

Two-Dimensional Simulation of Radio-Frequency Glow Discharges

by

Gilbert Lee Huppert

Bachelor of Chemical Engineering
University of Delaware (1987)

M. S. Chemical Engineering Practice
Massachusetts Institute of Technology (1989)

Submitted to the Department of Chemical Engineering in Partial Fulfillment of the Requirements
for the Degree of

Doctor of Science in Chemical Engineering

at the

Massachusetts Institute of Technology
February 1996

© Massachusetts Institute of Technology

Signature of Author _____

MIT

Certified by _____

Advisor

Professor of Chemical Engineering
Engineering and Computer Science, MIT

Certified by _____

Robert A. Brown, Thesis Advisor
Professor of Chemical Engineering, MIT

Accepted by _____

Robert E. Cohen, Chairman Departmental Committee on Graduate Students
Professor of Chemical Engineering, MIT

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY
MAR 22 1996

Science

LIBRARIES

Two-Dimensional Simulation of Radio-Frequency Glow Discharges

by Gilbert Lee Huppert

Submitted to the department of chemical engineering on October 30, 1995 in partial fulfillment of the requirements for the degree of Doctor of Science in Chemical Engineering

Abstract

Continuum, hydrodynamic modeling of the physics of radio-frequency glow-discharges at high pressures was investigated. One-dimensional simulations were performed to develop rigorous and efficient numerical methods appropriate to the extremely stiff and time dependent problem being examined. Simulations were capable of predicting the critical one-dimensional experimental phenomena. The numerical techniques developed in the one-dimensional simulations allowed efficient implementation of two-dimensional models. The two-dimensional simulations allow further comparison with experimental data collected in our laboratory, with reasonable semi-quantitative agreement. Two-dimensional simulations are shown to converge spectrally with refinement of the spatial grid, except at the junction of the insulator and the electrodes where a discontinuity in the ion boundary condition is enforced. The physics of the two-dimensional discharges is discussed at length, with particular attention paid to the role of the aspect ratio of the discharge on perturbing the solutions from a simple one-dimensional case. For both symmetric (equal area for powered and grounded electrode) and asymmetric (smaller grounded electrode) simulations, the perturbations extend a distance of one gap-spacing into the discharge radially. Simulations which have aspect ratios much greater than one are well represented by the one-dimensional solutions near the radial centerline. For the asymmetric geometries under consideration, the plasma is mostly confined within a region defined by the smaller electrode, and the larger electrode is not fully electrically active.

Thesis Supervisors: Professor Herbert H. Sawin and Professor Robert A. Brown

Acknowledgments

My sincere thanks are extended to all the members of the Materials Etching Technology Group (Professor Sawin's group at MIT) for their enlightening discussions on plasma physics, and their support during the long hours spent debugging the various versions of my simulation code. A special thank you is required to Reza Mehrabi, Howard Covert, and Todd Salamon of Professor Brown's group for their discussions on numerical methods, simulation techniques, and debugging methodology.

I thank Professors Herbert H. Sawin and Robert A. Brown for their discussions on plasma physics and modeling strategies, and for supporting me during my stay at MIT. I am grateful to the National Science Foundation for providing me a fellowship and to the Semiconductor Research Corporation for providing a Research Assistantship to me. This project would not have been possible without computational resources provided by grants at the Pittsburgh NFS supercomputer facility. Michael Mocella at DuPont graciously provided access to the DuPont supercomputer facility at the end of the project when other computer facilities were over-extended and time constraints were becoming important.

Finally, I thank my parents, my brother, and my good friend Amy Bernstein for providing an enormous amount of emotional support during the most difficult stages of my research. They were always there when I needed someone to kvetch to.

Table Of Contents

Abstract	2
Acknowledgments	3
Table Of Contents	4
Table Of Figures	6
List Of Tables	9
Chapter 1 : Introduction	10
1.1: Motivation	10
1.2: Previous Work	11
1.2.1 One-Dimensional Modeling	11
1.2.2: Two-Dimensional Modeling.....	15
Chapter 2 : The Continuum Model	18
2.1: A Description.....	18
2.2: The Weak Form	21
2.3: Numerical Methods For Solution	23
2.3.1: Spectral Element Discretization	23
2.3.2: Time Integration	26
2.3.3: Time Periodic Solution.....	28
2.3.4: Computational Cost	32
Chapter 3 : One-Dimensional Modeling Results	35
3.1: Introduction.....	35
3.2: Low Pressure Helium Plasmas	37
3.3: Numerical Convergence Results: Spectral Element Method.....	47
3.3.1: Temporal Accuracy.....	47

3.3.2: Spatial Accuracy	48
3.3.3: Comparisons Between Spectral Element Results and Finite Difference Calculations of Gogolides et al.	52
3.3.4: Conclusions.....	58
3.4: Comparison to One Dimensional Experimental Results: Argon Plasmas	59
3.5: Electronegative Gases: SF ₆	67
Chapter 4 : Two-Dimensional Modeling Results.....	77
4.1: Description of Two-Dimensional Geometry.....	77
4.2: Convergence Results in Two Dimensions	80
4.3: Extent of Two-Dimensional Perturbations	84
4.4: Symmetric Plasmas	88
4.4.1: Explanation of Plasma Physics	89
4.4.2: Experimental setup for comparison to simulation results.....	98
4.4.3: Comparison to Experimental Measurements	102
4.5: Asymmetric Plasmas.....	115
4.5.1: Explanation of Plasma Physics	117
4.5.2: Comparison to Experimental Measurements	130
Chapter 5 : Argon Plasmas In Complex Geometries	137
5.1: Symmetric Plasmas with Focusing Rings	138
5.2: Symmetric Plasmas with Bars.....	142
Chapter 6 : Conclusions and Future Work	147
6.1: Conclusions.....	147
6.2: Suggestions for Future Work	149
References	151

Table Of Figures

Figure 2-1 Time integration to a TPSS versus Newton shooting solution.	28
Figure 3-1 Helium plasma density for electrons and positive ions.....	39
Figure 3-2 Helium electron energy shown at 4 phases in the RF cycle.....	40
Figure 3-3 Power deposition in a Helium discharge at 300 mTorr.....	42
Figure 3-4 Current components for a Helium discharge	43
Figure 3-5 Spatially resolved currents in a Helium discharge at 300 mTorr	44
Figure 3-6 Time averaged a) ionization and b) emission rates for a Helium discharge while varying pressure.	46
Figure 3-7 Temporal accuracy of one dimensional solutions	48
Figure 3-8 Convergence of the spectral element method with both h-type and p-type mesh refinement.	50
Figure 3-9 Exponential convergence of the spectral element method with p-type mesh refinement for the peak energy per electron, which is the most sensitive simulation variable.....	51
Figure 3-10 Plasma simulations for a 2 cm argon discharge operating at 1 Torr and 13.56 MHz.54	
Figure 3-11 Time evolution of simulated electric field	55
Figure 3-12 Time evolution of simulated electron energy profile	55
Figure 3-13 Spatial inaccuracy of the finite-difference approach when upwind-differencing is included.....	56
Figure 3-14 Spatial inaccuracy of the finite-difference approach when upwind-differencing is included shown in the maximum electron energy in 1/2 the discharge.	57
Figure 3-15 The global physics for finite-difference and spectral-element techniques compares favorably for less sensitive variables.	58
Figure 3-16 Comparison of experimental and continuum model current-voltage characteristic for Argon Discharges at various pressures..	60
Figure 3-17 Comparison of experimental and continuum model Power characteristic for Argon Discharges at various pressures.	61
Figure 3-18 Model density predictions for an Argon discharge at various pressures.....	62
Figure 3-19 Comparison of the phase shift in an Argon discharge between current and voltage for experimental measurements and continuum model predictions at various pressures.	63

Figure 3-20 Comparison of experimental and continuum model Ion Flux for Argon Discharges at various pressures.....	64
Figure 3-21 Comparison of experimental and continuum model Ion Energy for Argon Discharges at various pressures.	65
Figure 3-22 Comparison of experimental and continuum model plasma emission for	66
Figure 3-23 Species concentrations in an SF ₆ discharge.....	69
Figure 3-24 Electron energy at various phases in the cycle for an SF ₆ discharge.....	71
Figure 3-25 Potential variation in an SF ₆ discharge.....	72
Figure 3-26 Components of the current in an SF ₆ discharge	73
Figure 3-27 Current through an SF ₆ discharge.....	75
Figure 3-28 Negative ion density for an SF ₆ discharge near the grounded electrode	76
Figure 4-1 Geometry for presentation of simulation and experimental results a) vector and contour plots b) surface plots.	77
Figure 4-2 Simulated regions for a) symmetric and b) asymmetric geometry.....	78
Figure 4-3 Positive ion density contours near the corner of the discharge where the insulating surface meets the grounded electrode.	82
Figure 4-4 Convergence of the Two-Dimensional Simulation with p-Type Refinement.....	83
Figure 4-5 Convergence of the electron energy with mesh refinement at different spatial positions within the plasma.....	84
Figure 4-6 Positive ion density for two-dimensional simulations.	85
Figure 4-7 Effect of plasma aspect ratio on continuum model results.....	86
Figure 4-8 Changes in the structure of the radial sheath with aspect ratio..	87
Figure 4-9 Effect of gap spacing on two-dimensional simulations.....	88
Figure 4-10 Current components and voltage for two-dimensional Ar plasma simulation at 1 Torr.....	90
Figure 4-11 Current density in a two-dimensional discharge at different radial sections.....	91
Figure 4-12 Total current in an Argon plasma at 2 phases in the RF cycle.	92
Figure 4-13 Electric fields for a two-dimensional Ar plasma.....	93
Figure 4-14 Total current for an Ar plasma with 2" gap.....	94
Figure 4-15 Electron density for symmetric Ar plasma..	95

Figure 4-16 Time averaged ionization rate for a two-dimensional Ar plasma simulation.	96
Figure 4-18 Ion Current to the powered electrode as a function of pressure.	98
Figure 4-19 Electrode geometry for a) symmetric and b) asymmetric reactor configuration.	99
Figure 4-20 Simulated and experimental plasma induced emission at high pressures	103
Figure 4-21 Simulated and experimental plasma induced emission at low pressures	104
Figure 4-22 Plasma induced emission for a wide gap 1 Torr plasma.	106
Figure 4-23 Langmuir probe characteristics for Argon Plasmas.	108
Figure 4-24 Electron Current characteristic for Langmuir probe showing fit to electron energy as a Maxwellian.....	109
Figure 4-25 Plasma potential in the symmetric Ar plasma at high pressures	110
Figure 4-26 Plasma potential in the symmetric Ar plasma at low pressures	111
Figure 4-27 Electron energy in the symmetric Ar plasma at high pressures.....	112
Figure 4-28 Electron Energy in the symmetric Ar plasma at low pressures..	113
Figure 4-29 Plasma density in the symmetric Ar plasma at high pressures.....	114
Figure 4-30 Plasma density in the symmetric Ar plasma at low pressures	115
Figure 4-31 Positive ion density contours for an asymmetric Argon plasma with an area ratio of 2.25.....	118
Figure 4-32 Positive ion density at the grounded electrode for the same conditions as Figure 4-31.....	119
Figure 4-33 Time averaged electric field for 1" gap, 1 Torr, asymmetric Argon plasma.....	120
Figure 4-34 Current for asymmetric 1" gap discharge.....	122
Figure 4-35 Current at the powered (larger) electrode for the asymmetric Argon Discharge at 1 Torr.....	122
Figure 4-36 Total current density on the powered electrode for an asymmetric Argon discharge at 1 Torr.....	123
Figure 4-37 Positive ion density contours for a highly asymmetric Argon discharge	124
Figure 4-38 Time averaged quantities for a highly asymmetric Argon discharge.	125
Figure 4-39 Simulated Plasma Emission for a highly asymmetric Argon plasma.....	125
Figure 4-40 Effect of plasma asymmetry on the positive ion	126
Figure 4-41 Positive ion density contours for a 1/2" gap, highly asymmetric argon discharge.	128

Figure 4-42 Positive ion density at the radial centerline of the discharge for various sized smaller electrodes. 128

Figure 4-43 Ion current to the grounded (smaller) electrode at various pressures and aspect ratios..... 129

Figure 4-44 Positive ion density for 1 Torr, 2” gap asymmetric Argon discharge. 130

Figure 4-45 Comparison of simulated plasma induced emission for asymmetric discharges ... 132

Figure 4-46 Plasma induced emission for wide gap, asymmetric discharges. 133

Figure 4-47 Plasma potential for asymmetric plasmas. 134

Figure 4-48 Time averaged electron energy for asymmetric discharges..... 135

Figure 4-49 Time averaged positive ion density for asymmetric discharges 136

Figure 5-1 Geometry for perturbation studies with a) ring b) angled bar on lower electrode... 137

Figure 5-2 Plasma emission for 1 Torr Ar plasmas with focusing rings on the electrode. 139

Figure 5-3 Plasma emission for 100 mTorr Ar plasmas with focusing rings on the electrode. 140

Figure 5-4 Electron energy as determined by Langmuir probe for the center of a perturbed Ar discharge. 141

Figure 5-5 Plasma perturbations by Al bars placed on grounded electrode at 1 Torr..... 143

Figure 5-6 Plasma perturbations by Teflon bars placed on grounded electrode at 1 Torr. 144

Figure 5-7 Plasma perturbations by Al bars placed on grounded electrode at 100 mTorr..... 145

Figure 5-8 Plasma perturbations by Teflon bars placed on grounded electrode at 100 mTorr... 146

List Of Tables

TABLE 3-1 Transport and reaction parameters for the continuum model37

TABLE 4-1 Computational requirements for convergence calculations.81

Chapter 1 : Introduction

1.1: Motivation

This work employs efficient and rigorous numerical methods for solution of the continuum model for plasma physics. Simulations are performed in one and two dimensional geometries, and compared to experimental measurements carried out in the lab at MIT. The main contribution of this work is a set of numerically rigorous solutions for the continuum model and suitable experimental measurements for comparison.

Plasma etching has become the most widely used pattern transfer mechanism as the critical feature length for micro-electronic devices continues to decrease and greater fidelity of pattern transfer has become necessary, as discussed by Graves (1989). For etching of large, shallow features it was not necessary to understand the plasma physics or chemistry to design etching reactors. However, for today's narrow and deep circuit components, and with the trend of using larger substrate areas to increase process throughput, a detailed understanding of the plasma physics which couples into the etching process is greatly desirable. The plasma physics, plasma chemistry, and surface evolution (etching) are intimately coupled. However, each process has a very different time scale, so it is most useful to understand each process separately.

One-dimensional modeling is very useful for understanding the basic physics of the plasma, and examining the differences between electronegative (electron attaching) and electropositive gases. The results of the one-dimensional model can be used as input for the simplest chemistry and surface evolution models, and are useful for preliminary design efforts. Two-dimensional models are necessary for more detailed understanding of the plasma physics in a real system, especially when the effect of radial non-uniformity of the plasma is to be incorporated into the surface evolution model. It is expected that the two-dimensional modeling efforts can contribute to the primary design process for new equipment, as well as suggesting modification to existing equipment for increased performance. By examining simpler cases and examining the two-dimensional physics in detail, two-dimensional modeling results will also be useful for determining the range of parameters to examine when testing out new equipment which cannot be directly modeled.

To apply the model results to equipment design, the model must be robust, computationally efficient, accurate, and contain sufficient physics. Without all of these features, results from simulations are of little value. If the simulation is not robust, too much effort will be wasted on achieving a solution, and generation of large sets of solutions would be too cumbersome for equipment design. If the model is not computationally efficient, the cost of the computations as well as the necessary computation time will exceed requirements for useful design efforts and plasma physics exploration. Accuracy is required to assure that the numerical methods are not perturbing the physics which the model encompasses. Finally, if the model does not include sufficient physics, then application to real systems is of no use, since the model is not, either quantitatively or qualitatively, correctly predicting the physics in the real system.

This work examines the geometry effects of parallel plate diode reactors for high pressure etching processes, such as the etching of silicon-dioxide. A specific suite of numerical methods were selected which are best suited to solving this problem, and the simulations were optimized for the choice of geometry, physics, and plasma model used. All of the numerical methods which are useful for simulations of diode reactors can also be used for simulation of next-generation high-density sources which use RF biasing on the wafer-chuck. The simulations also give the electron energy as a function of spatial position, which will be useful to couple into chemistry models used in plasma deposition simulations. Therefore, the present work can be considered to be not only a set of tools for computation of the physics of simple discharges in parallel plate reactors, but is also useful preliminary work for solving other problems which are just now becoming of interest.

1.2: Previous Work

1.2.1 One-Dimensional Modeling

Although the importance of plasma processing in microelectronics fabrication is well established, the physics and chemistry of plasma-assisted etching is just beginning to be understood. Many different tools are available for exploring the plasma physics, including Monte Carlo (MC) simulations, particle in cell (PIC), continuum models (also called fluid models), and Boltzman calculations. Boltzman simulations are the most rigorous, but are also extremely computationally intensive, and have generally been solved using very limiting

assumptions. Monte Carlo simulations are the simplest conceptually, but require simulation of large numbers of particles for convergent statistics. The PIC simulations are a modification of the MC models in which each particle represents an ensemble of particles, and each ensemble is allowed to move within a spatial grid. Continuum models treat each species as having a given distribution function (which is calculated in the Boltzman codes), and each species is considered concentrated enough that there is an equilibrium for that species at each spatial location. The connection between the continuum model and the Boltzman equations is discussed in detail by Gogolides and Sawin (1992). The PIC and MC codes are best applied at lower pressures where the gas is more rarefied and fewer particles have to be followed, while the continuum model is applicable at higher pressures where the equilibrium assumptions are not violated. All of these methods have been used to simulate DC discharges in one dimension. Since the RF simulations have a time dependent forcing function, the calculations are much more difficult and have long lagged behind the DC modeling efforts.

Continuum modeling shows promise for quantitative prediction of plasma processing in a computationally efficient framework. The first application of the continuum model was for prediction of the breakdown characteristics of SF₆ and N₂ discharges by Thompson and Sawin (1986), where it was shown that a simple continuum model is capable of predicting the breakdown characteristics quite well. Graves and Jensen (1986) used a more complete model to examine the discharge physics for DC and mixed DC/RF systems. This was the first attempt to model the complete plasma physics. The authors used a finite element spatial discretization and a Fourier series to handle the time-periodic forcing of the plasma. Graves and Jensen conclude that this approach is too computationally expensive to be widely used. There is also the added difficulty that the number of Fourier components must be chosen in advance, and it is difficult to estimate where to terminate the series.

Other simulation efforts have been based on finite difference methods and frequently have introduced artificial diffusion in the spatial approximations to help stabilize the numerical solution through upwind differencing or staggered grid approaches (Richards et al., 1987, Barnes et al., 1988; Gogolides and Sawin, 1992; Gogolides et al., 1992). Although these methods are numerically rigorous, the spatial approximations are low-order accurate and, as a result, fine meshes are needed to accurately resolve steep layers; otherwise, the approximations are overly

diffusive and predict solutions with gradients of the solution that are overly smooth. The only method for avoiding the application of artificial diffusion is to accurately resolve the solution structure and the true diffusive contribution to the conservation equations. This approach is extremely difficult with low-order accurate finite difference and finite element methods, because very fine meshes are needed to resolve the solution structure and the diffusive contribution. Therefore, the accuracy of these simulations is basically unknown, and results from simulations which use diffusive approximations are suspect. In fact, for the original work of Barnes et al., since a poor choice was used for the temporal convergence criterion, the solutions shown are not converged in time or space, and there is no error estimate available to indicate this. Boeuf (1987) used a Scharfetter-Gummel scheme for the spatio-temporal discretization, which links the upwind-differencing to the time and spatial discretizations. Therefore, this work also has an error of unknown magnitude in the solutions, although the discretization does guarantee that this error is minimal for a choice of grids.

Park and Economou (1990) simulated both electropositive and electronegative discharges using cubic-spline collocation on finite elements, and a strict convergence criterion for the temporal convergence. As with most modeling efforts no attempt was made to show that the spatial discretization is sufficient. There is adequate evidence in their publication that the plasma sheath electron energy is not properly refined, and this was ignored by the authors. Instead they attribute what might well be numerical error to the physics of the plasma in the sheath region. The physics for this version of the continuum model are also questionable, since the effective field approximation was not used, and the forms of the rate and transport parameters were inadequate.

Okazaki et al. (1989) refined the basis for the continuum model by adding in a separate relaxation time for each process in the energy and momentum balances. It is unclear whether this refinement made any real difference in the model results, although the authors stated that it was a necessary part of the plasma physics. However, there is little information on the numerical techniques that were used, and this work certainly did not examine convergence in any detail.

Paranjpe et al. (1990) was the first work to address the computational cost which comes with integrating the initial profiles until there is no longer a change over a cycle. The technique used to accelerate convergence was to examine the plasma on both long time scales, with time-

averaged equations, and on short time scales with the full equation set. This work presents very useful time averaged and reduced models for the plasma physics, and claims excellent computational speed. However, the solution technique is reported to have some problems converging to a solution, and indeed is not guaranteed to converge under any conditions. Gogolides et al. (1992) was the first work to address this issue by using a Newton-Raphson time shooting scheme to directly calculate the time-periodic state, although staggered-mesh finite elements were used which greatly limit the accuracy of the spatial discretization. This time shooting method is one of the basic numerical tools used in the present work.

