpolation sampling mechanism. Now, however, the PIC mesh is not simply a luxury that can be used to generate attractive contour plots. When running a full plasma simulation, it is necessary to ensure that the PIC cells are fine enough to resolve the local Debye length at any given point within the simulation. Finally, a large amount of extra species data is required by the Monte Carlo Collision routines. This includes collision cross sections for electron-neutral collisions, neutral-ion charge exchange collisions and the necessary parameters for Coulomb collision modelling (see §3.4). At this stage, a lack of data precludes the use of Krypton and Argon, since CEX collision cross sections are only known accurately for Xenon[82]. For this reason, since real data is available only for Xenon, the results use this gas. It is of course possible to run the code with the other two (Kr and Ar); in this case, the code will assume the CEX collision cross section is simply double the hard sphere cross section in the absence of real data; this approximation seems reasonable based on the data available for Xenon[60], for the ion energy range expected.

Parameter	Value	
cathode T_c	1680	K
insert ϕ_w	2.0	eV
start anode voltage V_a	100.0	V
anode current I_a	10.0	A
anode-cathode separation	5	mm

Table 6.1: Discharge: reference case parameters.

Additional parameters for the reference case discharge are shown in table 6.1. The flat plate anode is located on the downstream boundary. This of course alters the downstream neutral gas plume to some extent. Reference case anode and keeper electric parameters are set as shown in the table. According to some authors [88], the temperature is perhaps a high estimate for cathode temperature, but this is mitigated by the fact that the anode current is set to 10 A, so the cathode would tend to be hot. Of course, there is much debate over this, Malik and Fearn [56] estimate an internal surface temperature of approximately 2000 K based on analysis of thermal properties of the insert chemicals. The problem lies in the the fact that many analyses of cathode temperature are based on arguments relating to the current density predicted by the Richardson-Dushmann equation. Unfortunately (a) the work function is not known exactly and (b) the equation is highly non-linear in T. A combination of these factors make prediction of temperature a very difficult task. The only realistic way serious accuracy can be obtained would be for direct measurement to take place. A temperature of 1680 K is not required for running the cathode at lower current. The chosen startup voltage and default cathode-anode separation are as a result of the investigation into how to cause a rapid breakdown and hence equilibrium state of the plasma.

The final element of the discharge is the introduction of the neutral flow. Whenever the code is run, a restart file is written that, in addition to auxiliary information, simply contains details of each particle at the final time-step: position, velocity, species, weighting and excitation level (if required). In this way, incomplete simulations may be restarted. In addition, for the simulation of a discharge, the neutral flow can be modelled in the absence of the plasma at a large time-step, then the restart file of the neutral gas used when the full discharge simulation is run. This avoids the need to re-run the neutral simulation to equilibrium before electron emission is enabled.

Having completed the discussion of additional parameters that are set for the discharge above those set in the previous chapter, the next step is to investigate what happens when electron emission is enabled with no working propellant: this is the 'pre-breakdown emission' stage. Alternatively, we could have introduced the neutral flow while maintaining a discharge voltage below the first ionisation energy of the

propellant, so that now plasma would form.

6.2 Initialisation of the discharge

Before stepping directly into the analysis of the full, steady-state plasma discharge, it is necessary to briefly explain and present some details relating to the pre-breakdown transient state. In particular, what happens in the absence of neutral particles (an 'empty' discharge) and then what happens in terms of starting and sustaining the plasma.

6.2.1 Pre-breakdown thermionic discharge

It is of course possible to operate the cathode without any propellant: the body is heated and a thermionic (but space charge limited) current will flow from the cathode. In the absence of a plasma (any positive charge), thermionically emitted electrons will cause the the electric potential in the space upstream of the tip to be negative with respect to the cathode. This is simply the case of space charge limited current emission from the interior of a cylinder. It is worth examining this setup as this represents a data point for $\dot{m} = 0$, and it shows the nature of the vacuum emission.

Since we know that no plasma can form, it is futile to require that the anode demand 10 A as is set for the reference case. The cathode casing will emit mA of current, while regardless of the temperature of the insert, space charge limitation will retard the cathode current to the mA range as well. In other words, at a given cathode temperature, the casing gives up very little current, and the interior of the cathode is under space charge limitation: in the absence of a plasma, the anode will never receive 10 A (unless, for example, the temperature were set to an unrealistic value, millions of K, say). Instead, as was done for some of the verficiation cases of the PIC model (§4.5), the anode voltage is held constant and current received is recorded. The configuration behaves as expected and forms a space charge limiting sheath on the surface of the insert. As a result, under reference conditions, the anode receives only about 150 mA. Increasing the temperature does not increase the anode current above 150 mA: this means that under reference conditions in the absence of a plasma or source propellant, the discharge is Child-Langmuir law current limited. Of course, the anode current-voltage characteristic will not exhibit exactly the same trend as our original expression — eq. (4.35) — for one obvious reason: the cathode is no longer a flat plate located opposite the anode, it is now a confined cylinder.

Figure 6.1 shows time history plots of electron population and anode current. Here, electron population is defined as simply the total number of electrons present in the simulation volume at a given time. The electron population builds up very

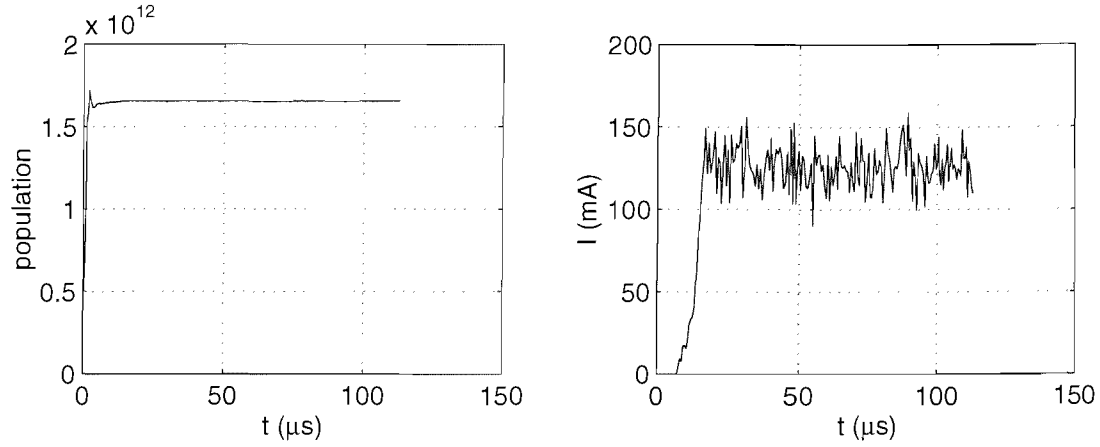

Figure 6.1: Time histories of electron population and anode current for the no-propellant discharge

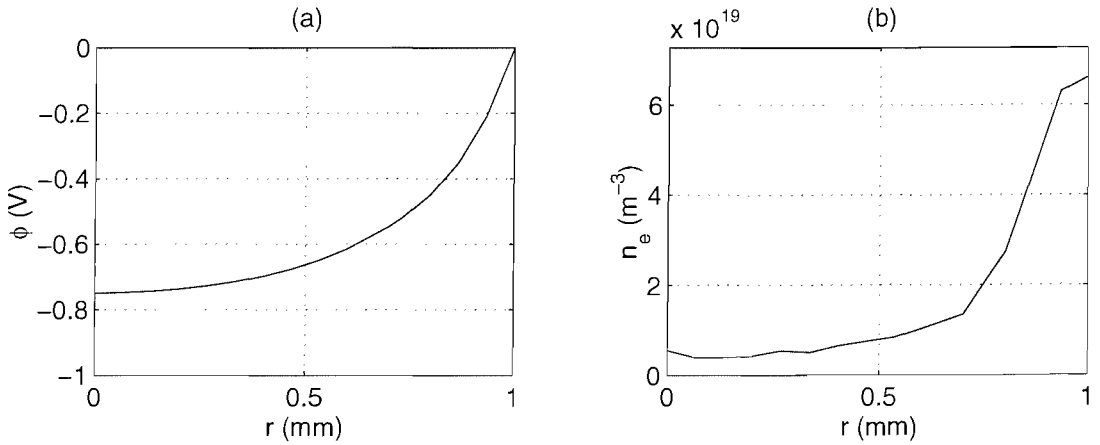

Figure 6.2: Radial cut showing plasma potential and electron density at $z = 10$ mm (mid way along the insert axially). The space charge limiting sheath can be seen.

rapidly, forming the space-charge limiting sheath on the surface of the insert. It then takes an additional 20 μs for the anode current to become steady. At this point the entire simulation volume is in equilibrium.

An impression of the radial profile of the limiting sheath in the cathode interior can be seen in figure 6.2. Here, fig. 6.2(a) shows that the presence of the electron gas in the interior cavity causes a negative electric potential and an electric field that tends to force emitted electrons back toward the cathode. This effect gives rise to the density profile seen in fig. 6.2(b). The on axis potential — approximately −0.8 V — is uniform axially for the entire length of the cathode cavity. At the tip, only the presence of the adjacent anode at a potential 5 V above the cathode serves to deflect the on axis plasma potential above -0.8 V. This data was generated using a permittivity modification that allows for a plasma of density 10^{20} m^{-3} to be simulated

with stability, on the reference mesh (150×60 cells in the area 15×3 mm). Recall that there would be little value in comparing the current here with the Child-Langmuir result, because a orificed cylindrical geometry C-L result is not known.

6.2.2 Breakdown

The process of breakdown in the hollow cathode is complex and not fully understood. In particular, some authors have found that discharge initialisation occurs at voltages lower than are predicted theoretically. Reasons for this include the possibility of Penning ionisation — the reader is referred to the thesis of Rudwan [88] for further details.

Initiating the discharge at low temperature (1000 K) — similar to that predicted experimentally — was, as expected, found to be difficult. This did not pose a problem however as the range of current values that are to be examined (5-20 A) demand a hot (2000 K) cathode. In this case, the cathode is set to the hot discharge temperature as an initial condition, and the discharge tends to form rapidly without exception. It would certainly be possible to conduct a future study of breakdown physics using this model where the temperature of the cathode is gradually raised from some initial low value until breakdown occurs. This would provide data on minimum breakdown temperatures. In all data presented here, however, we are interested in gaining a stable equilibrium discharge in the minimum time so that (a) sampling time of the equilibrium discharge is maximised and (b) compute time is minimised. In addition, since it is observed that the initial plasma forms outside the cathode, as is generally accepted, it was found that the speed of the breakdown was enhanced by raising the neutral density in the cathode exterior. This can be done in one of two ways. First, the neutral expansion can be confined radially by placing a reflecting boundary on the radial exterior limit. This is very similar to creating a 'closed keeper' type arrangement. Secondly, a fixed neutral density may be set on the downstream exit boundary and neutral gas introduced as per the method detailed in §3.3.3 (Bird's introduction of a free stream gas in the DSMC). It was found that the second method was the least intrusive and actually most similar to many experimental setups. This is because in a experiments, there is generally a non-zero 'vacuum pressure' present. In fact, Rudwan [88] states that for operation, the hollow cathode requires a background pressure of the order of 10^{-7} mbar, while the actual chamber reached pressures of 10^{-4} mbar (7.5×10^{-5} Torr). This is equivalent to densities of the order 10^{19} m^{-3} at room temperature. The result of replicating the experimental environment is that breakdown becomes easier in the model, and hence this simulated environment is used for all results presented here.

6.3 Numerical Validity

As was the case previously, it is necessary to first asses the impact of the theoretical assumptions made. Although it is possible to derive some confidence from the previous results section, *the code will not be free from non-physical numerical effects.* This is true for two reasons: first, components that have already been studied are now being used under different conditions, both numerical conditions such as time-step and physical factors such as increased temperatures. Secondly, there are elements present in the PIC code such as the PIC mesh, and numerical effects such as the electron mass modification.

In light of these concerns, several areas are now examined. In particular, the time-step dependence must be re-assessed as there are now plasma oscillations and high energy particles present, and because new charged particle collision models are introduced that will have different collision frequencies compared to the neutral flow. Secondly, the effect of the PIC mesh must be assessed. Although it was noted that long term numerical heating on poor meshes is likely to be minimal since the problem at hand is fundamentally a DC discharge (see §3.5.4, p. 75), it is still necessary to investigate the possibility in any case. Finally, the factor by which the simulation is accelerated computationally (the electron mass modification) must be considered.

6.3.1 Time-step

Holding all other parameters constant, including of course the real simulation time, calculation of the plasma is made using four different time-steps: 100, 200 (ref), 300 and 400 ns. Of course, a wider range of Δt would be desirable, and is generally expected in the case of conventional CFD numerical studies [35]. Here, a larger time-step, say by order of magnitude into the microsecond range cannot be tried due to reasons of PIC stability laid out in §3.5.4. Unfortunately, due to computational time constraints, it was also impossible to demonstrate the code running in the tens of nanosecond range. Although the range of analysis is far from what would be hoped for, it is still useful analyse the results, as a range between $\frac{1}{2}\Delta t$ and up to $2\Delta t$ is still possible. Data characterising the results are shown in figures 6.3 and 6.4; these data are represented as normalised plots of density and temperature taken from the cathode interior velocity distribution function sample point.

As would be expected, there is a dependence between the results and the time-step. Fig. 6.3 shows that in the case of neutral gas, there is now a slight density rise with time-step. This is different compared to the examination of time-step dependence for the neutral only code, which demonstrated very little numerical error with time-step at these values (fig. 5.4). It is logical to suggest then that the neutral dependence

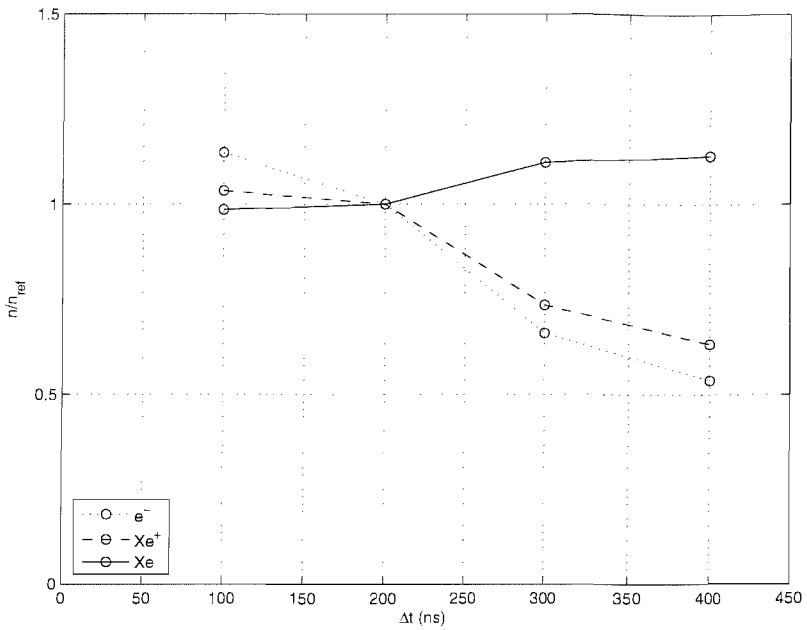

Figure 6.3: Normalised species densities as a function of simulation time-step, Δt. Values normalised to reference case value.

on time-step stems from the interaction between the neutral gas and the plasma, due to the plasma model exhibiting an error relationship with time-step. There is a strong relationship between simulation time-step and electron/ion density. Both plasma species show the same behaviour, and this is likely due to the fact that the at the sample point where the data is taken, the plasma is quasi-neutral, so the density of the two species is the same. The difference between electron and ion values plotted is then due to statistical scatter. Clearly, the density may be underestimated, although near 100 ns, the difference falls to 15%. In light of these findings, future results use only time-steps of 100 or 200 ns, and it is reasonable to suppose that values of density may be underestimated by up to 10%.

The second figure (6.4) shows the effect of time-step on temperature. Here we see that the neutral temperature varies very little, falling by perhaps 5% in the worst case (largest time-step). In the case of plasma species, the temperature tends to rise with increasing time-step, although again the error in the 100-200 ns range is minimal. It is likely that the rapid accumulation in error as Δt becomes large is due to numerical instability relating to the time-step-plasma frequency stability criterion of the PIC model, this is to be expected as values under 200 ns are chosen to be stable in this respect ($\omega_p \Delta t < 2$). The code remains stable at 400 ns, but the error is increased; it is of course reasonable to suppose that the error will increase as the stability point is approached (by increasing Δt). Indeed, it would be unexpected for such a numerical system to exhibit no error all the way until the maximum stability point, at which

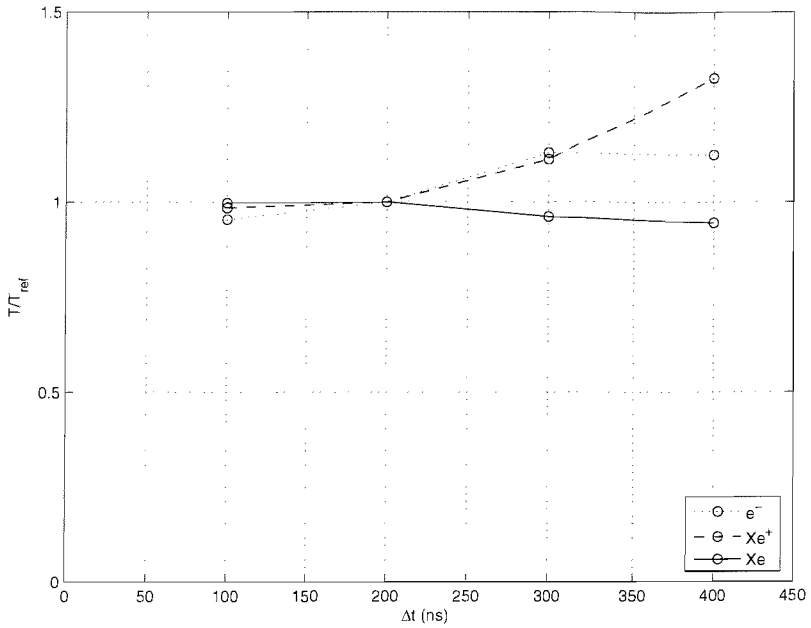

Figure 6.4: Normalised species temperature as a function of simulation time-step, Δt. Values normalised to reference case value.

the system would break down.

In conclusion, it has been shown that the numerical accuracy of the results is correlated with time-step, although statistical scatter and an inability to fully investigate very small time-steps must be considered. However, for the time-steps that are used in future results presented, the numerical error in density for the majority of the simulations is unlikely to exceed 20 percent, although in the worst case (peak density), this may rise even further. The error in temperature due to time-step size is generally smaller. *It should be remebered though that there is an absolute error introduced into the results as a direct consequence of being unable to use a small enough timestep.*

6.3.2 Mesh

There are really two meshes in the simulation at this point. The Monte Carlo mesh and the PIC mesh. It is possible to deduce that if we use a MC mesh similar to that presented in the neutral gas results, then it will be completely adequate for the plasma. This is because in light of the data on neutral density and the mean free path analysis in §2.4 it seems that the neutral-neutral collision rate will tend to be the highest of all collision types, due primarily to the fact that the neutral species has the highest density. Hence as long as the recommendations of §5.1.2 are adhered to, the MC mesh will be satisfactory for the plasma collisions.

Of course, the PIC mesh is in fact untested in the context of the hollow cathode simulation, and so must be examined now. The simulation is run at meshes both

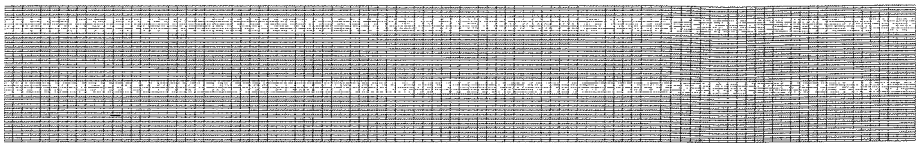

(a) PIC (100×60, uniform, aligned)

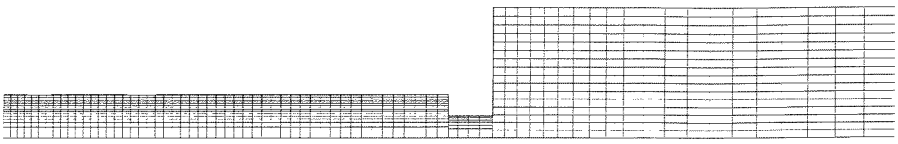

(b) MCC/DSMC (block structured, block aligned)

Figure 6.5: Reference PIC and MCC/DSMC meshes used in the plasma model.

smaller and larger than the reference case, on meshes of approximately 3, 6, 9 and 12 thousand cells. Specifically, 80×30 (2400 cells), 100×60 (6000 cells) and 120×90 (10800 cells), 150×90 (13500 cells). As an example, the standard mesh is shown in figure 6.5, alongside the DSMC/MCC mesh for reference. The PIC mesh is aligned along surface boundaries with the DSMC/MCC mesh and the blank area (cathode interior) in the PIC mesh is not calculated even though it is shown here. (This area is removed from calculation when the linear algebra problem is constructed).

Figures 6.6 and 6.7 show normalised data for the three species density and temperature respectively taken at the plume VDF sample point. Considering the plot of density first, there is a negative gradient in density with N, the number of PIC mesh cells. The final data point in the neutral (solid line) data is somewhat anomalous, but could arguably be due to statistical sampling error, as the value by which the point varies from the general trend is small compared to the trend itself. It is also true under these specific circumstances that the degree of statistical scatter increases for the higher N. This is because the number of sample time-steps available for sampling once a converged solution is reached is smaller if (as was the case here) all cases where run for equivalent compute times.

The plot of normalised electron and ion density in figure 6.6 is encouraging as it appears that there is very little deviation between the reference case and the highest N value. This indicates that in terms of mesh, the solution is changing by only a small extent for the chosen mesh. The plot also demonstrates that meshes such as case A should not be used.

The second plot (6.7, that of normalised temperature, shows almost no variation in neutral temperature with PIC mesh. This would be expected because there is no direct algorithmic connection between the two; only if the PIC mesh N were to

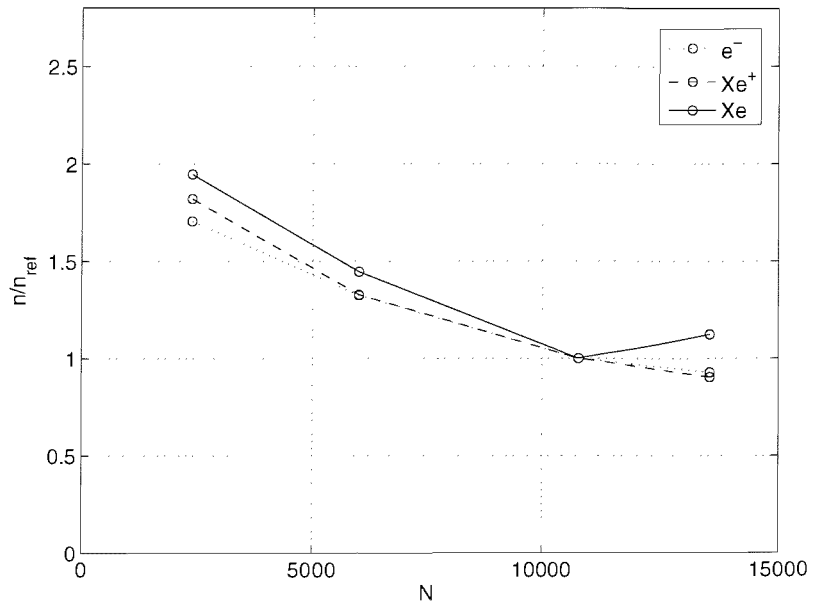

Figure 6.6: Normalised species densities in the tip region (sample point B) as a function of number of PIC mesh cells. Values normalised to reference case value ($N = 10800$).

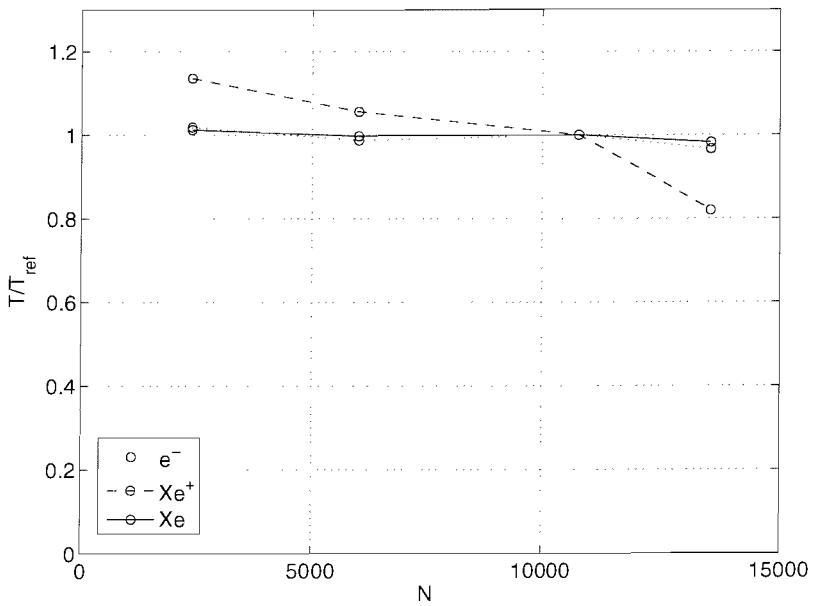

Figure 6.7: Normalised species temperatures in the tip region (sample point B) as a function of PIC mesh cells. Values normalised to reference case value ($N = 10800$).

significantly affect the plasma and hence then the neutral flow would any variation. It also appears that N has little effect on the ion species. The electron species does however show a slight negative gradient that shows little sign of flattening as N becomes large. Such an effect – a heating of the electrons on an over-sparse mesh – is well known in PIC codes, see for instance Birdsall and Langdon [12]. It was mentioned that some so-called "self-heating" of electrons may occur in section §3.5.4. It is highly likely that this is example of this phenomena. In the original discussion, the problem was dismissed due to the fact that in practical engineering DC discharges, the phenomena is rarely a major problem. In this case it is fair to say that self heating may be present, but is unlikely to have a major impact on the results. It should be noted however in discussion of electron temperature in future sections.

Another factor to consider is that the DSMC mesh remains the same for consecutive PIC meshes. Since the computational sample population is a function of the DSMC mesh density, this means that the sample rate per PIC cell decreases for comparatively denser PIC meshes. It is likely that this is a contributary factor to the error seen.

In summary, as expected, the density of the PIC mesh has some effect on the results. As usual, it is sometimes difficult to determine the error level amongst the statistical scatter (whose magnitude will be similar to that estimated in the previous chapter). In many cases, this effect is minimal, particularly at high N, where the code is generally operated. It is also likely that there may be some numerical self-heating of plasma species. This must be factored into the assessment of the accuracy of temperature predictions.

6.3.3 Electron Mass Adjustment (value of $\sqrt{f}$)

Here, we present results for the reference case electron mass correction factor, $\sqrt{f} = 480$, and two lesser values: $\sqrt{f} = 420$ and 360. These are referred to as the reference case (A), cases B and C respectively. Recall that a flux adjustment factor of $\sqrt{f}$ corresponds to an absolute multiplication of electron mass by f, so the reference case of 480 results in an artificial electron mass a little less than the mass of a Xenon atom. Of course, we could choose values of $\sqrt{f}$ to be more than 480, so making the artificial electrons heavier than the neutral particles. However, testing $\sqrt{f} > 480$ would surely highlight if the adjustment is affecting the results, we are more interested in any trends: recall that as $f \to 1$, computational electron dynamics become indistinguishable from real electron dynamics, so it makes sense to examine some values of $\sqrt{f}$ that are *less than* the reference value. It is difficult to predict if there will be a significant impact on the simulation due to electron mass modification. All of the measures discussed in §3.5.6 that correct for the modification are included, so provided that

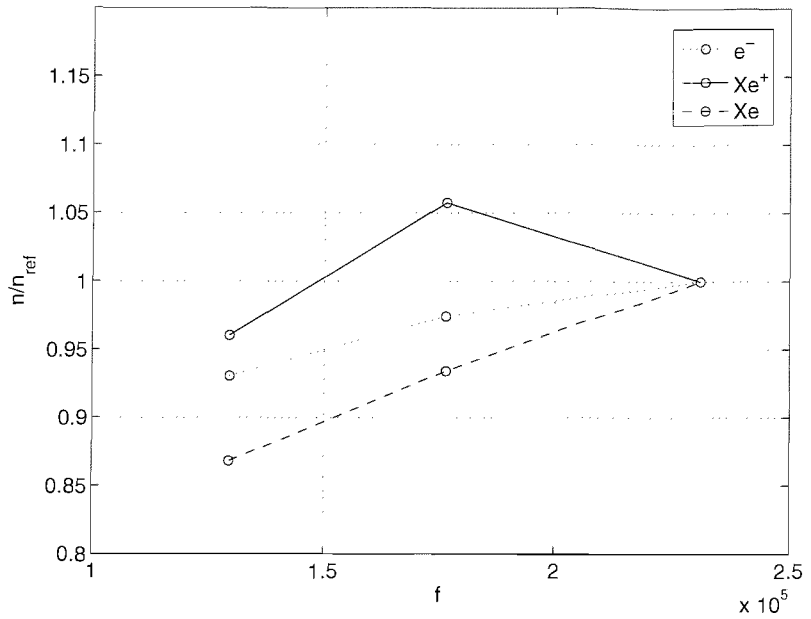

Figure 6.8: Normalised species densities as a function of electron mass adjustment factor, f. Values normalised to reference case value ($\sqrt{f} = 480$).

these modifications are (a) correct and (b) implemented correctly, it would be hoped that the impact would be minimal.

Figures 6.8 and 6.9 show plots of density and temperature at the sample point within the cathode. Considering the plots of density first, it is clear that the neutral density is not constant with electron mass factor: there is a positive relationship, and the variation between case A and C is around 10%. Interestingly, however, the ion and electron density variation falls within statistical scatter inherent in stochastic simulations. This outcome is quite unexpected. If anything, due to the comparative density and inertia of the neutral flow, and due to the fact that there is no *direct* link between the neutral dynamics and the PIC model, it could be suggested that changing f would in the best case have little impact on electron, ion dynamics and no impact on the neutral flow, while in the worst case, all three species would be affected. This is clearly not the case here.

The second plot, that of temperature (6.9) shows that changing the electron mass has little or no effect on the temperature of any of the three species. This is also unexpected in the light of the density plot since it could be proposed that neutral pressure remain constant as the electron mass factor is altered, so that temperature would change in an opposing way compared to the neutral density (assuming the neutral gas tends to obey an ideal gas law).

The trends shown are valuable in so far as they demonstrate that there is clearly no major correlation between f and the discharge. We can be confident that using the

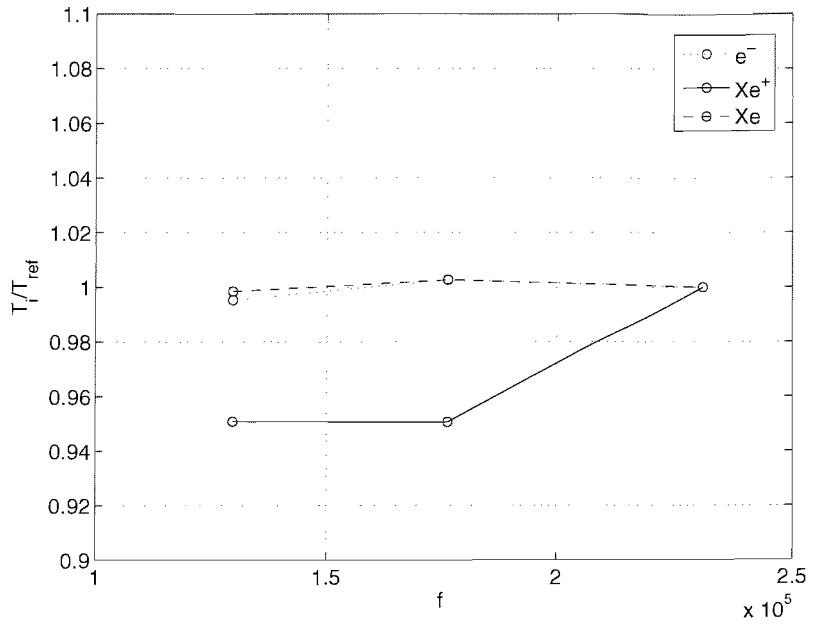

Figure 6.9: Normalised species temperatures as a function of electron mass adjustment factor, f. Values normalised to reference case value ($\sqrt{f} = 480$).

electron mass factor does not fully invalidate there results. It is however not understood why the neutral gas density varies as it does: this requires more investigation.

It true to say that changing f will alter the real time taken for convergence to equilibrium. This is trivially obvious since as $f \to 1$, so the absolute electron velocities are higher, and the transit time through the simulation drops. Of course, the motivation for using f in the first place is to bring the electron (plasma) equilibrium time closer to the neutral time so that it is not necessary to waste huge computational resources simulating an equilibrium plasma while waiting for the slow neutrals to come into equilibrium. Recall that for this numerical study, as must be the case, all computational and simulation parameters are fixed except f. This means that the simulation runs in all three cases for 60 ms real time (actually 12h compute time). This leads to the conclusion that for smaller f, the plasma tends to come into equilibrium faster (in real time terms, although of course in compute terms the simulation takes much longer) so although the plasma conditions may be very similar in all three cases, the slow neutrals may be at slightly different levels of convergence between A, B and C. Such an effect can be demonstrated by examining a time history of a critical convergence indicator for neutral particles.

Figure 6.10 shows neutral mass flux for cases A and C plotted against time. Although noisy, this plot tells us a lot, particularly within the first 15 ms. It is clear that in the low f case (dashed line), the mass flux is higher during the breakdown period. This indicates that the plasma discharge is likely to reach a state of equilibrium at a

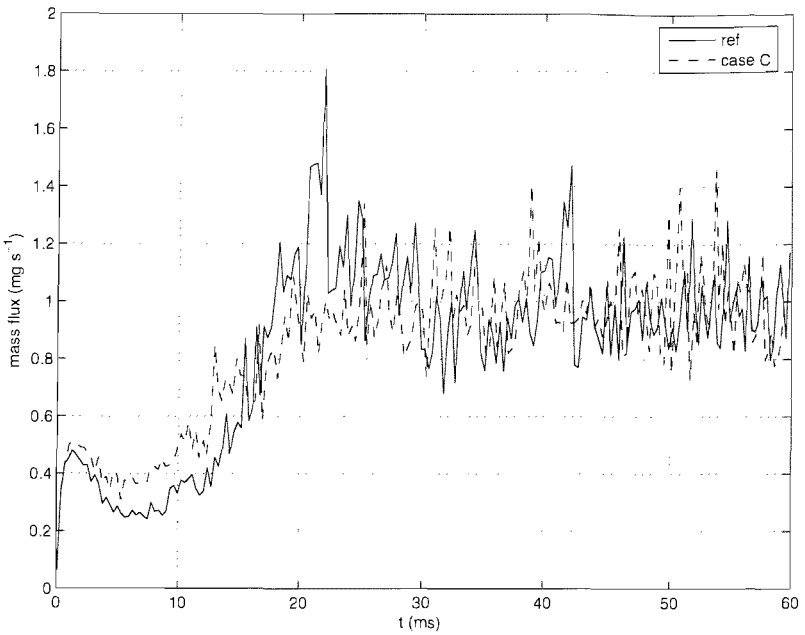

Figure 6.10: Neutral downstream mass flux time history for the reference and minimum (case C) electron mass acceleration factor.

different time compared to the neutral flow, or more importantly, that the conditions of the neutral flow at the point of plasma equilibrium will be different as f is changed.

It is highly probable that this could be compensated for by, for instance slowing the neutral flow evolution to allow for the varying the plasma behaviour. This approach of course leads to a circular argument as slowing the neutral in response to this situation has the same effect as setting f smaller, thus re-introducing the difference in τ_n compared to τ_p that we are attempting to avoid! In summary, two points may be suggested relating to the numerical study of the effect of f on the simulation. First, due to the fact that varying f necessarily varies the real plasma equilibrium time τ_p, we find slightly different neutral equilibria at different f. No immediate method presents itself by which this could be compensated for, but the difference in neutral equilibria is insignificant compared to that found due to variation in, e.g. cathode temperature, geometry, mass flux rate, etc. Put another way: although we know that slightly different neutral equilibria are reached at different f values, these are sufficiently insignificant that we will accept a potentially small error in some cases since the error is much smaller that the variation in neutral characteristics under other parameter changes. The second point to make is that fortunately, and perhaps most importantly, changes in f tend not to have any great impact on the plasma model. This is presumably due to the fact that the compensations to electron fluxes and velocities presented in §3.5.6 are scaling correctly. It is likely that the reason why the ion dynamics seem unaffected by altering f is due to the close relationship

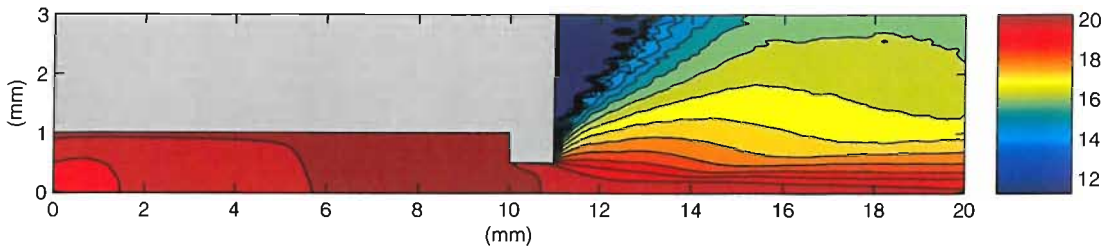

Figure 6.11: Contours of electron number density for the reference case discharge. Note log scale colourmap.

between electron and ion motion due to the PIC model.

This concludes the examination of the results of the numerical studies performed on the plasma discharge. The next stage of the analysis is to take a reference case and thoroughly examine every aspect of the output data.

6.4 Discharge structure

The structure of the equilibrium discharge produced by the code is now presented. The features are presented in three forms within this section; firstly, contour plots. These allow a general picture of the condition of the three species spatially within the discharge. While contour plots are attractive and allow for a macroscopic understanding of the physics, they necessarily contain too much information and are difficult to interpret and analyse for specific detail. Hence the other two subsections that present the discharge contain firstly axial centreline plots of key results, followed by some radial cuts. While axial centreline plots are perhaps the most powerful way to present these results, it must me remembered that due to the nature of the code, it was hoped that some sheath structure might be captured. This, of course, is primarily present on the cathode insert surface and so radial cut plots are of great value here, although the radial structure at the tip and in the plume are also of interest.

Contour plots of species densities are contained in figures 6.11, 6.12 and 6.13 for electron, ion and neutral species respectively. These plots serve to give the best qualitative view of what the computer simulation predicts for the discharge structure.

The first figure (6.11) shows electron density. The colour-map is on a log scale. The disturbance or noise in the region adjacent to the cathode outer face near the radial boundary is due to insufficient particle sampling statistics. In essence, the electron density is predicted to be so low here that the model tends not to provide many particles for data sampling. This is expected and a natural consequence of attempting to model a flow with very large density variations. The highest density occurs in the space charge limiting wall sheath attached to the cathode surface, but there is also a

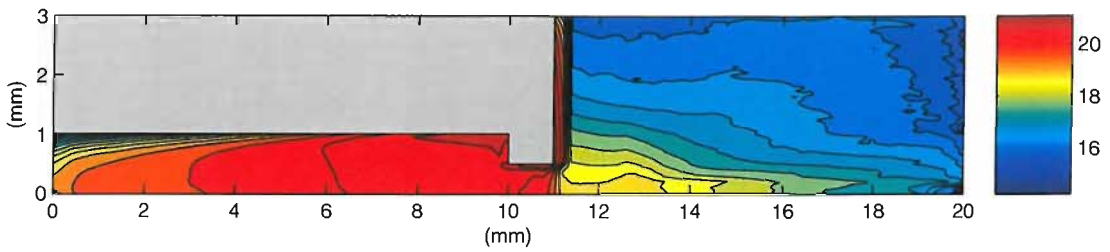

Figure 6.12: Contours of ion number density for the reference case discharge. Note log scale colourmap.

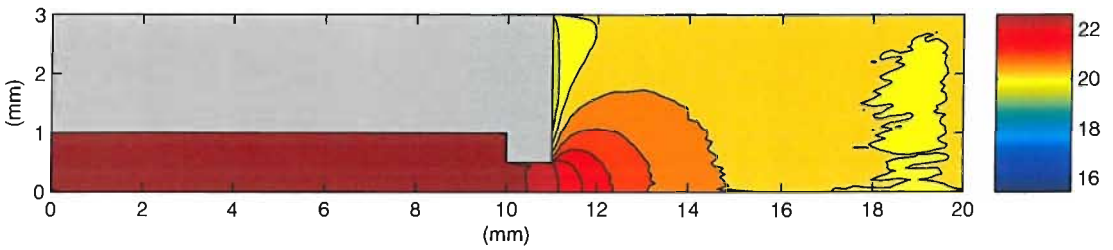

Figure 6.13: Contours of neutral number density for the reference case discharge. Note log scale colourmap.

region of high plasma density extending radially across the cathode interior just adjacent to the tip. This region of plasma density greater than 10^{20} m^{-3} extends around 5 mm inside the cathode. An area of similar size has been proposed by several authors [96], [34] although its existence has never conclusively been demonstrated previously. The presence in this plot of the so-called emission region, first introduced by Siegfried and Wilbur [96] as part of a theoretical model of the hollow cathode, can certainly be interpreted as being a successful resolution of the internal hollow cathode plasma. Comparison with experiment is difficult not least because measuring cathode internal plasma density experimentally is a difficult task. Even so, the plasma density within this region is within experimental error (the peak on-axis density, at around z=9 mm is greater than 2×10^{20} m^{-3}) of the predictions of Rudwan [88], although experimental predictions under these conditions are consistently a little higher, sometimes more than 10^{21}. The data of Rudwan was obtained by examining the cathode axially using spectroscopy. This means that the values given are actually integrated axially long the length of the cathode and adjusted. Additional uncertainty in experimental data is introduced by the fact that the value varies by up to a factor of 4 depending on the spectroscopic analysis technique used; in this case either local thermodynamic equilibrium (LTE) or collisional-radiative (CR). It is not particularly useful to make absolute comparisons such as these at this stage; particularly when the experimental data is so sparse and potentially inaccurate. Even so, it is a good sign that the computational results fall within the reasonable range set by experiment.

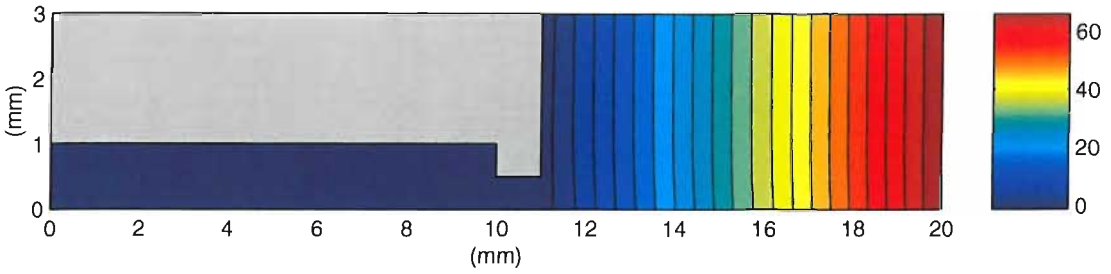

Figure 6.14: Contours of electric potential (V) for the reference case discharge.

The emission region can also be seen very clearly in the contour plot of ion density, fig. 6.12, again using a log scale colour-map. Here, the high density region can clearly be seen. Figure 6.13 shows the neutral gas density, from which, in comparison to the ion plot, it is possible to deduce that the ionisation fraction (n_i/n_n) will only be a few percent at the maximum point. This highest ionisation fraction region is located in the main cavity of the cathode, adjacent to the inner face of the tip: $z = 10$ mm.

Another feature of the contour plots is the dense ion sheath region located on the downstream face of the cathode. This is caused because this ion face has a low recombination fraction meaning that ions are not neutralised, but rather forced against the face by the exterior electric potential configuration, see contour plot fig. 6.14.

Figure 6.13 shows a plot of neutral Xenon density. Notice that qualitatively there is no significant deviation from the plots of neutral density shown in the previous chapter. From these plots it appears that the cathode interior plasma is nowhere near fully ionised; which is not to say that the plasma density is not similar to that predicted experimentally. From these results it could be suggested that authors who have asserted that the cathode interior plasma is nearly fully ionised simply underestimated the magnitude of the neutral gas density.

Finally, figures 6.15 and 6.16 show contours of current density magnitude $|\mathbf{j}|$ (although $|\mathbf{j}| \approx |j_z|$) and induced magnetic field. The data plotted are those directly calculated and used in the simulation to affect particle motion — *not* post processed from $\mathbf{u}_e, \mathbf{u}_i$ and n_e, n_i, rather using the real simulation data for $n_e, n_i, \mathbf{u}_e, \mathbf{u}_i$. Clearly, the fast moving electrons in the plume, but also to some extent in the tip, generate very high current densities. Bear in mind also that in the plume the electrons and ions have drift velocities in opposing directions further enhancing the current density. As expected then, a strong magnetic field is induced, the peak near the axis in the plume exceeds 3×10^{-3} T, 30 gauss, although it is also interesting to note the the magnetic field reaches around 10^{-3} T (10 gauss) just inside the cathode. Since the azimuthal field is negative (anticlockwise) and the electron drift velocity is predominantly positive in $\hat{\mathbf{z}}$, it is clear that what we observe in the plume is a classic z-pinch plasma. The

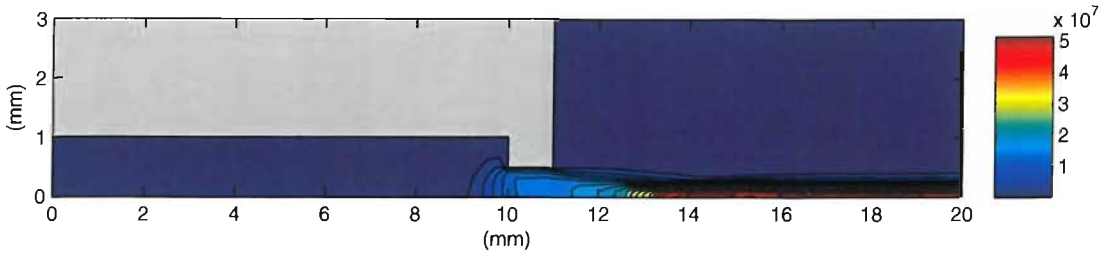

Figure 6.15: Contours of current density $|\mathbf{j}|$ for the reference case discharge.

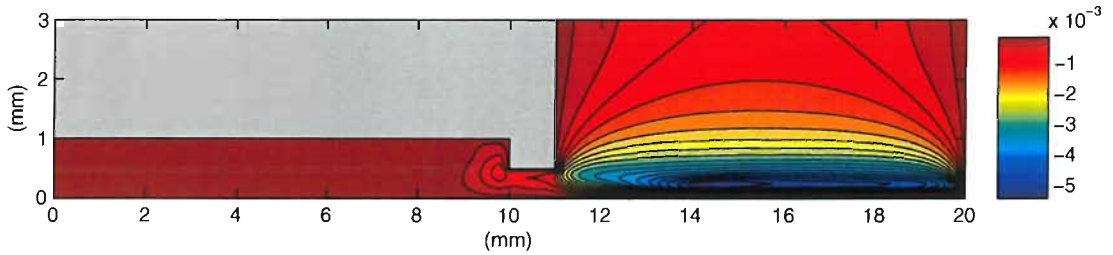

Figure 6.16: Contours of azimuthal induced magnetic field B_θ for the reference case discharge.

presence of the induced magnetic field tends to lead to a radial force on the electrons with positive u_z (and $u_z >> u_r$) in the direction of the z axis. The presence of the field pinches the electron flow toward the axis, increasing the on-axis density. This in turn leads to a larger current density j_z and correspondingly stronger magnetic field. It is also true that ions formed by ionising collisions outside the cathode will, due to the potential gradient, drift with a negative u_z, and experience the same pinching toward the axis. This is not as pronounced in the case of ion density since the ions are much heavier so the characteristic distance over which the ion dynamics would be affected by the magnetic field (effectively the Larmor radius) is far less.

6.4.1 Axial structure

The axial structure of the discharge is presented in plots containing data of the three species (electron, ion and neutral, respectively) in density, temperature, axial mean velocity, current density, neutral pressure and plasma potential. On all of these plots that follow, the two vertical lines indicate the position of the cathode tip; the inner and outer faces located at $z = 10$ mm and $z = 11$ mm respectively.

Figure 6.17 that shows the densities of the three species confirms the picture presented in the contour plots. Upstream of the cathode tip, the electron and ion densities are identical, indicating that we have a quasi-neutral plasma. There is also a positive density gradient in the plasma, including a peak density region extending close to the inner face of the tip, the density exceeds 10^{20} m^{-3} — the emission region.

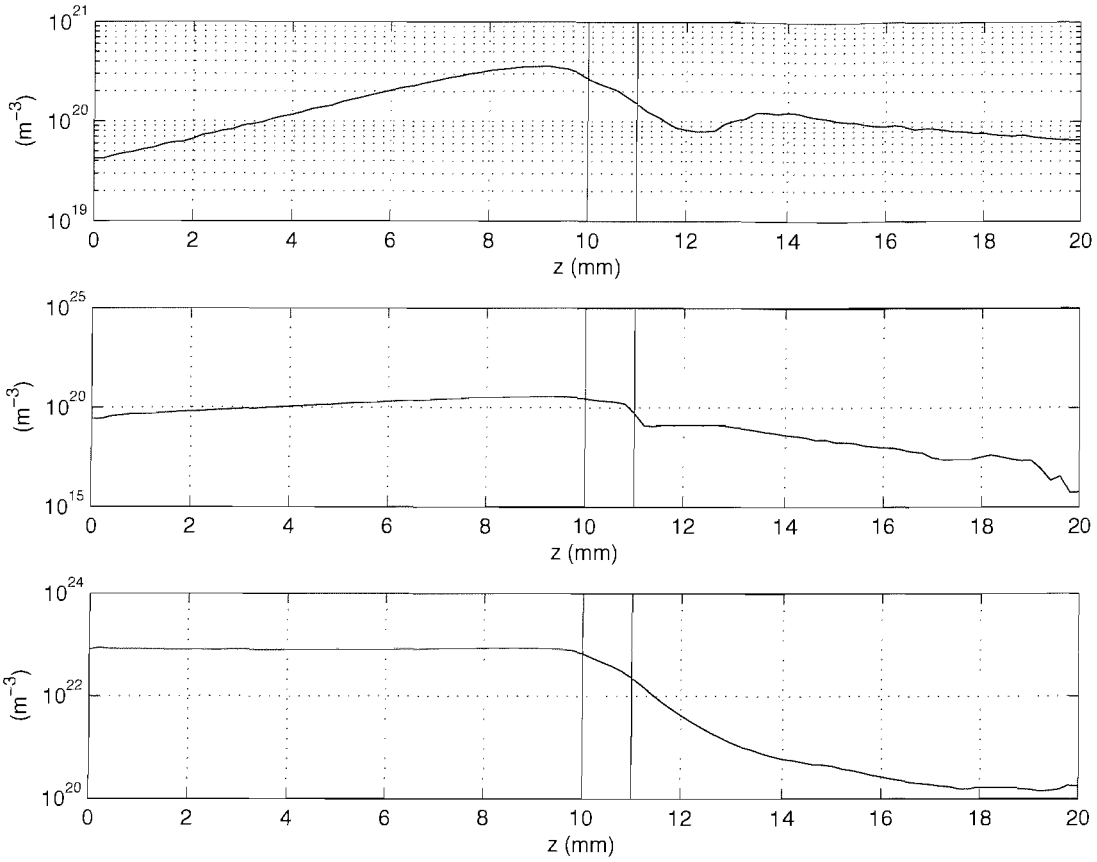

Figure 6.17: Axial centreline number density of the three species for the reference case discharge

In the far upstream the space charge limiting sheath exists as before, albeit somewhat weaker due to the presence of a low density plasma in the main channel of the cathode. Immediately downstream of the tip, the electron and ion densities diverge implying a loss of quasi-neutrality. This is the axial sheath that exists due to the presence of the anode: electrons are drawn from the internal plasma to the anode, while ions are repelled back into the interior plasma. The electron density profile in the tip is to some extent what would be expected due to the conclusions made previously regarding a z-pinch structure to the sheath. In the region of maximum magnetic field ($z > 15$ mm), the pinching effect leads to a rise in on-axis electron density. This cannot be seen in the ion profile, although based on the comparative ion/electron mass, it is quite reasonable that the electron dynamics be guided my the magnetic field while *over this scale*, there is little disturbance in the ion density.

As expected, due to its comparative magnitude, the neutral gas density is not altered significantly from non-discharge flow (compare to fig. 5.13 in the previous chapter). The gradual negative gradient exists as before.

The next triplet of axial centreline data plots (fig. 6.18) contains temperature

measurements for the three species. Recall from section 3.1.1, table 3.1, that temperature as directly sampled from a particle simulation is simply $(m(v^2 - v_0^2)/3k)$, where v is the mean (stream) velocity in a cell and v_0 is the root mean square of the velocity in a cell. This expression of course provides temperature in the correct way for the case of a purely Maxwellian distribution of particle velocities from which to sample. The exact meaning of the expression when a Maxwellian distribution is absent is not well defined, although of course for a non-Maxwellian velocity distribution, neither is the term 'temperature'. The degree by which any or all of the species velocity distributions diverge from the Maxwellian is examined later and further examination of this matter can be found in the discussion, later.

The electron energy is consistently close to 0.8 eV within the cathode, while the energy in the exterior (ionisation) region rises to a peak near 1.5 eV. These energies are certainly consistent with experimental estimates: for a 5 A Xenon discharge in a similar geometry, Rudwan [88] gives a value of 1.12 ± 0.05 eV. It must be remembered, however, that this value is an average integration of the electron energy from outside the cathode, so much include some contribution from the temperature in the plume. In the cathode exterior, adjacent to the anode, energies in the 1-2 eV range are predicted by the code. This is also consistent with experimental findings: the majority of reliable experimental data predicts electron energy in the 1-2 eV range within the cathode plume. Crofton and Boyd [27] for instance, measures the electron temperature in a T6 cathode plume to be 1.4 ± 0.1 eV. The ion species seem to be in thermal equilibrium with the neutral gas. This is not unexpected because the resonant charge exchange cross section is very high, while the ion density is at least an order of magnitude less than the neutral density. Due to the comparative importance of the CEX σ, ions will most likely spend most of their time colliding with neutrals while in the interior. In the neutral density region immediately outside the cathode, there is a very hot peak in ion temperature; it is reasonable to suggest this is where the majority of the ionisation is taking place, and that there is large energy transfer from electrons with large drift velocity. Of course, ionising electrons will have drift velocities in excess of 12.1 eV (Xe), so such energy transfer seems acceptable. The neutral temperature is largely as expected, although it is work noticing the heating of the neutral gas in the far plume ($z > 15$ mm), presumably due to the dense, on axis electron pinch. The neutral gas inside the cathode remains approximately in thermal equilibrium with the cathode surface. This is a somewhat important result as in important 1-D cathode models (Siegfried and Wilbur [95] for example), this is assumed to be the case.