Upwind and staggered mesh finite-difference methods inherently pollute the spatial discretization, leading to solutions which may be unacceptably perturbed. Spectral element methods, as introduced by Maday and Patera (1989), offer an attractive alternative for this approach. Here the spatial domain (a line segment for a one-dimensional problem) is divided into a few subdomains, or elements. Polynomial representations for the field variables are constructed in each element so that the fields are continuous across inter-element boundaries, as is usual for Lagrangian finite element approximations. Galerkin's method is used to discretize the differential equations and boundary conditions. This formulation of spectral elements therefore greatly resembles finite elements with the capability of using high-order basis functions. The difference comes when higher-order accuracy is required. Instead of increasing the number of elements, as is the approach in h-convergent finite element discretizations, the polynomial order of the approximation is increased with a given element discretization. The family of spectral element methods described by this procedure is part of the general class of p-convergent spectral techniques (Gottlieb and Orszag, 1977).

The major advantage of spectral element methods is rapid convergence to the exact solution for smooth problems. Analysis (Maday and Patera, 1989) shows that for linear elliptic and parabolic problems without singularities, the spectral element method converges exponentially to the solution with an increase in the polynomial order, compared to the algebraic convergence rate with decreasing mesh size that is expected for conventional finite element discretizations using fixed polynomial approximations.

1.2.2: Two-Dimensional Modeling

With the advent of faster computers, improvement in the numerical techniques, and a better understanding of the plasma physics which was being simulated, one-dimensional models became readily solvable. Even so, one-dimensional simulations remained expensive without an acceleration technique to find a time-periodic steady-state. The generalization to two-dimensions is trivial, except for changes in the boundary conditions. However, a two-dimensional simulation is equivalent to steady-state calculations of three-dimensional reaction-diffusion problems due to the harmonic forcing of the potential on one electrode. In fact, the two-dimensional modeling of the DC discharge is difficult enough that it was not attempted by Boeuf (1988) until after he had solved the one-dimensional RF problem. Even so, this first two-dimensional DC effort was not totally self-consistent since it did not include an energy balance, and treated the boundary conditions in a rather ad-hoc manner. Certainly, the difficulty of solving the simpler two-dimensional DC problem initially discouraged solution of the two-dimensional RF problem. This work addresses this issue by using rigorous techniques to achieve accurate solutions with reasonable computational resources.

Initial two-dimensional modeling efforts were forced to use a simple geometry with very restrictive assumptions on the boundary conditions and plasma physics to be simulated. All of the simulations reported to date are too computationally intensive to use for design work, and are in the category of preliminary work to demonstrate feasibility of achieving a solution. Tsai and Wu (1990) simulated electronegative discharges on very coarse meshes using a one-moment fluid model which was not self-consistent. The only effect examined was the supposed focusing of the plasma by the negative ions in the radial sheath, but the insulating wall was treated inconsistently, so any conclusions from the model are highly suspect. This work contributed to the field by showing that a continuum model is solvable in two-dimensions with very limiting assumptions. However, the physics in Tsai and Wu's model was inadequate since the local field approximation was used, which had previously been shown to be questionable at best. The main advantage of ignoring the bulk of the plasma physics for this study was that the resultant computation time was only a few XMP-hours. The spatial discretization used was a complex flux-corrected scheme which adds an unknown smoothing into the solutions, but also allows computation on coarser grids.

Young and Wu (1993,1993a) refined the model to include the full three moments of the Boltzman equation and therefore removed the limiting assumption of the local-field equation. The geometry simulated still incorrectly treats the insulating boundary as not perturbing the electric fields, and assumes that there is some type of field containing the plasma. This is incorrect, but is a reasonable first approach to the simulations in order to examine the physics of electropositive discharges. This study used a staggered grid approach to the spatial discretization which introduced an unknown amount of artificial diffusivity, but did allow computation on coarser grids. Even so, computation time on the order of hundreds of XMP hours were required for convergence since no time-acceleration scheme was used. The interesting feature of this simulation was that the charge densities were maximum along the axial center of the discharge, and had a global maximum near the insulating surface. The peak in the density is up to 20% higher than at the discharge center, and moves toward the center of the discharge decreasing in magnitude with decreasing pressure. The presence of the maximum in density from such a simple model is surprising at first, but is easily accounted for by the formation of the sheath on the insulating wall which creates a high field region and increased ionization.

Dalvie et al. (1993) noted a similar effect in simulating Argon discharges in two dimensions. They used a more realistic boundary condition on the insulating electrode, treating it as an ideal capacitor, but still did not account self-consistently for the charging of the insulator, and therefore would be slightly wrong at the corners. This work was done using a non-perturbed spatial discretization which should completely preserve the physics that are being simulated. An ad-hoc acceleration scheme was used which is supposed to save up to an order of magnitude of computation time. A maximum in ionization near the insulating wall was seen, which moved off the axial center for non-symmetric geometries (personal communication). The maximum in density was much less severe than reported by Young and Wu, and it is not clear if it exists for the case of an infinitely thick insulator, although this may be a function of the physics used to simulate the insulator.

The previous simulations were all done in close to symmetric geometries, but most commercial reactors are asymmetric due to a large grounded surface. Passchier and Goedheer (1993) examined a reactor with a more typical configuration with a large grounded area separated from the powered electrode by a thin insulator. This work used a Scharfetter-Gummel spatial

discretization on a very coarse mesh. No radial focusing effect was noted, although there was a maximum in ionization near the corners of the reactor. There was a large axial asymmetry noted in the profiles, which is interesting although of very little importance for plasma processing. More importantly, the ion flux increased toward the edge of the powered electrode before dropping precipitously across the insulating boundary. However, this simulation is still very computationally intensive (estimated 30 YMP hours), and the presence of the thin insulator between the powered and grounded electrode effectively introduces a singularity into the simulation. The electric fields must be singular near that junction, and a large amount of the observed physics may be due to numerical artifacts from the singularity. Since the Scharfetter-Gummel scheme effectively smoothes all of the profiles by upwind differencing, it is not clear how great an effect the singularity has on the results, or on what length scale the inaccuracies will propagate.

Chapter 2 : The Continuum Model

2.1: A Description

For plasma modeling, there is generally a trade-off between model accuracy and the computational requirements for solving the model equations. The simplest models treat the plasma as linear or non-linear circuit elements, and since this model requires next to no computational effort, the accuracy for calculating plasma physics is quite limited, and the information content in the model is very small. There also exist many simple analytic models which treat the plasma, usually only the bulk or sheath, in a very simple manner, and again the tradeoff is limited information content for reduced computational complexity. The continuum model and Monte Carlo simulation are the simplest methods which can predict the majority of the physics of interest for plasma processing. Monte Carlo techniques, generally implemented as Particle in Cell (PIC) codes, are efficient at lower pressures, and give more information than a continuum model would, but become very computationally intensive at higher pressures. The continuum model is most appropriately used at higher pressures, where ions have multiple collisions in traversing the plasma. At lower pressures, below 200 mTorr, there are not sufficient collisions in the plasma for the species to truly behave as a continuum fluid. The continuum model may be applied at lower pressures with adjusted transport parameters (which can be calculated in a Monte-Carlo simulation, resulting in a type of hybrid continuum/Monte Carlo code). Another problem exists in the plasma sheaths, where the species are rarefied and suffer few collisions while receiving large amounts of energy. This implies that the continuum model is not capable of correctly predicting the physics on the length scales of the sheath, but the detailed physics of the collisionality of the sheath seems to play very little part in the overall plasma trends, so using the continuum model is not too bad an approximation.

Different formulations of the continuum model will contain varying amounts of detail for the plasma physics. Each continuum equation can be viewed as a moment of the Boltzman equation, with the moments containing information on density, momentum, and energy for the

first three moments. This work uses a modified approach, where the momentum balances are simplified to yield directly the species fluxes, and the energy balance for ions is neglected as being unimportant. However, even the full continuum model does not yield the full information of a Boltzman calculation, which will yield the entire distribution functions for each species, but is extremely computationally intensive.

The continuum model equations used in this work are identical to those described by Gogolides and Sawin (1992) and Gogolides et al. (1992), and the derivation of the continuum equations is therefore not repeated here. For most plasmas of interest it is sufficient to examine only 3 species balances, electron density (N_e), positive ion density (N_+) and negative ion density (N_-), and the energy balance for the average energy per electron (ϵ). Most of the simulations in the present work are for electropositive gases, where the concentration of negative ions is insignificant, so the negative ion equation is not used. The continuum equations are then written in divergence form as:

$$\begin{aligned} \frac{d N_e}{d t} &= -\nabla \cdot \Gamma_e + R_{ion} - R_{attach} & \frac{d N_e \epsilon}{d t} &= -\nabla \cdot \Gamma_\epsilon - q \mathbf{E} \cdot \Gamma_e + R_{loss} \\ \frac{d N_+}{d t} &= -\nabla \cdot \Gamma_+ + R_{ion} - R_{recomb} & \frac{d N_-}{d t} &= -\nabla \cdot \Gamma_- + R_{attach} - R_{recomb} \end{aligned} \quad \text{Eq. 1}$$

where Γ_i denotes the species fluxes, q is the electronic charge and R_{ion} and R_{loss} are energy dependent reaction rates. From Eq 1 it is obvious that electrons are created by ionization, and can only be lost due to flux to the boundary, or through attachment to form negative ions (which is not present in electropositive plasmas). Positive ions are formed in the same manner as electrons, and are lost either through recombination with negative ions or through flux to the boundary. Note that there is no mechanism for bulk recombination of electrons and positive ions, since the cross-section for this reaction is small. Negative ions are created only when an electron is attached to a neutral, and are lost through recombination with positive ions or flux to a boundary. The energy balance is interesting because it shows that energy can be gained or lost through coupling to the electric field, or through flux to the boundaries, and is also lost through collisions.

The momentum balances are used in a simplified form, split into two components which describe the motion of each species due to diffusion and convection due to the electric field. The electrons have very small mass, and therefore respond directly to the instantaneous potential, but

ions are massive and cannot respond instantaneously. Note, that at the higher frequencies used in microwave discharges, neglecting the mass of electrons is no longer correct. The drift-diffusion equivalents of the momentum balances are then written as:

$$\begin{aligned} \Gamma_e &= -\frac{D_e}{\varepsilon} \nabla \varepsilon N_e - \mu_e N_e \mathbf{E} & \Gamma_\varepsilon &= \frac{5}{3} \varepsilon \Gamma_e - D_e N_e \nabla \varepsilon \\ \Gamma_+ &= -D_+ \nabla N_+ + \mu_+ N_+ \mathbf{E}^{\text{eff}} & \Gamma_- &= -D_- \nabla N_- - \mu_- N_- \mathbf{E}^{\text{eff}} \end{aligned} \quad \text{Eq. 2}$$

where D_i is the species diffusivity, with the electron diffusivity a function of the electron energy, and the μ_i are the species mobilities which account for the convection of electrons and ions by the applied potential including collisional effects. The electric field is calculated self-consistently, and the effective field is calculated as a function of the electric field as:

$$\nabla^2 V = \frac{q}{\varepsilon_0} (N_e + N_- - N_+) \quad \frac{\partial \mathbf{E}^{\text{eff}}}{\partial t} = \frac{q}{\mu_+ M_+} (\mathbf{E} - \mathbf{E}^{\text{eff}}) \quad \text{Eq. 3}$$

where ε_0 is the permittivity of free space and M_+ is the mass of an ion. For lower pressures and at higher frequencies the full momentum balances should be included for accurate simulation. For two-dimensional simulations, there will be two components for the effective field corresponding to the two components of the electric field (radial and axial), resulting in one more equation which needs to be solved.

The equation set (1) - (3) are formally parabolic, but the presence of the drift contributions caused by the fields leads to substantial convective character, measured by Peclet numbers that scale the magnitude of the drift to the diffusive contribution; as discussed in detail by Gogolides et al. (1992). The advantage of this formulation of the continuum model is that the statement as a convective-diffusive set of equations allows for solution using a Galerkin type method. Formulation of the continuum model in full three-moment form results in a formally hyperbolic equation set. Even though the present formulation is formally parabolic, the convective part of the equations is frequently dominant, making the equations seem hyperbolic on coarse discretizations. It is this hyperbolic character that compelled other workers to use upwind finite differencing for the discretization of the continuum equations.

The boundary conditions used to solve equations (1) - (3) are:

$$\begin{aligned} \Gamma_e &= \frac{N_e}{4} \sqrt{\frac{2\varepsilon}{M_e}} + \gamma_{\text{sec}} \Gamma_+ \cdot \bar{n} & \Gamma_\varepsilon &= \left[\frac{5}{3} \varepsilon \Gamma_e + \gamma_{\text{sec}} \varepsilon_{\text{sec}} \Gamma_+ \right] \cdot \bar{n} \\ \nabla N_+ \cdot \bar{n} &= 0 & N_- &= 0 \end{aligned} \quad \text{Eq. 4}$$

and are applied on all boundaries. The boundary condition for electrons and electron energy state that the flux to the surface is the equivalent to an ideal gas (isotropic) reaching the same surface at a given temperature. The positive ion boundary condition is used to remove a thin boundary layer which occurs for the correct boundary condition of $N_+ = 0$ (fast recombination on surfaces), which is not important. The plasma is driven by forcing the voltage to vary sinusoidally on one electrode, while holding the other electrode grounded. The two-dimensional asymmetric simulations are run with a capacitor between the electrode and the voltage source to allow computation of the self-bias potential. There is also an insulating portion which is simulated using Poisson's equation and requiring zero enclosed charge. The insulator introduces extra equations at the boundary to calculate the discontinuity in the electric field due to the charge buildup on the electrode:

$$\varepsilon_1 E_1 - \varepsilon_2 E_2 = C \quad \frac{dC}{dt} = q(\Gamma_e - \Gamma_+) \quad \text{Eq. 5}$$

Gogolides et al. (1992) forced the current to vary and enforced continuity of the current at one node in the finite difference grid by introducing a modification of Poisson's equation. This is correct and easily applied in a one-dimensional discharge as:

$$\frac{dE_1}{dt} = \frac{I_0}{\varepsilon_0} \sin(\omega t) + \frac{q}{\varepsilon_0} (\Gamma_e - \Gamma_+) \quad \text{Eq. 6}$$

where I_0 is the applied current to simulate the imposition of a forcing current. There is, however, not a proper two-dimensional analog since the current does not need to be constant within the discharge. Some computations were performed with the current boundary condition in order to complete a direct comparison to the finite-difference modeling.

2.2: The Weak Form

Since the spectral element technique is a Galerkin method, the weak equations must be formed. In general one has for a given field variable $U(x,t)$ and a flux $\Gamma_u(U)$:

$$\frac{\partial U}{\partial t} = \nabla \cdot \Gamma_u + \text{Lower - Order - Terms (LOT)}. \quad \text{Eq. 7}$$

The only terms which are required to be considered weakly are the second order spatial operators; the entire flux need not have a weak representation. The effective field has no second order operator, and is not strictly a PDE, but can rather be considered an ODE. Therefore, the weak form of the effective field equation will be just a representation of the effective field on the

spectral element grid, with the proper stiffness summation performed to account for the fact that the electric field need not be continuous at element boundaries.

The initial approach used to create the weak form was tailored specifically to allow use of a semi-implicit integrator with the diffusive portions split out from the convective portions, following the work of Maday et al. (1990). In order for this splitting to be computationally efficient, it is desirable that at least part of diffusive operator for each equation should be time-independent. This contribution is then used to stabilize the integrator, removing the Courant condition. This motivated representing only the diffusive portion of the positive ion flux weakly, and dividing the electron energy and electron density by the electron diffusivity (to maintain a constant portion in the electron weak forms). For example, for electron density, the result is:

$$\frac{1}{D_e} \frac{dN_e}{dt} = -\frac{1}{D_e} \nabla \cdot \left[\frac{D_e}{\epsilon} \nabla \epsilon N_e + \mu_e N_e \mathbf{E} \right] + \frac{1}{D_e} (R_{ion} - R_{attach}) \quad \text{Eq. 8}$$

where only the first term on the right hand side of Eq. 8 will be integrated by parts.

The weak form of each conservation equation is discretized by approximating the field variables using expansions of Gauss-Lobatto-Legendre polynomials of variable order.

The weak form of the general case in Eq. 7, or the electron density balance in Eq. 8, is derived by multiplying by a test function $W(x)$ that satisfies essential boundary conditions on $U(x,t)$, and then integrating-by-parts the contribution from the species flux. These operations yield

$$\frac{\partial}{\partial t} \int_A w u dA = - \int_A \Gamma_u \cdot \nabla w dA + \int_S \bar{n} \cdot \Gamma_u \nu dA + \int w LOT dA \quad \text{Eq. 9}$$

and

$$\int_A \frac{W}{D_e} \frac{dN_e}{dt} = \int_A \nabla \left(\frac{W}{D_e} \right) \cdot \left[\frac{D_e}{\epsilon} \nabla \epsilon N_e + \mu_e N_e \mathbf{E} \right] + \int_S \left[\frac{W}{\epsilon} \nabla \epsilon N_e + \frac{\mu_e W}{D_e} N_e \mathbf{E} \right] + \int_A \frac{W}{D_e} (R_{ion} - R_{attach}) \quad \text{Eq. 10}$$

Simplification of the first term in Eq. 10 yields two parts for the weak form, the first of which is equivalent to a constant, time-independent, parabolic term, and the second is nonlinear and time dependent, and the grouping under the second integral is equivalent to LOT's:

$$\int_A \nabla W \cdot \left[\nabla N_e + \left\{ \frac{N_e}{\varepsilon} \nabla \varepsilon + \frac{\mu_e}{D_e} N_e \mathbf{E} \right\} \right] - \int_A W \frac{\nabla D_e}{D_e} \cdot \left[\nabla N_e + \frac{N_e}{\varepsilon} \nabla \varepsilon + \frac{\mu_e}{D_e} N_e \mathbf{E} \right] \quad \text{Eq. 11}$$

After the gaussian quadrature is performed, the electron equations are again multiplied by D_e , yielding a form with a single term which looks like $(\nabla W \bullet D_e \nabla N_e)_{ij}$. This is the final version of the weak form used for electron density, even with integrators which don't use any type of semi-implicit method. For simulations which did not have constant mobility the full flux was integrated-by-parts for ease of calculation. It was observed that in constant mobility cases, changing the way that the convective flux was handled for positive ions did not affect simulations.

2.3: Numerical Methods For Solution

2.3.1: Spectral Element Discretization

For this work, it was desired that all the numerical method used be rigorous and accurate enough that it would be possible to show that the solution of the continuum model, as it was formulated, was not perturbed by the choice of discretization, time integrator, or solution technique to the time-periodic steady-state (TPSS). It was therefore imperative that the spatial discretization should be accurate, ruling out any type of upwind-differencing or Petrov-Galerkin method. The spectral element method of Patera (1989) was chosen as a flexible method with arbitrarily high spatial accuracy. The spectral element method is equivalent to a classical finite-element method, where the basis functions have been chosen so that the collocation points (including on the boundaries) and the node points coincide. The spectral element basis functions yields a diagonal mass matrix, and the choice of an arbitrary order for the basis function is facilitated. An advantage of this method is that convergence can be demonstrated by relatively small changes in the total number of degrees of freedom by raising the degree of the interpolating polynomial. As discussed by Maday and Patera (1989) for linear elliptic equations, the error in the solution, as evaluated in the L_2 -norm defined on the domain, converges exponentially, i.e.

$$\|\mathbf{u}(x, t) - \mathbf{U}_{exact}(x, t)\|_o = C e^{-N_p} \quad \text{Eq. 12}$$

where C is a constant that depends on the derivatives of the exact solution $U_{\text{exact}}(x,t)$. This exponential convergence with increasing N_p should be compared to the algebraic convergence of typical finite element discretizations, i.e. $\|U - U_{\text{exact}}(x)\|_2 = O(h^p)$, where p is the order of the polynomials used in the approximation for $U(x)$.

Given that a Galerkin weak form has already been selected for the continuum model equation set, the spectral element discretization can be formed. The solution field $U(x,t)$ is represented in a spectral element approximation (Maday and Patera, 1991) in one dimension as

$$U(x) = \sum_{i=1}^{N_{\text{elem}}} \sum_{j=1}^{N_p} U_{ij} \Phi_{ij}(x) \quad \text{Eq. 13}$$

where the summations are taken over all elements in the subdivision of the domain $0 \leq x \leq 1$, and over all nodes within each element. For the two dimensional representation, a tensor product formulation is used based on the one dimensional basis functions, which is exactly equivalent to the statement above, but now there are N_p^2 local basis functions. For the contributions to the nodal equations, this results in a direct stiffness sum of the components local to each element, with communication between adjacent elements only active at the element boundaries. The two dimensional forms are tensor products of the one dimensional basis functions, following the work of Maday and Patera (1991). As in finite element methods, the coefficients $\{U_{ij}(t)\}$ are determined from the discretized equation set and the basis functions are specified to interpolate the solution in the subdomain. In the spectral element approximation used here the basis functions $\{F_{ij}(x)\}$ are taken as the one-dimensional, Gauss-Lobatto-Legendre polynomials of degree (N_p-1) ; see Maday and Patera (1991) for details. The spectral element approximation defined by Eq. 13 is continuous across element boundaries, but has only piecewise continuous first derivatives at these intersections, as is typical of Lagrangian finite element methods. This degree of continuity is appropriate for solving the weak form of the equation 8 and leads to compact support in the discrete formulation, albeit the number of degrees-of-freedom that are coupled increases with increasing polynomial degree.