Another condition of all three species that is key to building a picture of the axial physics of the cathode is axial mean velocity. This is plotted for the three species in figure 6.19, as usual in vertically descending electron, ion, neutral order. Here we see

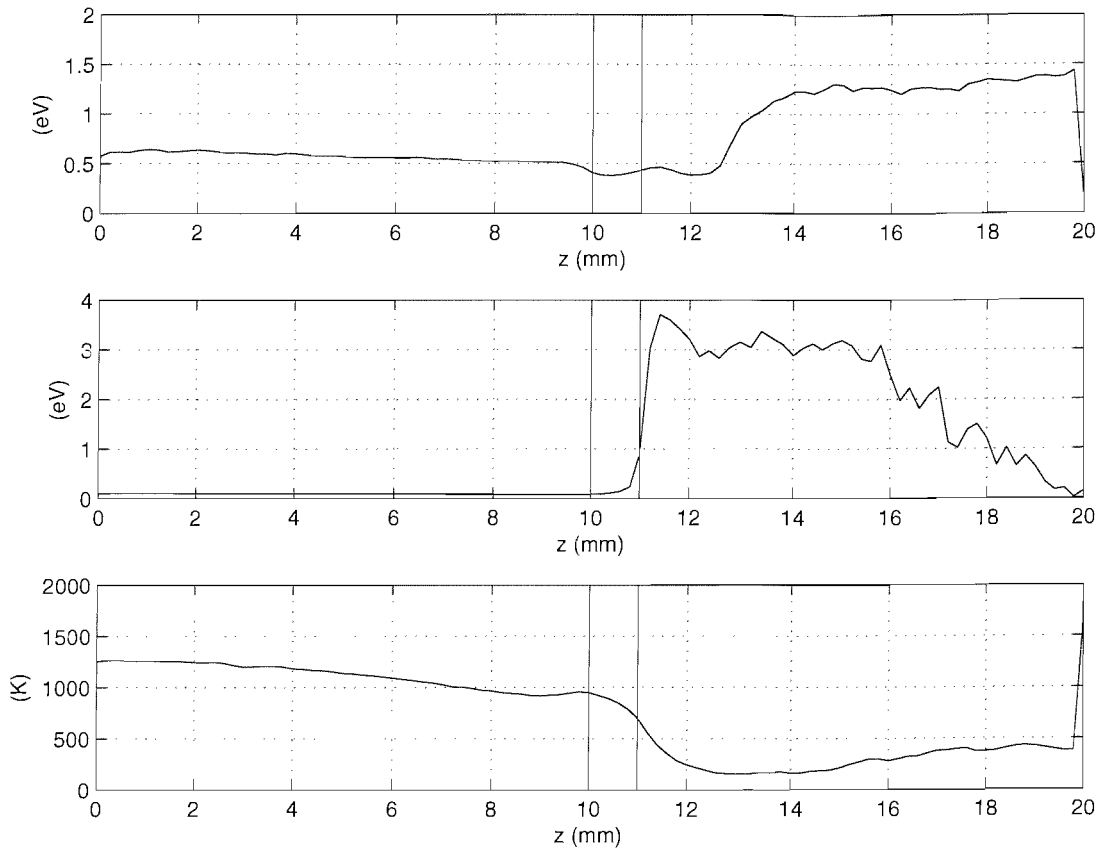

Figure 6.18: Axial centreline temperature of the three species for the reference case discharge, see note in text regarding calculation of temperature. The sudden drop on the downstream value represents the anode: no temperature is recorded here. The drop in temperature is not part of the modelled physics; it is a plotting anomaly.

very large electron and ion velocities associated with the acceleration region outside the cathode, near the anode. Although perhaps difficult to see on this scale, it is also true that there is a significant positive axial electron velocity throughout interior of the cathode. The ion gas within the cathode is almost stationary axially; there is no significant axial ion velocity within the heart of the emission region. This is expected: the ion velocity immediately downstream of the tip is large and negative. The ions are confined to the cathode by exterior conditions. Finally, we see a sustained positive axial velocity in the neutral gas within the cathode, followed by the customary expansion in the tip. The expansion is rapidly curtailed in the exterior because of the presence of the anode, and the introduction of neutral boundary gas.

The final triplet of axial plots show some more derived characteristics of the axial structure: pressure, current density and potential. Thankfully, in subplot 6.20-1, a negative pressure gradient can be seen within the cathode (the plot uses a logarithmic y-axis scale). This means that that the high interior axial neutral velocity seen in fig. 6.19-3 can be sustained.

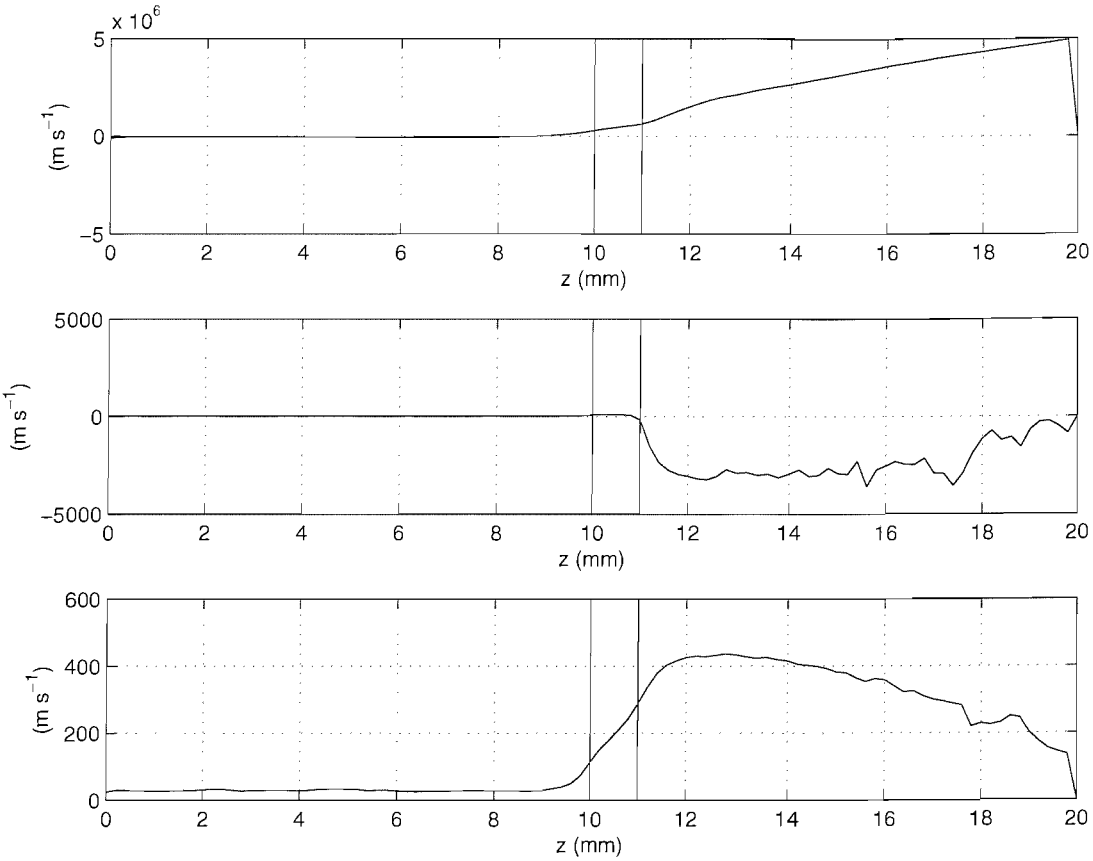

Figure 6.19: Axial centreline mean velocity in the axial direction of the three species for the reference case discharge. The sudden drop on the downstream value in the upper plot represents the anode: no data is recorded here. The drop in temperature is not part of the modelled physics; it is a plotting anomaly

The second subplot of figure 6.20 shows current density, plotted on the centreline assuming zero radial and circumferential velocity ($u_r = u_c = 0$) so that $\mathbf{j} = j_z\hat{\mathbf{z}}$ only. This is an important plot to show as it demonstrates the region of maximum induced magnetic field due to the axial current flow. Within the cathode, for the duration of the emitting region and tip, the current density increases monotonically. Next, the *on-axis* j_z falls a little, presumably due to the tip expansion. Note that the induced magnetic field close to the tip is weak compared to that in the fully developed plume. Hence, as we proceed downstream and the z-pinch plume develops, the on axis current density increases. Such deviations are not inconsistent with the net current flow in the simulation. If the current density is integrated radially at any point along the z-axis, the result is consistent in terms of the net current flow in the simulation. Finally, the lower plot shows the plasma potential, demonstrating that the potential in the cathode interior is negligible compared to the gradient in the exterior, generated by the anode. To better examine and understand the interior plasma potential, it is useful to now examine radial profiles of data.

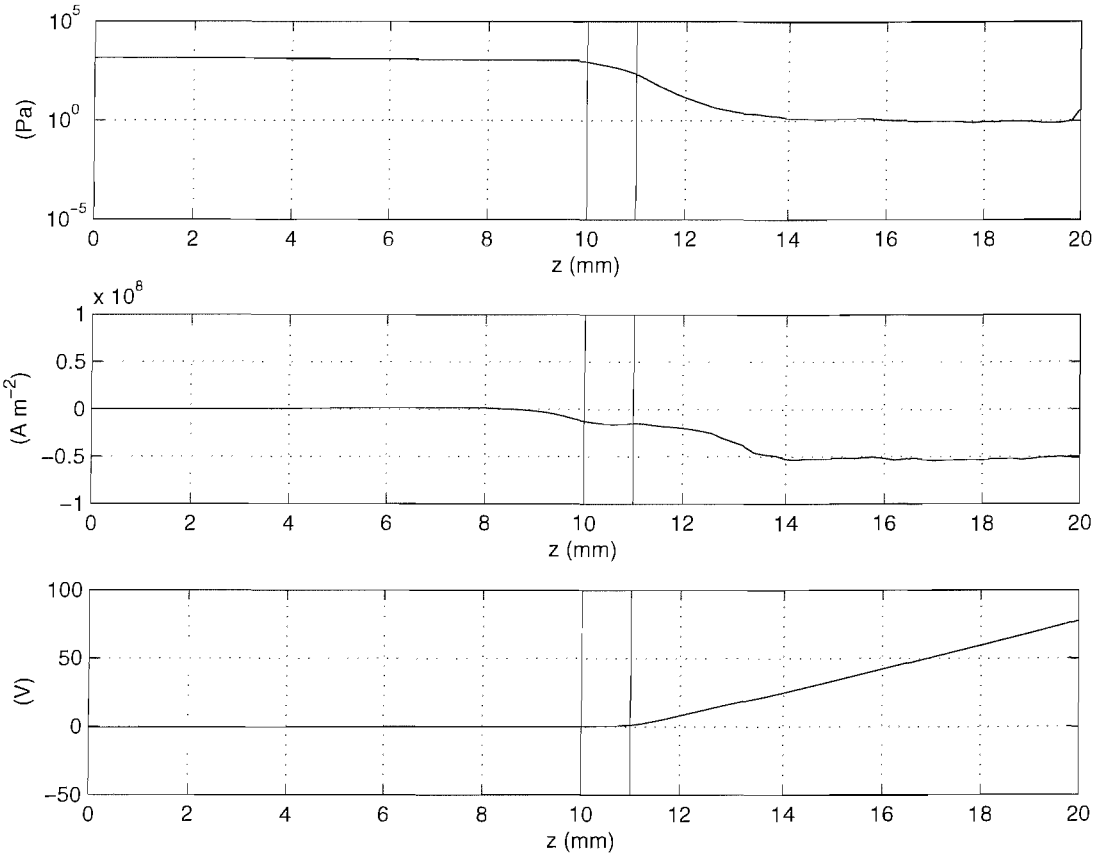

Figure 6.20: Axial centreline plots of current density — $e(n_i u_i - n_e u_e)$, neutral scalar pressure and plasma potential for the reference case discharge

6.4.2 Radial profiles

Radial profiles are particularly useful in the case of these results. Not only can the provide a 'cross-section' of areas such as the plume, more importantly, they allow for the sheath structure of the plasma attached to the cathode insert and tip to be examined. Remember however the comments made previously about the modifications to the permittivity in the stability section. This means that the sheath sizes will not be physically realistic, even though we anticipate that the fundamental physics is retained.

Figure 6.21 shows radial cuts of electron density at z positions listed above the individual plots (in mm) radially, while figure 6.22 shows radial cuts of ion density. Near the wall, we expect these values to differ significantly; this is of course how the sheaths are formed in any case. Notice, the plot at 10.8 mm extends only 0.5 mm on the x-axis; this is because $z = 10.8$mm is in the tip, of radius 0.5 mm. For ease of reference, plots are referred to as subplot 1-6, where plots 1,2,3 are the top three reading left to right, 4,5,6 are the lower three in a similar way.

The first plot is very similar to those seen previously in the space-charge limited electron-only discharge. As could be deduced from the contour plots, near the upstream boundary, space charge limitation is still present. Here, for instance, comparing subplots in figures 6.21-1 and 6.22-1 we see that the ion density is the same as the electron density away from the $r = 1$ mm wall sheath, but the sheath is still space-charge limiting. This can be further confirmed by examining the plasma potential; recall that if the potential is negative then current flow is inhibited so we have space charge limitation. Figure 6.23-1 shows just this. There is a sharp drop in potential within the sheath, followed by a largely flat region in the main channel. Notice that this does differ somewhat from a pure electron discharge in a vacuum, which tended to have a monotonically decreasing plasma potential moving away from the wall ($-r$ direction). In 6.23-1, the plasma potential is somewhat altered due to the presence of the low density quasi-neutral plasma in the main channel.

Subplots 2 and 3 are within the emitting region of the cathode. Here, the plasma potential (6.23-1,2) rises sharply adjacent from the wall, serving to extract the current emitted without allowing electrons to return to the cathode. Near the centreline, a flat potential profile prevails, accompanied by a quasi-neutral plasma. This sheath is generated by the charge difference near the wall, figs. 6.21-2,3 and 6.22-2,3. Notice that the peak ion density (on the wall) exceeds 10^{21} m^{-3}, while the peak electron density does not. The insert wall is unable to provide enough electrons to neutralise the sheath as they are immediately drawn away from the wall and downstream through the tip. It is then possible to conclude that the emission at a point like this is *saturated*. In other words, the size of the emitting region seems as though it must

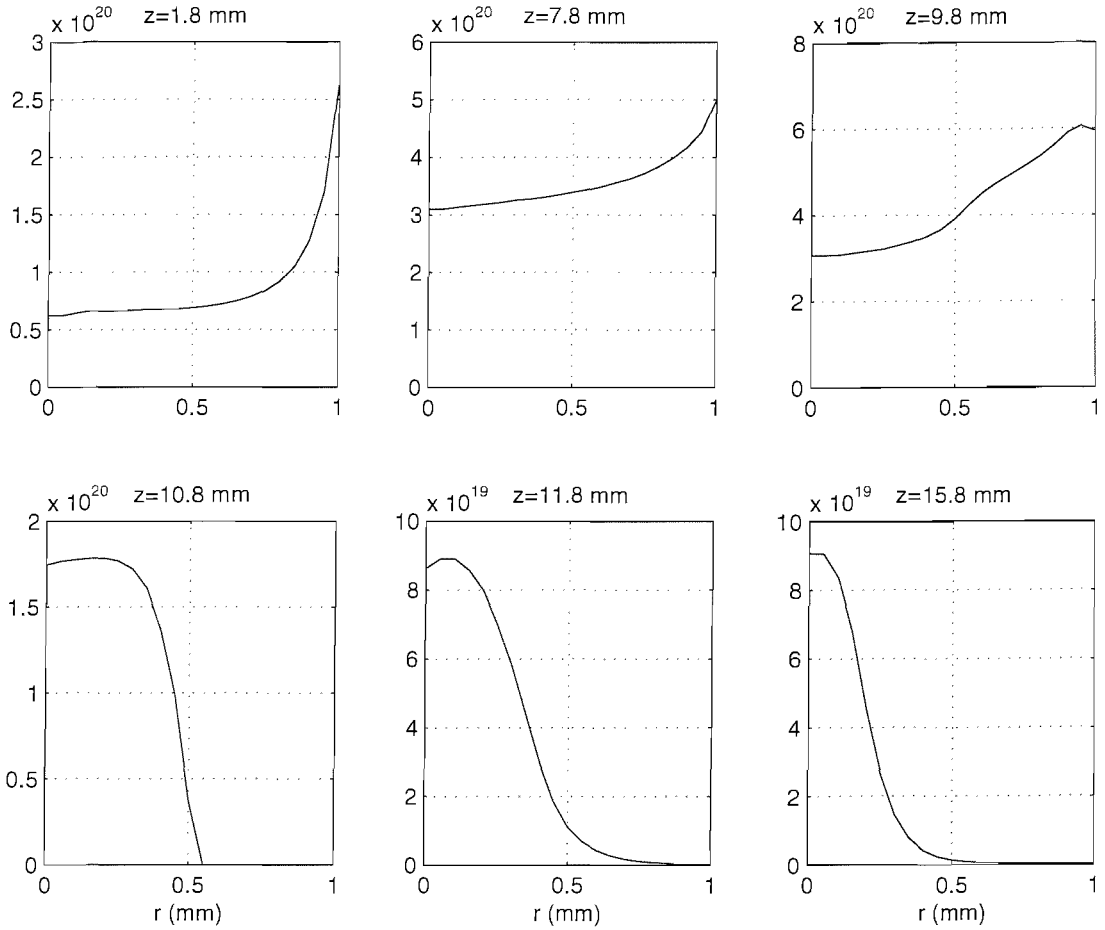

Figure 6.21: electron density: Radial cuts at various points axially (see individual plot title) showing electron density, reference case discharge.

expand upstream, completely saturating the insert emission current, until it is able to meet the anode current. The validity of this hypothesis is further considered in the discussion.

Further downstream, in the tip, the sheath type is again very different. Although the tungsten tip emits some current, the work function is more than 4.5 eV. This means that at the same temperature and under similar electric fields, the tip will emit orders of magnitude less current compared to the insert material. Hence, the electron density on the tip is effectively zero. Strangely, however, the ion density is quite high near the tip. It is possible to conclude that the reason for this is that there is a large supply of ions (as is the case due to the intense ionisation rate nearby: just outside the cathode) and also because the plasma potential in this region is strongly influenced by the proximity of the anode. See figure 6.23-4 showing plasma potential. Here, what looks like an emitting sheath has formed, but on the surface of the tungsten tip. It is more likely that the high external potential caused by the

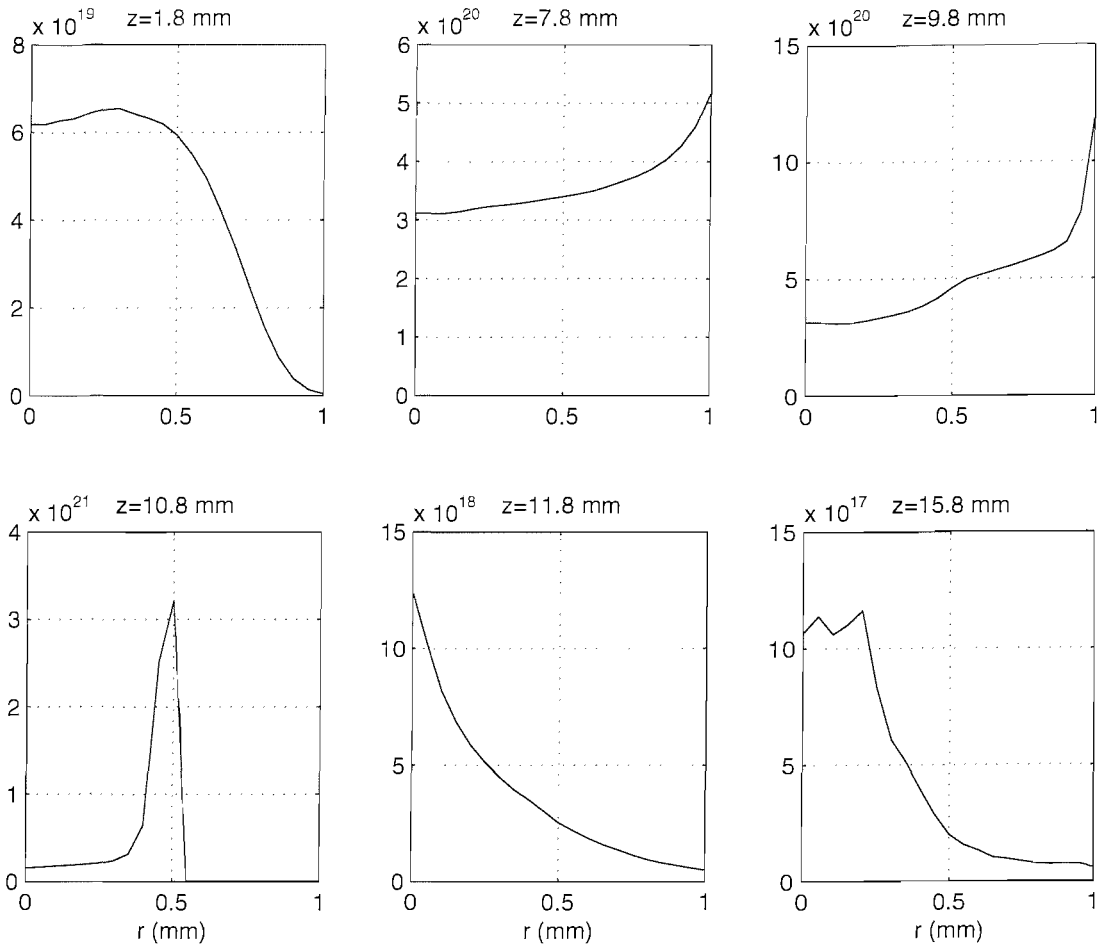

Figure 6.22: ion density: Radial cuts at various points axially (see individual plot title) showing ion density, reference case discharge.

anode is creating this structure. The result is that ions are rapidly drawn toward the cathode creating a region of intense ion bombardment on the tip surface. Such an effect can be qualitatively said to exist in experiment because authors frequently note that the cathode tip glows under some circumstances [87].

Within the plume, away from solid surfaces, there are no significant radial sheaths to observe; the plasma is close to quasi-neutral throughout. Indeed, far from the cathode, the electron and ion densities are equal (figures -6), but in the fifth figures of density, this is not quite true. Examining figures 6.23-5,6 showing plasma potential shows a dip in potential toward the centreline. This could be attributed to a number of effects, although a likely explanation is that due to the z-pinch effect show previously, there is little opportunity for radial electron diffusion in this region: there is an electron beam emitted from the cathode. In fig. 6.22-5, it can be seen that there exists a radial gradient at $r = 0$. This is not expected and should not be the case. The most likely explaination of this relates to an interation between the magnetic and

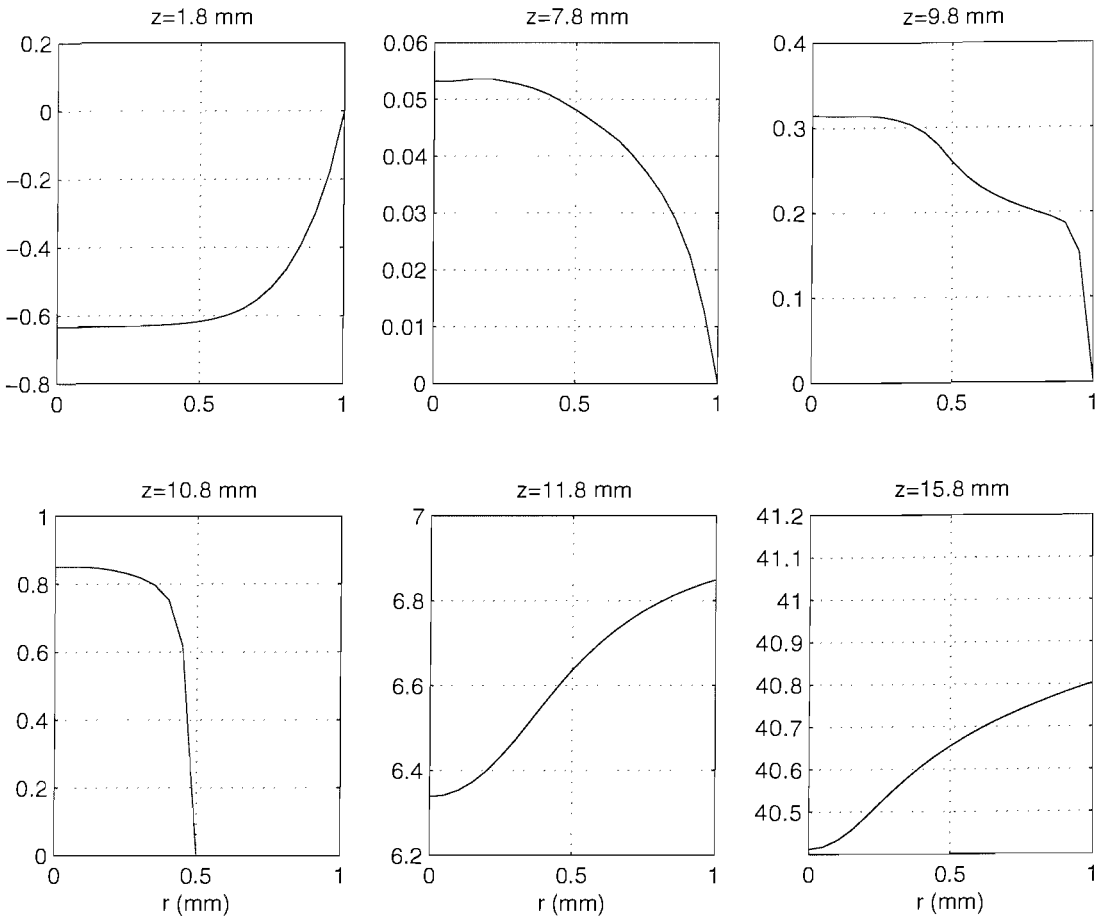

Figure 6.23: Plasma potential: radial cuts at various points axially (see individual plot title) showing plasma potential, reference case discharge. Cathode (at $z = 1$ mm, plots 1-3 and $z = 0.7$ mm, plot 4) is at 0 V.

electric field solutions near $r = 0$. A solution to this may be to use a higher order scheme for the magnetic solution on the axis, or different meshes for magnetic and electric solvers.

This completes the presentation of radial cuts through the cathode. While these plots go some way to explaining the structure of the radial sheaths, there still remains much to be understood. Recall that one of the original justifications for using a particle model was that some portions of the flow may be non-Maxwellian in velocity distribution; such regions are now investigated.

6.4.3 Velocity Distribution Functions

Finally, as was the case with the analysis of the neutral gas flow, it is possible to actually examine the particle velocity distributions. The analysis of contour plots, axial and radial cuts has been somewhat conventional in traditional CFD terms until

this point. It is however very unusual to be able to examine the particle velocity distributions in this way; this is a unique advantage of direct particle simulation over continuum models. Recall that the plots show two pieces of information: the distribution sampled directly from particles in the simulation (data points), the Maxwellian distribution as per the mean velocity and temperature sampled *for all particles*. Starting top-left, in a clockwise direction, the four plots show axial, radial and azimuthal velocity, followed by the energy distribution. Above the plots are shown various pieces of auxiliary information such the stream velocity. In the lower-left plot, that of azimuthal velocity, stream velocity is shown and is generally non-zero. This is counter-intuitive since it has already been stated that the problem under consideration is assumed axially symmetric. In fact, non-zero circumferential velocity is simply a statistical phenomenon: notice that the u_θ is generally at least an order of magnitude lower than the minimum of (u_z, u_r). To within statistical sampling accuracy then, the azimuthal velocity *is* zero in all species.

The following figures show electron and ion velocity distributions in three dimensions plus velocity magnitude distribution at two sample points: one within the cathode in the main emitting plasma (at $z = 5$ mm) and another at $z = 12$ mm, $r = 0.2$ mm. This second point is approximately at the point of maximum current density, located in the heart of the plume and ionisation region.

The plot of electron distribution within the cathode is shown in fig. 6.24. In all cases, the distribution of velocity agrees well with the classical definition. Deviations from the Maxwellian are only observed near zero, particularly in the radial case. Recall that the emitted energy distribution (§2.3.1 is an odd tail of the Fermi-Dirac distribution, with few electrons with very small energies. This means that although Coulomb collisions, electron-neutral collisions and the existence of a quasi-neutral plasma in the PIC model have tended to pull the electron vdf into a true Maxwellian, the influence of the emission energy remains in the low numbers of very low energy electrons. This effect can be most clearly seen in the axial distribution. The radial distribution contains deviations presumably to to the electron-sheath interaction.

Figure 6.25 again show electron velocity distribution, only this time in the plume. It is immediately obvious that the axial distribution is significantly shifted to the right, due to the natural drift velocity of the electrons toward the anode. It is also true that the axial simulated temperature is lower than the classically predicted temperature. This can be inferred by the fact that the computed distribution is 'thinner' than the classical one. Another interesting feature is the absence of higher energy electrons in the azimuthal direction. The only natural explanation to this is that the magnetic field is tending to restrict swirl in the model. Thankfully, this is a symmetric effect, because where it not, significant swirl instabilities may have occurred.

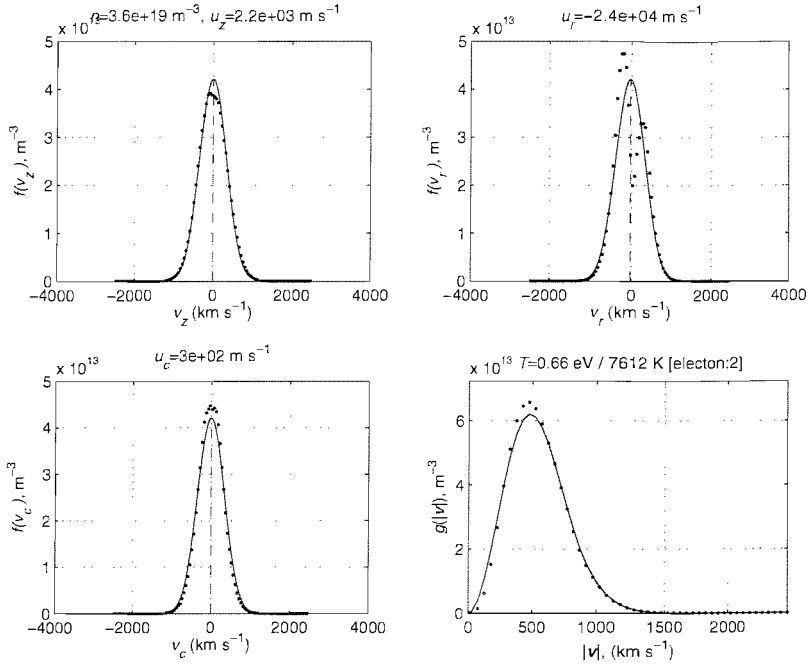

Figure 6.24: Velocity distribution plots of electrons in the cathode interior, at $z = 5$ mm.

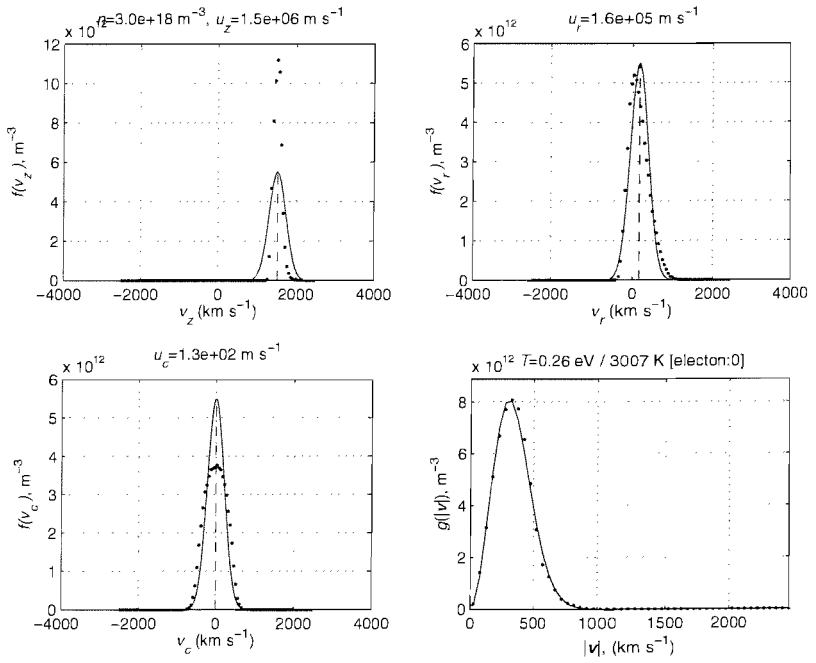

Figure 6.25: Velocity distribution plots of electrons in the cathode plume, at $z = 12$ mm, $r = 0.2$ mm.

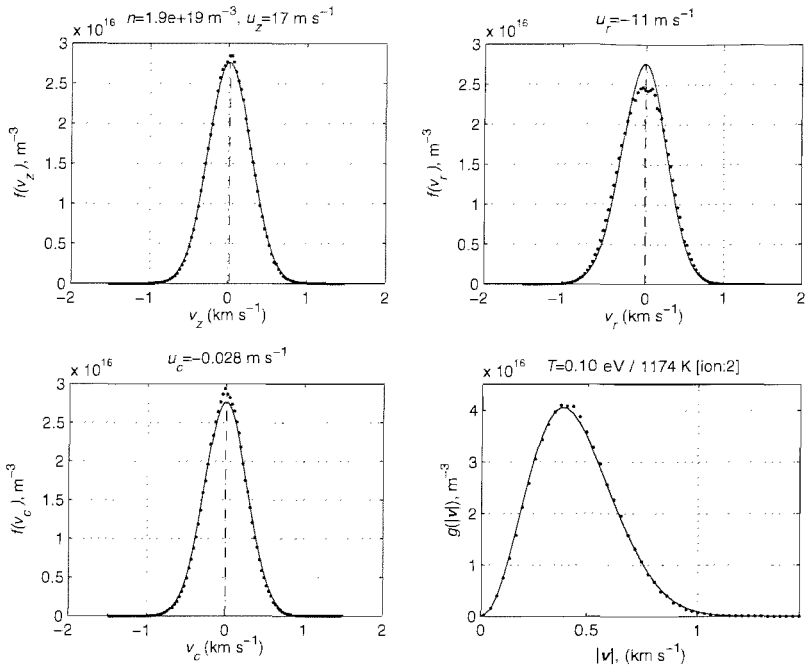

Figure 6.26: Velocity distribution plots of Xenon ions in the cathode interior.

Considering the ions, fig. 6.26 shows ion velocity distributions for the cathode interior. Here, there is little drift velocity, only -17 m s^{-1} axially. The numerical simulation fits well with the classical Maxwellian distribution plots. This is to be expected because the collision rate between neutrals and ions is high and the neutral gas is close to being a continuum fluid in Knudsen number. This means that it is likely that due to ion-neutral collisions, the two will be of a similar temperature and tend toward a classical Maxwellian. The only deviation from the natural case occurs in the radial, most likely due to the presence of the radial sheath.

Finally, fig. 6.27 shows velocity distributions of the ions within the plume. This plot is completely different compared to the interior case. Because the neutral density is very low in this region there is little neutral ion interaction. In addition, many ions sampled at this point may be newly created ions due to ionising collisions between electrons and neutrals. The axial plot is very interesting. It is clear that the ion drift velocity is very strongly toward the cathode; this is due to the positive potential gradient present in this region. In addition, the distribution is non-Maxwellian.

The majority of the ion gas is observed with a strong negative axial velocity: these are the ions heading upstream toward the cathode. A final note when considering the axial distribution is that the temperature — in qualitative terms the 'width' — of the distribution is very large axially. The most reasonable explanation for this is that most of the momentum involved in the ionisation is axial, so a large fraction of the resulting redistribution of momentum that occurs during ionising collisions is

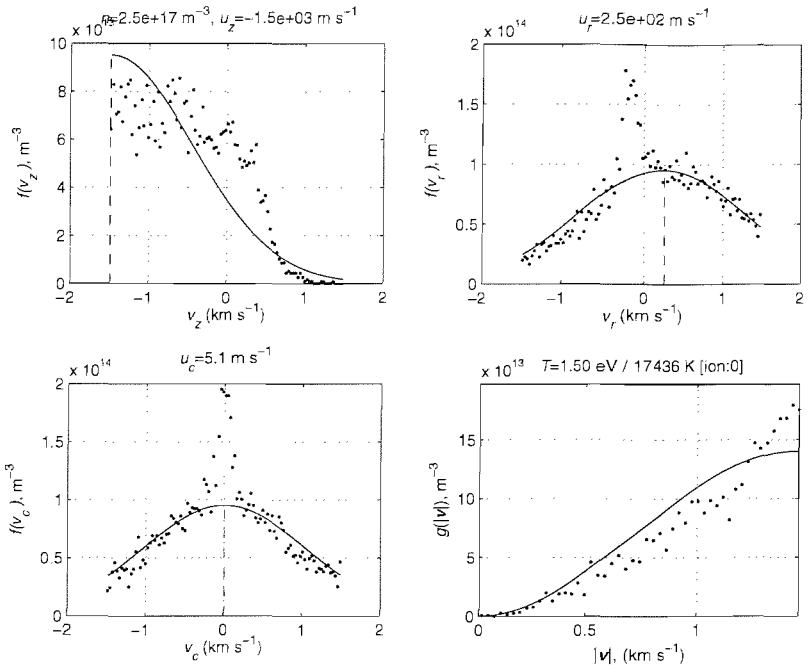

Figure 6.27: Velocity distribution plots of Xenon ions in the cathode plume.

found in the axial direction. Hence, the large, hot distribution in the negative axial velocity is due to newly formed ions from ionising collisions. Similar trends are seen in the radial and azimuthal distributions. Considering the radial plot, we see that the distribution is broadly shifted positive, implying, as expected, that the ion gas created by ionisation is expanding radially. There is however a peak region for slow ions that is shifted negative radial. First, we must ask: what is the origin of these cold ions? The answer must of course be that they are the product of CEX collisions. But what of their motion? This is naturally due to the presence of the magnetic field. Recall that ions with low radial velocity but very high axial velocity will tend to be pinched toward the central axis. This explains the negative radial velocity in the figure. In the final plot, that of azimuthal velocity, we see that there is again a peak non-Maxwellian density region corresponding to low u_θ ions. Again, the origin of the peak of cold ions is most likely due to charge exchange collisions. This is most likely due to the magnetic pinching effect on low radial velocity ions and the interaction between velocities in the axially symmetric particle mover.

6.5 Current and Mass flow rate

When conducting hollow cathode experimentation, once the cathode is working, the next step is always to run a voltage current characterisation. This indicates that the device is operating as expected, although the results can often be obscured by spot

to plume transition phenomena. It seems logical then that the first this to do with a numerical model of a cathode is to run a characterisation at a couple of different propellant flow rates across a range of currents.

Under some conditions, it is difficult to associate a particular mass flow rate as being either spot or plume mode (since all very high current discharges tend to be spot in any case), a 'high' and 'low' mass flow rate are chosen to generally represent 'spot' and 'plume'. This is a common assumption. Of course, it may be the case that the numerical model operates entirely in one mode; as has been stated, then cause for the mode transitions in cathodes is still not understood.

In a very general sense then, the two mass flow rates are chosen refer to values at which plume and spot modes have been observed for the propellant under consideration[88]. Hence, 1 mg s^{-1} and 3.29 mg s^{-1} Xe; the reference case presented previously used $\dot{m}$ = 1 mg s^{-1} Xe. The current range is chosen to be similar to that which experimental researchers might use when characterising a hollow cathode, namely 5 to 20 Amp in 5 Amp intervals.

This section now proceeds to examine the current voltage relationship before a presentation of the variation in plasma conditions with current and mass flow rate. First however, a brief note on the mechanics of recording current from an inherently stochastic simulation.

On the accuracy of recording current in stochastic simulations

A note is required to cover the accuracy and error in the values of current and voltage. This may seem strange in the case of a numerical model: on face value, if the current is set to 15 A, say, then exactly 15 A might be expected to be drawn. This is not true for stochastic particle based models! Figure 6.28 shows a time history of both anode current (left axis) and anode voltage (right axis) for the reference case conditions (20 A, 1 mg s^{-1}). This plot is also an excellent judge of the level of equilibrium present in the discharge. The final values recorded in this case where 32 V and 19.98 A. The noisy trace is that of current. This line rises rapidly during the startup to a current of nearly 30 A after less than 0.1 ms. This is the phase where the cathode interior is filled with ions and the internal emitting plasma forms. The internal density is then regulated: since the anode is receiving too much current, the voltage drops from the initial condition of 100 V to as little as 25 V. The discharge comes broadly into equilibrium in terms of macroscopic currents after around 0.2 ms. Of course, flow conditions with longer characteristic periods such as temperature may take longer than this to reach equilibrium.

The noisiness of both lines in fig. 6.28 is due to the stochastic nature of the simulation. An increase in computational particles would have the effect of smoothing

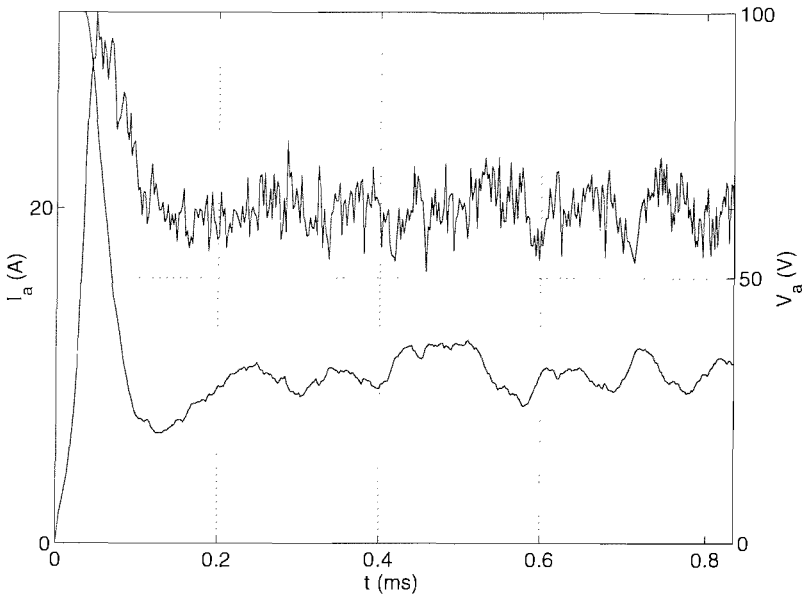

Figure 6.28: Time history of anode current and voltage for the reference 20 A discharge case, Xe gas at 1 mg s^{-1}.

the trace. Regardless of this, it is possible to measure the average current and voltage and write the result. Sampling of this average is generally done during the last 20 % of the run time available in the simulation, which, in the context of the figure means that sampling begins at around 0.6 ms and continues to the end. In the case of this specific run, this period consists of 10,000 time-steps with samples taken on each time-step. Clearly, this is the origin of the noise. We could simply take the precaution of sampling anode current and voltage data at wider intervals.

6.5.1 Current-Voltage

Figure 6.29 shows a current voltage plot for the two mass flux conditions. At low mass flow rate ('plume'), we see a positive current-voltage characteristic; voltages are similar to and well within the range expected based on experimental data. The positive gradient is also expected for the low mass flow plume condition. In the high mass flow rate ('spot') condition, a slightly negative current-voltage characteristic is observed, and the anode voltage remains lower than the plume mode case. This is all consistent with cathode behaviour under experimental conditions. In the case of a real hollow cathode, results such as these would tend to suggest a valid setup.

Another feature of the IV plot is the absence of any unexpected trends. Often, in experimental IV plots of cathode behaviour, spot to plume transition may occur within one current sweep. This is intriguing: on the one hand, it may simply be that there is insufficient data points in fig. 6.29 to show spot to plume transition. Perhaps

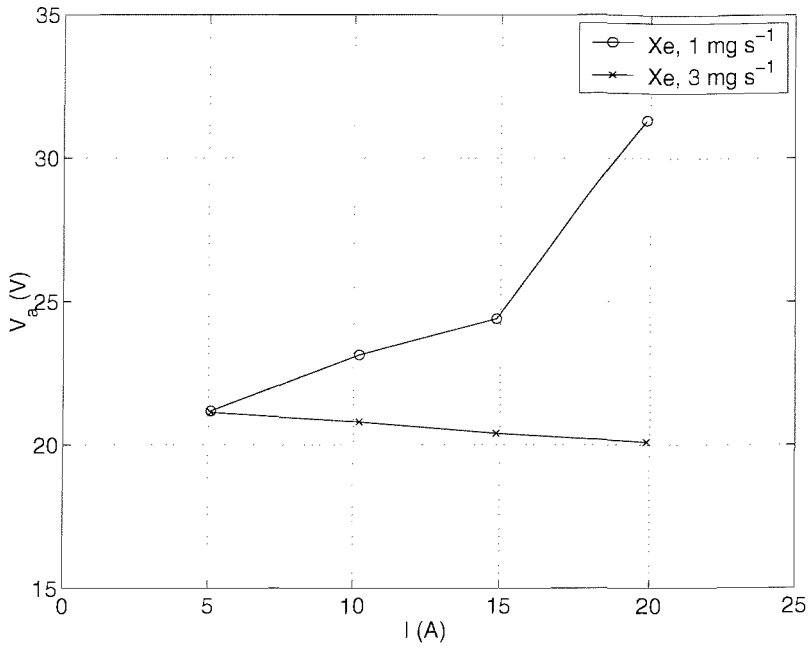

Figure 6.29: Current-voltage behaviour of the cathode using Xenon at two different mass flow rates ('plume' and 'spot' conditions).

the two $\dot{m}$ cases are simply well within spot or well within plume. This seems unlikely however because the current range is very large in cathode terms. If one subscribes to the hypothesis of Rudwan [88], however, then we should expect that with this configuration, there should be no spot to plume transition at all. This is because Rudwan proposes that transition phenomena are a function of large scale exterior plasma effects that form when using a hollow anode located at least a few centimetres from the cathode. Hence we should definitely not see transition. Unfortunately, as was discussed in the introductory chapters, such matters are still poorly understood so it is difficult to draw firm conclusions about the validity of the code based on the absence or otherwise of transition.

This completes the presentation of the current-voltage characteristic of this simulated cathode. It was included for completeness since a reader of any hollow cathode research often expects the ubiquitous I-V plot; however from the perspective of this research it provides little useful information bar a cursory indication of the validity of the model. What it is possible to do now is to investigate how the plasma conditions *inside* the cathode scale with current and mass flow rate.

6.5.2 Plasma conditions

Using the same results then to those that provided the I-V characterisation (fig. 6.29), it is now possible to show the effect current and mass flow rate changes have upon the

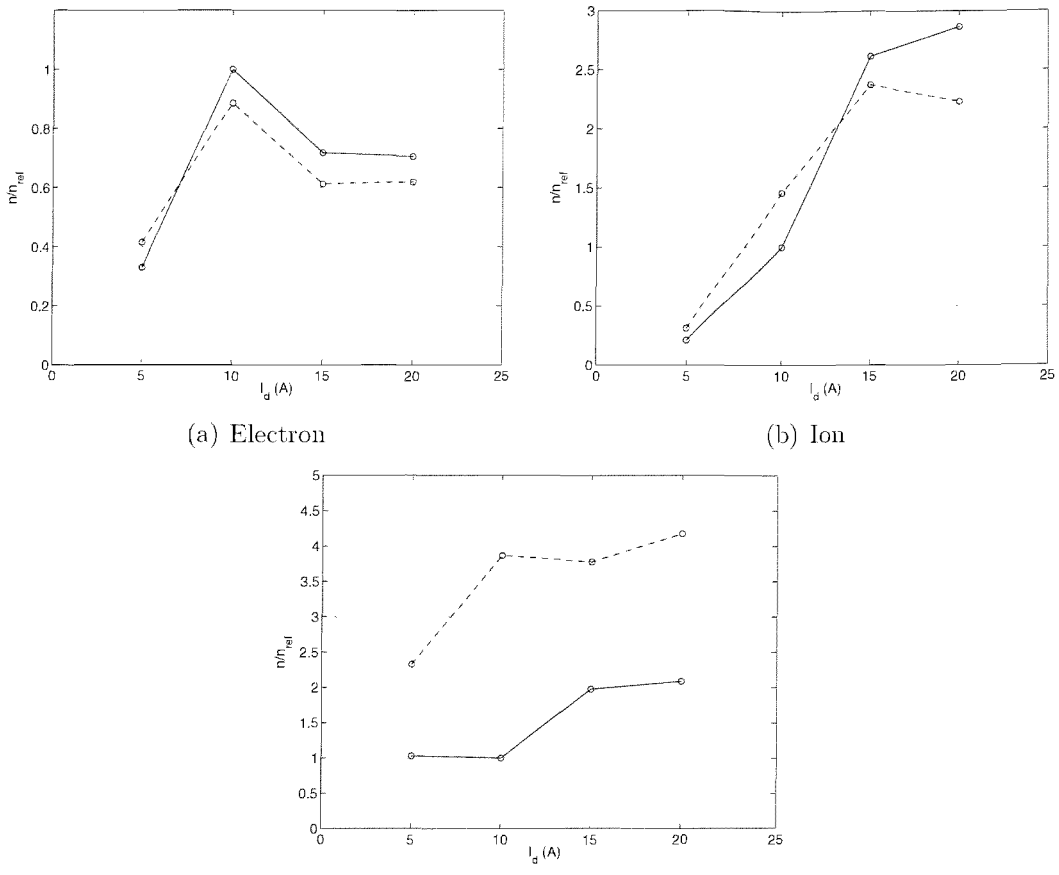

(a) Electron

(b) Ion

(c) Neutral

Figure 6.30: Cathode internal (i.e. reference position A) plasma density as a function of current for two mass flow rates, 1 mg s^{-1} (−) and 3.29 mg s^{-1} (−−). The data are normalised to the 10 A, 1 mg s^{-1} case for each species.

plasma itself. Considering first the interior plasma, in terms of density 6.30 shows the species densities. The general trend for heavy particles is for the density to increase with mass flow rate. Considering the first two plots (a) and (b), (that show electron and ion densities), it is clear that the internal density does not vary significantly between the two mass flow rate cases. This is unexpected: it would seem logical for the plasma density to rise in the case of a higher neutral flow rate. Plot 6.30(c) shows that the neutral density is generally at least a factor of two higher for the 3.29 mg s^{-1} case compared to the 1 mg s^{-1} case. There is also an anomalous result in the electron density, as the peak value occurs at the 10 A point, rather than (as is the case for the heavy particles) as expected at 20 A. Remember that the internal sample cell is in the maximum density emitting region near the emitting sheath, so if the sheath is changing size with current (as would be expected), this might explain the anomalous electron density.

Figure 6.31 shows plots of temperature as a function of discharge current. Con-

sidering first the electron temperature, fig. 6.31(a), there is a strong positive gradient in temperature. At both $\dot{m}$ conditions, the electron temperature rises with discharge current. This seems consistent based on the analysis of the previous figure showing electron density: since the density seems not to vary to a great extent in this region, higher currents must then be due to larger electron drift velocities. A higher drift velocity would then naturally imply a greater level of electron heating. The variation in temperature is over a wider range than might have been expected; the range is between a minimum of approximately 0.6 eV up to a maximum of around 1.5 eV at 20 A. Such a gradient is unexpected as it is generally assumed from experimental studies that the electron temperature remains at around 1 eV. Considering next the ion temperature, it is clear that higher discharge currents lead to cooler ions. Such a result is quite counter-intuitive, but there are certainly explanations. A possible argument is that since the ions are typically at a temperature higher than the neutrals there is a heat flow from the ions to the neutral (that is particularly true due to the frequency of charge exchange collisions). Hence given that (a) the heat capacity of the neutral gas is very large compared to the ions because of the neutral density and (b) there is a considerable rise in neutral density with discharge current, it is possible that the neutral cooling effect on the ions actually increases with discharge current. Such a hypothesis is of course strongly based on the assumption that the heat flux between the ion and neutrals is closely related to neutral density. The small rise in neutral temperature with discharge current shown in fig. 6.31(c) would support such a hypothesis.

Before considering the plume conditions, let us first recall the key findings in terms of the internal plasma. The electron density shows an unexplained peak at 10 A, while there is a strong positive gradient in electron temperature with discharge. The heavy particle densities rise monotonically with current at both mass flow rates, while there is an anomalous cooling of the ions with current density and a heating of the neutrals. This is possibly explained by considering the heat flux between the species and the comparative importance of CEX collisions.