The specification of the discrete equations 9 is completed by using the $\{F_{ij}(x)\}$ as the test functions, i.e. $W(x) = F_{ij}(x)$, $i = 1, \dots, N_{\text{elem}}$ and $J = 1, \dots, N_p$, so that Eq. 9 is equivalent to a $N_{\text{elem}} \times (N_p - 1) + 1$ set of ordinary differential equations. The test functions are calculated using a set of subroutines provided by Einar Ronquist at MIT.

The Gauss-Lobatto-Legendre polynomials have several unique properties that make them particularly well suited for the solution of parabolic initial-value problems if the integrals in Eq. (9) are evaluated at the quadrature points defined by the polynomials; that is at the Gauss-Lobatto-Legendre points. First, the time derivative on the left side of Eq. 9 gives rise to the typical mass matrix contribution for finite elements (Strang and Fix, 1973) with terms

$$\mathbf{M}_{ij} \frac{d}{dt} U_{ij} \equiv \left[\int_A \Phi_{ij} \Phi_{km} dA \right] \frac{d}{dt} U_{ij} \quad \text{Eq. 14}$$

Using Gauss-Lobatto-Legendre quadrature these terms have the particularly simple form $\mathbf{M}_{ij} = c_0 \delta_{ij}$, where δ_{ij} is the Kronecker Delta function and c_0 accounts for the quadrature weights and element sizes. In other words, the mass matrix is diagonal, as discussed previously. Also for the contribution of a constant diffusive flux, i.e. $\Gamma_u = D \nabla U$, the first integral on the left side of Eq. (9) becomes

$$\int_A D \nabla \Phi_{ij} \cdot \nabla \Phi_{km} dA U_{ij} \quad \text{Eq. 15}$$

which is the stabilizing contribution to the equation set; see Strang and Fix (1973) for a discussion of this point. These integrals are evaluated exactly by Gauss-Lobatto-Legendre quadrature using the quadrature rule for $\{2(N_p - 1) - 1\}$ th degree polynomials.

The equations discretized by the spectral element formulation for the plasma model can be written compactly as a set of Differential-Algebraic Equations (DAE's) of the form

$$\mathbf{M} \frac{\partial \mathbf{u}}{\partial t} = \mathbf{R}(\mathbf{u}) \quad \text{Eq. 16}$$

where \mathbf{M} is the mass matrix, with either the constant weights c_0 or zero on the diagonal and zeroes elsewhere, $\mathbf{R}(\mathbf{u})$ is a nonlinear vector function and \mathbf{u} is the vector of coefficients in the spectral element expansions. The variables \mathbf{u} are assumed to be of dimension N , so that \mathbf{M} is an $N \times N$ matrix. The zeroes on the diagonal of \mathbf{M} arise from discretization of Poisson's equation for the electric potential and lead to algebraic, not ordinary differential equations, in the Eq. 16. With the completion of Eq. 16, the spectral element spatial discretization is complete, and this set of equations will be referenced as the starting set for plasma simulation in the next subsection. Fine boundary and internal layers exist in the solution fields, leading to a problem notoriously difficult to solve (Finlayson, 1980).

2.3.2: Time Integration

The DAE's have index one and are solved using an Adam's Moulton integrator that is second-order accurate in the time step Δt , as described by Brenan et al. (1989). The nonlinear algebraic equations that arise at each time step are solved by Newton's method with the Jacobian matrix computed using one-sided finite difference approximations for efficiency. The algorithm is programmed with the capability to use full Newton iterations, but this is generally inefficient since it saves only at most 2 iterations but requires inversion of the Jacobian of Eq. 16 multiple times per time step (which is the most expensive part of the calculation). In order to further lower the computational cost, a heuristic was developed whereby the same Jacobian is saved and used until, at a single time step more than a preset (usually 9) number of iterations are required for convergence, or until the magnitude of the residuals increases after a correction. This results in a Jacobian being used for an average of 5-6 timesteps, and can approximately halve the computational time. For the one dimensional simulations, the LINPACK routine SGBCO was used to solve the system of linear algebraic equations. For two dimensional simulations a memory-resident solver would require too much memory overhead, so a frontal-type solver based on the work of Irons (1970) was used. The frontal solver was provided by Professor Brown's research group at MIT, and has been optimized to run on several platforms. This also has the advantage that more of the sparse structure of the Jacobian is incorporated into the solution instead of simply assuming that the Jacobian is fully banded (the Jacobian for two-dimensional systems is actually a series of blocks with N_p diagonals, where N_p is the degree of interpolating polynomial, and is empty off the block-diagonals).

Other time integration schemes were also investigated. It was found that the Euler integration method was too inaccurate for our purposes, and the problems with starting up a higher order method make their application inefficient. A variable-order, variable time-step scheme was also investigated, in particular the package DASSL, and was found to be unacceptable. It was noted that the time-periodic forcing of the problem interacts with the error predictors for the variable-order integrator, and results in an instability which causes the integrator to stop converging after a few cycles of integration when following transients. This occurred because the heuristics used in DASSL were constantly adjusting the integrator order and timestep, and eventually the accumulated errors required a time step small enough to make the

simulations intractable. It was also not clear whether DASSL could be properly applied to integration of the linearized system for calculation of the TPSS.

The form of the initial guess for solution of the nonlinear system was also investigated. If the Adams-Bashforth second order predictor is used, an estimate of the integrator accuracy is readily available. However, it is known that the second order guess is unstable at the desired time step for the implicit integration scheme, and occasionally the error in the guess will make the Newton iterations unstable. The first order guess performed better, and for most computation (especially in the two-dimensional simulations), the solution at the previous time step was used as the guess for the following time step. When a prediction was within the range of convergence of the Newton iterations, the accuracy of the prediction did not have any noticeable effect on the number of iterations required for convergence of the overall scheme.

For the two-dimensional simulations the LU decomposition of the system Jacobian for solution of the Newton iterations of the implicit time integrator consumes over 65% of the computation time. The remainder of the computational cost results from the other back-substitutions against a Jacobian which has already been LU decomposed (for calculation of the TPSS, the LU decomposition importance with respect to the other inversions is reduced in importance, but the overall importance of these matrix computations is always over 95%). Since the two-dimensional simulations were still quite computationally intensive, even with the heuristic applied to the solution of the implicit integrator, a method for further reducing the computation time was implemented.

It was originally hoped that a semi-implicit method could be used, which would decouple the implicit parts of each equation and lower the total number of unknowns for the matrix which needed to be inverted by a factor of 6, as well as allow computation of the LU decomposition only a single time. However, the nonlinear boundary conditions on the electron density and electron energy equations seem to be an insurmountable problem in applying this type of integrator. However, it was noted by Graves (1994, personal communication) that since the ions move extremely slowly, the Courant number for the ion equations should allow one to follow their concentration with an explicit integrator. Since the ions move so slowly it is indeed noted that the error in the ion equations is always much lower than for other equations, so if the integrator is stable for a high ion time step, it will be accurate enough. The procedure was

therefore to split the integrator into a fully implicit part for the electron concentration, electron energy, and potential (since the electron energy is exponentially dependent on the electric field), and a fully explicit part for the ions and the effective field. Since the Jacobian inversion are completely dominant in the calculation, and since the LU decomposition of the matrix scales as $N B_w^2$, where B_w is the Jacobian half band-width, it was expected that this reformulation of the integrator could result in up to a factor of eight speed enhancement for the simulations. Due to the lowering of the vector length concomitant with the lowering of the bandwidth of the matrix, a real speed enhancement of three to four was actually observed.

2.3.3: Time Periodic Solution

As described by Gogolides et al. (1992), the disparate time scales for electron and ion diffusion make time integration to the TPSS very inefficient. The difficulty with direct time integration is illustrated in Figure 2-1, with the ratio between the fast and slow time scale lower than in a plasma for illustrative purposes (the actual ratio is closer to 10^4 , while the figure uses 10). To find the TPSS it may be necessary to integrate for hundreds or thousands of cycles, while using a shooting technique only a few shooting iterations should be required for convergence. It is clear that in the case where the cost of calculation of the TPSS shooting step is on the order of a few cycles of integration, the simulation speed is enhanced by at least an order of magnitude.

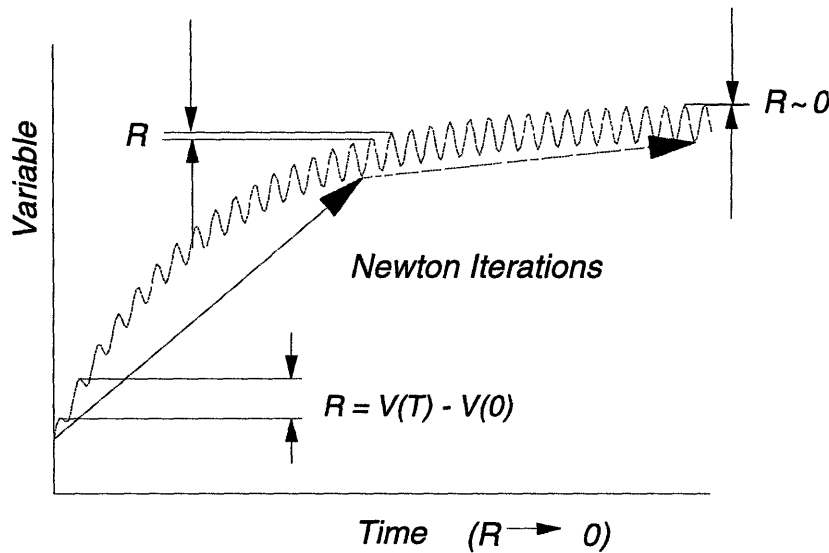


Figure 2-1 Time integration to a TPSS versus Newton shooting solution. For plasma simulations, the short time scale is at least an order of magnitude smaller than shown here.

If it were possible to completely decouple the electron and ion equations in a time averaged sense, an efficient acceleration technique could be formed. This was the approach of Paranjpe et al. (1990). However, the ions depend too significantly on the electron relaxation because of the strong coupling through the potential field, and no rigorous splitting is available. Instead, we used the Newton shooting algorithm described by Gogolides et al. (1992) for direct calculation of a TPSS; the Newton shooting scheme is only briefly described below; see Gogolides et al. (1992) for more details. Here Eqs. 16 are integrated over a single period of the RF cycle with the initial condition updated by a Newton iteration to force the solution to be time periodic, i.e. so that the solution at the end of the period is equal to the initial condition. This results in definition of the residual equations:

$$\mathbf{r}(\mathbf{u}_o) \equiv \mathbf{u}(t = T, \mathbf{u}_o) - \mathbf{u}_o \quad \text{Eq. 17}$$

which is the mathematical statement of the TPSS. Newton's method is used to find the correct initial conditions for integration over a single period to force the residual to zero. The algorithm is a version of the shooting algorithm for solving boundary value problems described by Keller (1976) and has been used by Doedel and Heinemann (1983) to compute time-periodic solutions in a CSTR.

The time-periodic shooting technique simply takes an initial value problem in time, which would otherwise be integrated directly until the residual met some pre-specified tolerance, and changes the problem into a boundary value problem in time specified by Eq. 17. The TPSS calculation then is independent of the way that the residuals are formed, and it is possible to solve Eq. 17 using any desired method. Newton's method is generally applied because it converges quadratically and usually has a large radius of convergence.

Using Newton's method does require the Jacobian of the system 17, hereafter referred to as the TPSS Jacobian, which is quite computationally intensive. The mathematical statement of Newton's method for the TPSS shooting algorithm is:

$$(\underline{\underline{\mathbf{A}}} - \underline{\underline{\mathbf{I}}}) \Delta \mathbf{u} = -\mathbf{r}(\mathbf{u}_o) \quad \text{Eq. 18}$$

where $\underline{\underline{\mathbf{A}}}$ is the Monodromy matrix, $d\mathbf{u}(T)/d\mathbf{u}_0$. It is possible to calculate the Monodromy matrix by perturbing each variable separately, and see what the resultant perturbation is in the residual one cycle later, thus yielding the sensitivity matrix after N cycles. However this method of

calculating the Monodromy matrix ($\underline{\mathbf{A}}$) is very inefficient. The Monodromy matrix is simply a fundamental matrix used to represent the TPSS, and Gogolides et al. (1992) showed that another method of computing the TPSS Jacobian is to follow the linearized version of Eq. 16 with an initial condition of the identity matrix:

$$\underline{\mathbf{M}} \frac{d\underline{\mathbf{A}}}{dt} = \underline{\mathbf{J}}(t) \underline{\mathbf{A}} \quad \underline{\mathbf{A}}(t=0) = \underline{\mathbf{I}}. \quad \text{Eq. 19}$$

This computation results in a fundamental matrix representation of the TPSS Jacobian since the rows are linearly independent perturbations as is discussed by Keller (1976). This can be viewed as calculating the result of independently perturbing each variable, and seeing how this effects all the others if the perturbation is small enough (and therefore the equations of evolution are linear). Although this method still computes by solving N independent perturbations as above, efficiency is increased because the most expensive part of the computation need only be done one time (the inversion of the system Jacobian, $\underline{\mathbf{J}}$). Since the equations are generally stiff, the linearized set is as well, and a first order implicit Euler integrator is used to calculate the time evolution of $\underline{\mathbf{A}}$.

For the one-dimensional simulations, it is most efficient to compute and invert the TPSS Jacobian to calculate the corrections for solution of Eq. 17, even though this is a full matrix since its size is relatively small. The linearized equation set is time-integrated using the Euler implicit algorithm; since the nonlinear equation set is unstable with respect to explicit time-integration, the linearized system is expected to be unstable as well. The accuracy of the TPSS Jacobian is not critical, since it is primarily determining the relative magnitudes of the corrections. A larger time-step for calculation of the TPSS Jacobian is therefore acceptable, and it is not necessary to use a high-order time-integrator for these calculations. The time-splitting scheme which was used in the two-dimensional calculations was also applied to the one-dimensional calculations (the TPSS Jacobian was also computed using the split integrator), and was found to produce equivalent results with equivalent time-steps. This shows that the ion equations are capable of using a very large explicit time-step.

For two-dimensional simulations, the size of the TPSS Jacobian prohibits its simple calculation and inversion since it is a full matrix. However, if one examines how a single column of the TPSS Jacobian is formed, the calculation consists of a series of sparse matrix

solves which originate from the implicit time-integration scheme in Eq. 19. This is clear examining Eq. 19 after the implicit integrator has been applied:

$$\underline{\underline{\mathbf{A}}}(t) = \left[\prod_{j=1}^{N_{steps}} \left(\frac{\underline{\underline{\mathbf{M}}}}{\Delta t} - \underline{\underline{\mathbf{J}}}(t_0 + j\Delta t) \right)^{-1} \right] \underline{\underline{\mathbf{A}}}(t=0). \quad \text{Eq. 20}$$

If the number of time-steps required in the calculation is small, then it is simple and efficient to store the LU decomposition of these system Jacobians, and view the time-integration of the TPSS Jacobian as a set of successive back-solutions. The difficulty of calculating and inverting the full TPSS Jacobian is then avoided by applying the GMRES algorithm to iteratively solve for the Newton step, since GMRES requires only the products of the TPSS Jacobian and a vector. GMRES solves a set of linear algebraic equations by forming the residual, $\underline{\underline{\mathbf{W}}}\mathbf{x}-\mathbf{b}$, and then projecting unto a Krylov subspace which expands with every iteration. The result for the TPSS Newton shooting algorithm using GMRES is that only products of $\underline{\underline{\mathbf{A}}}$ and the correction vector Δu need to be calculated. The algorithm is guaranteed to converge in less than N iterations, where N is the rank of the matrix, and the residuals are guaranteed to decrease monotonically. For general systems, it is observed that GMRES will converge with only a few percent of N iterations, where the exact number of iterations is highly dependent on the convergence criteria which is chosen. The cost of each iteration is very small, since we are merely performing a few back-substitutions against a known right-hand-side. In fact, this is exactly equivalent to solving the Newton iterations for the time integrator for the full system. In the nonlinear case an average of 6-7 back-substitutions are required for each time step. If we use an equivalent number of timesteps for GMRES, the cost of the TPSS Newton step will be equivalent to calculating the residuals as long as the number of iterations is on the order of the number of timesteps used to follow the linear system for a cycle. The cost of using GMRES is minimal for our simulations because the matrix-vector multiply (the back-substitutions which are used in Eq. 20) are very sparse compared to the full matrix-matrix multiply used in the general implementation of the TPSS Newton shooting algorithm.

It is readily seen that the residual required for GMRES is computed by using the previous GMRES vector as an initial condition and applying the time-stepping scheme for the linearized system equivalent to Eq. 20. This method is very efficient because GMRES generally converges to a sufficient accuracy in less than 200 iterations (typically 50-70 iterations), while the full

matrix has rank or order 25,000. This is a property of the algorithm, since the Krylov space heavily weights the highest eigenvalues for the system, and eliminates the largest components of the residual first. It is possible to further enhance this effect by preconditioning the matrix so that the system appears closer to symmetric positive-definite. This is generally achieved by multiplying by a matrix consisting of a very limited band around the diagonal (usually either diagonal or tridiagonal) which can be efficiently inverted. For the two-dimensional simulations this is not possible, since we are using GMRES mainly to avoid computing the total Monodromy matrix and the diagonals are therefore not available. GMRES was also tested on the one-dimensional system, and was found to converge with similar properties to the direct calculation and inversion of the TPSS Jacobian, and with approximately a 25-30% savings in computation time.

2.3.4: Computational Cost

The computational cost of varying the discretization is assessed to demonstrate the difference in using p-type versus h-type refinement for a smooth, well-behaved problem. The main costs associated with solution of the continuum model for one-dimensional solutions are:

A)	B)	C)	D)
computation of the matrix for the nonlinear DAE's Jacobian (Eq. 16)	LU decomposition of Jacobian matrix in A)	computation of the matrix for iteration to a TPSS (Eq. 20)	LU decomposition of the matrix in C)

These four computations make up more than 99% of the total calculation cost. The computation of pieces A-C is performed at each time-step, so the ratio of their costs is most meaningful.

Computation of piece D, however, is done once per TPSS Newton iteration, so this computation is always negligible for more difficult computations which require many time-steps per integration cycle. For two-dimensional calculations, the costs are slightly different, because the relative costs change as the number of variables increases. In two-dimensional simulations, the cost of part D is estimated to be at least an order of magnitude higher than the rest of the simulation, while part C is expected to be nearly two orders of magnitude more expensive. There is also the difficulty of storing the large full matrix, \mathbf{A} , which is estimated to be large enough that it would not fit in core memory of most of the larger supercomputers. These difficulties are

removed by using the GMRES algorithm. This computation results in the calculation of an additional N_{tpss} system Jacobian inversions equivalent to piece B, and N_{GMRES} iterations which require N_{tpss} back-solves of the equation set each time. In practice, the result is that GMRES requires nearly double the computation time as calculating the change of the solution over one cycle. If it were possible to precondition the iterations, a computational savings for GMRES of up to a factor of 5 might be possible, but this would result in no more than a doubling of the overall simulation speed.

For the one-dimensional calculation, the cost of computation A is approximately $N_{\text{nodes}}B_w C_1$, where C_1 is the cost of computing an entry for the system Jacobian matrix and B_w is the half-bandwidth of the Jacobian matrix. C_1 averages 8-10 operations, but does depend on the discretization to some extent. The cost of computation B is $2N_{\text{nodes}}B_w^2$. The cost of computation C is due almost completely to the solution of the N_{nodes} independent linear systems, since the LU decomposition has already been performed to solve the nonlinear system in computation B. This yields a cost for the formation of \mathbf{A} of $3 N_{\text{nodes}}^2 B_w$. The cost of computation D is $1/3 N_{\text{nodes}}^3$, but the ratio to the other costs requires a factor of $1/N_{\text{steps}}$ since the other computations are performed at every timestep. The LU decomposition of \mathbf{J} is performed only once per integration over an RF cycle.

The tradeoffs in computation cost between p-type and h-type refinement are considered by examining these costs in terms of the more primitive variables, N_p and N_{elem} , where N_p is the degree of the polynomial for the discretization and N_{elem} is the number of spectral elements in the discretization. This leads to the cost estimates:

A)	B)	C)	D)
$C_1 N_p^2 N_{\text{elem}} N_{\text{fl}}^2$	$2 N_p^3 N_{\text{elem}} N_{\text{fl}}^3$	$3 N_p^3 N_{\text{elem}}^2 N_{\text{fl}}^3$	$1/3 N_p^3 N_{\text{fl}}^3 N_{\text{elem}}^3 / N_{\text{steps}}$

Comparing these computations with the LU decomposition of Jacobian matrix of the DAE system (computation B), yields the following relative computation costs:

A)	C)	D)
$C_1 / (2N_p N_{\text{fl}})$	$3/2 N_{\text{elem}}$	$1/6 N_{\text{elem}}^2 / N_{\text{steps}}$

This indicates that computation A becomes less important as the degree of the interpolating polynomial increases, computation C increases with the number of elements, and computation D

increases as the number of elements squared. Using p-type refinement only decreases the relative cost of computation A, while h-type refinement emphasizes calculations C and D.

For one-dimensional simulations, the calculation of the Jacobian matrix of the DAE system is usually negligible with respect to its LU decomposition, and the importance of back-solving for a previously decomposed Jacobian can be manipulated by saving the matrix for a variable number of iterations. Both calculations required for the Newton iteration to find the TPSS are usually more expensive than the LU decomposition of the DAE Jacobian matrix, but the most important cost overall is the formation of $\mathbf{\Delta}$, except for very fine spatial discretizations. Vectorization efficiency of the calculations can greatly affect the relative costs of the different calculations, although this fact is not apparent from the scalings on operation counts, and only becomes obvious from timing estimates. For the one-dimensional calculations in this work, computations A and D vectorize most efficiently, and are not of significant computational cost. The scaling between computations B and D still does seem to hold.

For two-dimensional simulations, the algorithms have been chosen in such a way to minimize the overall cost of the system Jacobian inversion, and the calculation of the TPSS shooting step. Vectorization efficiency seems to keep the relative computation cost approximately constant with increasing mesh size, and use of the GMRES algorithm assures that the TPSS calculation will always have similar relative cost to the calculation of the residuals. The cost of calculating the system Jacobian is quite minimal because of the way the calculations have been performed.

For both the one- and two-dimensional computations, all functions were hand-inlined for efficiency. Loops over the equation type (to discriminate between positive ions, electron concentration, electron energy, voltage, and effective field) were generally written out explicitly. Furthermore, to increase vector efficiency, multiple loops were chained where ever it was possible. For example, to do computations for all the local contributions, one can either loop over the degree of the polynomial and the number of elements separately, or these loops can be chained using a set of arrays which contain information necessary for unchaining the loops. This results in extra computations for the indices, but results in very long vector lengths (especially for the two-dimensional simulations), and increases computation speed by up to a factor of 10. Some optimization was also done on the vectorization for the calculation of $\mathbf{\Delta}$ for the one-

dimensional case. Instead of treating the calculation as an inversion versus N independent right-hand-sides, the order of the calculations was inverted so that the computation was performed as an inversion versus all N right-hand-sides simultaneously. This loop inversion results in a speed enhancement of over a factor of 10, since the number of equations, N , is much greater than the bandwidth of the matrix being inverted.

The optimization of the simulations results in very efficient computation of a TPSS for both the one and two dimensional simulations. For the one dimensional case, less than 8 MW of memory and approximately 1 minute computation time are required on a CRAY C90 to find a TPSS from an arbitrary starting profile for Argon discharges. The SF_6 case requires a smaller timestep and a much finer discretization, and can be solved in approximately 10 minutes on the same platform. For the symmetric two-dimensional cases less than 16 MW of core memory are required, as well as 1 GB of temporary storage for the frontal solver. Computation of one TPSS shooting step requires approximately 10 minutes of CPU time to calculate the residuals, and approximately 20-30 minutes of CPU time to calculate the GMRES iterates. For the asymmetric cases 1-2 GB of temporary storage are required and slightly more than 16 MW of core memory. There is a concurrent increase in computation time to 20-25 minutes for calculating the residuals and 45-60 minutes to calculate the GMRES iterates. For the most expensive asymmetric simulations 10 Newton steps were required to achieve convergence, resulting in a total cost of approximately 14 CPU hours on a C90. This is not the minimal cost for the simulation since the discretization was not optimized.

Chapter 3 : One-Dimensional Modeling Results

3.1: Introduction

As a demonstration of the physics, accuracy, and capabilities of the continuum model, this Chapter will present results from the continuum model in one dimension. The first section demonstrates the general physics which are included in the continuum model for electropositive discharges. The results shown are based on bench-mark modeling work for the 44th annual Gaseous Electronics Conference, and compare very favorably with other works in the field, as is evident from the comparison presented by Surendra (1995).

The second section shows why the methods used for this modeling effort are superior to other works which use artificial-diffusion in the spatial discretization, and presents motivation for using accurate representations of the spatial operators for plasma modeling. Another advantage of the techniques used for this work is the inherent speed of the computations which is a result of the higher coupling of the spectral-element basis functions, and hence longer vector-chain lengths.

The third section shows a brief comparison of the continuum model results to experimental measurements performed in the lab at MIT. As will be shown in the next chapter, except when one is examining the physics relatively close to an insulating wall, the plasma is very nearly one-dimensional. Therefore, one expects that zero- and one-dimensional measures of the plasma physics, such as the total current versus voltage or the power input to the plasma, and spatially averaged quantities will compare well with experimental results. The perturbations on the one-dimensional solution are examined in detail in the next chapter.

The last section in this chapter examines the changes which occur when the plasma is highly electron-attaching. A detailed description of the physics of an SF₆ discharge is examined, and the motivation for confining initial two-dimensional studies to electropositive discharges is discussed.

The model used for all of the simulations is identical, except that for the benchmark Helium simulations the positive ion transport parameters are allowed to vary as a function of the effective field to allow direct comparison to other modeling efforts. The transport and rate

parameters for all the simulations, both one and two dimensional, in this work are presented in Table 3-1. For experimental comparison all parameters are taken to be fixed: no attempt is made to achieve a better fit by adjusting any of the simulation parameters.

Table 3-1 Transport and reaction parameters for the continuum model

	Argon Parameters	Helium Parameters	SF₆ Parameters
Ion Diffusivity (cm ² -Torr/ s)	40	$7.8\sqrt{E/P} + 378.7$	13
Ion Mobility (cm ² -Torr/ V-s)	1444	$\frac{15148}{\sqrt{1+0.16E/P}}$	500
Electron Mobility (cm ² -Torr/ V-s)	3×10^5	9.42×10^5	1.4×10^5
Electron Diffusivity (cm ² -Torr/ s)	$2/3 \mu_e \epsilon$	$2/3 \mu_e \epsilon$	$2/3 \mu_e \epsilon$
Ionization Rate Constant (1/cm ⁶ -s)	$8.7 \times 10^{-9} * (\epsilon - 5.3) * e^{\frac{-4.9}{\sqrt{\epsilon - 5.3}}}$	$2.0 \times 10^{-7} * \sqrt[3]{\sqrt{\epsilon - 1.8}} e^{\frac{-4.7}{\sqrt[3]{\epsilon - 1.8}}}$	$2.1 \times 10^{-9} * (\epsilon - 5.0) * e^{\frac{-6.2}{\epsilon - 5.0}}$
Electron Energy Loss Rate Constant (1/cm ⁶ -s)	$8.1 \times 10^{-9} * (\epsilon - 5.3)^2$	$\frac{2.0 \times 10^{-7}}{\sqrt[3]{\epsilon - 1.8}} e^{\frac{-4.7}{\sqrt[3]{\epsilon - 1.8}}}$	$1.37 \times 10^{-8} * (\epsilon - 5.0)^{1.6}$
Electron Attachment Rate Constant (1/cm ⁶ -s)	0	0	8.0×10^{-10}
Positive/Negative Ion Recombination Rate Constant (1/cm ⁶ -s)	0	0	7.9×10^{-8}

3.2: Low Pressure Helium Plasmas

Much previous work has been devoted to the study of the physics of Argon plasmas, and the two-dimensional simulations in this work are limited to Argon as well. Surendra (1995) organized a detailed model benchmark comparison for both continuum type models and PIC/Monte-Carlo simulations to compare speed and accuracy. Only one other group doing continuum simulations was able to demonstrate any accuracy in the code, and the continuum model compares quite favorably with the PIC/Monte-Carlo simulations even at 30 mTorr where the fluid assumptions are dubious at best. The results are also compared qualitatively to some experimental measurements and are shown to be in reasonable agreement. One of the difficulties in the comparison is that even though all of the transport parameters are provided, each model requires slightly different inputs and therefore has to tailor the parameters to specific needs. The parameters used for simulation of Helium discharges are presented in Table 3-1.

Helium simulations were run at 300 mTorr, 100 mTorr, and 30 mTorr. The continuum model was expected to work best at the highest pressures, while the PIC scheme was expected to be more efficient and more accurate at the lower pressures. The results for Helium are similar to results which were obtained for Argon discharges. The first measure of interest is the plasma density, which is shown in Figure 3-1. The shape of the curve is typical for high pressure electropositive discharge simulations, and occurs because there is no bulk loss mechanism to offset the production of positive ions other than drift-diffusion. Ions must therefore constantly decrease in concentration from the center of the discharge outwards. The positive ions are always modulated in time at a scale much smaller than is visible on the plots, and the electrons move visibly only in the sheath regions adjacent to the plasma boundary. The sheath becomes larger as the discharge pressure is lowered, and the overall density of both species decreases, although not linearly with pressure.

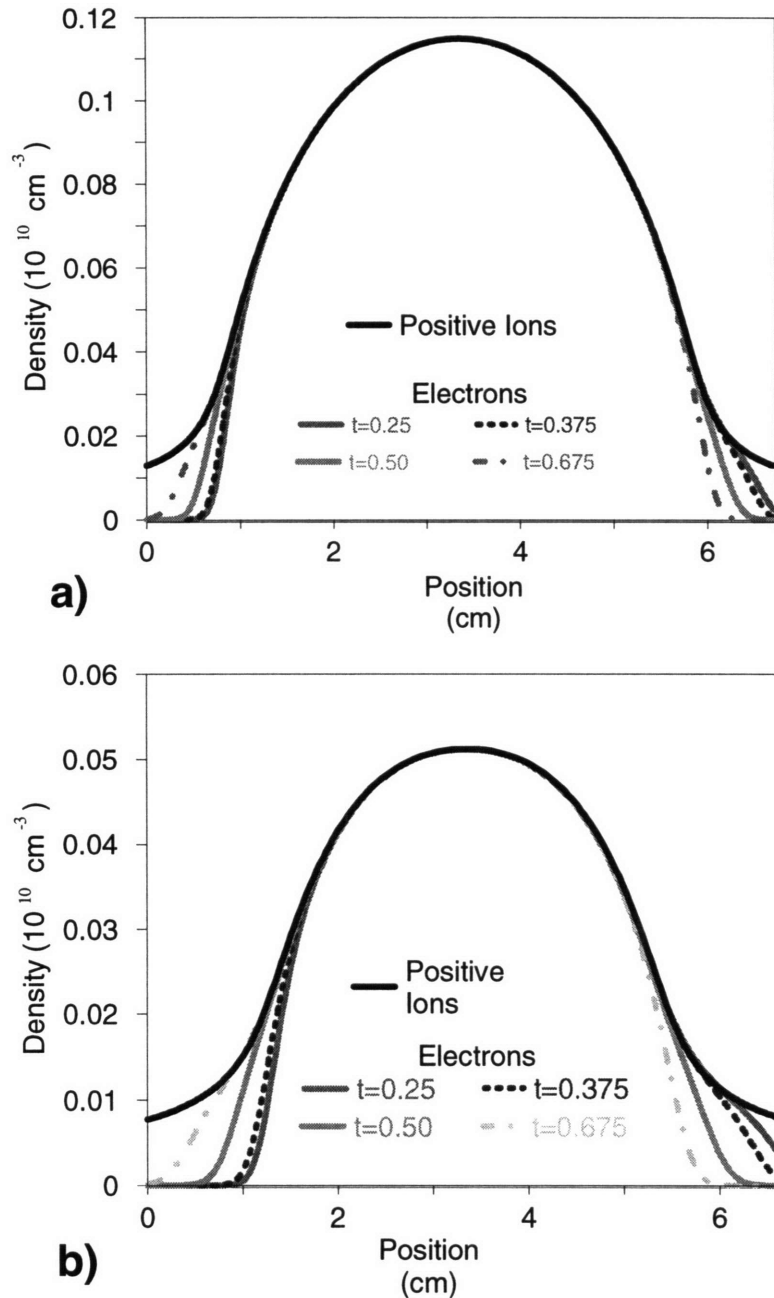


Figure 3-1 Helium plasma density for electrons and positive ions. Electrons are shown at various phases in the RF cycle, while positive ions are stationary in time. 13.56 MHz, 6.5 cm gap, 1 mA/cm² current. a) 300 mTorr; b) 100 mTorr. Sheath formation is evident at both pressures, and the plasma behaves similarly independent of pressure.

It is interesting to note that the electron concentration becomes vanishingly small in the plasma sheath region, and that the details of this physics is accurately captured by the spectral element method. In fact, one of the drawbacks of using a highly accurate method is the necessity

of resolving all the physics in the model even if it is of little consequence. This is highly evident in the plot of electron energy presented in Figure 3-2 where the leftmost plots contain a small amount of secondary electron emission to stabilize the model. The goal was to solve the model with a zero secondary electron emission coefficient, to compare to the other modeling efforts. This was not tenable because the electron density was vanishingly small on the electrode, forcing the electron energy equation to be nearly singular. The difficulty with a singular electron energy balance is evident in Figure 3-2 b), where there is a slight oscillation in the solution near the bulk-sheath interface. This oscillation is purely numerical, and can be removed through sufficient mesh refinement.

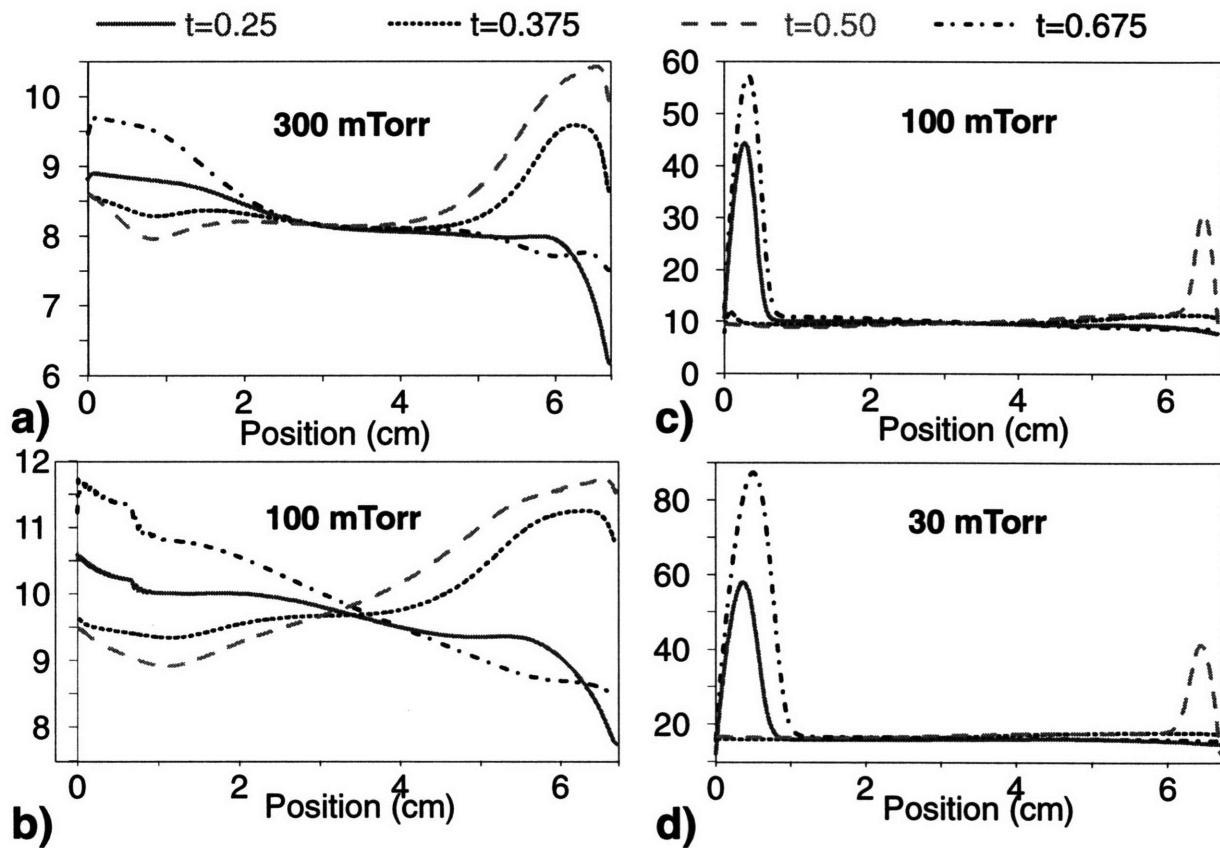


Figure 3-2 Helium electron energy shown at 4 phases in the RF cycle. 13.56 MHz, 6.5 cm gap, 1 mA/cm² current. a) 300 mTorr; b) 100 mTorr no secondary electron emission, c) 100 mTorr 10^{-5} secondary electron emission, and d) 30 mTorr with secondary electron emission. The addition of a small secondary electron emission stabilizes the numerical methods by maintaining a small electron concentration in the sheaths, changes the energy in the sheaths, and has very little effect on the global plasma physics.

Another approach proved much more valuable in investigating the plasma physics, and stabilized the solutions by removing the cause of the problem. Providing only 10^{-5} secondary electrons per ion striking the electrode greatly changes the electron energy in the sheath, as shown in Figure 3-2 c), but is only changing the total electron population in the sheath by an inconsequential amount. Indeed, in any real plasma, there will be secondary electron emission from all surfaces. Other modeling efforts do not show a problem with the electron energy because they are not correctly capturing the physics, and therefore ignore this problem. Even though this small change does not affect the overall discharge model, it emphasizes the importance of knowing how the numerics affect the overall solution. Note that even though Figure 3-2 c) and d) show very high electron energy in the sheath, this is for a vanishingly small number of electrons which probably exist under the plasma conditions simulated. The other interesting observation is that the secondary electron physics which we are adding is something that the PIC codes do not model because of the difficulty of accounting for electrons so far out in the distribution function.

It is also interesting to examine where the plasma power is deposited. If one looks at the time-dependent power deposition for ions and electrons, as shown in Figure 3-3 , it is clear that the electrons couple in the majority of the power near the sheath-edge, and lose significant energy to the plasma during the part of the cycle when the sheath is collapsing as they diffuse against the field. It is also apparent that this point in the cycle is the most difficult for the simulations, as is apparent from the small oscillations in the electron power. These oscillations occur because the power is a product of the electron flux and the electric field, both of which are not required to be continuous across element boundaries, and converge to smooth solutions with mesh refinement. The ions gain energy by falling across the time-averaged sheath potential, and will deposit this energy in the form of heating of the electrodes. Note that the ion power is not constant in time, and that this is mainly due to a modulation in the effective field, not in the ion current.

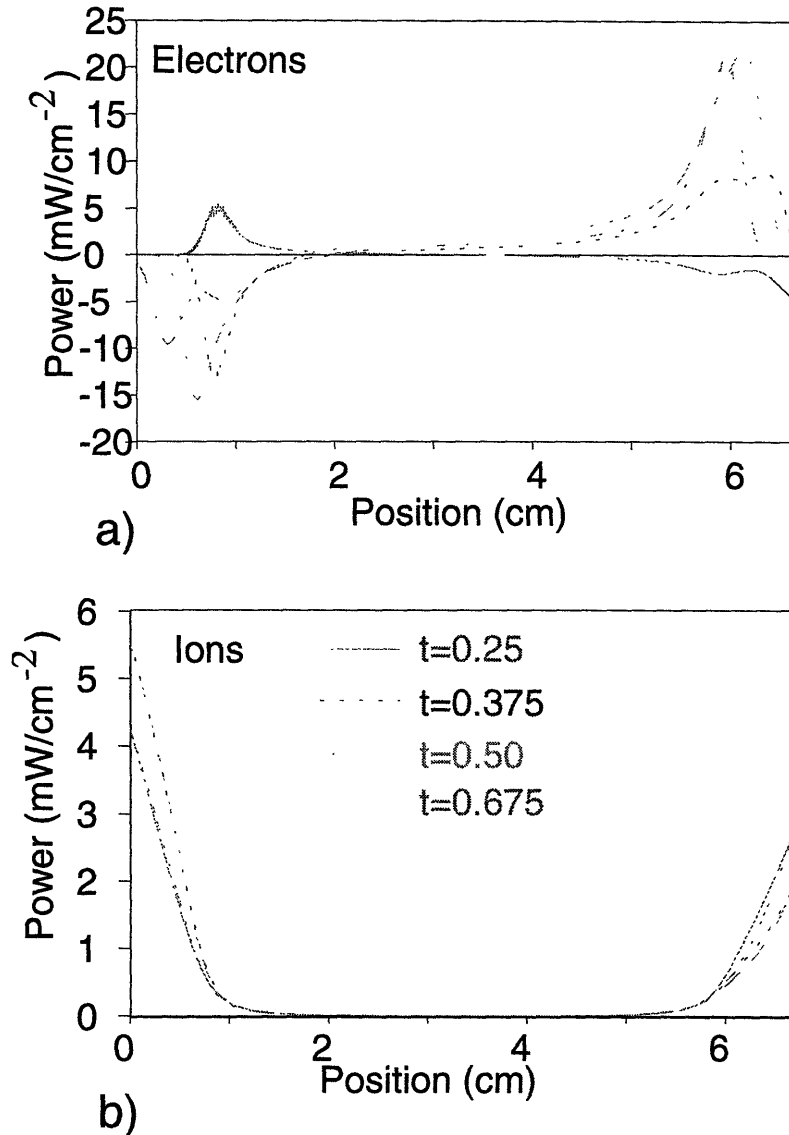


Figure 3-3 Power deposition in a Helium discharge at 300 mTorr. Shown at 4 phases in the RF cycle for a) Electrons and b) Ions. Ions are always heated by falling through the plasma sheath, while electrons are cooled for a significant part of the RF cycle while diffusing against the sheath field.

The simplest measure which can be made of the plasma physics is to simply measure the way the current changes with applied voltage, giving the plasma impedance. This is shown in Figure 3-4, and for the Helium discharge it is obvious that at this level of detail, there really is very little difference in the discharge with respect to pressure changes. The discharge is highly capacitive, and most of the current is due to the capacitance of the sheaths. The electrons flow into the electrode during a very small part of the cycle when the field has collapsed, and the ions are a very small component of the current which changes very little over time. As is evident by

the rich physics in the rest of the discharge simulation, a simple model which just deals with predicting current and voltage will be overly simplistic, and matching just the current is not a good measure of the plasma physics.