It is interesting to note that the trends in plasma conditions for the internal plasma did not seem to be particularly perturbed by the variation in mass flow rate (with the exception of the neutral density, of course). When considering the plume density, figure 6.32, we observe a monotonic rise in density with current for all species at all mass flow rates. It is also shown that for all species, the density in the high mass flow rate case is higher than the low mass flow rate (recall that the data are normalised to the 10 A, 1 mg s^{-1} case) — again this is expected. What is perhaps not expected is that the ion density is generally a factor of 1.5 greater in the high mass flow rate case, while the electron density is only around 1.2 greater. This implies that the axial

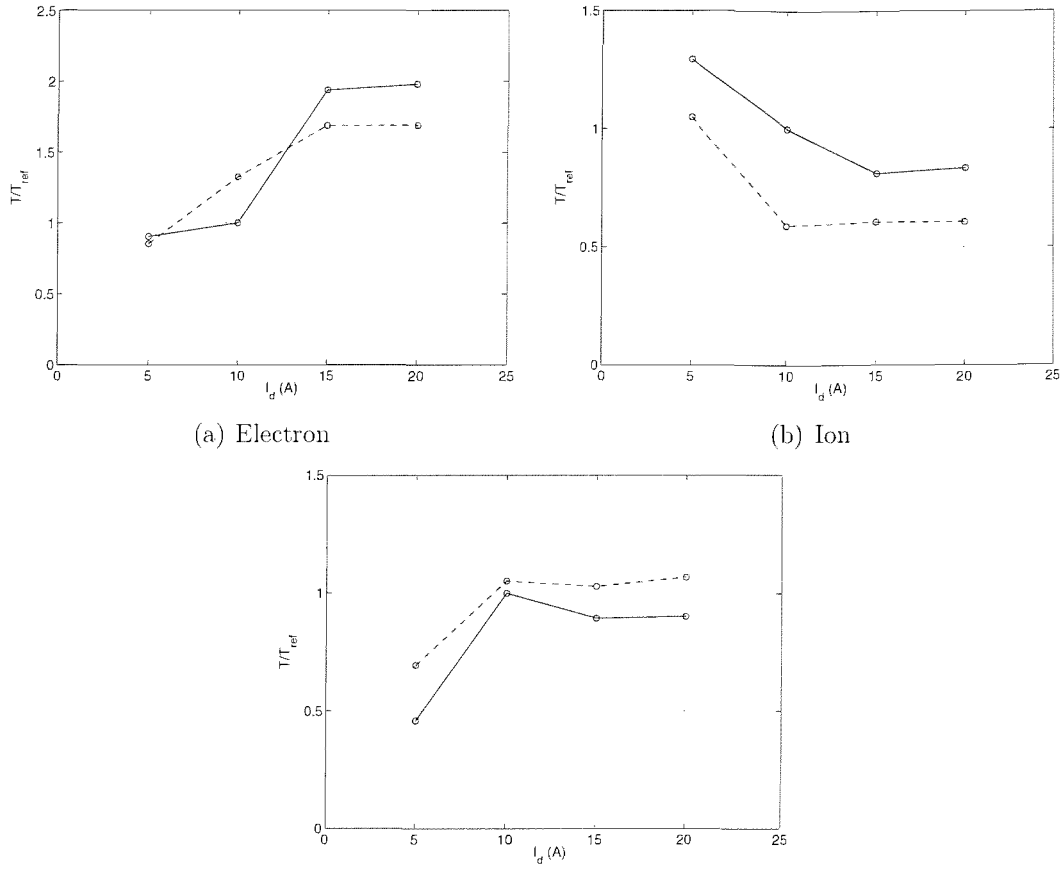

(a) Electron

(b) Ion

(c) Neutral

Figure 6.31: Cathode internal (i.e. reference position A) species temperature as a function of current for two mass flow rates, 1 mg s^{-1} (−) and 3.29 mg s^{-1} (−−). The data are normalised to the 10 A, 1 mg s^{-1} case for each species.

sheath in the plume is strengthening with increase in mass flow rate. In the case of the neutral flow, a similar trend to that of the internal gas is seen, with the high mass flow rate case rendering a neutral plume around a factor of 2 denser compared to the low case. The neutral density also increases with current, implying that the neutral containment near the cathode, presumably due to ion collisions, occurs in the plume as well as within the cathode. It seems that a denser neutral plume is caused at high current. This is of course important because the ionisation rate in the plume is a function of neutral density, so in some senses, at high current, the ion compression of the plume may lead to higher a ionisation rate.

Figure 6.33 shows plots of temperature in the plume as a function of mass flow rate. In all cases, the variation in temperature is very little, particularly in the high mass flow rate case. In the heavy particles, the higher mass flow rate is associated with an increase in temperature as a function of discharge current, although this trend is at most around 10%. Interestingly, however, in the case of the electrons, the

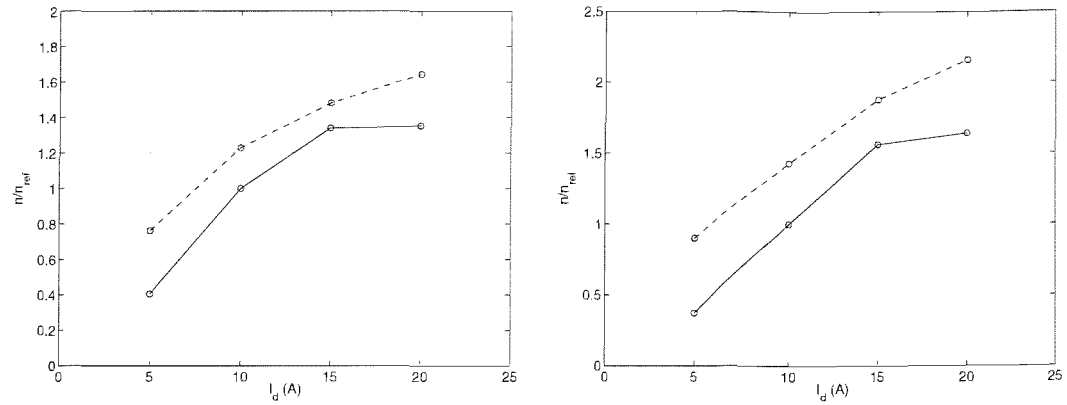

(a) Electron

(b) Ion

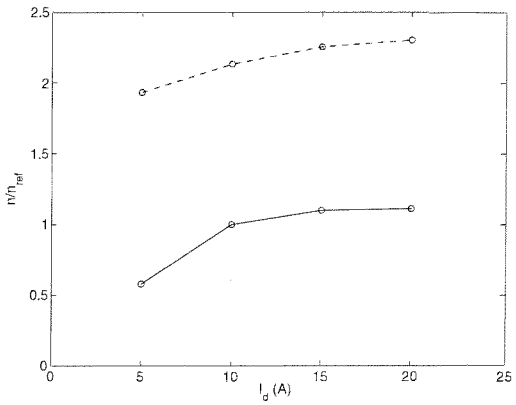

(c) Neutral

Figure 6.32: Cathode plume plasma density as a function of current for two mass flow rates, 1 mg s^{-1} (−) and 3.29 mg s^{-1} (−−). The data are normalised to the 10 A, 1 mg s^{-1} case for each species.

temperature is slightly higher at the lowest current setting at both mass flow rates. This contradicts some experimental results, Rudwan [88], and is hence considered in greater detail during the experimental comparisons found in the following chapter.

In summary, the plume plasma at low mass flow rate behaves much as expected[27]: there is a rise in density associated with higher currents, while there is very little temperature change. No spot-plume transition was found to exist, although this was not unexpected due to the geometry of the simulation.

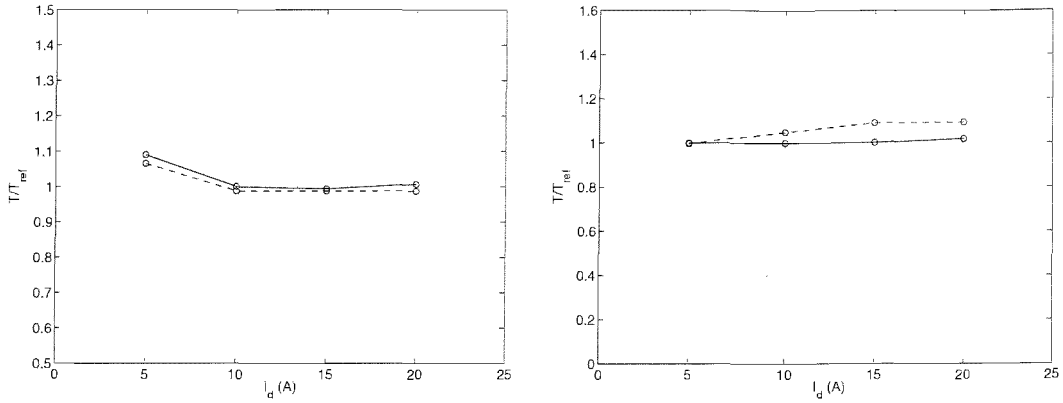

(a) Electron

(b) Ion

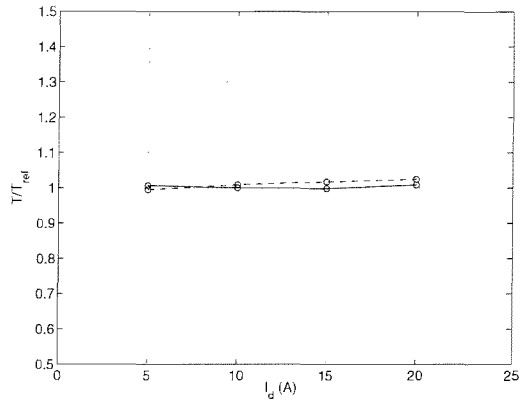

(c) Neutral

Figure 6.33: Cathode plume species temperature as a function of current for two mass flow rates, 1 mg s^{-1} (−) and 3.29 mg s^{-1} (−−). The data are normalised to the 10 A, 1 mg s^{-1} case for each species.

6.6 Tip radius

As was the case for the neutral flow, it is expected that the tip radius of the cathode will have a significant effect on the discharge characteristics. Four different values are run, ranging from the minimum realistic radius of 0.3 mm, 0.5, and 0.7 mm, and finally 0.9 mm. Cathodes are not typically run at 0.9 mm as this is very close to a no-tip case.

In all cases, the simulation converged to an equilibrium state. Figures 6.34 and 6.35 shows plots of density in the cathode interior for all three species as a function of tip radius. This may also be compared to a similar plot in the case of a neutral only flow, fig. 5.25. For neutral Xe atoms the trend is very similar to that of the neutral-only flow. The ion and electron trend is also of the same general character as the neutral plot. In all species, the trend is for the density to decrease monotonically with increasing tip radius. This is of course qualitatively expected since a smaller tip radius implies less geometric containment of the internal plasma. For a stable discharge, a higher internal neutral density then implies a higher plasma density: it appears that the internal plasma density scales with tip radius as the neutral density. Such a relationship is not unexpected since it is quite reasonable to suggest that the plasma density in the emitting region rises to a point where the emitting region is supplying sufficient current for the discharge. This can be further explained: first of all, suppose that the size of the emitting region is strongly related to a pressure balance between between the internal neutral flow and ion flow from the cathode exterior into the emitting region. This means that it could be the case that there is simply a direct relationship between the emitting region density and the neutral density. Such a hypothesis is of clearly an oversimplification, particularly in the case of a any temperature variation with tip radius, which it is necessary to examine next.

Species temperature in the interior is shown in fig. 6.35. The first point to notice is that as with the plot of density, temperature decreases monotonically with increasing tip radius. The neutral temperature does not change significantly at all, the change being a few percent only. This supports the hypothesis that the neutral temperature is strongly related to the cathode body temperature, and that neutral-neutral collisions and neutral-wall collisions dominate the neutral temperature compared neutral-plasma species collisions (at least in the cathode interior). The electron temperature also does not deviate by a large amount, particularly between the more enclosed configurations (0.3 and 0.5 mm). The temperature does however drop by a factor of nearly a half in the 0.9 mm case. Such behaviour could be explained in terms of the lower density. Perhaps the most interesting effect is the ion temperature, which is greatly increased. Clearly, in the case of a small tip, and correspondingly a higher density plasma, the ion temperature increasingly tends toward the electron

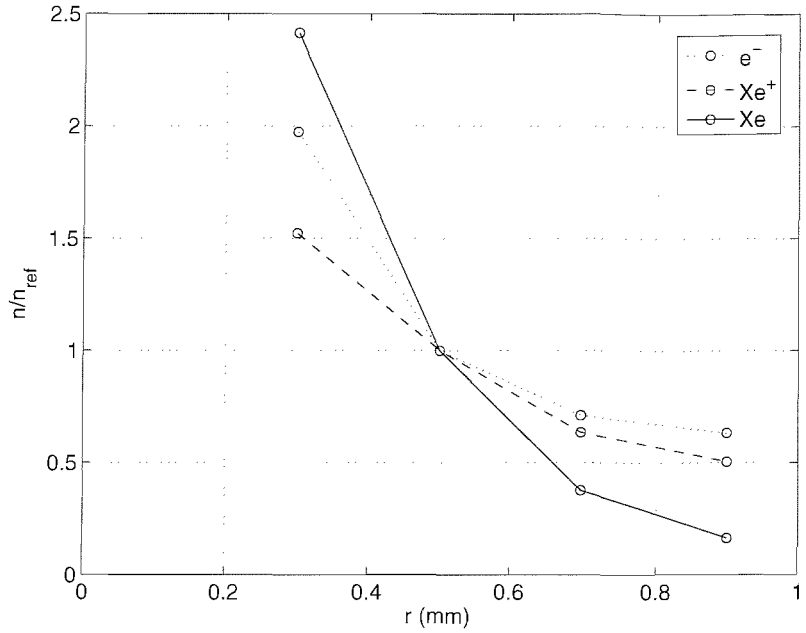

Figure 6.34: Interior species density as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm.

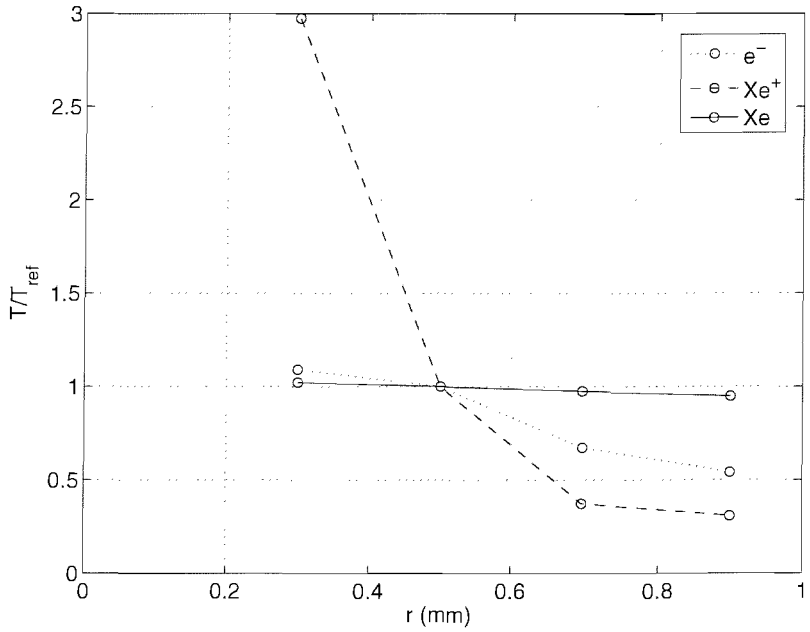

Figure 6.35: Plume species temperature as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm.

temperature, while in the lower density case (wide tip), the ions are much cooler. Such a pronounced trend is unexpected and important because ion temperature (and hence ion heat transfer to thruster components) is a critical factor in thruster lifetime and reliability. The validity of this result, particularly since the model assumes an isothermal cathode, needs to be considered carefully however.

6.7 Anode position

It is currently a popular hypothesis that the position, type and/or geometry of the anode in a diode configuration cathode discharge has a defining effect on the spot/plume characteristics (§2.1, fig. 2.2 and [88]. For this reason, it is worth examining the effect of the anode position has on the plasma conditions. There is however a reason why this model may not show such an effect in detail due to the compromises made in choosing parameters. Effectively, the model has been designed and verified to generate the most accurate reproduction of the hollow cathode *internal plasma* as possible. This was because the least is known about the internal plasma, and hence the greatest utility derived from examining it. Unfortunately, due to the dramatic difference in density between the internal plasma and plume plasma, coupled with the fact that particle codes do not tend to scale well over very large ranges of density, means that we are less confident about the precise plume plasma. In essence, the code is optimised for the internal plasma, but we suppose that changing the anode position most greatly affects the external plasma. Finally, due to computational constraints, the range over which the anode position can currently be varied is not great, compared to that found in experimental studies.

Recall that the anode is simply set to exist on the downstream boundary (i.e. on the line $z = L$), so changing anode position is akin to changing the 'length' of the simulation. The default anode position in the reference case is at $z = 20$ mm (this is 10 mm from the inner face of the tip). In order to investigate the effect of the anode upon the results, it is moved in increments of 2 mm both closer to and further from the original position; the data points for anode position are: 8, 10, 12, and 14 mm, expressed as a distance from the inner face of the cathode tip. In general, however, anode position is expressed in plots in terms of deviation from the reference case Δz. There are no particular computational concerns over this range, although as the wider separations are modelled, a larger simulation volume is required and compute time rises accordingly.

All runs of the simulation where found to converge to an equilibrium state. Figure 6.36 shows a plot of anode voltage with anode position. This plot is unexpected both in the trend, and in the absolute values shown. The trend is simple enough: the

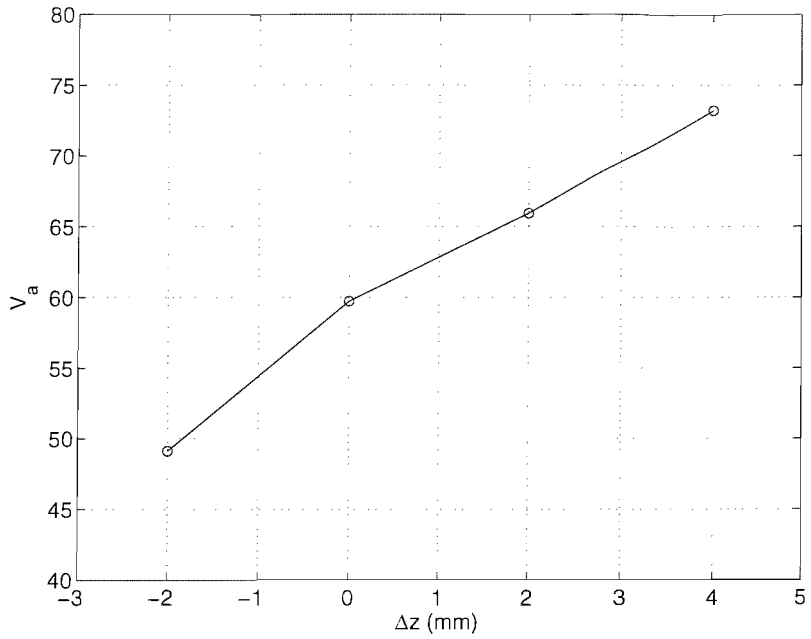

Figure 6.36: Anode voltage as a function of anode position at a fixed current.

voltage appears to be proportional to anode displacement. This is unexpected as one would assume that as the anode-cathode separation approaches zero or ∞, the voltage would tend in an exponentional fashion to either zero or a very large value. What might be expected would be a much shallower gradient, in which case the plot would seem more reasonable. In any case, although such a result is worrying, it certainly demonstrates that something significant is happening as the anode is moved. The next question naturally is: how does the plasma vary with anode position?

A plot showing the variation of species density in the cathode interior and plume are shown in figures 6.37 and 6.38 respectively. As usual, data are normalised to the reference case value, and all three species are plotted.

Considering first figure 6.37 it is clear that the variation in internal plasma conditions with anode position is considerable. The neutral density drops as the anode is moved away from the cathode. This is natural since the extent by which the neutral expansion is impeded is reduced as Δz increases; in other words, a closer anode implies a higher back pressure so that the internal density will tend to be higher. There is a less pronounced change in internal plasma density, although in the case of most distant anode, electron and ion densities drop in a similar fashion to the neutral gas. As before it is natural to construct an argument based on the existence of an emitting region requiring a constant plasma density to neutral density ratio (or equivalently ionisation fraction), so if the neutral density drops, a lesser plasma density is required to generate the same discharge current.

What is now interesting is that the anode position has a comparatively smaller im-

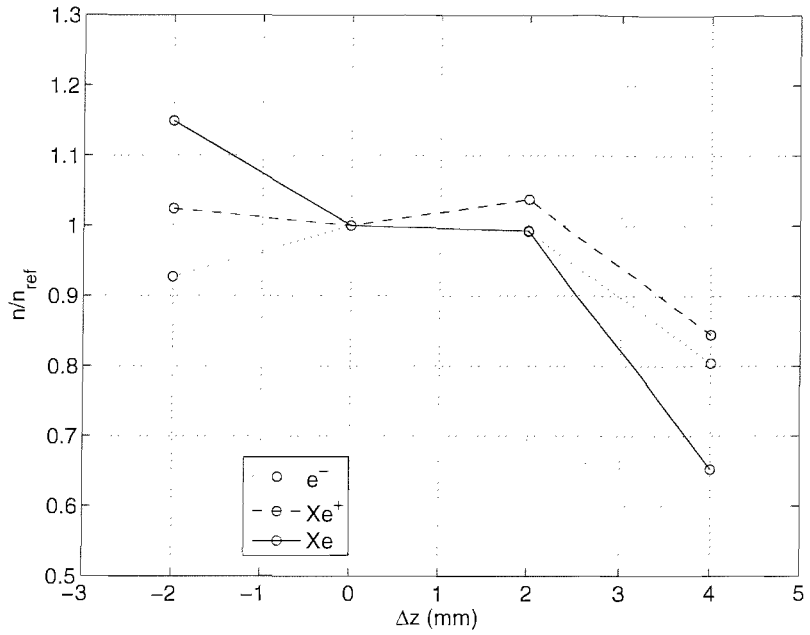

Figure 6.37: Interior species density as a function of anode position. Anode position of 0 represents reference case.

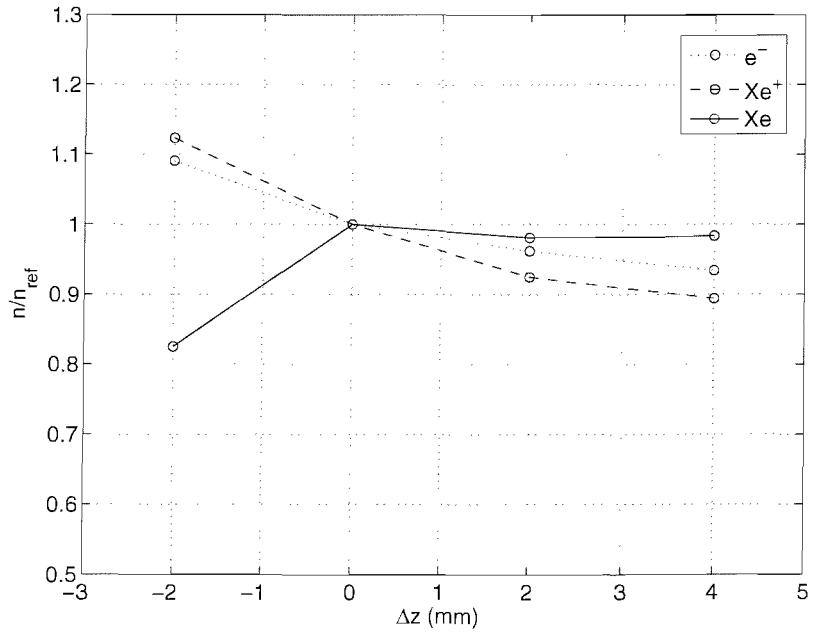

Figure 6.38: Plume species density as a function of anode position. Anode position of 0 represents reference case.

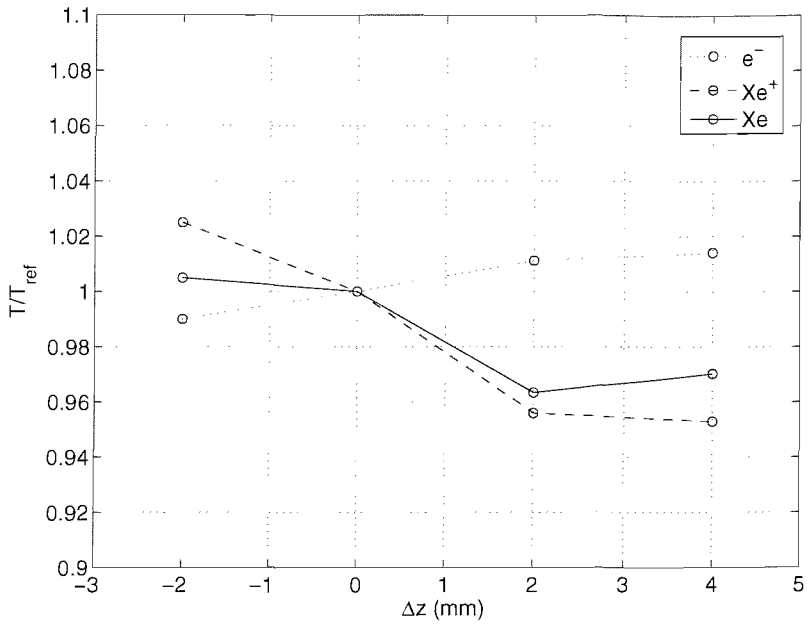

Figure 6.39: Plume species temperature ratio as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm.

pact on flow density in the plume (fig. 6.38). Apart from in the minimum case, there is no change in neutral density, while for the plasma species, density drops monotonically with increasing Δz. The trend in plasma density is clear (apparently little statistical scatter) and shows a difference in ion density of $\pm 10\%$ over the range considered; the plasma plume is thinning as the anode is moved away from the cathode.

Finally, species temperature in the plume is shown in 6.39. The extent by which temperature varies is very little, a few percent only rising to a maximum of 5%. The electron temperature rises slightly with Δz, perhaps due to the increasing anode voltage, while anode and neutral temperature both fall. It is significant to learn then, that it seems as though the anode position apparently affect the voltage by a very great degree (at least in this range of Δz), that the density certainly changes, while there is very little change in temperature.

Figures 6.40 and 6.41 show plots of the variation in plume pressure and velocity for the three species. In both plasma species, there is a consistent drop in pressure as the anode plate is moved away from the cathode. At the closest position, the neutral pressure is low, while it tends not to change a great deal at the other positions. It is difficult to explain why the neutral pressure should be lower in the closest position, particularly as this corresponds to the highest ion and electron pressures. It is possible this relates to the exact placement of the plume sample point (i.e. it is moving relative to the plate at different conditions).

The second figure, that of axial stream velocity, shows that there is a gradual rise

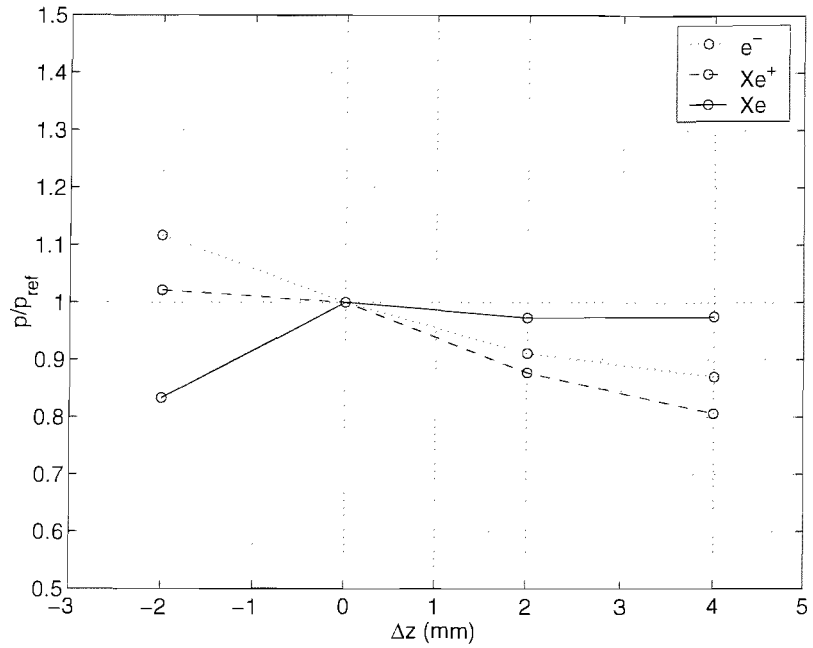

Figure 6.40: Plume species pressure ratio as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm.