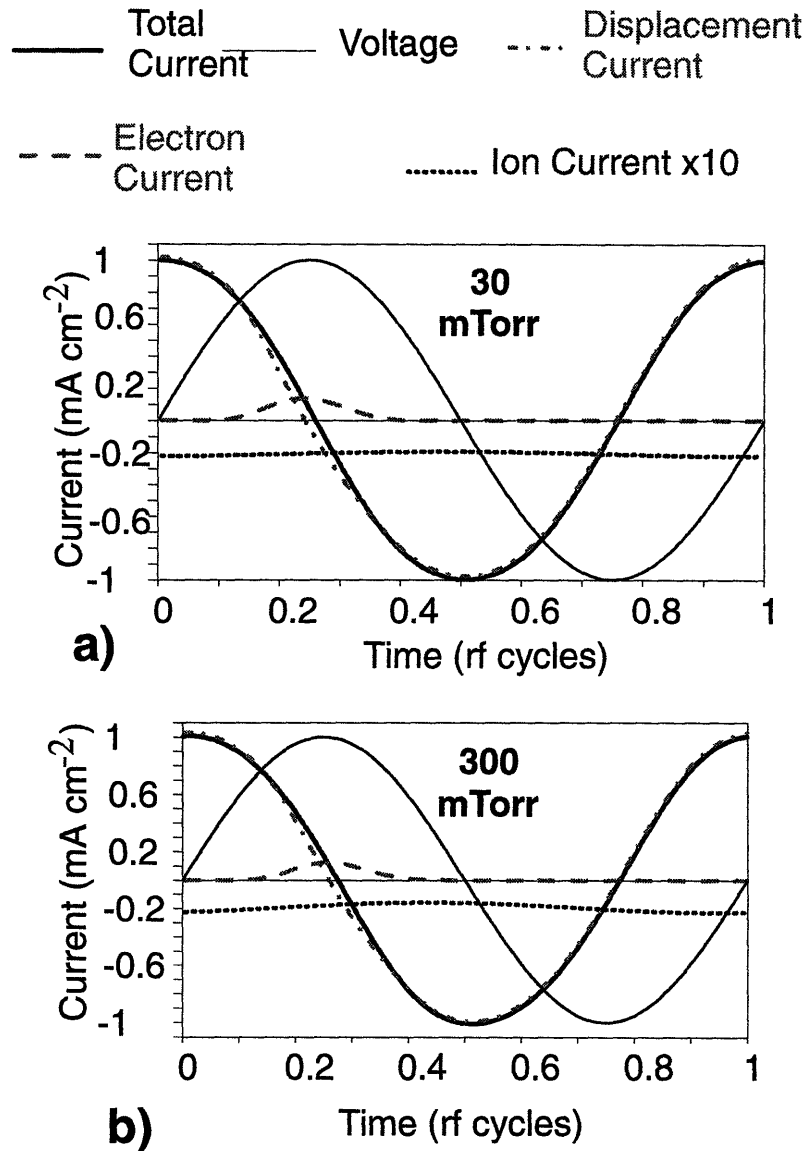
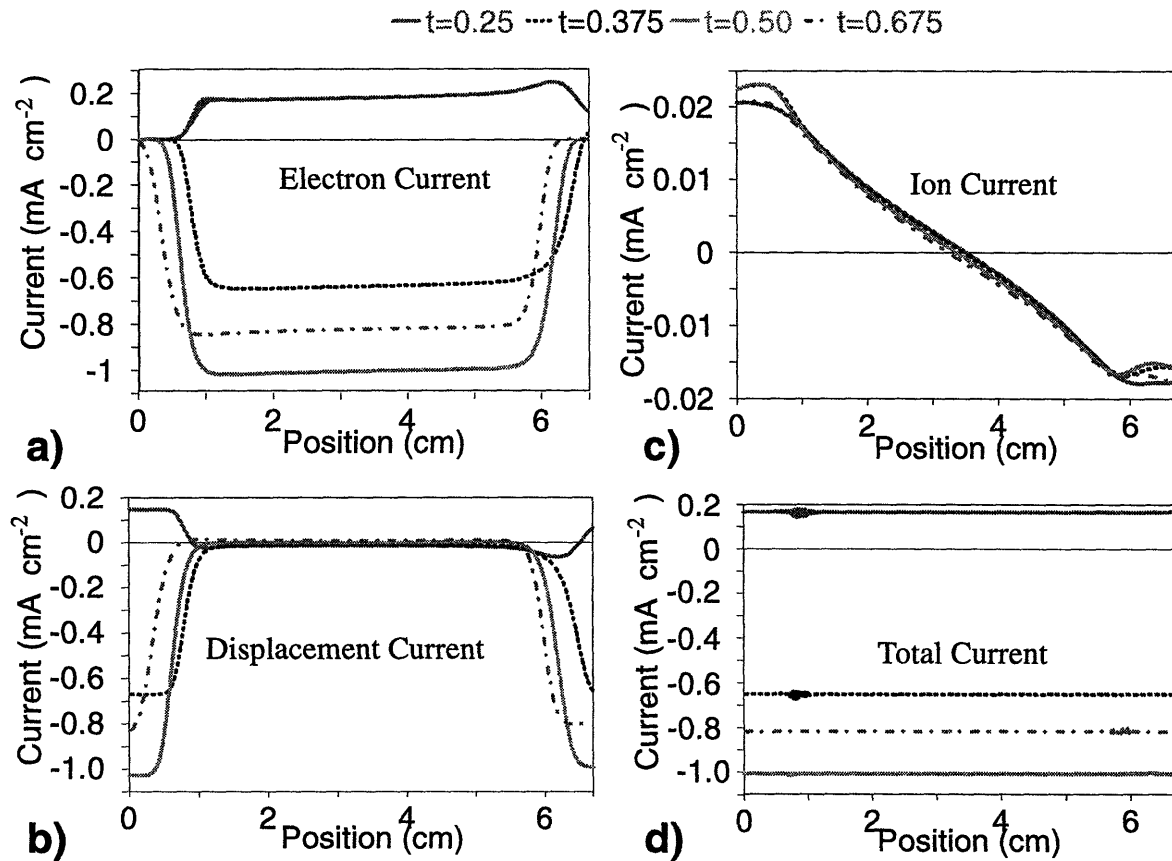



Figure 3-4 Current components for a Helium discharge. a) 30 mTorr and b) 300 mTorr. The electropositive discharge is highly capacitive, and the current is not affected by changes in the discharge pressure.

Although simply examining the current is not a good measure of the plasma physics, the components of the current contain much more detailed information, especially when viewed

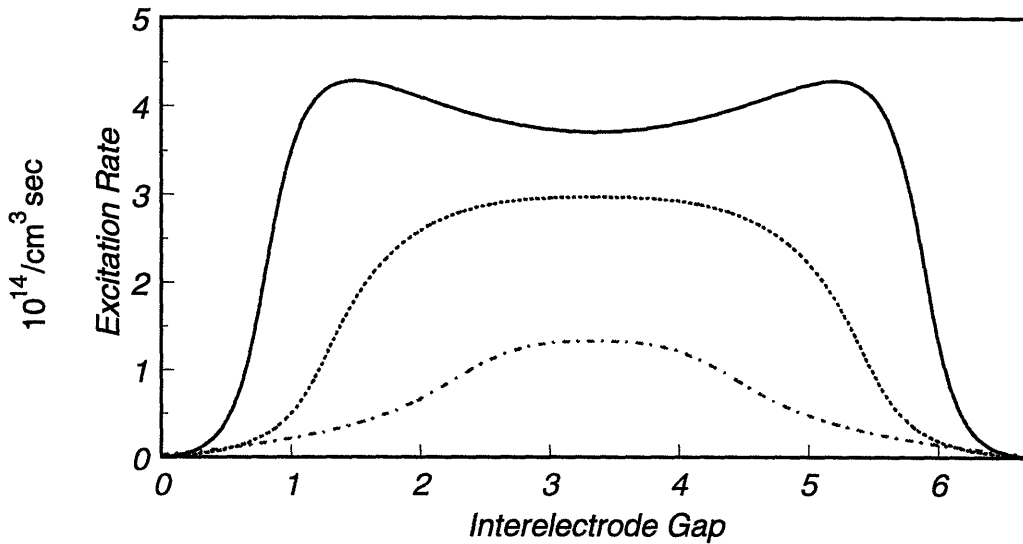
spatially resolved. It is then obvious that although the electrons do not carry the majority of the current to the electrode, it is the electron current which dominates the bulk of the discharge, Figure 3-5. The capacitance of the sheaths is very evident in the displacement current, as is the modulation in the sheath thickness. One can also see that ions generally move outward from the center of the discharge, and reach some limiting velocity in the sheath, which is slightly modulated in time. The total current is constant across the discharge since the simulation is one-dimensional and there is only one path for current to flow (and indeed, one can simulate using a forcing current instead of a forcing voltage). Here again, the small oscillations in the total current are due to the numerical approximation for the electron equation: they disappear with mesh-refinement, and are not significant to the simulation.



When viewing a discharge, the most apparent feature is a bright glow which is due to the light emitted when species excited by electron impact relax to their ground state. For this work, the general approach to modeling the light emission is to look at the ionization rate as equivalent to excitation, although the threshold for the two processes are different. For the Helium discharge, however, the excitation cross-sections were given in order to calculate part of the energy loss rate, so a more direct measure of the glow was available, and can be readily compared to the ionization rate. One expects to see two brightly glowing regions at the sheath edges, and a duller glow in the bulk of the discharge for most electropositive discharges. Both the ionization profile and the excitation profile show this trend, as seen in Figure 3-6. However, the sheaths span the discharge at 100 mTorr, which is not experimentally observed. The excitation rate does show the glow decreasing with pressure, while looking at the ionization rate seems to suggest that the glow should increase as the pressure decreases. To correctly model the plasma glow, it is evident that the simple energy loss process modeled in the excitation is not correct, and the ionization rate does not decrease quickly enough with pressure. A population balance on the metastables should be enforced to rigorously examine the plasma glow, but that is beyond the scope of this modeling effort.



a)



b)

Figure 3-6 Time averaged a) ionization and b) emission rates for a Helium discharge while varying pressure. Both profiles show a maximum in intensity at the sheath edge at the highest pressure, but neither ionization nor emission can correctly capture the bright glow at the sheath which is experimentally observed at lower pressures.

3.3: Numerical Convergence Results: Spectral Element Method

The accuracy and computational efficiency of the spectral element, Newton shooting method is established by a series of calculations with varying time step and spatial discretization. These calculations were carried out using the properties of the argon plasma described by Gogolides et al. (1992) for an electrode spacing of 2 cm, a pressure of 1 Torr, and a frequency of 13.56 MHz for varying applied current or voltage.

3.3.1: Temporal Accuracy

High accuracy of the time integration is crucial to take advantage of the exponential convergence of the spectral element spatial discretization. This point is emphasized by comparing the accuracy for the integration of the solution for a fixed time step and an arbitrarily small value for the calculation of the residuals after integration over a single period for varying values of Δt ; these results are shown in Figure 3-7 with u_{exact} defined by integration for $\Delta t = 1/6400$. The quadratic convergence of the time integrator is obvious. Note that the error in the integrator for the time step $\Delta t = 1/200$ is approximately 10^{-5} . Hence, it is expected that the error in the spatial discretization cannot be smaller than this limit.

The minimum value of this residual at the end of the Newton iterations for each value of Δt is plotted in Figure 3-7 and is uniformly smaller than the error in the residual due to the time integrator. The residual does not attain machine zero due to inaccuracy in calculation of both the residual and Jacobian. The minimum in the residual at Δt near 1×10^{-3} is caused by the limitation in the accuracy of the LU decomposition of the Jacobian matrix for the TPS Newton iteration, which introduces round-off error of approximately 10^{-8} when LU factorization is performed using a LINPACK algorithm; the condition number for this matrix is usually in the range of $10^7 - 10^9$.




Figure 3-7 Accuracy of solutions with respect to: (+) trapezoid rule time integrator, integration for one full cycle, (O) calculation of the minimum for the TPSS residual, Eq. (14). The trapezoid-rule integrator converges as $(\Delta t)^{-2}$ while the minimum TPSS residual converges as $(\Delta t)^{-3}$.

The error in the residual seems to be converging as $O(\Delta t^3)$, instead of $O(\Delta t^2)$ as expected from the accuracy of the integration over a single period. The increased accuracy of the TPSS solution is linked to the accuracy of the integrators for both the nonlinear equation set **16** and the linear sensitivity equations that are used in the Newton shooting method. There appears to be an aggregate effect of using higher-order integration for solving both equation sets. This result suggests that low-order integration of the sensitivity equations may lead to results with the same accuracy, but at lower cost.

3.3.2: Spatial Accuracy

The major advantage of the spectral element approximations over conventional finite element approximations is the higher-order accuracy and the resultant increase in computational efficiency. The higher-order accuracy will be apparent by comparison of results computed with different finite element meshes. The increased computational efficiency is a result of the smaller number of differential equations (smaller N) that are needed to achieve a given level of accuracy, compared to the discretizations created by finite difference and finite element methods. Since the cost of the calculations is directly related to the reduced number of degrees-of-freedom (N_{nodes}) in the discretization, as described in Chapter 2, reducing N_{nodes} is an efficient means of limiting calculation cost.

The discretization error computed for the TPSS solution for the argon plasma is shown in Figure 3-8 with the error defined as the difference between the solution computed for a given mesh and the solution computed using 23 elements and 11th degree polynomials, which is taken as the exact solution. Calculations are shown for holding the degree of the polynomial constant and increasing the number of elements, as is typical for a finite element or h-type approximation, and for spectral or p-type refinement where the number of elements is held constant and polynomial degree is increased. Calculations are shown with two different values of Δt to verify that the accuracy of the time integrator is not affecting the evaluation of the accuracy of the spatial discretization.

The spectral refinements demonstrate exponential convergence as plotted on this semi-logarithmic graph, except for the data points corresponding to the coarsest and the most refined approximations. Evidently the coarsest approximation is too crude for the asymptotic error estimates to be valid; theoretically spectral convergence is only achieved in the limit of a vanishing error. The finest approximation is too close to the solution defined as the exact result, so that the computation of the error is no longer reasonable. Spectral convergence also is demonstrated by the error in the average energy per electron at a point that represents that maximum value of this variable for the coarsest discretization; this is plotted in Figure 3-9 as a function of the degree of polynomial used in the spectral approximation with 23 non-uniform elements. The error in this variable is less than 10^{-3} even for the coarsest discretization used.




Figure 3-8 Convergence of the spectral element method with both h-type and p-type mesh refinement. p-type refinement used 23 nonuniform elements and h-type refinement used 23, 46 and 92 elements. For h-type refinement the time-step was set to $\Delta t = 0.005$, and the symbols refer to refinement with: (□) 4th order polynomials, (○) 3rd order polynomials, (△) 2nd order polynomials. For p-type refinement, two different time-steps were used on the same elemental mesh, with the symbols representing: (+) $\Delta t = 0.0025$, (x) $\Delta t = 0.005$.




Figure 3-9 Exponential convergence of the spectral element method with p-type mesh refinement for the peak energy per electron, which is the most sensitive simulation variable.

Because both h- and p-type refinements are produced by changing N_{elem} and N_p , respectively, the differences in accuracy and efficiency of the two methods can be easily assessed. For a set polynomial order, the Galerkin approximation with Lobatto-Gauss-Legendre polynomials reduces to a classical finite element method that should converge as h^p , where p is the polynomial order. Calculations for $p = 4$ are shown in Figure 3-8 with halving of the finite element mesh. The error appears to decrease by 2^4 , as expected. Results for $p = 3$ also show the expected behavior.

Results also are shown in Figure 3-8 for comparable h- and p-type spatial discretizations. Starting with fourth degree polynomials on a grid of 23 elements and increasing the number of elements to 46 gives the same number of degrees-of-freedom as 8th degree polynomials on the original 23 element mesh. The error in the spectral element discretization is nearly 5 times smaller than the h-type reduction achieved by decreasing the element size. The computational cost of the spectral discretization, measured in CPU time on the Cray XMP-464 at MIT is nearly

double for the p-type discretization, as expected because of the increased bandwidth of the Jacobian matrices caused by the use of the higher-order polynomials. This increased computational cost is evident from the scaling arguments in Section 2.4. Lowering the order of the polynomial in the h-type discretization results in more dramatic results. For discretizations with the same N , the error in the spectral element results is over 10 times smaller with only a factor of two increase in computation cost. Clearly, the spectral element approximations have a distinct advantage if very high accuracy is needed. The need for this high accuracy in the modeling of rf-plasma discharges is the focus of the discussion in the next section.

We demonstrate the cost of using low-order finite elements for calculations with more than minimal accuracy by examining the computational cost of finding solutions with an error of 10^{-5} . For the p-type refinement, using 23 elements and 6th order polynomials suffices. This yields a total of 139 nodes, and relative to the base case of 4th order polynomials on 23 elements, a relative cost of 3.1. For h-type refinement with cubic polynomials the same calculation requires more than 300 nodes, with a relative cost greater than 11. For quadratic polynomials, the savings are even more impressive: the calculation requires more than 2000 nodes and the relative cost would be more than 480. In other words, using spectral elements instead of quadratic finite elements saves more than a factor of 150 in computation cost.

3.3.3: Comparisons Between Spectral Element Results and Finite Difference Calculations of Gogolides et al.

As described in Chapter 1, many of the previous calculations for rf-plasma discharges have used low-order accurate finite difference discretizations that apply upwind differencing schemes to stabilize the discretized equations. However, the upwinding scheme may introduce artificial diffusion into the approximation that incorrectly smooths sharp gradients in the solution field and hence gives inaccurate results for fluxes that depend on these gradients. Accurate computation of these gradients and the resulting fluxes requires higher-order accuracy in the spatial discretization and is one of the major advantages of spectral element methods. To demonstrate this point we compare calculations using the p-type spectral element method

described here with calculations using the upwind finite difference approximations of Gogolides et al. (1992).

The calculations described below demonstrate the effect of the artificial diffusivity introduced by the upwinding in the finite difference approximations. We have attempted to mimic this effect in the spectral element calculations by artificially increasing the diffusion coefficients for electrons and ions to 1.5 and 60 times their proper values, respectively. This calculation represents a crude attempt at approximating the behavior of the finite difference approximation, because the effective diffusivities in the upwind approximations should behave as (Thomasset, 1981)

$$D_{eff} = D + \Delta x v = D + \Delta x \mu_e \Delta V \quad \text{Eq. 21}$$

or

$$= D + \Delta x \mu_+ \bar{E}^{eff}$$

where Δx is the mesh spacing in the finite difference discretization. Corrections of the form of Eq. 21 would be highly nonlinear because of the temporal and spatial dependence of the potential fields.

Calculations of the current-voltage (I-V) curve and the maximum density of charged species in the plasma are shown in 3-10 for calculations with the spectral element and finite difference methods and for computations with the spectral method using modified diffusivities. For both variables there is a significant difference between the results of the spectral and finite difference calculations. The IV curves differ by as much as 30% for low values of the voltage, and this difference is not very sensitive to the value of the diffusivity used in the spectral element simulation. The difference in the maximum density is also approximately 30% and is much more sensitive to the values of the diffusivity, This difference is approximately modeled by using the increased diffusivities given in Eq. 21. This agreement is not unexpected because the maximum density of the plasma is set by the loss rate of ions from the center of the plasma, where the contribution of ion drift is extremely small. We believe that only a fraction of this difference is a result of the difference in boundary conditions used in the calculations (current forcing in the

finite difference calculations and voltage forcing in the spectral results) because the difference is still 15%, when compared to the spectral calculations performed with current.

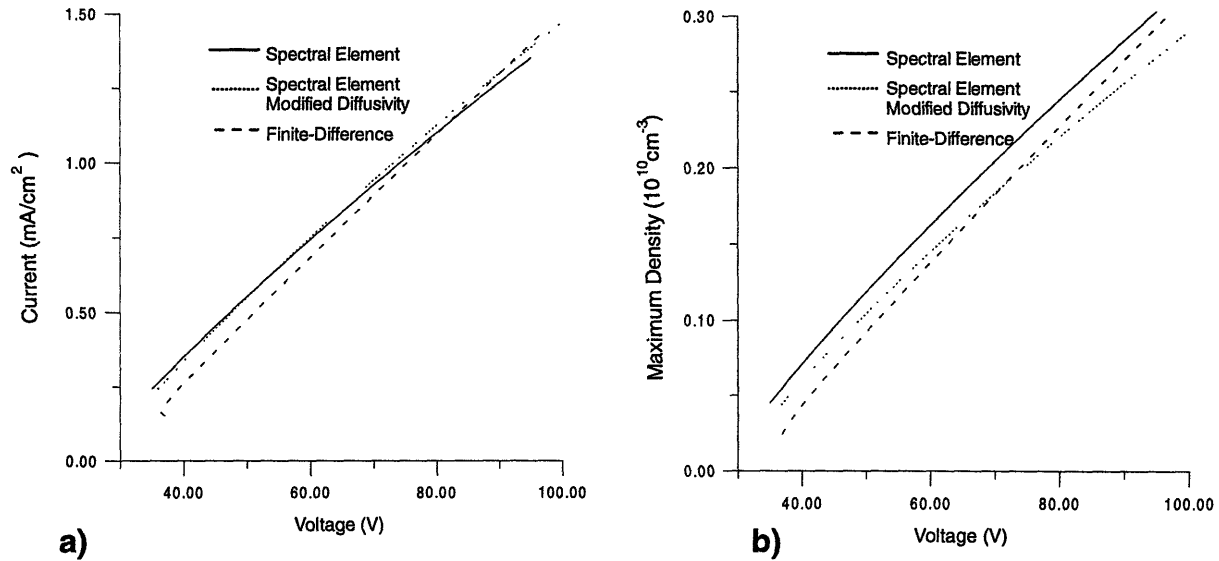


Figure 3-10 Plasma simulations for a 2 cm argon discharge operating at 1 Torr and 13.56 MHz; a) current-voltage characteristic, and b) the maximum plasma density as a function of voltage.

The differences between the spectral element and finite difference calculations are small for some variables, as is demonstrated in Figure 3-11 for the variation of the electric field, $\partial V/\partial x$, as a function of position at three times during the rf-cycle. The difference between the spectral element and finite difference results is generally decreased by artificially increasing the diffusivity in the spectral element calculations, indicating that the artificial diffusivity is responsible for much of the error in the finite difference results. The three calculations have very similar spatial and temporal behaviors, indicating that the qualitative physics of the plasma is preserved in all three numerical methods.

The most significant errors caused by the upwind differencing appear in the calculation of the average electron energy, as shown in Figure 3-12 by the spatial dependence of this variable at a particular time in the period. The spectral calculations resolve a very thin boundary layer in the energy density near the electrode at $x = 0$, which is not present in the finite difference calculations. Increasing the diffusivity in the spectral element calculations does not explain this difference.

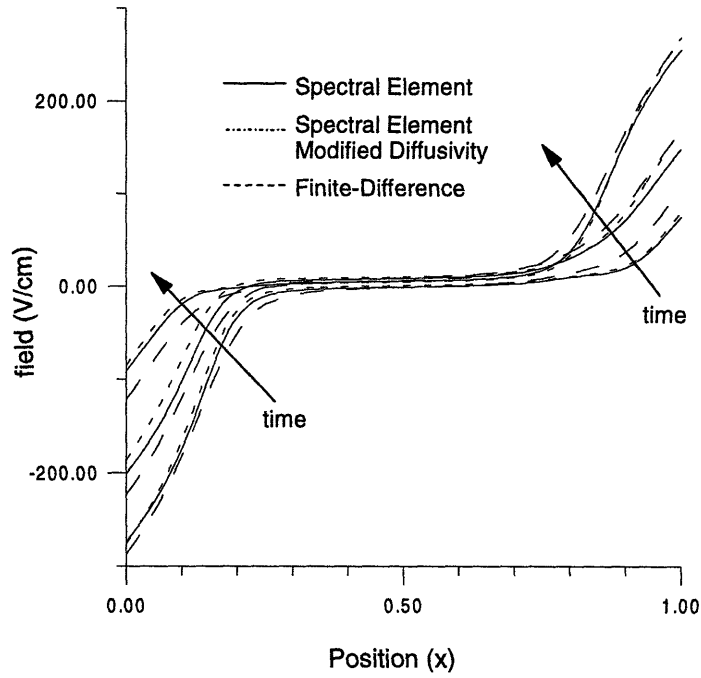


Figure 3-11 Time evolution of simulated electric field profile using the conditions in Figure 3-10 with 60 V forcing voltage.

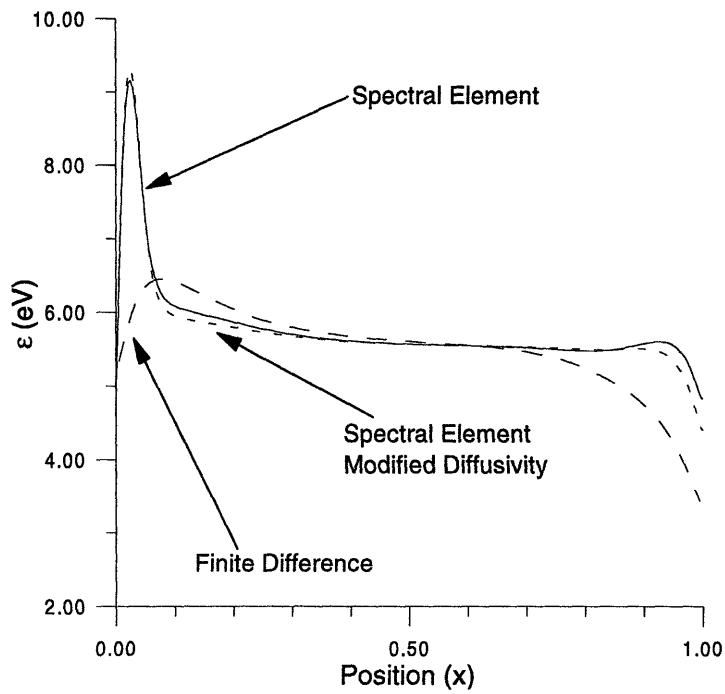


Figure 3-12 Time evolution of simulated electron energy profile using the conditions in Figure 3-10 with 60 V forcing voltage.

The differences in the calculations of the average energy per electron are seen further in the three-dimensional contours of the entire field as a function of both space and time, which are shown in Figure 3-13 for calculations using current forcing for the plasma. The calculations with the spectral element discretization (Figure 3-13 a) predict higher peaks at the electrodes during the cycle and resolve additional sub-structure near the electrodes during the middle of the period. Three distinct peaks are clear in Figure 3-13 a, two broad peaks centered around a very sharp peak. Three distinct peaks are clear in Figure 3-13 a, two broad peaks centered around a very sharp peak.

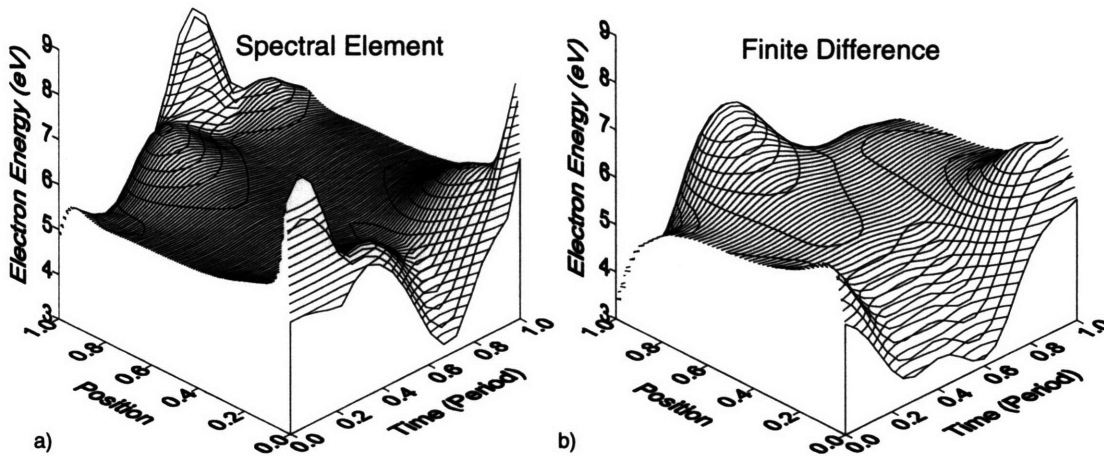


Figure 3-13 Spatial inaccuracy of the finite-difference approach when upwind-differencing is included. Average electron energy simulated using a) spectral-element technique; b) finite-differences with upwind differencing.

The finite-difference calculation (Figure 3-13 b) only shows one very broad peak, which is most likely a result of smoothing-out the three peaks evident in the spectral-element calculation. This structure is due to secondary electron emission during the portion of the cycle when the electric field is highest. Since the finite difference calculations do not correctly predict the very small electron fluxes away from the electrode during this phase of the cycle, these calculations are not expected to capture the highly peaked structure that appears in the spectral element calculations. The differences between the two approximations are made more apparent by plotting the maximum of the average electron energy as a function of time in the cycle, as shown in Figure 3-14 , where differences of 50 percent are seen at the beginning and the end of the period.

Even though the upwind difference approximation leads to large variations in the electron energy, the overall predictions of the physics of the discharge are largely unaffected by these approximation errors. Because the electron concentration is small at the point where the peak average energy occurs, the error in the total energy in the discharge is small and the peak in the ion density is changed only slightly: this difference is linked more closely to a larger loss rate of ions in the upwind difference approximation. The voltage is perhaps the best macroscopic measure of the discharge physics. It is obvious from Figure 3-15 that the peak voltage may be in error by as much as 15%, but the details of the voltage profile as a function of space and time are relatively accurate.

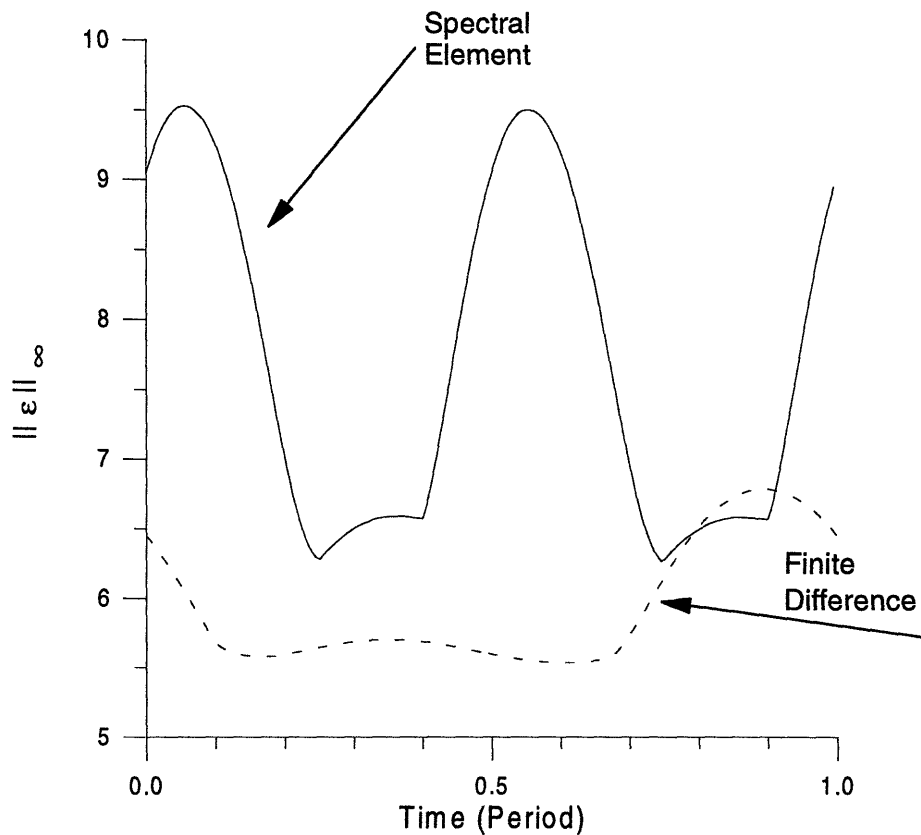


Figure 3-14 Spatial inaccuracy of the finite-difference approach when upwind-differencing is included shown in the maximum electron energy in 1/2 the discharge.

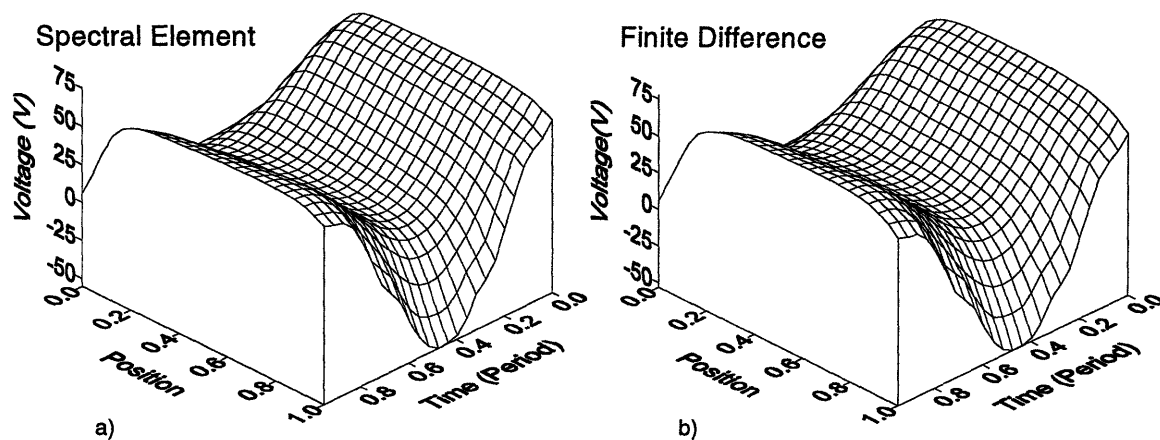


Figure 3-15 The global physics for finite-difference and spectral-element techniques compares favorably for less sensitive variables. Plasma potential simulated using a) spectral-element technique; b) finite-differences with upwind differencing.

3.3.4: Conclusions

Spectral element approximations for solution of a continuum model for rf-glow discharges are demonstrated here to be more accurate than simulations with traditional calculations based on upwind finite difference and low-order accurate finite element methods. Highly accurate simulation is necessary for modeling of the typical operating conditions for industrial plasmas, which are known to cause thin boundary layers near the electrodes. The exponential convergence expected of the spectral approximations is demonstrated.

Direct comparison between calculations with the spectral element approximation and the upwind finite difference method applied by Gogolides et al. (1992) demonstrates the need for the increased spatial resolution of the spectral approximation in order to resolve boundary layer structure that develops in the energy fields adjacent to the electrodes. Errors of over 50% are seen in the finite difference predictions of the average electron energy and details of the solution structure are unresolved. For the case of an argon plasma the inaccuracy of the finite difference method is not overly important to the global discharge physics, since the electron concentration is very small in the area where the energy is highest.

The cost is prohibitively high for either finite difference or finite element approximations with equivalent accuracy to the spectral element discretization. As discussed in Section 3.3.2,

the spectral element calculations are expected to be less costly by a factor of 150 for solutions that are accurate to one part in 10000. This difference is exacerbated when simulations are generalized to two spatial dimensions, as would be needed to model the geometries of plasma reactors used in the microelectronics industry.

3.4: Comparison to One Dimensional Experimental Results: Argon Plasmas

A great deal of attention has been paid to assuring that the physics in the continuum model is correctly maintained with the solution techniques used for this work. It is also very important to show that the equations we are solving have some connection with the real plasma physics; ideally the model should be fully quantitative. This section demonstrates that the plasma physics simulated using the continuum model is at least qualitatively correct. The two-dimensional version of the model incorporates more physics, and therefore yields more quantitative comparison.

A reactor configuration with a one inch gap and two 6.5 inch electrodes surrounded by a glass surface, with the remaining ground surfaces far away from the reactor, was used to compare to the one dimensional simulations. This configuration has a high enough aspect ratio that we expect the slender body approximation to apply over 70-85% of the electrode area. All simulations were run on the same mesh, and the voltage was varied until the solutions became unstable. The current is generally predicted to be approximately a factor of two too low, Figure 3-16. One explanation for this put forward by Gogolides et al. (1992b) that the positive ion mobility was incorrectly modeled in the sheath as a constant. The line in Figure 3-16 labeled 1 Torr modified uses the correct functionality for the ion mobility as Gogolides et al. suggested, but results in very little modification in the I-V characteristic. Most of the difference that was previously seen in the computations with the modified mobility was probably due to the inaccuracy inherent in the spatial discretization for the finite-difference methods.




Figure 3-16 Comparison of experimental and continuum model current-voltage characteristic for Argon Discharges at various pressures. The continuum model captures all the important trends, and is in reasonable agreement with the experimental measurements.

Since we know that at most 85% of the plasma is likely to be one-dimensional, a systematic error of more than 20% in the current is easily explainable. The comparison of the model results to the experimental data is better at lower voltage and lower pressure, where the discharge is more uniform, more one-dimensional, and noticeably fills the reactor more fully. Therefore, for the experimental results, the actual area that the plasma sees is not very well known, so calculating the current based on the electrode size probably accounts for the remainder of the differences in the results. Liu (1992) noted that the current does not scale linearly with the electrode area, which supports the supposition that the plasma is not maintaining a constant area. This is consistent with the current in the plasma being mainly capacitive.

The power deposition in the plasma is represented in Figure 3-17, and it is obvious that the model results match the experiments quite well. Since the power is mainly coupled into the electrons at the bulk-sheath interface, as well as some bombardment energy from the ions hitting the electrodes, scaling the power by the electrode area works properly. These are both resistive

losses, and the non-uniformity in the plasma sheath with varying radius tends to decrease the plasma density, and increase the resistance, resulting in power deposition which is more uniform across the discharge than the current to the electrodes.




Figure 3-17 Comparison of experimental and continuum model Power characteristic for Argon Discharges at various pressures. The continuum model agrees quantitatively with the experimental measurements. Note that there are no adjustable parameters contained in the model.

The power input to the plasma decreases with pressure, since the plasma density decreases, increasing the resistance less quickly than the current decreases. For the lower pressure plasmas, power into the ions in the sheath becomes more important, and increasing the voltage does not affect the bulk density as significantly as at higher pressures, where increasing power results in a large increase in the plasma density, as seen in Figure 3-18. In fact, it is obvious that the power and the density are intimately linked. The lower pressures have a totally different power deposition mechanism than at higher pressures, and Liu (1992) discussed in great detail the difference between bulk and sheath power coupling. Since it appears that the change of

power with pressure and voltage is modeled correctly with the continuum assumptions, at least for a first approximation the major power-coupling mechanisms are correct in the modeling.




Figure 3-18 Model density predictions for an Argon discharge at various pressures. Note that density varies almost linearly with voltage, and therefore with applied power.

It is not surprising to see that the phase difference between the current and voltage, shown in Figure 3-19, is correctly predicted, since the power has already been shown to match correctly. Correctly predicting the phase lag also indicates that the physics in the continuum model is reasonable, since the phase lag indicates the relative importance of the resistive and capacitive portions of the plasma impedance. It is difficult to sustain a discharge for very low voltages, as the plasma begins to extinguish. The model result for one Torr shows the extinguishing point for the plasma quite clearly, as the phase shift rapidly heads for the pure capacitive solution at -90° .




Figure 3-19 Comparison of the phase shift in an Argon discharge between current and voltage for experimental measurements and continuum model predictions at various pressures. The model predicts the experimental phase shifts quantitatively. Note that the 1 Torr simulation shows the turning point where the plasma is extinguishing.

For surface evolution simulation, the most important input parameters are the ion flux and the average ion energy, shown in Figure 3-20 and Figure 3-21. The experimental measurements are made with a three-grid retarding energy analyzer which is accessed through a pin-hole in the electrode. Measurement of the ion energy and flux is not as simple as measuring the current and power, and is quite easily capable of perturbing the discharge, although the analyzer design used here has been shown to be relatively non-invasive by Liu et al. (1990).




Figure 3-20 Comparison of experimental and continuum model Ion Flux for Argon Discharges at various pressures. The ion flux is not a strong function of pressure, and varies approximately linearly with power for both the experimental data and the model results.

The model predicts the flux quite well to within the experimental accuracy, while the model ion energy is considerably low. The breakdown of the model for ion energy is understandable, since the continuum assumption for the positive ion breaks down on the length scale of the plasma sheath. The ion energy is calculated by examining the velocity (mobility times effective field) at which the ions reach the electrode. However, the assumption that the ions can be described by a mobility flux in the sheath is quite poor, since ions can gain considerable energy between collisions, and will not be in thermal equilibrium except at the highest pressures. The continuum model does predict the general trends in the ion energy, and for most applications it is expected that a Monte-Carlo simulation of the sheaths will be necessary to correctly predict the ion energy. The trends in ion energy are still qualitatively captured by the continuum model, so if some model for a modified mobility were available for the plasma sheaths, the average ion energy could be correctly predicted.




Figure 3-21 Comparison of experimental and continuum model Ion Energy for Argon Discharges at various pressures. The continuum model is not capable of predicting the ion energy because there are not sufficient ion collisions in the sheath to maintain an equilibrium distribution.

It would be optimal if one could measure the electron density and the electron energy across the plasma, and indeed Langmuir probes provide a reasonable means for doing so. Interpreting the results of these measurements are not straight-forward, and the probe will not yield useful results in the plasma sheath, so for one-dimensional simulations the utility of Langmuir probe measurements in RF discharges is limited. Measurement of the light emission from the discharge does not perturb the plasma, and is easily interpreted if one examines only a single wavelength. The emission is a convolution of the electron density and the electron energy, so correctly predicting the plasma emission indicates that these variables are correct, although one is not able to state anything about each variable separately. For this study, the plasma emission is modeled as the convolution of the ionization rate with the relaxation time for the excited species. This model is used because the concentration of the excited species is not

modeled, and excitation and ionization should have similar cross-sections, albeit with different thresholds.

The results for the model plasma emission are compared with the experimental measurements in Figure 3-22. The plasma emission is measured both spatially and temporally resolved, and the model results are computed on the same basis. At 1 Torr, as shown in Figure 3-22 a) for the experimental data, and Figure 3-22 b) for the simulation results, the model predicts the trends for emission as both a function of time and space. The experimental results include an arbitrary offset for the baseline emission which is not added in the model, and is most likely due to the dark current of the photodiode and ambient light. The sheath thickness, as measured by the maximum in the glow intensity, is predicted to be about 50% larger in the model than is observed, but part of this error is due to the difficulty in imaging the plasma near the electrode, and the limited spatial resolution of the optical system.




Figure 3-22 Comparison of experimental and continuum model plasma emission for Argon a) Experimental result 1 Torr, b) continuum model 1 Torr, c) experimental result 100 mTorr, d) continuum model 100 mTorr.