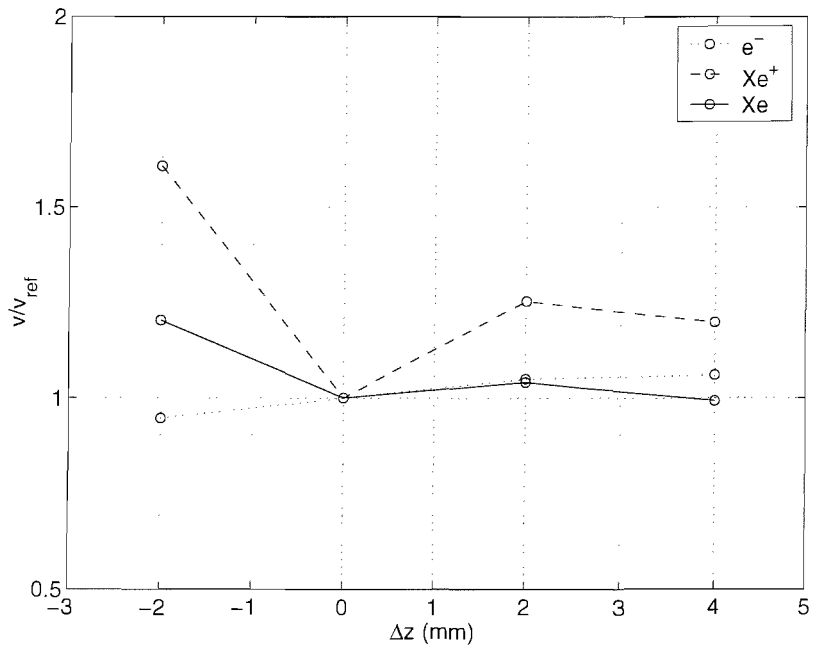

Figure 6.41: Plume species axial stream velocity ratio as a function of cathode tip radius. Data shown for tip radii of 0.3-0.9 mm in steps of 0.2 mm.

in electron velocity with anode position. The neutral velocity tends to be unchanged at all but the closest position. Here, the velocity is elavated. This is consistent with the plot of pressure: lower pressures corresponding to higher velocities. The trend in ion velocity is very difficult to explain. This is likely due to a combination of factors, perhaps relating to the ion sheath attached to the downstream wall of the cathode.

Summary

The results produced by the model that characterise the hollow cathode plasma have been presented and examined. This concludes the main results section, although the comparison between the numerical model and experimental data is contained in the following chapter.

Chapter 7

Discussion

This chapter aims to assess how successful the research has been in fulfilling the objectives outlined in the introduction. This is done by discussing the results presented previously in the context of the goals that were set. Additionally, the implications of the presented results are considered.

7.1 Neutral Gas

The rarefied gas flow within the hollow cathode was examined in detail. The data presented included a set of numerical studies designed to build confidence in the results, followed by an examination of the flow structure. Several key parameters were then examined in turn, including mass flow rate, cathode tip radius (open area fraction), the working gas and the cathode temperature. Finally, some comparisons to experimental data were given.

The discussion of the results initially follows the order of the objectives set out in the early stages of this thesis. The degree by which the objectives were met are considered and further questions discussed.

7.1.1 Rarefaction

The objectives set out for the investigation of the neutral flow were three-fold. First, it was supposed that the condition (rarefaction) of the flow could unique due to a possible transition between flow regimes. In several of the plots shown in the results chapter, Knudsen number was plotted for a given parameter, or for different points within the flow. The code indicates that such a transition is present, as hypothesised by the first objective of the research. This can be seen most clearly in figure 5.14. The plot shows the three regions clearly, demonstrating that it is valid to assume a near continuum flow within the cathode. The flow in the vicinity of the tip — and area that

should be considered to be very important in providing an accurate description of the flow — is shown to be transitional in Knudsen number ($0.1 < \mathrm{Kn} < 1$). In this region, slip flow is expected to be significant. Finally, the plume of the cathode is certainly rarefied. In the case of an expansion to a hard vacuum, it has been shown that within a 1-2 cm downstream of the cathode, the plume expansion is such that $\mathrm{Kn} \rightarrow \infty$ — the definition of free molecular flow. When a downstream pressure condition is applied, even that of the interior of a typical laboratory vacuum chamber, the plume rarefaction is reduced significantly.

The rarefaction transition that has be shown to exist here is important for future cathode models. In particular, a model that assumes one level of rarefaction throughout the flow is liable to error. It is however fair to state the the error introduced by such an assumption is comparatively small and only relevant for models striving for high levels of overall accuracy. The recommendation that can be drawn from this finding is that simple models assume a continuum (preferably slip-flow) models, unless the cathode is (a) exposed to a hard vacuum downstream and (b) operating a very low mass flow rate. The justification being clearly that the combination of (a) and (b) can lead to a rarefied condition even upstream of the tip. A final note is that the rarefaction is amplified if the mass flux per unit area remains constant while the cathode dimensions are reduced.

A final implication of the flow rarefaction is the possible extension into microfluidics. Indeed, this was the subject of Crawford and Gabriel [25]. It is true say that since the Knudsen number is the ratio of mean free path to characteristic length, regardless of gas conditions, if the length is sufficiently small, then the flow must become rarefied. DSMC simulations *room temperature, atmospheric pressure* air flowing inside MEMS micro-channels have been completed, for example [73]. A natural conclusion is then that micro hollow cathodes may be devices whose fluid flow is entirely rarefied. This is potentially very important because as was seen in the plasma results, then conventional scale cathode seems to heavily rely on the neutral pressure to contain the main emitting plasma. A conclusion of this work would be that building a very small cathode would necessitate an increase in mass flow rate *per unit area* to retain the internal pressure. Hence, if it assumed that in all other respects, the cathode performance scales linearly with size, but the normalised (to cathode dimension) mass flux *increases* then it must be concluded that the specific impulse must fall comparatively. Hence this work predicts that cathode micro-thrusters may under-perform compared to their full scale counterparts.

7.1.2 Comparison to Experimental Results

Regrettably, there exists very little experimental data that describes neutral gas flow within hollow cathodes. The most common data is a measurement of pressure in the propellant feed pipe to the cathode. This can be compared to the upstream pressure the code produces. The other candidate for experimental comparison is the results of Gessini et al. [38], who aims to measure the thrust produced by the cathode.

Upstream pressure gauges

Fearn and Patterson [33] conducted some experiments on a cathode with similar geometry to the reference case (the cathode used in the later T5 and current T6 ion thruster). While the experiments presented in [33] were primarily aimed at investigating stability under high currents, the data are useful because a commercial pressure transducer was attached to the Xenon feed pipe to the cathode. The authors then provide results for cathode pressure. They find that the pressure in the feed pipe is proportional to mass flow rate at a given discharge current. As the current is increased (they examine currents of 5, 10, 15, 20 and 25 A) the pressure remains proportional to mass flux, although at a higher discharge current, the same mass flux results in a higher pressure. It is not unreasonable to suggest that the positive correlation between discharge current and pressure is due to a heating of the cathode (and hence the gas). It has already been shown in section §5.7 of this work that an increase in cathode surface temperature yields an increase in pressure.

Although many aspects of the experiment can be reproduced in the numerical model (the geometry of the cathode, for example), there are some factors that are unknown that could significantly affect pressure measurements. In particular, in light of the discussions previously regarding the difficulty of measuring insert temperature, it is not known what the insert and feed pipe temperature is. An additional point to consider is the geometry of the feed pipe where the pressure is measured. It is not possible to use the reference case data as it is because there is one very important difference between the geometry of Fearn and Patterson and that used above: in the experiments, the tip diameter was 0.75 mm. Given the strong dependence of cathode interior density with tip radius (covered in §5.5), a new set of data is required with the new tip radius (i.e. 0.375 mm, representing an open area fraction of 0.14 compared to the cathode interior). This data set covers cases for two mass flow rates (0.63 mg s^{-1} and 1.0 mg s^{-1}, and a range of three cathode temperatures (1300, 1400, 1500 K). This range is chosen because it is not known what the internal wall temperature is, and such a range covers the extremities of the expected temperature. A final adjustment to the numerical data is required. The standard T6 cathode has a 20 mm long insert. The numerical data here is obtained using a 10 mm long insert, so the

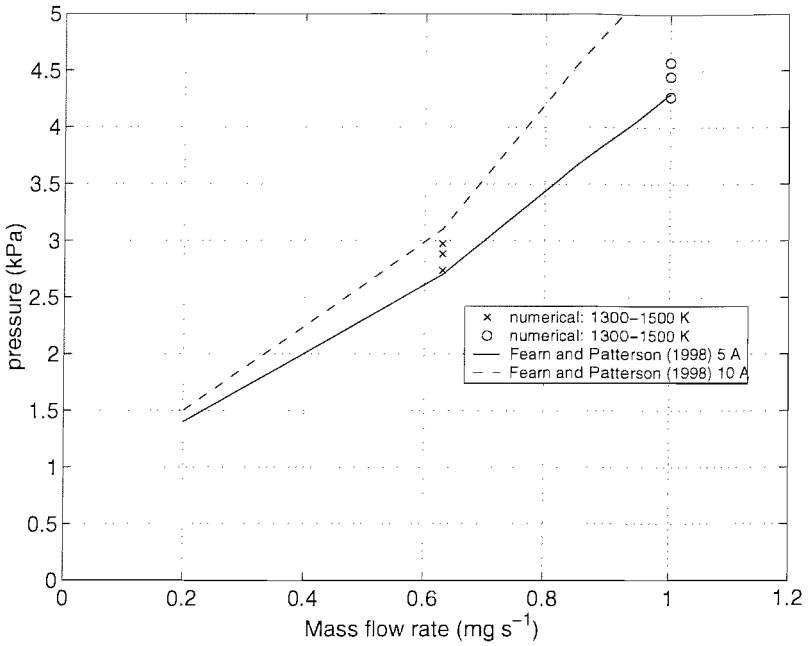

Figure 7.1: Comparison of the experimental data of Fearn and Patterson for hollow cathode upstream pressure to numerical results for the same cathode geometry. Numerical data is plotted for a range of insert temperatures and experimental data is shown for 5 and 10 A discharge sweeps. In the simulation, the pressure monotonically increases with cathode temperature, so the larger pressure data points are at the higher temperatures.

pressure should be extrapolated upwards to correspond to the additional 10 mm not modelled. This extrapolation is valid for two reasons: firstly, pressure exhibits a direct proportionality with axial position, with a negative axial gradient, obviously. This can be seen is several figures contained previously in this chapter, see figure 5.27 for example. Secondly, when the full 20 mm of the insert is modelled, it has been found that extrapolated pressure line follows the real results very well. The full 20 mm is is not modelled in the data presented in this section simply because the upstream 10 mm tells us very little, does not significantly affect the rest of the flow and is very computationally intensive to capture correctly.

Figure 7.1 shows the numerical results plotted alongside the experimental data. The numerical results compare well with the experimental data, particularly for the cooler 5 A discharge. It is not expected that the results should be identical because the numerical data is neutral gas only, while the experimental data is taken during discharge. The similarity between the results tells us that it is possible that the presence of the plasma near the tip of the cathode is not severely affecting the neutral gas dynamics near the cathode entrance. This would support the commonly held hypothesis that the plasma only extends a few centimetres upstream of the tip. At the higher discharge current, the experimental data shows pressures higher than are

achieved even for a 1500 K insert. This does not necessarily imply that the insert is hotter than 1500 K in the experiment, as the increase in pressure is most probably due to a stronger heating of the downstream gas as a larger current passes through it. Fearn and Patterson give data for discharge currents all the way up to 25 A in [33] (the paper primarily covers high current hollow cathode characterisation). At higher discharge currents still, the pressures are increased, although even at 20 A, the upstream gauge does not read much higher than 4 kPa (at 0.63 mg s^{-1}).

Plume gas dynamic measurements

Of interest downstream is how the thrust varies with $\dot{m}$, because there are some analytical models against which we can compare the results. Gessini et al. [38] also provides data from an experimental thrust measurement system; see appendix D for details of the experimental method. It is easy to compare results with the setup of Gessini because the configuration of the reference case broadly follows that of Rudwan [88], and Gessini uses the same experimental rig, vacuum chamber and cathode setup (in both cases the cathode from the T6 ion thruster). The setup is then nearly identical to the reference case. The exceptions are as follows. Firstly, the Gessini experiments use Argon as a propellant. Next, cathode temperature is set to 300 K; this is the thrust calibration test of Gessini, where the cathode is *not* heated. Finally, the range of mass flux data taken is 0.207, 0.414, 0.621, 0.734 and 1.01 mg s^{-1} (0.207 mg s^{-1} Argon is 0.5 Aeq). The data can be found in `data/neutral/gessini` including the raw experimental data (`thrustreadings.dat`) on the CD.

The simplest way of deducing thrust in the numerical model is to measure the plume exhaust velocity. If the gas is assumed to be stationary on the upstream boundary (a reservoir condition), then the thrust is simply the product of the plume velocity and the mass flow rate,

$$T = \dot{m} v_{ex}. \tag{7.1}$$

The question is: what is the exhaust velocity? Figure 7.2 shows a plot of axial velocity along the centreline of the cathode at reference conditions. After having negotiated the main expansion, the flow actually slows slightly at a position adjacent to the keeper. After this point, the flow accelerates to a plateau. It is the mean axial component of the flow velocity in this plateau region that is equivalent to the exhaust velocity. Radial velocity is 'lost' in terms of contribution to thrust as due to axisymmetry, only net axial forces may exist.

Given these statements, it is possible to use the velocity measured by the sample point in the plume as a good approximation to v_{ex} in equation 7.1. In addition, it is possible to compare the code and experimental results to the ideal analytical result for continuum expansion and for free molecular expansion. Specifically, from

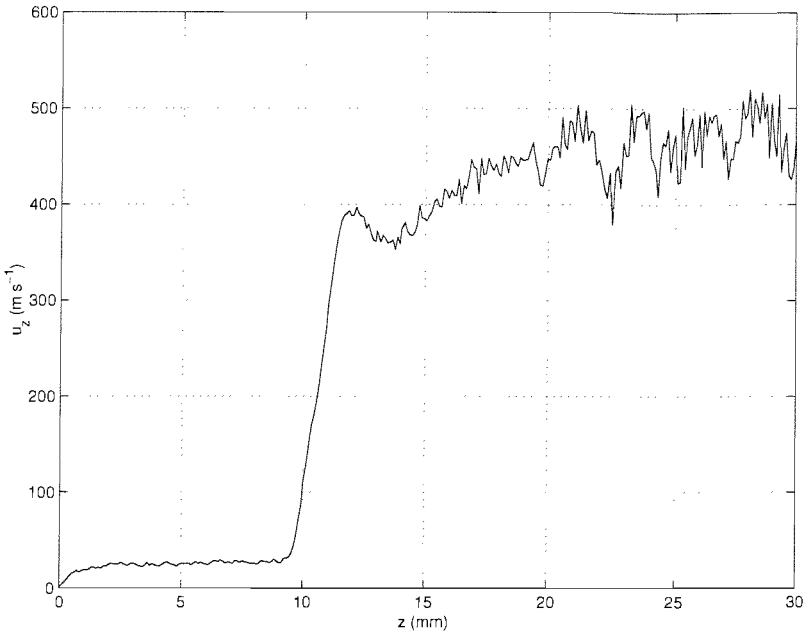

Figure 7.2: Axial centreline plot of axial stream velocity.

the expression for exhaust velocity (1.12), the thrust becomes

$$T = \dot{m}\sqrt{2\frac{\gamma}{\gamma-1}\frac{kT_0}{m}} \tag{7.2}$$

assuming a vacuum downstream conditions so that $p_e/p_0 = 0$. Gessini et al. [38] also quotes the thrust of a quasi 1-D free molecular expansion as

$$T = \dot{m}\sqrt{\frac{\pi}{2}\frac{kT_0}{m}}. \tag{7.3}$$

It is expected that any results should lie between these two results: the continuum result should act as a maximum thrust (if we find data points above the continuum results we should be worried!), while the free molecular result should act as a minimum. Of course, since the free molecular prediction is a theoretical maximum thrust for an expansion of a collisionless gas, it is still possible to have a poorly designed nozzle that produces less thrust than the collisionless maximum while operating in a continuum regime. In summary: it should be expected that thrust data lie under the continuum prediction (because the hollow cathode is by no means a perfect nozzle), and that numerical results from the code are similar to those measured experimentally.

Figure 7.3 shows both of the above expressions plotted with both the data of Gessini and numerical results. Both the experimental results and numerical results show a similar trend to the analytical prediction: this is not surprising; it would be expected that thrust should be proportional to mass flow rate. There is in fact very

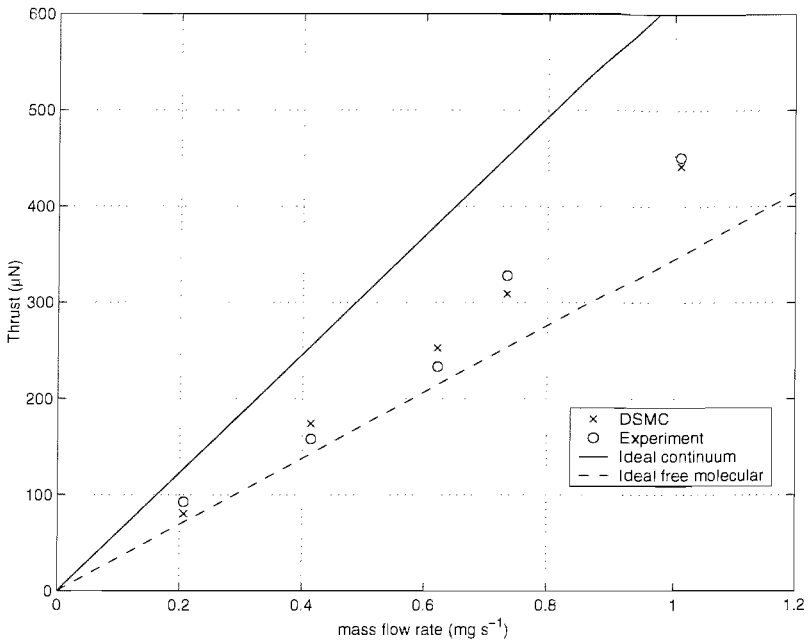

Figure 7.3: Comparison of numerical, experimental and analytical results for the thrust produced by cold gas flow in the hollow cathode.

good agreement between experiment and computation. Both seem to indicate a similar level of underestimation of thrust compared to the continuum prediction. The data points from the numerical model are consistently within $\pm 5\%$ of the experimental data. Error bars for the experimental data are unavailable at this time, although initial investigation indicates that they are of the order of 40 μN — equivalent to approximately $\pm 8\%$ for the 1 mg s^{-1} case. Based on the error estimate for the experiment, it is possible to state that the numerical results lie within experimental error. This serves as an excellent verification of the numerical model.

Conclusions

The second major objective of the neutral gas analysis was to build a numerical model that is accurately able to predict forces produced by (and within) the cathode. The data were compared to the data of Gessini for absolute thrust produced by a cold gas flow through a cathode. Since good agreement was found between experimental and numerical data, it is fair to say that the model achieves the second objective. While the evaluation of internal forces is possible, this is often not useful in the case of cathode engineering since the design limits of the cathode are generally set to withstand thermal conditions: the device is always heavily over-engineered with respect to gas dynamic and structural forces present.

Collaboration with Rudwan: hollow cathode breakdown

During the period of the research, extensive use was made of the neutral gas model data by Rudwan [88]. The specific data used can be found in [88, appendix B]. As part of the research, the mininum voltage for breakdown of the cathode was investigated. The data were used to plot Paschen curves for the theoretical breakdown voltage. The use of the data provided by this research by Rudwan was found to be very helpful in crystallising the arguments put forward for the anomalous breakdown characteristics of the cathode. The ability use of neutral gas data provided by this model to construct Paschen curves was important because experimental evidence indicated a 'severe departure' from the classical behaviour: this was demonstrated by comparison of the curves to actual breakdown behaviour. Briefly, Rudwan hypothesises that the primary reason for the departure from classical theory is due to the Penning ionisation effect. This effect is occurs in a gas mixture and allows for ionisation to occur at energy lower than would be expected for the primary gas species (the inert gas). The hypothesis that Barium evaporation from the cathode insert causes a gas mixture conducive to the Penning effect seems to adequately explain the departure from classical theory. The reader is referred to the thesis [88] for further information, in particular the appendix that covers how Rudwan used the data produced in this research to construct the Paschen curves. Some additional modelling of electric field penetration into the cathode pre-breakdown was also conducted. This is discussed in greater detail in the relevant (following) subsection.

7.1.3 Cathode geometry

It is clear that the primary design element to be changed in the cathode is the tip radius. This is because in any nozzle, the throat area is of critical importance to the flow. Several plots were shown as a function of area ratio; this is simply the ratio of tip radius to insert radius, and so ranges from near zero for a very enclosed cathode to 1 for a cathode with no protruding tip. The results show that the density at the tip point drops following a e^{-x} form where x is the area ratio, such that for area ratios larger than around 0.7 the density changes very little. A similar trend is seen for tip velocity. Indeed, for area ratios larger than 0.5, very little change is seen in the axial tip speed. That the relationship between area ratio and upstream/tip flow properties was consistently predictable for all area ratios indicates that under no value of A_t/A_i does the flow become supersonic. This was confirmed. A final point to raise is that very little change was observed in the plume of the cathode. As has been mentioned, taking experimental measurements of the plume is frequently done partially because interior experimentation is so difficult. This means that there exists a large body of plume

data. Fearn and Patterson [33], for example, provide data points for cathode radii of 0.375 mm and 0.65 mm (area ratios of 0.1406 and 0.4225 respectively) and analyse cathode operation mode transitions for the two values. What this research indicates is that if the neutral gas dynamics are affecting the transition characteristics of the cathode at various tip radii, then the transition phenomena is either (a) occurring outside the cathode but comparatively independent of neutral gas effects, (b) processes that occur inside the cathode that are in some way related to neutral gas interaction or of course (c) a combination. At this point, the discussion clearly leads toward the implication of these results with respect to the various hypothesis regarding the mechanisms of HC mode transitions. It is perhaps wise, however, to discuss the findings based on this research relating to operation modes once the full plasma results have been presented. The reason for the delay of the discussion is clear: the neutral gas results presented thus far are always susceptible to the argument that the gas flow may be significantly perturbed by the discharge.

7.1.4 Neutral gas parameters

To conclude the analysis of the neutral flow in the cathode, the effects of altering various parameters were investigated. The results from these studies led to further understanding of the onset of rarefaction within the cathode, as well as some detail regarding the thermal modelling and heat transfer within the flow.

Mass flow rate

A very clear relationship between interior gas density and mass flow rate was found (fig. 5.20): it was discovered that interior density is proportional to injected mass flux. The confirmation of this result is important as it can be used to help mass flux dependence into the basic 1-D models being developed[50 86]. In relation to the rarefaction of the gas, it was shown that the flow is rarefied for Xenon below 2 mg s^{-1} (1.48 Aeq) *in the tip of the cathode*. It had been suggested by some authors [38] that the gas flow could always be assumed to be continuum (or at least slip-flow) within the cathode and rarefied/collisionless outside. For low flow rates, particularly less than 1 mg s^{-1} (0.74 Aeq Xenon), the slip and rarefaction in the interior was shown to be significant. Of course, such observations are only of relevance to low mass flux cathode operation: generally when the cathode is used as an ion engine plume neutraliser.

Propellant species

Results were provided for flow using three different propellants, Argon, Krypton and Xenon in ascending mass order. The results matched the rather obvious predictions very well; in reality, since the *mass* flow rate was being held constant, but the mass of the individual species varied between sets, predictions such as higher densities for the less massive species were found to be correct. Since the cathode temperature was fixed, and the channel was sufficiently long to heat all the species to the same temperature, the resulting data showed a fixed fixed temperature, but variable density and pressure. Higher density and thus pressure resulting from the lighter species. It was surmised that the power input to the gas from the cathode would be different if the different species where to reach the same temperature. A plot of wall heat flux showed this to be true and revealed a common difference with particle based codes when compared to traditional models. Setting absolute temperature in this way was found to be simply the most robust, foolproof method in terms of the stability of the code, despite the apparent disadvantages we see in some results.

Cathode body temperature

The results for neutral flow with different cathode body temperature revealed perhaps the greatest weakness in the neutral model: that the cathode temperature must be fixed and uniform. Ironically, however, the necessity during the neutral gas analysis to fix the temperature led to an opportunity to directly validate the code by comparing to basic compressible nozzle theory, and observing the $v_{ex} \propto \sqrt{T}$ behaviour (figure 5.29). Upon completion of the results for neutral gas flow it became clear that this was an area of weakness, so it would clearly be a recommendation for future work to implement a thermal model. This is of course considered in greater detail later.

7.2 Plasma Discharge

Following the work done to characterise and investigate the neutral flow in the hollow cathode, the second phase of the research was to add a plasma model so that the discharge could be simulated. It was first necessary to assess the degree of stability present in the numerics.

7.2.1 Numerical Validity

Although it was possible to have some confidence in the validity of the model based on the study of the neutral model, it was necessary run some addition numerical studies to build confidence in the results. These studies allowed either for the numerical factor to be shown to have no effect on the results, or to show the extent by which the factor might affect the results — an estimate of error.

Although the time-step had already been analysed for the accuracy of the neutral model, it was of course necessary to re-assess time-step dependence since the PIC model is intimately linked to it. It was found that the numerical accuracy of the results is strongly correlated with time-step. However, for the time-steps that where used for all the results presented, the influence of time-step can be considered minimal. It was unfortunately impossible to practically run the code at a time-step that we could be sure of no influence. It was concluded then that the analysis showed that the peak error in density due to time-step dependence was unlikely to exceed a few percent, while the error in temperature is smaller than the statistical scatter.

Four different PIC meshes were tried, and it was found that density increased with larger meshes, although the variation decayed so that for the mesh used to generate the results presented, there was little difference in density compared to a denser mesh. In terms of temperature, there was some evidence of so-called PIC 'numerical heating', although again this was mitigated by using a finer mesh. It is possible to conclude that compared to the other studies completed for the plasma model, the mesh seems to have the least severe numerical impact on the final results. The magnitude of any variation in other cases is greater.

Although several compensations were made to the model to effectively remove the effect of electron mass modification (§3.5.6) it was necessary to test for electron mass adjustment for two important reasons. First, to verify that the compensations are correct and second, to judge whether even given correct compensations, the mass modification would still affect the results. Although it was found that variation in the data was minimal due to changes in mass modifications, further investigation did show that the transient behaviour of the model was affected by electron mass modification. It is possible then to tentatively assume that for steady-state studies, the electron

mass modification has little adverse effect. However, we know now that it has some effect, and this must be considered when discussing the findings of later results. That the transient behaviour is affected by f means that the potential applicability of the model to the study of hollow cathode breakdown or instability is limited.

7.2.2 Discharge Structure

The structure of the discharge was presented in four sections: contour plots, axial centreline plots, radial cuts and velocity distribution functions. These allowed the hollow cathode internal plasma to be examined in great detail. It was found that the the proposed higher density emission region exists, and that two different types of sheath exist within the insert section, with a different sheath again attached to the tip. As expected, it was shown that the plume of the hollow cathode in diode configuration forms a $z-$pinch plasma, so that there is a high on axis plasma density and a sharp drop far from the axis. A significant magnetic field is generated by such a configuration. It was also found that the plasma species tended to form Maxwellian velocity distribution functions in the high density cathode interior, while highly non-Maxwellian distributions where observed in the plume. This was also expected, see Crofton and Boyd [28] for examples of experimentally measured non-Maxwellian plume distributions. It is fair to say that there were no major problems or inconsistencies with the detailed analysis of the results. The major question that remains is: to what extent are the absolute values presented similar to experimental data?