The comparison with the continuum model breaks down at 100 mTorr, shown in Figure 3-22 c) for the experimental data and Figure 3-22 d) for the model results, and this can not be explained by the failure of the continuum model. At 100 mTorr, the model is at the edge of applicability, but should still produce reasonable results. The model ionization occurs almost completely in the center of the plasma, and the sheaths have seemingly collapsed. Examining Figure 3-22 d) carefully one can see that the peak in ionization occurs slightly off center, indicating that there is still a spatial separation between the peaks in the ionization. In fact, the maximum in the electron energy still occurs in the sheath, but the electron density is so small that the high energy has no effect on the ionization rate. The experimental plasma emission does show a shape similar to the model results in the bulk of the plasma. The breakdown in the model is most likely due to the difference in the thresholds for ionization and excitation. If the excitation is modeled analogous to ionization, except with a lower threshold energy, the effect of the high energy in the sheath becomes pronounced, and glow at the edge of the discharge will result.

This section has covered the simple one-dimensional physics which can be compared directly to experimental data. The comparisons are overall quite promising, and most of the differences between the experimental and model results are relatively easy to explain; the only inherent flaw in the model is the inability to correctly account for the ion energy at the electrode, but this result is not unexpected.

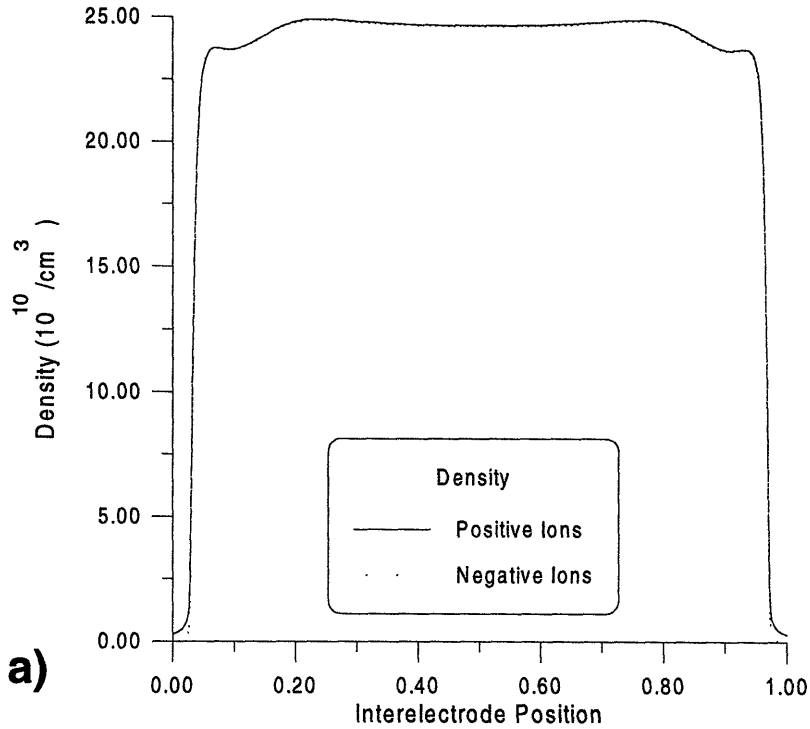
3.5: Electronegative Gases: SF₆

Since many of the gases used for microelectronic fabrication are electronegative, examining only electropositive discharge physics is only part of the necessary modeling effort. Fortunately SF₆ has been studied nearly as well as the noble gases and sufficient data is available to use for modeling. Note that SF₆ is on the other end of the spectrum from Ar and He, in that it is so highly electron attaching that most of the negative charge is held by the negative ions. The treatment of the electronegative gas is no different from electropositive gases, except for the additional equation describing the negative ion concentrations. In theory, solution of the continuum model for an electronegative gas should be equivalent in computational cost to


solution of the model for electropositive gases. In practice, this is not the case since the difference in the plasma physics, especially the thinner sheathes, makes the electronegative gas much more complex.

The negative ions behave similarly to both electrons and positive ions. Like the electrons, they carry a negative charge and therefore balance the positive ions to allow the bulk of the plasma to remain quasi-neutral. Like the positive ions, the negative ions are very massive and only respond to a time-averaged potential. The predicted concentration of the two ion species is shown in Figure 3-23 a). Neither species shows significant modulation in time, and except in the very thin sheath region, the two ions very nearly balance. Note that a much larger forcing voltage (300 V) is required to strike the electronegative discharge with respect to the electropositive discharge (30 V). The concentration of the ions is much higher for the electronegative discharge, although the higher density can be attributed to the higher applied power.

The electron concentration is not as simple in shape for SF₆ as the He discharges. The electron profiles show very strong variation in both space and time, Figure 3-23 b), especially in the sheath regions. The electrons are created during the phase of the cycle when the sheath is locally imparting energy to the plasma, and once the sheath is fully formed are quickly lost through attachment to form negative ions, or as current to the electrodes. This results in the interesting features in the electron concentration profiles. During the anodic phase of the cycle, all the electrons in the sheath are depleted for a large portion of the discharge, which is significantly thicker than the region holding most of the space-charge. During the cathodic phase when the electrons are being formed they are created in this region and have a steep concentration gradient much closer to the electrode, balancing out the positive charge since the negative ions can not move into the sheath on this time-scale.



a)



b)

Figure 3-23 Species concentrations in an SF₆ discharge. 1 Torr, 300 V, 1 inch gap, 13.56 MHz, at various phases in the cycle for a) positive and negative ions; b) electrons. Note that the concentration of negative and positive ions is nearly equal except in the sheath regions, and that the electrons are two orders-of-magnitude lower in concentration.

The concentration profiles for all the variables are no longer sinusoidal, as is generally seen in electropositive discharges, because the mechanisms for creation and loss have changed quite significantly. The addition of the negative ion equation adds a homogeneous loss rate for the positive ions, which is balanced by the ionization rate. Electrons are mainly lost to attachment in the bulk, and as current in the sheaths. The power coupling changes quite significantly for electronegative discharges, as is evident from the modulation of the electron energy in Figure 3-24. The energy does not exhibit a sharp peak at the bulk-sheath interface, although it does rise in the sheath during the cathodic phase to create a pulse of electrons. The energy is not pinned near the ionization threshold in the plasma bulk, as is the case in electropositive discharges, indicating that power couples into the electrons everywhere in SF₆, whereas power is coupled to the electrons mainly at the bulk sheath interface in electropositive discharge.

The bulk modulation of the electron energy shows that ionization in the bulk is not insignificant, and is expected to be a primary means of maintaining the discharge. This is also evident in Figure 3-25 a) showing the plasma potential which varies quite significantly across the discharge. In fact, using the typical definition of the sheath as a region where the potential rises rapidly has little meaning for this discharge, although Figure 3-25 b) shows that the fields in the sheath are significantly higher than in the plasma bulk. The most interesting feature of the plasma is hidden when looking at the potential, but is readily seen in the plot of the electric field. The field is never monotonic, and in fact has a reversal (due to the accumulation of electrons) right after the cathodic part of the RF cycle. It is this field reversal which traps electrons and allows them to form a large number of negative ions at the bulk-sheath interface. Note that the negative ions don't see the field reversal, since they are responding only to the time-averaged field.

It has been shown that the simulations for electropositive gases predict the discharge to be mainly capacitive, and it is obvious from Figure 3-25 that although the sheaths are still capacitive and have a significant field, the bulk of the discharge also has a large potential drop. It is obvious from Figure 3-26 that the current is very close to a sine wave with zero phase-shift, which is the form of the applied potential. The current in the plasma bulk is now much more complex than in the electropositive case, with a large displacement current in the bulk of the plasma balanced by the ion currents, and the electrons still carry the majority of the current. For SF₆, the current at the electrode resembles the electropositive case except for a phase shift toward

zero degrees. The displacement current comprises the majority of the current in the sheath, with a small electron current being allowed through the sheath at end of the cathodic part of the cycle. In the sheath, the negative ion concentration is negligible, while the positive ion current is small and nearly constant.




Figure 3-24 Electron energy at various phases in the cycle for an SF₆ discharge. The energy is highly modulated in the plasma bulk, and is not limited to the ionization potential, as opposed to the energy in electropositive discharges.

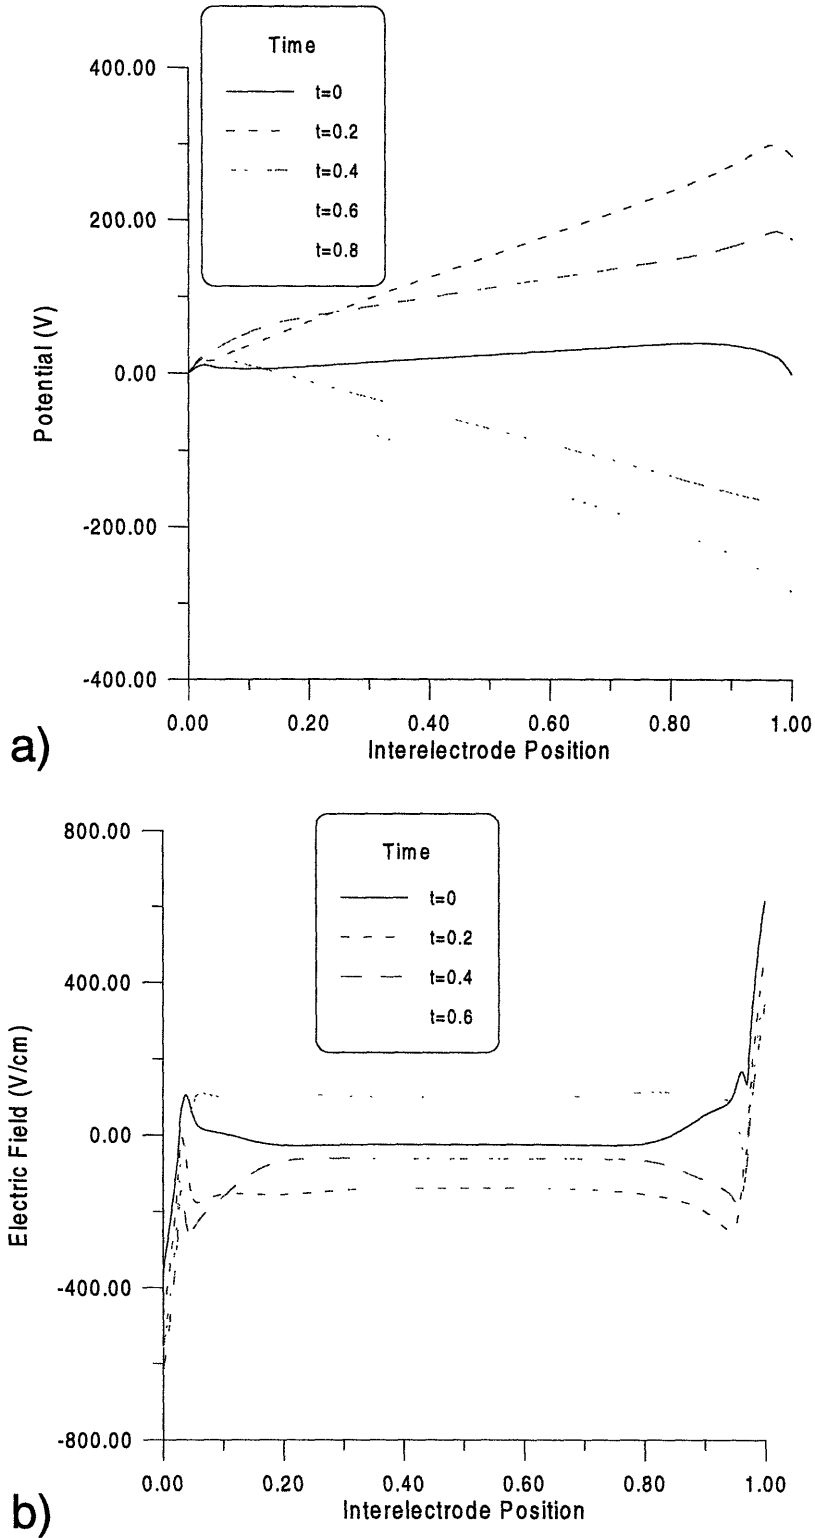


Figure 3-25 Potential variation in an SF₆ discharge. a) potential and b) electric field profiles at various times in the cycle. The bulk of the plasma is highly resistive because the negative ions are much heavier than electrons, and energy can therefore couple to the plasma both in the bulk and in the sheath.

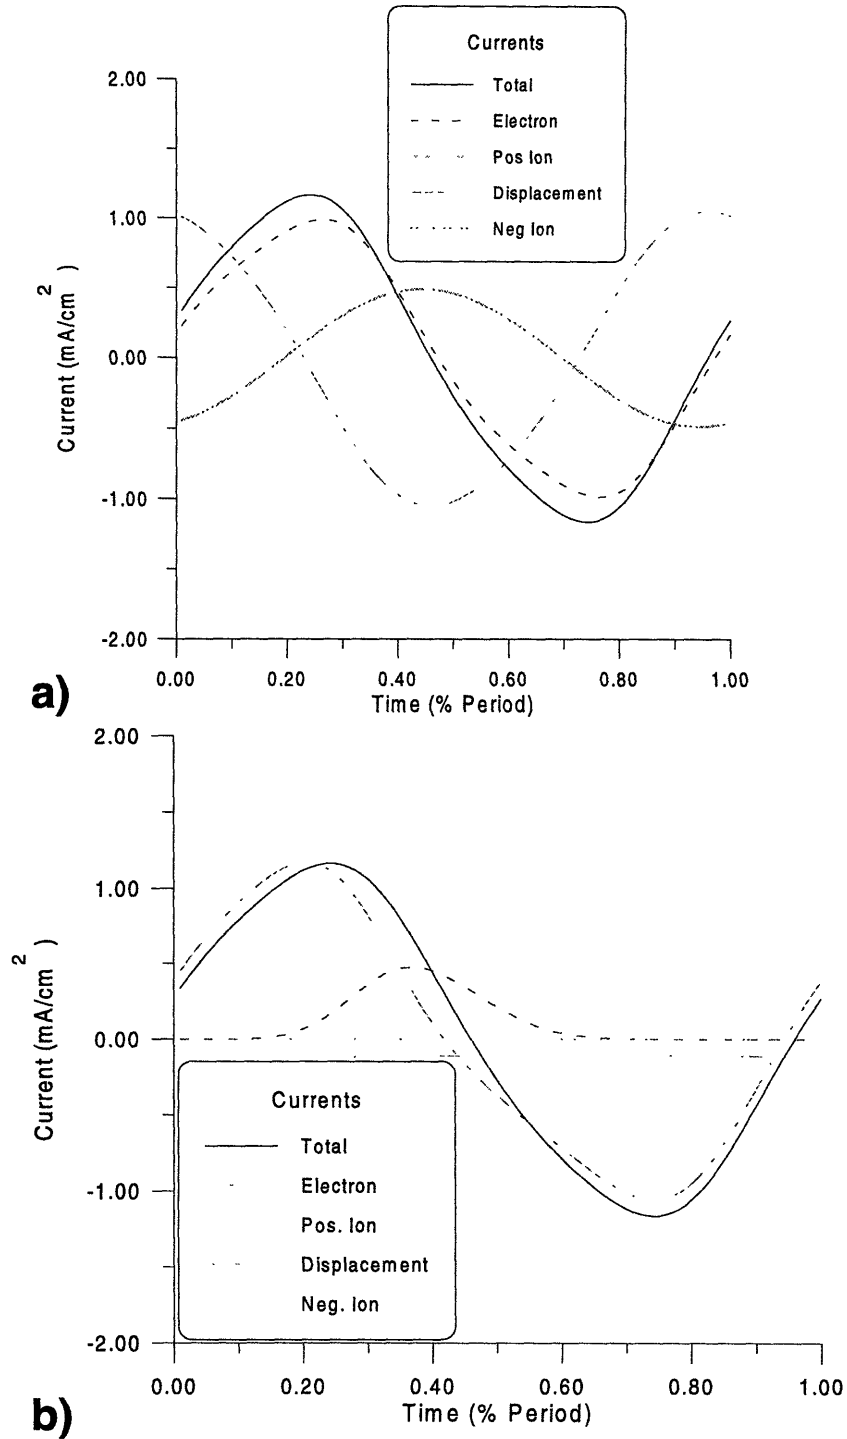


Figure 3-26 Components of the current in an SF₆ discharge. 1 Torr for a) the center of the discharge and b) the powered electrode. The bulk current is still mainly carried by the electrons, despite the reduction in concentration, while the ion current is balanced by the displacement current.

The variation of the current components across the discharge is also quite complex compared to the electropositive case, Figure 3-27, but this is expected from the spatial variations

in the electric field. The total current across the discharge is shown to be constant, as is required by current continuity for a one-dimensional discharge. The bulk of the discharge contains a small, but non-negligible displacement current which is balanced by the ion currents, and most of the current throughout the bulk of the discharge is carried by the electrons, even though they are the minority charge carrier. The change in sign of the electric field within the sheath requires a non-monotonic displacement current. The displacement current is balanced by the ion currents, with the positive ion current even changing sign in the sheath (or at the bulk-sheath interface) at this phase in the cycle.

The physics of the SF₆ discharge are quite different than for He or Ar, but there is nothing inherently difficult to simulate in the profiles presented up until now. The thinner sheaths merely mean a small increase in the mesh density, as is required for the higher fields and field reversals. In fact, it should be easier to simulate SF₆ than the electropositive gases because the electron concentration does not vary as much in the plasma sheath, and the energy is not sharply peaked in the sheaths. However, these simulations are quite difficult to perform. Even though the electron balance appears more facile to solve in SF₆, the negative ions present an unexpected problem.

Since the negative ions are massive, they should respond only to a time-averaged effective field, and not to the instantaneous potential across the discharge. However, since the field is so high in the sheath and the density gradient so steep, the negative ions are modulated on a very small length scale as shown in Figure 3-28. Although this result should not be unexpected, since the negative ions are not infinitely massive and have to respond to the applied field on some length and time scale, the strong modulation of the negative ions near the electrode is problematic. The full solution for the ions, including the detail of how the concentration approaches zero, is required in order for the spectral-element method to remain stable.

For the abrupt change in slope evident in Figure 3-28, a very large number of elements need to be concentrated in this region, even though the physics preserved is probably of very little importance to the overall simulation. The requirement for small elements also placed a smaller limit on the maximum time-step allowed for the simulations, making the overall cost extremely high compared to Ar and He simulation. The magnitude of the negative ion density near the electrode is small and the overall contribution to the charge balance and current balance must therefore be proportionately small, so the overall importance to the simulation is

demonstrably small as well. However, there is no simple change in the negative ion boundary condition which removes this oscillating boundary layer. It is conceivable that some reasonable change in the physics, such as addition of a field-dependent mobility, would remove the difficulty with the negative ions. It was decided that this work would remain rigorous to the plasma physics, and therefore two-dimensional simulations were limited to electropositive discharges.




Figure 3-27 Current through an SF₆ discharge. 1 Torr near the peak in the current phase (20% of the RF cycle). The ion currents are nearly equal except in the sheath, and are due to ion drift. Electrons carry the majority of the current in the plasma bulk.




Figure 3-28 Negative ion density for an SF₆ discharge near the grounded electrode. Expanded view at varying phases in the cycle. Note that the negative ions respond to the highly modulated field in the same way that electrons do, but at a much smaller length scale. Capturing the negative ion sheath requires very fine meshes, resulting in excessive computation time.

Chapter 4 : Two-Dimensional Modeling Results

4.1: Description of Two-Dimensional Geometry

Two different geometries were used for both the experimental measurements and the simulation results. The geometries are termed symmetric, because the powered and grounded electrodes are the same size, and asymmetric, with a smaller grounded electrode. The two-dimensional effects are limited to the perturbation from the insulating boundary for the symmetric cases, and are a convolution of the insulating boundary and the difference in electrode areas for the asymmetric cases. For standardization, all simulation results and experimental measurements are presented on the same basis. Three different types of graphics are used: vector plots, contour plots, and surface plots. For all plots, the bottom and top axis represent the grounded and the powered electrode, the left axis represent the insulating wall, and the right axis is the center of symmetry, shown in Figure 4-1. For the asymmetric simulations, the smaller

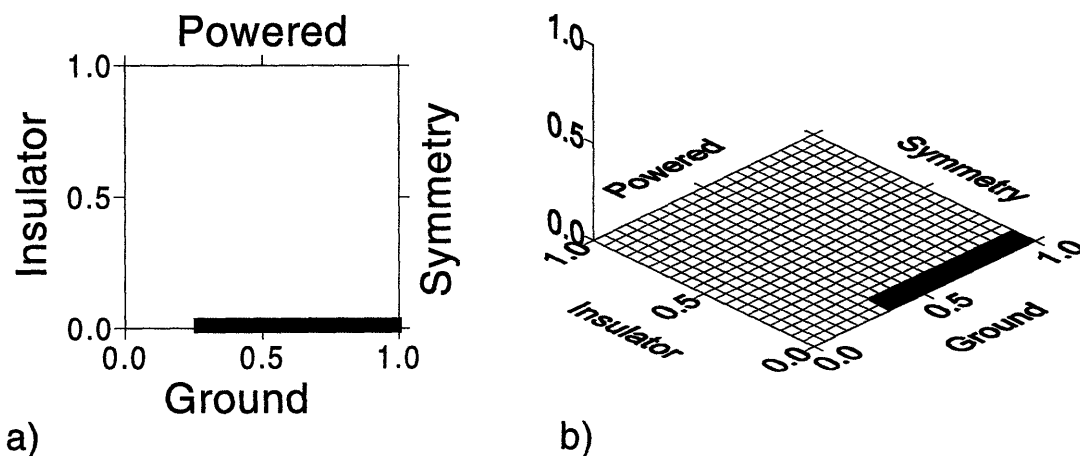


Figure 4-1 Geometry for presentation of simulation and experimental results a) vector and contour plots b) surface plots.

grounded electrode starts at the center of symmetry and extends part way to the insulating boundary, as represented by the dark bar in Figure 4-1. The rest of the lower boundary is an insulator which extends for a few gap spacings to another ground plane parallel to the grounded electrode. The exact geometry used for the simulations, including the regions which do not contain plasma, is shown in Figure 4-2. Note that the model infinitely large electrodes by forcing the radial field to vanish at the boundary labeled 5 in Figure 4-2, and that the insulating surfaces

labeled 7 and 8 are allowed to charge up self-consistently. The discretization in the confinement insulator is taken to be coarse since the fields in this region are not of primary importance.