7.2.3 Comparison to Experiment

The reference case geometry is effectively that of the cathode from the T6 ion thruster and is a very popular cathode used by many researchers in experimental work. Perhaps the best way to gain an initial impression of the similarity between experimental work and this research is to extend table 2.1 (p. 14) to include the numerical results. Such comparison tables are often very informative and useful when judging the validity of a numerical model. This is the case here, although it is also true to say that experimental data between different authors tends to vary as much as the results from the numerical model vary compared to experiment. This can be seen in table 7.1.

Considering first the internal conditions, the best match between experimental setup and numerical setup is that of Rudwan [88]; it is this data that is compared in more detail to the simulation results later. For the range of data considered, both the plasma density and energy fall well within the selection of experimental data given.

	Propellant	Cathode	$\dot{m}$ (Aeq)	I_d (A)	T_e (eV)	n_e (m^{-3})			V_p (V)
Internal Plasma									
Fearn and Phillip (1973)	Hg	UK-10 (T5)	0.014–0.12	1.5	0.1-0.6	10^{16}	–	10^{18}	4.5–7.5
Siegfried (1978)	Hg	SERT-II	0.1	2.0	0.6	1.5×10^{16}	–	4×10^{20}	4.0–8.0
				6.0	0.4–0.6	10^{17}	–	3×10^{20}	4.0–6.0
Ahmed Rudwan (2002)	Xe, Ar, Kr	UK-25 (T6)	0.1–4	1–5	0.8	10^{20}	–	6×10^{22}	–
Crawford (2004)	Xe	UK-25 (T6)	1–3	5–20	0.8	10^{20}	–	2×10^{21}	1
External Plasma									
Csiky (1969)	Hg	SERT-II	0.05	0.3	1.5-2.0	10^{16}	–	10^{17}	11
				2.0	0.5	10^{17}	–	10^{18}	11
Siegfried (1978)	Hg	SERT-II	0.1	2.0	2–5	5×10^{16}	–	3×10^{18}	12–28
				6.0	0.6	3×10^{17}	–	5×10^{18}	11–12
Singfield (1990)	Xe	UK-10 (T5)	0.175–0.475	1.1	0.4–2.2		no data		8–13
Friedly (1992)	Xe	High I_d	0.37	20	2.0			1.5×10^{19}	12
				40	2.9			5×10^{19}	13–15
				60	3.8			6×10^{19}	15–17
Monterde (1997)	Xe	UK-25 (T6)	1.67–3.11	5–10	0.3–0.4	1×10^{17}–	5×10^{17}	$\approx 10^{20}$	12
Edwards (1997)	Xe, Ar, Kr	UK-25 (T6)	0.3–1.2	2–15	1–5	5×10^{17}	–	3×10^{18}	16–30
Crofton (2003)	Xe	UK-25 (T6)	0.2–2.4	4–8	1.15		no data		10
Crawford (2004)	Xe	UK-25 (T6)	1–3	5–20	1-2	10^{18}	–	10^{19}	15–30

Table 7.1: Plasma properties reproduced for comparison with this research. Some of this table was originally complied by Edwards [31] and has been adapted and extended to include several other, more recent results. Cathodes with names such as UK-25 refer to the those that accompany RAE/DERA/Qinetiq gridded ion thrusters where the number refers to the main discharge chamber diameter in centimetres. SERT refers to the US 'Space Electric Rocket Test' program that evolved to provide the NASA DS-1 ion thruster.

Bear in mind however that experimental readings of plasma density in the cathode interior are typically associated with error bars of up to two orders of magnitude [87].

One thing that is always immediately apparent from glancing at a table such as this is that the majority of the data is external data. In fact, the table could be further expanded with external plasma data on hollow cathodes; the ones that are included are simply a representative set. The numerical data for the external plasma compares very favourably with experimental data. Perhaps the only major discrepancy is that the numerical values for plasma density tend to be quite high. This could be explained due to the fact that the numerical data are for comparatively high discharge currents and also due to geometry considerations. Virtually all of the experimental data shown were obtained using diode configuration, as of course was the numerical data. The difference is that in many of the experiments, a hollow cathode located in some cases several cm downstream of the cathode were used. The numerical anode is comparatively very close to the cathode, so geometric containment enabling higher densities is to be expected. It will be seen in the following comparisons of experimental/simulation data that this effect may in fact be significant.

Plume magnetic field

Patterson and Fearn [77] provide a plots of current density and magnetic field radially in the plume of a cathode of similar geometry to that presented in this document. The data presented in [77] are derived from experimental measurements of current density radially in the plume. For a discharge current of 15 A, the profiles obtained are very similar to those generated in the code, while the peak magnetic field, located at r =0.4 mm is given as 6.4 mT. This compares very favourably with the prediction given by this model that the peak magnetic field is again located at approximately r =0.3 mm, while the peak magnetic field is 3.8 mT, this can be seen in figure 7.4.

The magnetic field profile compares very well with that of Patterson and Fearn, although the simulation predicts a slightly weaker field that is marginally closer to the axis. There are several reasons for this. It should be noted that in Patterson's analysis, some assumptions are made to generating the plot of magnetic field from current density data. In particular, it is assumed that the cathode has an effective orifice of 1 mm, and further that the magnetic field is zero at the tip (*'there is no pinching in the orifice'*). When factors are considered, it seems that the degree of agreement simulation and the results of Patterson are very good.

These quality of the agreement of the results of Patterson and Fearn and the simulation are particularly important for two reasons. Firstly, they demonstrate that modelling induced magnetism in the plume of the cathode is critical to a correct cathode model. There was little doubt over this issue for high currents, but it seems

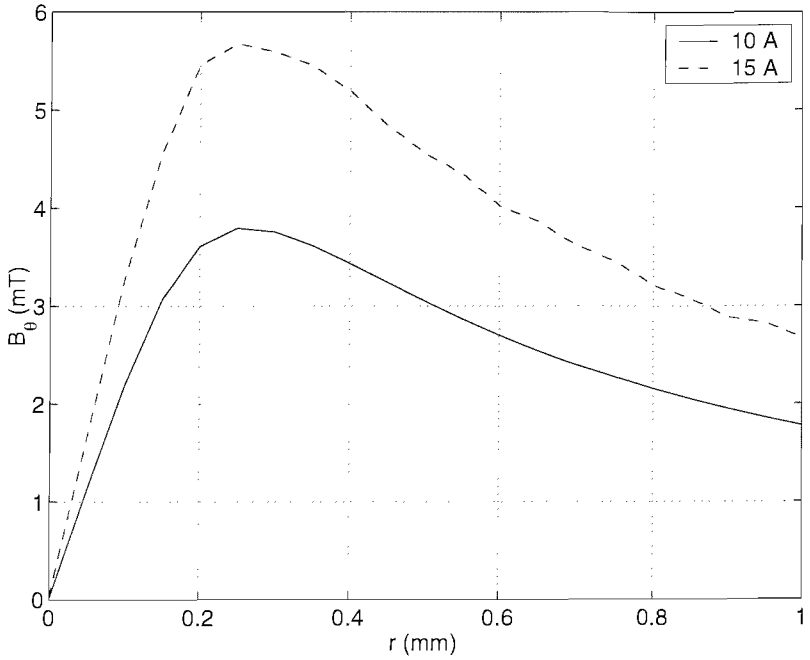

(a) Simulation, 10 A, 15 A

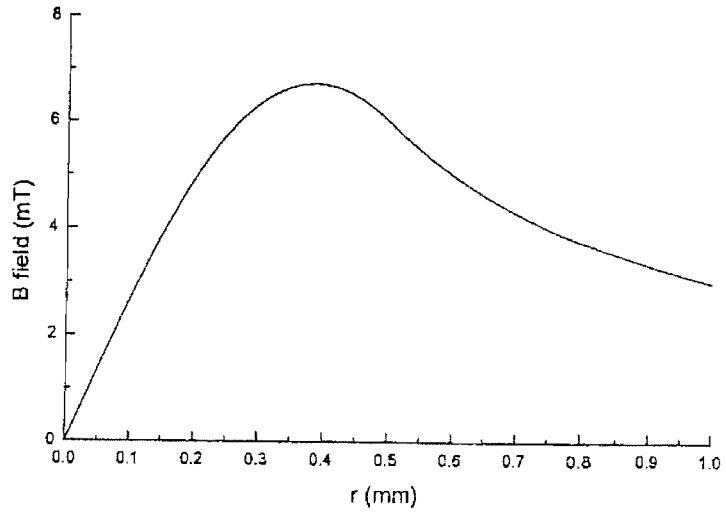

(b) Patterson and Fearn [77], 15 A

Figure 7.4: Radial magnetic field profile in the plume. In the simulation, the radial cut is made at $z = 14$ mm. Plot (b) is from Patterson and Fearn [77]

that even at 5 A the presence of the magnetic field significantly dictates the plume structure. The second reason why this comparison is very important is that it acts as a validation of the code's magnetic field solution. This is in fact an excellent validation case because for the exact magnetic field structure to be reproduced requires that the magnetic field equations are solved correctly, but also that the $\mathbf{v} \times \mathbf{B}$ particle rotation is implemented correctly.

Axially integrated electron temperature

In the recent experiments of Rudwan [88], the electron temperature was measured by examining the plasma along the central axis, thus resulting in an average integration of data. Such a result is very important because it represents non-intrusive experimental data from within the cathode. The only other example of this is Monterde et al. [63], and the current work of Pottinger and Gabriel [80]. Here, the data is compared to values in the simulation. Simulation data is processed so as to come as close to the experimental method as possible; i.e. the effect of the axial integration is taken into account. The comparison would tell us very little if a random point on the axis were to be compared to the data of Rudwan [88]. In fact, since the simulation and experimental data intersect, a judicious choice of sample point would render an exact match, but with little scientific value. Hence, the mean of the simulation data is taken axially. It is expected then that the simulation result will predict a temperature *lower* than that of the experiment for one important reason. Specifically, it appears that the higher temperatures (above 1eV) are found primarily in the plume, and further that the plume temperature seems to vary little as a function of anode position. In the experiments, Rudwan has a nominal cathode-anode separation of 40 mm, a factor of 4 greater than in the simulation. Further, the experiment uses a hollow anode so the plasma could extend as far as a quartz shield included to protect the optical diagnostic equipment, located 55 mm from the cathode. It could be suggested then that the volume of higher energy electrons would be larger in the experiment compared to the simulation, so correspondingly the experimentally measured temperature due to axial integration could be higher also.

Figure 7.5 shows a plot containing both the experimental data and simulation data for electron temperature at a 5 A discharge for a range of mass flow rates (1.0, 1.5 and 3.29 mg s^{-1}). It is important to note also that Rudwan observed plume to spot transition between the 1.5 mg s^{-1} data point and 3.29 mg s^{-1} point. As expected, the simulation result is a temperature slightly lower that the experimental result. Although in the high mass flow rate the simulation data falls within the experimental error bar for temperature. To assess the validity of the suggestion that the hot plume might contribute to this underestimation, the simulation data is adjusted so that the

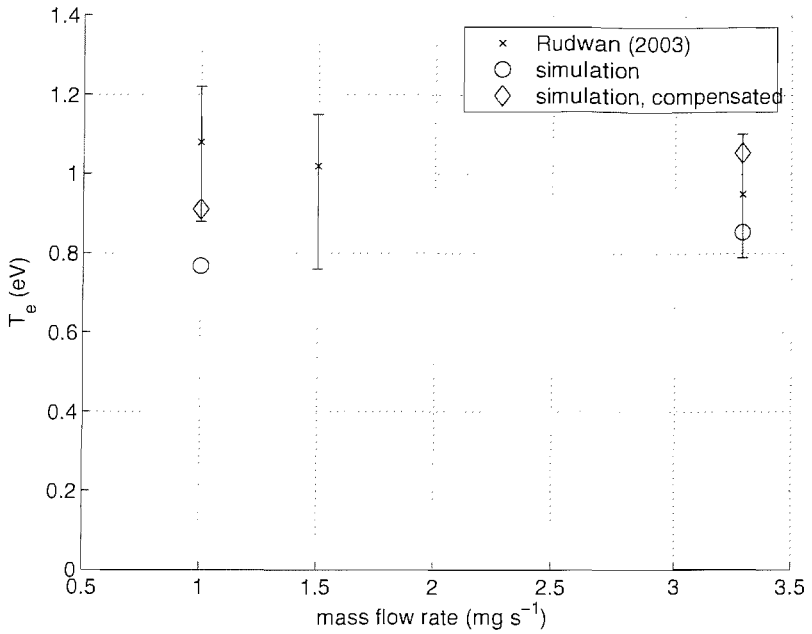

Figure 7.5: Comparison of experimental data of Rudwan [88] to simulation results for integrated axial electron temperature. Simulation data is given as average axial temperature and compensated average axial temperature, where the plume temperature is weighted to account for the experimental configuration.

plume region is temperature is weighted as if the data were continued for the volume present in the experiment. To elaborate: in the experiment, the ratio of external axial plume extent to internal plasma length is around 4, while in the simulation, this ratio is 1. Hence, given N data points in the simulation, the compensated numerical results are produced by multiplying the N_{ext} external data points by 4, adding the internal values and dividing by $(N_{int} + 4N_{ext})$. Although crude, this method should give a result similar to that produced by a spectroscopic examination of the plasma, as per the experiments. These data are also plotted as 'compensated' simulation results. In this case, the data lie within experimental error bars at both mass flow rate levels.

The degree of similarity with experimental results of Rudwan [88] is an excellent indication of the accuracy of the results. This is particularly true because the experimental results used are generally considered the most accurate data on cathode internal plasma properties due to the fact that the acquisition technique is non-intrusive.

7.2.4 Cathode Geometry

Two factors where examined to assess the impact geometry changes have on plasma conditions. These were the tip radius and anode position. Tip radius is a critical factor in all cathode studies; it is typical, for instance, for experimenters to progressively drill out the cathode tip to gain an understanding of the tip behaviour. Hence it was

logical to examine this. In terms of the anode position, recent findings by Rudwan [88] suggest that external experimental configuration can significantly affect the cathode behaviour.

Tip Radius

As expected, it was found that reducing the tip radius leads to a hotter, denser plasma. Plasma density was found to rise by a factor of 4 as the r is reduced from 0.9 to 0.3 mm. In terms of temperature, although neutral temperature remained very similar, the model predicted a very significant rise in ion temperature, by a factor of more than 3 when comparing the extreme cases. If such a trend is correct, the implications are important, as it implies that cathodes with small tips produce very hot ions. It is assumed that much of the damage caused to ion thruster components is due to high energy ions. Hence, the recommendation based on the results of this work would be that small cathode tips should be avoided if possible. Of course, one does not design the cathode geometry purely on the premise that ion temperature be reduced, and there will be performance trade-offs if the tip radius is increased by too large an amount.

Anode position

It is difficult to summarise whether the analysis of anode position produced the results that where expected or not. This is of course because in one sense, the results support current thinking on the subject, while when taken in isolation it is difficult to come to terms with the absolute values in some cases. Rudwan [88] proposes that one can understand the majority of cathode mode transition and behaviour due to the existence and interaction with an anode. Hence, if these hypotheses are to be believed, we would expect a strong dependence between the macroscopic cathode characteristics and anode position. This is certainly what was found, but the variation in voltage with anode position was so large that doubt must lie over the result. Further, as has been noted before, the model is optimised and designed to understand the internal plasma first and foremost, even at the cost of decreased confidence in (low density) external phenomena.

What was further confusing about these results was the fact that the plasma conditions did not alter as violently as the anode voltage. Indeed, the results for both internal and external plasma conditions were as expected, the most obvious trend being for a drop in density corresponding with a more distant anode. It is fair to say that the results presented for anode position had a negative effect on the confidence we can have regarding the accuracy of the code, predominantly in studies of the cathode exterior. It is certainly true that a trade-off was made to capture the internal plasma

behaviour at the expense of the external. Was this necessary? The answer is that within the constraints of this project, yes. However, the problems associated with particle codes and large density variation are not so fundamental that it cannot be imagined that in the future, a method could be devised whereby both regions could be captured correctly in a single model.

7.2.5 Parallelism

The results presented in this document were almost exclusively produced by running the serial version of the code. However, a computationally parallel version was developed and used to produce some of the neutral flow results, including some contained in references [25] and [26]. The issues relating to the code performance in parallel are now discussed. The hardware available for parallel computation was a 'cluster' machine. In parallel computing terms, this is a 'distributed memory' architecture, which means that each compute node has direct access only to local memory, and must communicate with other nodes if access to data located elsewhere is needed. In this case, the primary task when developing parallel software is to reduce both the volume and frequency of communication between nodes.

A simple approach was taken with the neutral gas code. The mesh (see, for instance fig. 5.5) is divided axially between N nodes, so that each node has approximately the same number of cells. Since the particle data is stored and accessed on the DSMC mesh, this means that by domain decomposing the mesh, the particle data is also automatically distributed. When particles pass from one node's mesh area to another, the particle data is sent to the correct node. This scheme proved fairly effective. The volume of communication was always small, as only a few particles would pass between nodes. The frequency of communication was quite high, however, because particle sending occurred on every time-step. This tended to adversely affect performance because time-steps are completed in a very small period of time (tens of milliseconds). Indeed as N grows, each node has only a very small volume to model. In addition, no overlapping of communication with computation was possible because the DSMC cells must be filled correctly for the algorithm to proceed, and a node receiving particles cannot know where those incoming particle need to be. It was concluded that parallelisation of the DSMC via domain decomposition is moderately successful. In fact, the efficiency is much better with larger numbers of particles. The upshot of this is that this method leads to a parallel code that will compute a more complex problem in the same time as the serial code. Unfortunately, on a distributed memory architecture, computing the same problem in $1/N$ the time taken for a serial run to complete is very difficult.

Of course, once the particle simulation and DSMC are parallel, it is trivial to add

all kinds of Monte Carlo based collision routines. Hence, addition of the MCC is very easy. Key numerical parts of the code (PIC, DSMC, MCC) are written in serial and are insulated from the parallelism: particles are accessed on the DSMC grid and the grid is segmented across nodes; each node 'sees' a serial problem. Solution of the Poisson equation efficiently in parallel under the hardware constraints present was found to be impossible. This is primarily because the Poisson problem is so simple and solved so rapidly (because a previous solution that is *almost* correct is always available) that the solution time in serial is of the order of the time to complete a single trivial communication! Hence, for this kind of code, solving this problem size, on this architecture, it was difficult to see how the performance of the full PIC-DSMC-MCC could become efficient.

In summary, the problem with parallelisation of codes such as this one stems from the fact that many millions of very short (in computation) time-steps are completed, with a necessity to communicate on every time-step. It is the opinion of the author then that communication/computation overlapping would yield a parallel neutral DSMC that may achieve good efficiency. However, on todays cluster architectures, PIC codes will suffer from very poor efficiency.

Chapter 8

Conclusions

Three topics are now summarised: first, a general overview of the findings of the research. This is followed by a critical analysis of the novelty of the research itself, the results produced, and the discussion. Finally, a summary of the key points regarding future work is provided.

8.1 General Remarks

A fairly general particle based plasma physics code has been developed to aid understanding of the hollow cathode internal fluid dynamics. The choice of a particle model rather than a continuum one is based on the prediction that a significant fraction of the gas/plasma collisions may be rarefied and result in non-Maxwellian velocity distribution functions. This was shown to be the case and justifies the use of particle models in low mass flow rate or vacuum downstream conditions.

The code was validated in several ways. This included component verification for the random number generator and a test of the iterative solver against a PDE with analytical solution. Standard examples such as a rarefied supersonic Taylor-Couette flow and a flat plate discharge were then used to build confidence in the model. In addition to these component and standard case tests, numerical studies were undertaken to judge the degree of dependence between the data produced and numerical factors such as time-step and mesh spacing. Finally, with real results from a hollow cathode to hand, then model was compared to experimental data and was found to show correct trends in almost all cases and come within experimental error of many readings.

A full characterisation of the cathode neutral flow was presented including analyses of DSMC collision model, variation in mass flux, cathode geometry, propellant and cathode temperature. These data were used in collaboration with the research of Rudwan [88]. It was shown that the predictions regarding pressure and thrust agree

well with experimental data from the hollow cathode.

An attempt was made at full characterisation of the internal plasma during steady-state discharge. The method degree of numerical error was assessed via some numerical studies and although the error was small, it was impossible to fully eradicate such error due to computational constraints. The analysis allowed for an error estimation to be made. In terms of the plasma itself, several features that had been theorised for the internal structure of the plasma where reproduced, in particular the existence of an emitting area region.

In addition, the plume region was reproduced accurately in terms of comparison to experimental data. It was found by study of the plume structure that even at low current (5 A), the plume structure was related strongly to the induced magnetic field. This is an important result as previous PIC-DSMC codes have tended not to implement steady-state induced magnetic solvers (Szabo Jr et al. [101], Crofton and Boyd [28] for example).

The plasma data was compared to the experimental data including that for induced magnetic field and plasma temperature. In both cases, the simulation fitted well with experimental predictions, providing increased confidence in the predictions of the model. Two elements of cathode geometry — variation of tip radius and anode position — were examined using the full plasma code.

8.2 Novelty of the Research

The research can be considered novel due to the following:

- This research represents the first complete and comprehensive analysis of the rarefied neutral gas dynamics within the hollow cathode. The data produced have already been used in more than one case to aid the analysis of experimental data.

- It was demonstrated that the hollow cathode neutral flow was indeed transitional in terms of the degree of rarefaction present. This justifies the use of a particle-based code to capture non-Maxwellian velocity distributions, particularly within the plume.

- Due to the nature of the rarefied gas dynamics and scaling of the cathode geometry, based on the new results presented here, the construction of micro-hollow cathodes for thrust purposes would be discouraged. This is because the structure of the discharge requires a certain gas pressure to contain the emitting region within the cathode: smaller cathodes would need to scale up the mass flux per unit area.

- The results presented here are the first complete 2D model of the cathode internal plasma physics. This is of great importance because although much has been written about cathode interior physics, due to the difficulties associated with experimental investigation, there is almost no data available.

- In many of the investigations, it was found that the position of the anode had a significant effect on the condition of the plume plasma and discharge voltage. Although there are doubts regarding these findings, such a result does confirm the recent suggestion that hollow cathode mode of operation depends primarily on exterior discharge configuration.

- Rather than set a fixed magnet in the solution of the PIC simulation, an experimental steady-state induced magnetic field solver was used. This proved to be very effective in reproducing the plume structure, and critical to obtaining results that compared well to experiment plume data.

- It was predicted that internal ion temperature is strongly related to cathode tip radius. This has very important implications as hot ion bombardment is considered one of the most destructive elements of ion propulsions. Due to these findings it is suggested that greater design pressure be placed on a wider cathode tip since this may have the effect of increasing reliability and lifetime.

8.3 Recommendations for Future Work

With any numerical code, the scope for improvements and additional features is inevitably enormous. Rather than provide an exhaustive list of all possible additions, this section seeks to examine a smaller set of improvements and research areas that would be most beneficial to the role of the model in cathode physics research. The improvements and recommendations presented here broadly represent those that would have been included if time had been available.

Thermal model

It is easily deduced from reading the discussion of heat transfer between cathode and gas/plasma that the fixed temperature wall model is a poor assumption. Although the fixed model does not seem to introduce significant problems, self-consistent determination of cathode temperature would be an important advantage, due to the fact that temperature is a critical engineering parameter with respect to cathode operational lifetime. This means that the first enhancement to be considered would be a thermal model of the cathode (and possible the electrodes). This could take the

form of a simple model whereby each surface has uniform temperature based on heat fluxes to and from the surface all the way to a 2-D axisymmetric solution of the heat flow problem within the cathode body. A problem that would be anticipated is that the characteristic time for thermal dynamics (the thermal "capacitance") is in reality much larger than the plasma breakdown time, say, and indeed the physical time it is possible to model on todays computers. This would mean that an artificially high thermal conductivity would need to be used to allow the thermal and plasma dynamics to come into equilibrium. The relationship and feedback between the plasma conditions and thermal conditions would need very careful considerations as the cathode electron emission scales as $T^2 \exp(-T^{-1})$ (see eq. (2.4)). It is highly likely that with an artificially low thermal time, instabilities could exist due to the coupling between the electron emission and cathode temperature. It is possible, although improbable, that the thermal time required for stable simulation could be prohibitively large in a computational sense.

Hybrid continuum-particle CFD

Authors such as Boyd et al. [14] have recently been involved in research to create hybrid particle-continuum models. Clearly, such an approach would be useful in this case since to the problems encountered with the computational requirements of the particle code at high plasma density would be mitigated by using a continuum model instead. Unfortunately, such methods are currently not well understood and the accuracy and stability of the transition between a particle and continuum model are poor. It seems likely, however, that in the near future the best approach to hollow cathode modelling would be such a hybrid model. Indeed, it is not hard to imagine the addition of a continuum (MHD) model to the code presented here, once the numerics of merging the codes are developed fully.

Other applications

There are two natural extension to the application of the cathode model. First, it is clear that the ion thruster chamber could be modelled using the code in its existing form. Minor extensions that would be required might include a method to efficiently handle the extractor grid, and additional code to implement non-azimuthal externally applied magnetic fields.

Secondly, it would be interesting to further investigate the downstream structure of the discharge, perhaps to try to prove the hypothesis of Rudwan [88] regarding spot-plume transition. As has been noted previously, it was found that the requirements for accurately modelling the inside of the cathode differ to those for the exterior. The emphasis in this work was the fundamental interior physics. It is easy to imagine,

however, generating a suitable boundary condition at the tip exit and modelling the plasma only in the exterior. An additional application here could include examination of the cathode-ion thruster plume coupling plasma.

Species extension

Finally, it is often supposed that some trace species (such as Barium) are of great importance in both cathode breakdown and operation. Deposition of harmful contaminant species is also a major concern in the reliability of ion thrusters. For these reasons the model could be trivially extended to model more species and predict deposition rates and damage rates within the ion thruster. Similar studies have been undertaken in the case of HETs.

Appendix A

The Software

This thesis should be accompanied with a CD containing a copy of the code that produced all of the results found within, a copy of all of the raw data, the processed data and associated scripts used to construct the plots. In addition. a copy of this thesis and associated scripts is included. The struct of the CD is as follows:

- **src** This contains the source code of the numerical model under the subdirectory `PICMCC/main`. In this directory there are auxiliary scripts located under `matlab`, while the code is built using a `GNU Make` makefile. Architecture specific compiler flags are contained in the specific `.mk` files, so that for instance `linux-gcc.mk` contains configuration for a Linux based machine compiling using GCC. Also in the src directory is the cvs repository of the code. This can be used to track and analyse changes made to the model during development.

- **data** Here, both raw and processed data is contained, along with scripts used for generating plots. To save space, only data directly used in the thesis is provided here. As such the arrangement broadly follows the classification by which data is presented in the results section(s).

- **thesis** This directory contains a copy of the LaTeXsource used to generate this thesis, plus the postscript files of the various images and plots.

Appendix B

Field solver Algorithms

Algorithms for the three solver types compared and validated in section 4.4 are given in full here. They are the trivial Jacobi (algorithm 1), and the BiCG (algorithm 2) and BiCGStab (algorithm 3). The absence of the standard CG algorithm from the triplet is due to the fact that it is only applicable to symmetric matricies. Given that only cylindrical axially symmetric descretisations are considered in this work, only asymmetric matricies are used. The most basic useful CG variant is therefore the BiCG.

Algorithm 1 Jacobi

Choose initial $x^{(0)}$
for $k = 1, 2, ...$ **do**
  for $i = 1, 2, ...$ **do**
    $\bar{x}_i = 0$
    for $j = 1, 2, ..., i-1, i+1, ..., n$ **do**
      $\bar{x}_i = \bar{x}_i + a_{i,j} x_j^{(k-1)}$
    end for
    $\bar{x}_i = (b_i - \bar{x}_i)/a_{i,i}$
  end for
  $x^{(k)} = \bar{x}$
  check convergence
end for

Iteratively approaches a solution for x given $Ax = b$. Note that the number of operations completed in one iteration of the Jacobi method is approximately equal to one Matrix-Vector product [5].

Algorithm 2 Bi-Conjugate Gradient

Compute $r^{(0)} = b - Ax^{(0)}$ for some initial $x^{(0)}$
$\tilde{r} = r^{(0)}$ or some other estimate
for $i = 1, 2, ...$ **do**
  solve $Mz^{(i-1)} = r^{(i-1)}$
5: **solve** $M^T \tilde{z}^{(i-1)} = \tilde{r}^{(i-1)}$
  $\rho_{i-1} = z^{(i-1)^T} \tilde{r}^{(i-1)}$
  if $\rho_{i-1} = 0$ **method breaks down**
  if $i = 1$ **then**
    $p^{(i)} = z^{(i-1)}$
10: $\tilde{p}^{(i)} = \tilde{z}^{(i-1)}$
  else
    $\beta_{i-1} = \rho_{i-1}/\rho_{i-2}$
    $p^{(i)} = z^{(i-1)} + \beta_{i-1} p^{(i-1)}$
    $\tilde{p}^{(i)} = \tilde{z}^{(i-1)} + \beta_{i-1} \tilde{p}^{(i-1)}$
15: **end if**
  $q^{(i)} = Ap^{(i)}$
  $\tilde{q}^{(i)} = A^T \tilde{p}^{(i)}$
  $\alpha_i = \rho_{i-1}/\tilde{p}^{(i)^T} q^{(i)}$
  $x^{(i)} = x^{(i-1)} + \alpha_i p^{(i)}$
20: $r^{(i)} = r^{(i-1)} - \alpha_i q^{(i)}$
  $\tilde{r}^{(i)} = \tilde{r}^{(i-1)} - \alpha_i \tilde{q}^{(i)}$
  check convergence
end for

The **solve** statements (lines 4 and 5) represent the precondition phase where M is the preconditioner matrix. This algorithm requires two matrix-vector products (lines 16,17) and two inner products. The is equivalent to around one matrix-vector product plus two inner products work per iteration. The CG methods derive their speed from typically needing an order of magnitude fewer iterations to reach the same residual compared to stationary methods.

Algorithm 3 Bi-Conjugate Gradient Stabilized

Compute $r^{(0)} = b - Ax^{(0)}$ for some initial $x^{(0)}$
$\tilde{r} = r^{(0)}$ or some other estimate
for $i = 1, 2, ...$ **do**
  $\rho_{i-1} = \tilde{r}^T r^{(i-1)}$
5: **if** $\rho_{i-1} = 0$ **method breaks down**
  if $i = 1$ **then**
    $p^{(i)} = r^{(i-1)}$
  else
    $\beta_{i-1} = (\rho_{i-1}/\rho_{i-2})(\alpha_{i-1}/\omega_{i-1})$
10: $p^{(i)} = r^{(i-1)} + \beta_{i-1}(p^{(i-1)} - \omega_{i-1}v^{(i-1)})$
  end if
  solve $M\tilde{p} = p^{(i)}$
  $v^{(i)} = A\tilde{p}$
  $\alpha_i = \rho_{i-1}/\tilde{r}^T v^{(i)}$
15: $s = r^{(i-1)} - \alpha_i v^{(i)}$
  check norm of s; if small set $x^{(i)} = x^{(i-1)} + \alpha_i \tilde{p}$ and stop
  solve $M\tilde{s} = s$
  $t = A\tilde{s}$
  $\omega_i = t^T s/t^T t$
20: $x^{(i)} = x^{(i-1)} + \alpha_i \tilde{p} + \omega_i \tilde{s}$
  $r^{(i)} = s - \omega_i t$
  check convergence
  if $\omega_i = 0$ **method breaks down**
end for

Again, **solve** statements show the preconditioner points. Notice that in addition to the BiCG breakdown on $\rho_{i-1} = 0$ (line 5), there is an additional breakdown possibility at line 23, where $\omega_i = 0$. BiCGStab requires two matrix-vector products (13 and 18) and four inner products (lines 4,14 and 19). This is two more inner products than the BiCG, but it should be noted that the BiCGStab will take fewer iterations than the BiCG method under most circumstances.

Appendix C

Differential Operators in Cylindrical Coordinates

As a reference, vector calculus operators in cylindrical polar co-ordinates are now provided. Most of the references in the text refer to the axially symmetric form of the following, where $\partial\psi/\partial\theta = 0$. See Kreyszig [55] or Goldston and Rutherford [40] for further derivation.

Gradient

$$\nabla\psi = \left(\frac{\partial\psi}{\partial r}, \frac{\partial\psi}{r\partial\theta}, \frac{\partial\psi}{\partial z}\right) \tag{C.1}$$

Divergence

$$\nabla \cdot \mathbf{A} = \frac{\partial(rA_r)}{r\partial r} + \frac{\partial A_\theta}{r\partial\theta} + \frac{\partial A_z}{\partial z} \tag{C.2}$$

Curl

$$\nabla \times \mathbf{A} = \left(\frac{\partial(A_z)}{r\partial\theta} - \frac{\partial A_\theta}{\partial z}, \frac{\partial A_r}{\partial z} - \frac{\partial A_z}{\partial r}, \frac{\partial(rA_\theta)}{r\partial r} - \frac{\partial A_r}{r\partial\theta}\right) \tag{C.3}$$

Laplacian

$$\nabla^2\psi = \frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial\psi}{\partial r}\right) + \frac{1}{r^2}\frac{\partial^2\psi}{\partial\theta^2} + \frac{\partial^2\psi}{\partial z^2} \tag{C.4}$$

Laplacian of a vector

$$\nabla^2\mathbf{A} = \left(\nabla^2 A_r - \frac{2}{r^2}\frac{\partial A_\theta}{\partial\theta} - \frac{A_r}{r^2}, \nabla^2 A_\theta + \frac{2}{r^2}\frac{\partial A_r}{\partial\theta} - \frac{A_\theta}{r^2}, \nabla^2 A_z\right) \tag{C.5}$$

Appendix D

Experimental setup of Gessini *et al.*

Gessini contents, as several other authors do currently, that the hollow cathode is a viable candidate as a stand-alone microthruster. This is due to the presence of anomanlously high energy ions that have been found to be emitted from the cathode under certain circumstances. If this high energy ion emission could be optimised, then the cathode may exhibit high specific impulse and become a viable thruster design. In order to record the thrust produced by the cathode so that the high energy ion emission may be characterised in terms of specific impulse, Gessini devised a thrust measurement system. This consists of a cantilever beam target (CBT) and a Laser optical Level (LOL). The plume of the cathode strikes the target and deflects the beam. A laser is directed at the target and the reflected beam deflection is measured. Since the structural properties of the beam are deliberately designed to be well defined, it is possible to relate a the beam deflection to the force required to bend the beam. This force is simply the force exerted due to the plume impinging upon the target.

Bibliography

[1] Ramesh K. Agarwal, Keon Young Yun, and Ramesh Balakrishnan. Beyond Navier-Stokes: Burnett equations for flows in the continuum-transition regime. *Physics of Fluids*, 13(10):3061–3085, October 2001.

[2] John David Anderson. *Modern compressible flow: with historical perspective.* McGraw-Hill, 2nd edition, 1990.

[3] K. Asano, N. Ohno, M. Y. Ye, S. Fukuta, and S. Takamura. 2-d PIC simulation on space-charge limited emission current from plasma-facing components. *Contrib. Plasma Physics*, 40(3–4):478–483, 2000.

[4] C. Barbeau and J. Jolly. Spectroscopic investigation of energetic atoms in a DC hydrogen glow discharge. *J. Physics D: Applied Physics*, 23(9):1168–1174, 1990.

[5] R. Barrett, M. Berry, T.F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijhout, R. Pozo, C. Romine, and H. Van der Vorst. *Templates for the solution of linear systems: Building blocks for iterative methods.* SIAM, Philadelphia, PA, 1994.

[6] G. A. Bird. Approach to translational equilibrium in a rigid sphere gas. *Physics of Fluids*, 6:1518–1519, May 1963.

[7] G. A. Bird. Monte Carlo simulation of gas flows. *Annual Review of Fluid Mechanics*, 10:11–30, 1978.

[8] G. A. Bird. Monte Carlo simulation in an engineering context. *Progress Astronautics Aeronautics*, 74:239–255, 1981.

[9] G. A. Bird. *Molecular Gas Dynamics and the Direct Simulation of Gas Flows.* Number 42 in Oxford Engineering Science Series. Oxford University Press, 1994.

[10] C. K. Birdsall. Invited paper: Particle-In-Cell charged-particle simulations, plus Monte Carlo collisions with neutral atoms, PIC-CIC. *IEEE Trans. on Plasma Science*, 19(2):65–85, April 1991.

[11] C. K. Birdsall and D. Fuss. Clouds-in-clouds, clouds-in cells physics for many-body simulation. *J. Computational Physics*, 3:494–511, April 1969.

[12] C. K. Birdsall and A. Bruce Langdon. *Plasma Physics via Computer Simulation.* Institute of Physics Publishing, 1991.

[13] J.P. Boris. Relativistic plasma simulation - optimisation of a hybrid code. *Proc. 4th Conf. Num. Sim. Plasmas*, pages 3–67, NOV 1970.

[14] Iain D. Boyd, J. Vanden Beukel, and Y. Jafry. Particle simulations of Helium microthruster flows. *J. Spacecraft and Rockets*, 31(2):271–277, 1994.

[15] Iain D. Boyd, M. Keidar, and W. McKeon. Modeling of a pulsed plasma thruster from plasma generation to plume far field. *J. Spacecraft and Rockets*, 37(3):399–407, May 2000.

[16] George R. Brewer. *Ion Propulsion Technology and Applications.* Gordon and Breach, Science Publishers Ltd., 1970.

[17] G. Carter. *Ion bombardment of solids.* Heinemann Educational, 1968.

[18] M. Celik, M. Santi, S. Cheng, M. Martínez-Sánchez, and J. Peraire. Hybrid-PIC simulation of a Hall thruster plume on an unstructured grid with DSMC collisions. In *Proceedings of the 28th International Electric Propulsion Conference, Toulouse, France*, March 2003.

[19] P. R. Chalise, Watanabe M., Okino A., Ko K., and E. Hotta. Characteristics of an ion induced secondary emission electron gun. *Plasma sources science and technology*, 12:235–243, 2003.

[20] Edward H. Chao, Stephen F. Paul, Ronald C. Davidson, and Kevin S. Fine. Direct numerical solution of Poisson's equation in cylindrical (r,z) coordinates. Technical report, Princeton Plasma Physics Laboratory, 1997.

[21] S. Chapman and T. G. Cowling. *The mathematical theory of non-uniform gases.* Cambridge University Press, 2nd edition, 1952.

[22] J.D. Cobine. *Gaseous conductors.* Dover, 2 edition, 1958.

[23] Sylvain Coulombe and Jean-Luc Meunier. Thermo-field emission: a comparative study. *J. Physics D: Applied Physics*, 30:776–780, 1997.

[24] Francis T. A. Crawford and S. B. Gabriel. Modeling small hollow cathode discharges for ion microthrusters. *AIAA paper 2002–2101*, May 2002.

[25] Francis T. A. Crawford and S. B. Gabriel. Microfluidic model of a micro hollow cathode for small ion thrusters (invited). *AIAA paper 2003–3580*, June 2003.

[26] Francis T. A. Crawford and S. B. Gabriel. Numerical simulation of the hollow cathode plasma using a PIC-DSMC code. In *Proceedings of the 28th International Electric Propulsion Conference, Toulouse, France*, March 2003.

[27] Mark W. Crofton and Iain D. Boyd. Plume measurements and modeling results for a Xenon hollow cathode. *AIAA Paper 2002-4103*, July 2002.

[28] Mark W. Crofton and Iain D. Boyd. The T6 hollow cathode: Measurements and modelling. *AIAA Paper 2003-4171*, June 2003.

[29] G. Csiky. Investigation of hollow cathode discharge plasma. *AIAA paper 69-258*, 1969.

[30] M.T. Domonkos, A.D. Gallimore, and G.J. Jr. Williams. Low-current hollow cathode evaluation. *AIAA Paper 99–2575*, 1999.

[31] C.H. Edwards. *Discharge characteristics and instabilities in the UK-25 ion thruster operating on inert gas propellants.* PhD thesis. University of Southampton, August 1997.

[32] D.G. Fearn. Electric propulsion of spacecraft. *J. British Interplanetary Society*, 35:156–166, 1982.

[33] D.G. Fearn and S. W. Patterson. Characterisation of the high current hollow cathode for the T6 ion thruster. *AIAA paper 98–3346*, July 1998.

[34] D.G. Fearn and C. M. Philip. An investigation of physical processes in a hollow cathode discharge. *AIAA Journal*, 11(2):131–132, 1973.

[35] J. H. Ferziger and M. Perić. *Computational Methods for Fluid Dynamics.* Springer-Verlag, 2nd edition, 2002.

[36] Nikos A. Gatsonis and Xuemin Yin. Hybrid (particle-fluid) modeling of pulsed plasma thruster plumes. *J. Propulsion and Power*, 17(5):945–958, October 2001.

[37] P. Gessini, S.B. Gabriel, and D.G. Fearn. The hollow cathode as a micro-ion thruster. *Proceedings of the 27th International Electric Propulsion Conference, Pasadena, California*, October 2001.

[38] P. Gessini, S.B. Gabriel, and D.G. Fearn. Hollow cathode thrust measurements using a target: initial results and some issues. *Proceedings of the 28th International Electric Propulsion Conference, Toulouse, France*, March 2003.

[39] Sergey Gimelshein, Gennady Markelov, and Marc Rieffel. Collision models in the `hawk` DSMC implementation. Technical report, California Institute of Technology, Computer Science department, July 1996.

[40] R.J. Goldston and P.H. Rutherford. *Introduction to Plasma Physics.* Institute of Physics Publishing, 1995.

[41] W.A. Harrison. *Elementary electronic structure.* World Scientific, 1999.

[42] John K. Harvey and Michael A. Gallis. Review of code validation studies in high-speed low-density flows. *J. Spacecraft and Rockets*, 37(1):8–20, January 2000.

[43] Philip G. Hill and Carl R. Peterson. *Mechanics and Thermodynamics of Propulsion.* Addison-Wesley, 1970.

[44] E.H. Hirsch. Abnormal electron temperatures and electron reflection in the cylindrical thermionic diode. *J. Physics D: Applied Physics*, 35:2766–2771, October 2002.

[45] R. W. Hockney and J. W. Eastwood. *Computer Simulation Using Particles.* Institute of Physics Publishing, 1988.

[46] K.G. Hutchinson and S.B. Gabriel. Conceptual design of a micro ion thruster. *Proceedings of the 27th International Electric Propulsion Conference, Pasadena, California*, October 2001.

[47] Siegried W. Janson. IEPC-97–070: Batch-fabricated resistojets: Initial results. In *Proceedings of the 25th International Electric Propulsion Conference, Cleveland, Ohio*, pages 427–434, August 1997.

[48] Manish Jugroot and J.K. Harvey. DSMC neutral and charged particle simulation of a Kaufmann-type ion thruster. In *Proceedings of the 10th International Congress on Plasma Physics, Joint 42nd meeting of the APS Divison conference, Quebec, Canada*, 2000.

[49] Manish Jugroot and J.K. Harvey. Particle simulation of the main chamber of a Kaufmann-type ion thruster. *Proceedings of the 3rd International Conference on Spacecraft Propulsion, ESA SP–465, Cannes, France*, 2000.

[50] Ira Katz, John R. Anderson, James E. Polk, and John R. Brophy. a model of hollow cathode plasma chemistry. In *Proceedings of the 38th Joint Propulsion Conference*, July 2002.

[51] Ira Katz, John R. Anderson, James E. Polk, and Dan M. Geobel. model of hollow cathode operation and life limiting mechanisms. In *Proceedings of the 28th International Electric Propulsion Conference, Toulouse, France*, March 2003.

[52] J. W. Koo, I. D. Boyd, J. M. Haas, and A. D. Gallimore. Computation of interior and near-field flow of a 2 kW class Hall thruster. *AIAA Paper 2001–3321*, 2001.

[53] K. Koura and H. Matsumoto. Variable soft sphere molecular model for inverse-power-law or Lennard-Jones potential. *Physics of Fluids A*, 3:2459–2465, 1991.

[54] Horst Kraemer. Personal communication, May 2003.

[55] Erwin Kreyszig. *Advanced Engineering Mathematics*. Wiley, 7th edition, 1993.

[56] A. K. Malik and D.G. Fearn. The study of the physics of hollow cathode discharges. In *Proceedings of the 23rd International Electric Propulsion Conference, Seattle, WA*, 1993.

[57] M. J. Mandell and I. Katz. Theory of hollow cathode operation in spot and plume modes. *AIAA paper 94–3134*, 1994.

[58] Manuel Martinez-Sanchez and J. E. Pollard. Spacecraft electric propulsion — an overview. *J. Propulsion and Power*, 14(5):688–699, September 1998.

[59] M. Matsumoto and T. Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudorandom number generator. *ACM Trans. on Modeling and Computer Simulation*, 8(1):3–30, January 1998.

[60] J.S. Miller, S.H. Pullins, D.J. Levandier, Y. Chiu, and R.A. Dressler. Xenon charge exchange cross sections for electrostatic thruster models. *J. Applied Physics*, 91(3):984–991, February 2002.

[61] M. Mitchner and C. H. Jr. Kruger. *Partially Ionized Gases*. Wiley, 1973.

[62] A. Mondinos. Theory of thermionic emission. *Surface Science*, 115:469–500, 1982.

[63] M.P. Monterde, M.G. Haines, A.E. Dangor, A.K. Malik, and D.G. Fearn. Kaufman-type xenon ion thruster coupling plasma: Langmuir probe measurements. *J. Physics D: Applied Physics*, 30:842–855, 1997.

[64] E.L. Murhpy and R.H. Jr. Good. Thermionic emission, field emission and the transition region. *Physical Review*, 102(6):1464–1473, June 1956.

[65] D.M. Murray, O.R. Tutty, and S.B. Gabriel. Numerical modelling of ion thruster hollow cathode interior flow. *AIAA paper 97–0793*, January 1997.

[66] D.M. Murray, O.R. Tutty, and S.B. Gabriel. Numerical modelling of the gas flow in an ion thruster hollow cathode. In *Proceedings of the 2nd European Spacecraft Propulsion Conference*, pages 217–233, August 1997.

[67] Y. Nakamura and M. Kurachi. Electron transport parameters in Argon and its momemtum transfer cross section. *J. Physics D: Applied Physics*, 21:718–723, 1988.

[68] K. Nanbu. Probability theory of electron-molecule, ion-molecule, molecule-molecule, and Couloub collisions for particle modeling of materials processing plasmas and gases. *IEEE Trans. on Plasma Science*, 28(3):971–990, June 2000.

[69] K. Nanbu. Theory of cumulative small-angle collisions in plasmas. *Physical Review E*, 55(4):4642–4652, April 1997.

[70] K. Nanbu and Y. Kitatani. An ion-neutral species collision model for particle simulation of glow discharge. *J. Physics D: Applied Physics*, 28:324–330, October 1995.

[71] K. Nanbu and S. Kondo. Analysis of the three-dimensional DC magnetron discharge by the particle-in-cell / Monte Carlo method. *Japanese J. Applied Physics*, 36:4808–4814, 1997.

[72] Kenichi Nanbu, Tamotsu Morimoto, and Masaharu Suetani. Direct simulation Monte Carlo analysis of flows and etch rate in an inductively coupled plasma reactor. *IEEE Trans. on Plasma Science*, 27(5):1379–1388, October 1999.

[73] Choong K. Oh, Elaine S. Oran, and Bohdan Z. Cybyk. Microchannel flow computed with the DSMC-MLG. In *AIAA paper 95–2090*, June 1995.

[74] Y. Okawa and H. Takegahara. Particle simulation of ion beam extraction phenomena in an ion thruster. In *Proceedings of the 26th International Electric Propulsion Conference, Kitakyushu, Japan*, pages 805–812. JSASS, October 1999.

[75] Elaine S. Oran, Bohdan Z. Cybyk, and Choong K. Oh. Direct Ssimulation Monte Carlo: Recent advances and applications. *Annual Review of Fluid Mechanics*, 30:403–442, 1998.

[76] C. A. Ordonez and R. E. Peterkin Jr. Secondary electron emission at anode, cathode, and floating plasma-facing surfaces. *J. Applied Physics*, 79(5):2270–2274, March 1996.

[77] S.W. Patterson and D.G. Fearn. The generation of high energy ions in hollow cathode discharges. In *Proceedings of the 26th International Electric Propulsion Conference, Kitakyushu, Japan*, pages 695–702, October 1999.

[78] A. V. Phelps. Electron collision cross section data, `ftp://jila.colorado.edu/collision_data/eletrans.txt`. Online, 2003.

[79] C. M. Philip. A study of hollow cathode discharge characteristics. *AIAA Journal*, 9(11):2191–2196, 1971.

[80] S. Pottinger and S.B. Gabriel. Optical measurements on the internal plasma of a T6 type hollow cathode. *Proceedings of the 28th International Electric Propulsion Conference, Toulouse, France*, March 2003.

[81] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipies in C: The Art of Scientific Computing*. Cambridge University Press, 2nd edition, 1992.

[82] V. Puech and S. Mizzi. Collision cross sections and transport parameters in Neon and Xenon. *J. Physics D: Applied Physics*, 24:1974–1985, 1991.

[83] Donald Rapp and W.E. Francis. Charge exchange between gaseous ions and atoms. *J. Chemical Physics*, 37(11):2631–2645, December 1962.

[84] Zoran M. Raspopović, Sava Sakadžíc, Svetlan A. Bzenić, and Zoran Lj. Petrović. Benchmark calculations for Monte Carlo simulations of electron transport. *IEEE Trans. on Plasma Science*, 27(5):1241–1248, October 1999.

[85] D. Reichelmann and K. Nanbu. Monte Carlo direct simulation of the Taylor instability in a rarefied gas. *Physics of Fluids A*, 5:2585–2587, 1993.

[86] Paola Rossetti, Fabrizio Paganucci, and Mariano Andrenucci. A hollow cathode model for application to the electric propulsion. *AIAA Paper 2002-4239*, 2002.

[87] I.F.M.A. Rudwan. Personal communication, 2004.

[88] I.F.M.A. Rudwan. *Physics of hollow cathode breakdown and steady-state operation with several inert gas propellants*. PhD thesis, University of Southampton, 2003.

[89] Wilhemus M. Ruyten. Density-conserving shape factors for particle simulations in cylindrical and spherical coordinates. *J. Computational Physics*, 105:224–232, 1993.

[90] Gilberto M. Sandonato, Joaquim J. Barroso, and Antonio Montes. Magnetic confinement studies for performance enhancement of a 5cm ion thruster. *IEEE Trans. on Plasma Science*, 24(6):1319–1329, December 1996.

[91] Gilberto M. Sandonato, P.E. Lima, H.S. Maciel, and C. Otani. The influences of magnetic field strength on the magnetic confinement of primary electrons in an ion source. *Contributions to Plasma Physics*, 39(3):187–195, 1999.

[92] J. Schou. Transport theory for kinetic emission of secondary electrons from solids. *Physical Review B*, 22(5):2141–2174, September 1980.

[93] H. Shelton. Thermionic emission from a planar Tantalum crystal. *Physical Review*, 107(6):1553–1557, 1957.

[94] D. E. Siegfried and P. J. Wilbur. An investigation of Mercury hollow cathode phenomena. *AIAA paper 78-705*, 1978.

[95] D. E. Siegfried and P. J. Wilbur. Phenomnological model describing orificed, hollow cathode operation. *AIAA Journal*, 21(1):5–6, January 1983.

[96] D. E. Siegfried and P. J. Wilbur. A model for Mercury orificed hollow cathodes: theory and experiment. *AIAA Journal*, 22(10):1405–1412, October 1984.

[97] James S. Sovey, Vincent K. Rawlin, and Michael J. Patterson. Ion propulsion development projects in the U.S.: Space electric rocket test I to Deep Space 1. *J. Propulsion and Power*, 17(3):517–526, May 2001.

[98] S. Stefanov and C. Cercignani. Monte Carlo simulation of the Taylor-Couette flow of a rarefied gas. *J. Fluid Mechanics*, 256:199–213, 1993.

[99] Kay Sullivan, Manuel Martínez-Sánchez, Oleg Batichev, and James Szabo. Progress on particle simulation of ceramic-wall Hall thrusters. In *Proceedings of the 28th International Electric Propulsion Conference, Toulouse, France*, March 2003.

[100] J. J. Szabo. *Fully Kinetic Numerical Modeling of a Plasma Thruster*. PhD thesis, Massachusetts Institute of Technology, February 2001.

[101] J.J. Szabo Jr, Manuel Martinez-Sanchez, and O. Batishchev. Particle-In-Cell modeling of thruster with anode layer (TAL). In *Proceedings of the 26th International Electric Propulsion Conference, Kitakyushu, Japan*, pages 581–588. JSASS, October 1999.

[102] S. Takamura, M. Y. Ye, T. Kuwabara, and N. Ohno. Heat flow through plasma sheaths. *Physics of Plasmas*, 5(5):2151–2158, May 1998.

[103] V.G. Terpigor'yev. The energy spectrum of thermal electrons emitted from an oxide cathode. *Radio Eng. Electron. Phys. (translation of Radiotekhnika i Elektronika (USSR))*, 18(2):319–321, February 1973.

[104] W. Wagner. A convergence proof for Bird's direct simulation Monte Carlo method for the Boltzmann equation. *Akademie der Wissenschäften der DDR, P-Math-223-90*, July 1990.

[105] J. Wang, D. Brinza, and M. Young. Three-dimensional particle simulations of ion propulsion plasma environment for Deep Space 1. *J. Spacecraft and Rockets*, pages 433–440, May 2001.

[106] Wells and Harrison. Hollow cathode spot to plume transition. Unpulished, 1980s.

[107] K. Weltner, J. Grosjean, P. Schuster, and W.J. Weber. *Mathematics for Engineers and Scientists: translation of Mathematik für Physiker (1961)*. Stanley Thornes Ltd, 1986.

[108] Richard Wirz, Jesse Escobedo, Patrick Sheehan, James Polk, Coleen Marrese, and Juergen Mueller. Development and testing of a 3cm electron bombardment micro-ion thruster. *Proceedings of the 27th International Electric Propulsion Conference, Pasadena, California*, October 2001.

[109] J. Woods and D.A Wright. Secondary electron emission from barium oxide. *Br. J. Appl. Phys.*, 3:323–326, April 1952.

[110] G.J. Yashko, G.B. Giffin, and D.E. Hastings. IEPC-97–072: Design considerations for ion microthrusters. In *Proceedings of the 25th International Electric Propulsion Conference, Cleveland, Ohio*, pages 443–449, August 1997.