Figure 4-2 Simulated regions for a) symmetric and b) asymmetric geometry. The regions in the figure are as follows: 1) Powered electrode, 2) grounded electrode, 3) Confinement insulator, 4) radial symmetry centerline, 5) edge of insulator, set to symmetry to simulate infinite electrode, 6) plasma region, 7) insulating wall confining plasma, and 8) insulator adjacent to smaller electrode to enforce asymmetry.

For convenience, all of the simulation results are plotted as non-dimensional variables, except when directly comparing to experimental data. The characteristic length is chosen to be one inch, so all coordinate axes are equivalent to physical dimensions in inches. All plots are produced maintaining the correct aspect ratios between the radial and axial dimensions. The lower axis indicates the distance from the insulating (glass) boundary, and the radial position is this number subtracted from the total radius of the reactor. For vector plots the arrowheads are of constant dimension with the length of the vectors proportional to the its magnitude. Time is always referenced with respect to the percent phase in the RF cycle, so that a time of 0.5 is equal to 180° , or π radians. The brightness of the boundary between shaded regions for the PIE plots is proportional to the emission intensity, while the shading follows the opposite trend (darker shading is more intense emission).

For the one dimensional case simply setting a tolerance on the TPSS residual norm is sufficient to achieve convergence. For the two dimensional cases Determining when a simulation has converged to a TPSS is more difficult, since variables such as the electron density

vary over six orders-of-magnitude both in space and time. The difference occurs because the sheaths intersect at corners (for example, the boundary between 1 or 2 and 7 in Figure 4-2), resulting in a greatly reduced density. Although the exact magnitude of these variables is not important to the overall simulation, they may have relatively large TPSS residuals. The effective field is high in this region, and converges slowly to a TPSS, although this has a negligible effect on the simulations due to the change in the ion boundary condition. It was noted that, for reasonable initial guesses, the TPSS residual and the structure of the solution is mainly defined after 3-4 TPSS Newton iterations. The residual norm does continue to decrease, at a much slower rate, but the changes in the solution are unimportant to the overall physics, and generally fall well below the error inherent in the model.

Strict convergence criteria were used where possible. From the one-dimensional simulations it is known that there will always be a minimum in the TPSS residual for a given timestep and discretization, due to inaccuracy in calculating both the TPSS Jacobian and the TPSS residuals. For the two-dimensional convergence results in the next section, the initial profiles are already very close to converged, and are allowed to 3 full Newton shooting steps to reach a steady state. The result is that for all simulations the TPSS residual reaches its minimum value, which is different for each grid. The effect of this criteria on the resultant TPSS is discussed further with the convergence results. For the symmetric simulations the TPSS residual norm falls below 10^{-2} within a few Newton shooting steps (usually 3-4), and further simulations change the TPSS residual very slowly. All symmetric simulations are converged to a final TPSS residual of less than 10^{-3} . The resultant bulk plasma quantities are then converged to an error of less than 1%, and have TPSS residuals on the order of 10^{-8} . All of the large residuals are in the sheath where the simulation results are not sensitive to the exact values of the variables. For example, reducing the TPSS residual from 5×10^{-3} to 8×10^{-4} in the 0.1 Torr simulation resulted in no more than a 5% change in the ion flux to the electrodes, and no significant change in the overall plasma physics (peak density, current, power, and other important parameters). For the asymmetric simulations, the singularity in the field at the insulator/electrode boundary results in a high effective fields which converge slowly. The convergence criterion for these cases is relaxed to 10^{-2} for this reason, and all variables not near the insulator/electrode boundary are much more strictly converged. This is further discussed in the section on the asymmetric simulation results.

4.2: Convergence Results in Two Dimensions

Since the major contribution of this work is a set of rigorous computational tools for solving the continuum model, it was necessary to study the convergence properties of the two-dimensional simulations to assure that the simulations were accurate. The geometry chosen for this study was a 1 inch by 1 inch cylindrical plasma with equal area electrodes and an insulating radial boundary with a pressure of 1 Torr, frequency of 13.56 MHz and 100 V_{pp} applied potential. It is expected that the corner where the insulating boundary and the electrodes meet will introduce discontinuity in one of the higher derivatives for many of the variables, so spectral convergence is not guaranteed; the solutions are not expected to be infinitely differentiable. This is especially true for positive ions, since the charging of the insulating boundary leads to a reversal in the effective field near the corner (where boundaries 1 or 2 meet boundary 7 in Figure 4-2). The change in sign of the effective field necessitates a change in the boundary condition for the positive ions to exclude aphysical leakage of ions into the plasma from the corner. The original boundary condition is changed from zero gradient to weakly imposed zero concentration:

$$\bar{\mathbf{n}} \cdot \nabla N_+ = 0 \quad \text{becomes} \quad |\bar{\mathbf{n}} \cdot \nabla N_+| + |CN_+| = 0, \quad \text{Eq. 22}$$

where C is a large constant which sets the ultimate magnitude of N₊. The change in boundary conditions on positive ions essentially introduces a discontinuity in the ion fluxes near the corner, and may destroy spectral convergence of the simulations.

Simulations were performed on a non-uniform grid of 17 axial and 9 radial elements, with a greater concentration of elements in the radial and axial sheaths than in the bulk region. A solution was already available for 4th order polynomial basis functions, and was interpolated to provide a guess for the other polynomial orders. Convergence was tested by examining the results for 3rd, 4th, 5th and 6th order basis functions, using 8th order basis function results as the exact solution. The number of degrees-of-freedom for each simulation is shown in Table 4-1. Simulations were run for 3 Newton iterations, which results in a minimum in the TPSS residual for all grids. Since the TPSS Jacobian was computed using a large time step, and since there is significant error in the TPSS residuals from time integration over one cycle, quadratic convergence of the TPSS shooting method was not observed. It was noted that the minimum error in the TPSS residuals was the smallest for the lowest order basis functions and higher for higher order basis functions. The increase in error is attributed to the increase in the Courant

number within the element as the order of the interpolating polynomial increases, decreasing the minimum separation of nodes. Convergence of the TPSS shooting method with time step for the integrator or for the calculation of the TPSS Jacobian was not explored due to the expense of calculating solutions with very small time steps. These calculations scale linearly with the timestep, and would require a minimum of the cases requiring 2,4,8 and 16 times the cost of the base case simulations shown in Table 4-1.

Table 4-1 Computational requirements for convergence calculations.

	Degree of Polynomial				
	3	4	5	6	8
DOF	9645	16872	23579	37412	66024
CPU time (hours)	0.2	0.5	1.4	2.5	6.6

The convergence of the spectral element method is tested by plotting the L_2 norm of the error in the variables versus the degree of the interpolating polynomial; spectral convergence is indicated by a straight line. Spectral convergence is achieved for the two-dimensional simulations, despite the corner discontinuity, as shown in Figure 4-4. Two different error norms are presented: the first contains all of the plasma variables, represented by the filled circles, and the second ignores errors at the corner elements where a discontinuity is expected. Exponential convergence is only guaranteed for sufficiently fine meshes. For the 3rd order polynomials the mesh is not fine enough to achieve spectral convergence, and the result is a high error. Both norms are exponential for the 4th, 5th and 6th order polynomials.

If one considers all of the variables, the discretization error decreases very slowly with increasing polynomial order, although spectral convergence is achieved. The second norm ignores the effect of the discontinuity in the corner, and shows a greatly enhanced convergence rate. Indeed, examining where the largest errors are shows that the effective fields in the corner have by far the highest errors, and overall 90% of the error is in 1% of the equations in the corner elements. A high error in the corner is indicative of the discontinuities there, and suggests that local refinement of the elemental mesh would be much more effective than increasing the polynomial globally. Indeed, by examining the contours of positive ion density in the corner in

Figure 4-3, it is obvious that there is a sharp drop in positive ion density which is limited to the corner of the discharge. This is a direct result of the boundary condition we are using, and if one does not wish to resolve this physics, it is possible to change the physics in such a way to smooth out or delete the change in ion boundary condition.




Figure 4-3 Positive ion density contours near the corner of the discharge where the insulating surface meets the grounded electrode. Note that the spatial scales and the ion density scales are multiplied by 100 to show the fine detail.

The discontinuity produced by the ion boundary condition propagates into all the other variables through the coupling of the boundary conditions, and through Poisson's equation. However, the effect of the discontinuity is very localized. The error in the electron energy, which is a very sensitive plasma variable, illustrates how the overall convergence is affected. Three different energy norms are plotted in Figure 4-5: the energy at the corner of the discharge where the discontinuity occurs, the error for the maximum in the energy, and the error for the corner element, at the node furthest from the corner of the discharge. Except for the energy in the corner, all the other norms on the energy appear to exhibit spectral convergence. The error in the maximum energy is higher because it is not normalized to the magnitude of the energy, and the maximum energy is a factor of 4 higher than in the corner.




Figure 4-4 Convergence of the Two-Dimensional Simulation with p-Type Refinement. • All simulated variables, + 2 error from corner elements ignored. The majority of the error is contained in the nodes near the corner discontinuity, and this error is not affected by using higher-order interpolating polynomials.

It is interesting to note that the discontinuity does not affect the convergence of the energy on the node in the first element diagonally opposite from the discontinuity, although all the variables in the first element are coupled through the spectral discretization. Some of the error in the corner variables is attributable to the TPSS shooting method, since the variables which have the highest final TPSS residual are always the ones near a corner of the discharge, and the spatial and temporal errors may therefore not fully decouple.




Figure 4-5 Convergence of the electron energy with mesh refinement at different spatial positions within the plasma. Except at the corner discontinuity, the energy converges spectrally.

4.3: Extent of Two-Dimensional Perturbations

Both the one and two dimensional models are available, and a comparison between the two is of interest to understand when the much simpler one-dimensional model will be of use. This is best done by examining the positive ion concentrations since they are neither steep nor changing in time on the scale of the RF cycle. The positive ion density for the one dimensional simulation is compared in Figure 4-6 to the density along the radial line of symmetry for 3 different two-dimensional simulations which are different only in the width of the plasma. The two dimensional simulations were performed with 1, 2.25 and 4 inch radius and a 1 inch gap, giving an aspect ratio for the plasma of 1, 2.25 and 4. There is very little difference in all of the simulation results, indicating that even at 1 inch radius the plasma behaves close to one-dimensional near the radial center of symmetry.




Figure 4-6 Positive ion density for two-dimensional simulations. Contours along the centerline for aspect ratio 1, 2.25 and 4 plasmas compared to the results from one-dimensional simulations.

The profiles for the two plasmas with aspect ratios of 2.25 and 4 are nearly identical at the centerline, indicating that the effects of the insulating wall no longer perturb the solution above a certain radius. In fact, both two-dimensional simulations match the results for the one-dimensional simulation nearly exactly. However, the simulations of a plasma with aspect ratio 1 exhibit nearly a 10% increase in the bulk density. This effect is easily explained by looking at the full two-dimensional density plots shown in Figure 4-7. All of the two-dimensional simulations exhibit a maximum in the density not at the centerline, but near the insulating wall, due to an enhancement in the ionization rate from the presence of a radial electric field. Therefore, the density increases with increasing radius at any fixed axial position. The loss rate of ions through diffusion is lower in the two-dimensional simulation because of the second spatial degree of freedom, since it is partially balanced by the radial flux. Since the one

dimensional simulation cannot model the radial ions flux, the centerline density must decrease as the aspect ratio increases.




Figure 4-7 Effect of plasma aspect ratio on continuum model results. Positive ion density for 3 different aspect ratio plasmas, with constant electrode gap spacing. The perturbation due to the confining insulator is similar in all cases.

Even in the simulation of the widest two-dimensional plasma, there is a non-zero radial gradient in the positive ion density, although the gradient is very small. The two-dimensional spatial effects are easily seen by looking at the positive ion density at the center of the gap as a function of the radius, Figure 4-8. As the plasma becomes wider, the density at the center of symmetry decreases slightly, since the gradient has become closer to vanishingly small. This appears as a small difference between the density at the center of the widest two-dimensional simulation and the one-dimensional simulation. The peak in the ion density also decreases, more dramatically, with the width of the plasma. For the aspect ratio 1 simulation, the radial density

gradient is large enough that the peak density increases by 10%. However, at the center of symmetry, the ion density is much closer to the prediction from the one-dimensional simulation. It is also apparent that for the majority of the plasma assuming a one-dimensional solution results in a very small error. The solution near the insulating wall is identical in every case, suggesting that the plasma may be divided into 3 radial regions: the one-dimensional uniform region, a region with a thickness on the order of 1/2 of the plasma gap-spacing due to perturbations from the insulator (a radial sheath due to ambipolar diffusion), and a matching region for the two solutions which has a thickness on the order of the plasma gap-spacing.




Figure 4-8 Changes in the structure of the radial sheath with aspect ratio. Positive ion density as a function of the radius of the simulated plasma. The simulations approach the one-dimensional solution one gap-spacing from the insulating wall.

The two-dimensional effects for these 1 Torr plasmas do not extend very far radially into the plasma. The critical dimension appears to be the gap spacing, since the perturbations die out close to one gap spacing radially into the plasma. Even the aspect ratio one plasma is very close

to one-dimensional at the centerline, and the majority of the perturbation occurs for all the simulations within a 1/2 gap spacing of the insulating wall. Examining the density contours for a 1/2" gap with 2.25" radius (aspect ratio of 4.5), as shown in Figure 4-9, it is apparent that the figure of merit is indeed the aspect ratio. The perturbation due to the insulator extends almost exactly 1/2" into the plasma radially, or $O(1)$ of the gap spacing. For a simulation with 2" gap and 2.25" radius (aspect ratio of 1.125), the radial perturbations extend the entire width of the plasma, again demonstrating the scaling of the perturbations with aspect ratio. In this high gap plasma the physics have changed enough that there is no longer a radial maximum in the ion density and a different radial sheath structure is present.




Figure 4-9 Effect of gap spacing on two-dimensional simulations. Positive ion density for an Argon plasma at 1 Torr with 1/2" gap spacing. The density profiles are similar to those in the 1" gap case at a similar aspect ratio.

It is interesting to note that the perturbation for the 1/2" gap is almost exactly the same as for the 1" gap at the same aspect ratio. The functionality in the radial direction for the slender plasmas is similar enough to suggest that the radial density functions are separable from the axial density functions. A one-dimensional model could then be used to predict the density across the gap, with the radial variation expected to be similar for all the simulations.

4.4: Symmetric Plasmas

The following two sections present a synopsis of the physics for the two-dimensional simulations using equal area electrodes. The general plasma physics have already been discussed in section 3.1, so the focus of these sections is how the insulating boundary perturbs the one dimensional solution, and how the simulation results compare to the experimental measurements.

Here the geometry includes an insulating portion at $x=0$, with dielectric constant 3, which completely fills a gap between the electrode extending 2.5 gap lengths beyond the plasma as illustrated in Figure 4-2 a). Charge accumulates along the insulating surface, resulting in a jump in the electric field between the insulator and the plasma. A small surface diffusion coefficient was added into the charge balance to increase numerical stability for the TPS Newton iterations, but has no noticeable effect on the overall simulations.

4.4.1: Explanation of Plasma Physics

The base case simulation for the results described in this section are from a two-dimensional simulation with a 1" gap spacing, a 2.25" diameter, 1 Torr pressure, 13.56 MHz forcing frequency with the transport and rate parameters for Ar. The secondary electron emission coefficient was set to 0.1 to simulate an oxidized surface, and the secondary electron energy was arbitrarily set to 1 dimensionless energy unit (5.93 eV). The standard grid consisted of 4th degree polynomials using 23 axial and 12 radial non-uniform elements, with the greatest concentration of elements in the sheath regions.

One change in the boundary conditions was found to be necessary for all two-dimensional simulations. Since the accumulated charge on the boundary leads to a negative time averaged potential in the corner, the effective field close to the corner of the discharge points away from the electrode for at least some portion of the RF cycle. The standard boundary condition for ions results in a ion flux leakage from the corner into the plasma. The most obvious change in boundary condition would be to set the flux to zero at this point, but this was found to be numerically unstable for typical cases. A weakly enforced zero concentration on the wall was used instead. Although this condition results in a stable simulation, the change from zero gradient to zero concentration introduces a discontinuity in the positive ion boundary condition which is noticeable in the convergence results.

The current and voltage waveforms for the plasma are one measure which is readily available in most plasma systems. However, this measurement is inherently zero-dimensional, and therefore is not likely to be a good indication of perturbations due to the more complex geometry in the symmetric two-dimensional system. Indeed, examining the components of the current and the voltage, shown in Figure 4-10, there is very little difference between the two-dimensional simulation for Ar and one-dimensional simulation results. The discharge current is

almost completely capacitive, with a small electron current flowing to each electrode right after the cathodic part of the RF cycle. The ion current is very small and is nearly constant in time.




Figure 4-10 Current components and voltage for two-dimensional Ar plasma simulation at 1 Torr.

Although there is very little information in the current examined in this way, it is instructive to examine how the current varies across the electrode as one approaches the insulating boundary. The current density as a function radius divided into 4 sections is shown in Figure 4-11. The current density decreases by more than a factor of two as the insulating boundary is approached, and the phase shift increases, indicating that the plasma is becoming more capacitive. This trend is expected since the ion density is slightly peaked near the insulating wall, and rapidly decreases as the insulator is approached. The electron density must also decrease at least as fast to limit the electron current to the surfaces. The result is both an increased sheath thickness with radius, and a decrease in the plasma conductivity. The increase in plasma impedance leads to a lower current as the radial wall is approached, and since the bulk of the plasma is disappearing as the resistance increases, the plasma has a higher capacitance near the insulator.




Figure 4-11 Current density in a two-dimensional discharge at different radial sections. Note that the current is more capacitive close to the confining insulator, and the current decreases as the insulator is approached.

The flow of the total current through the plasma, Figure 4-1 , illustrates quite well how the radial insulator perturbs the plasma physics. At 75% of the RF cycle, the upper electrode is a momentary anode, and the current flows toward the cathode. For most of the discharge width, the current is essentially one-dimensional. Near the insulating boundary the electrons flow either inward or outward radially depending on the magnitude of the electric field. The bounding insulator is perturbing the electron flow in this area in order to keep the plasma confined. The steep radial gradient in electron concentration allows the electrons to diffuse against the field near the lower electrode, while the high electric field near the upper electrode reverses the electron flow. The current will flow in the opposite directions at 25% of the RF cycle.




Figure 4-12 Total current in an Argon plasma at 2 phases in the RF cycle. a) $T=0.75$, b) $T=1.00$. The electron current is repelled from the insulator when the electric field is high and the current is low. For the majority of the RF cycle, the current appears mainly one-dimensional.

At the beginning of the RF cycle, the electric fields have mainly collapsed, and the potential on both electrodes is zero, while the insulating boundary has a complex potential distribution due to the competing effects of charging and field penetration. The overall effect is that the insulator barely perturbs the current flow, and the current is essentially one-dimensional everywhere. The electric field are shown at the same times in the RF cycle in Figure 4-13. In Figure 4-13 a), the field is a maximum at the upper electrode, while in Figure 4-13 b) the magnitude of the field is relatively small everywhere. Small changes in the electric field do have the capability of producing large changes in the current, since the drift component of the flux is in close balance with the diffusive flux. Figure 4-13 c) shows the time averaged field in the plasma, showing that the electrons are confined by the plasma sheath both radially and axially. There is a point of zero average field which corresponds to the maximum in electron and ion density. All charge carriers, on the average, move away from the density maximum which was described earlier. The maximum in density must occur because the interaction of the radial and

axial sheaths produces a point where there is zero time-averaged electric field at which the electrons accumulate.




Figure 4-13 Electric fields for a two-dimensional Ar plasma. a) $t=0.75$, b) $t=1.00$, c) time average. The fields are always highest in the sheath regions. The time averaged field clearly shows the regions where the fields vanish, allowing electrons to accumulate.

The effect of the radial/axial sheath interaction is more pronounced in the case where the gap spacing is wider and the overall density is lower, as shown in Figure 4-14. When the upper axial sheath is nearly fully formed, as in Figure 4-14 a), the current forms what appears to be a vortex close to the insulator. Again, as in the case with the higher aspect ratio plasma, the insulator is mainly acting at this point to confine the plasma, and limits the flux of electrons to the bounding insulator surface. Since the aspect ratio is high for this case, the effect of the insulator is greatly pronounced. However, when neither sheath is dominant, as in Figure 4-14 b), the current flows directly between the electrodes, and is almost completely capacitive.

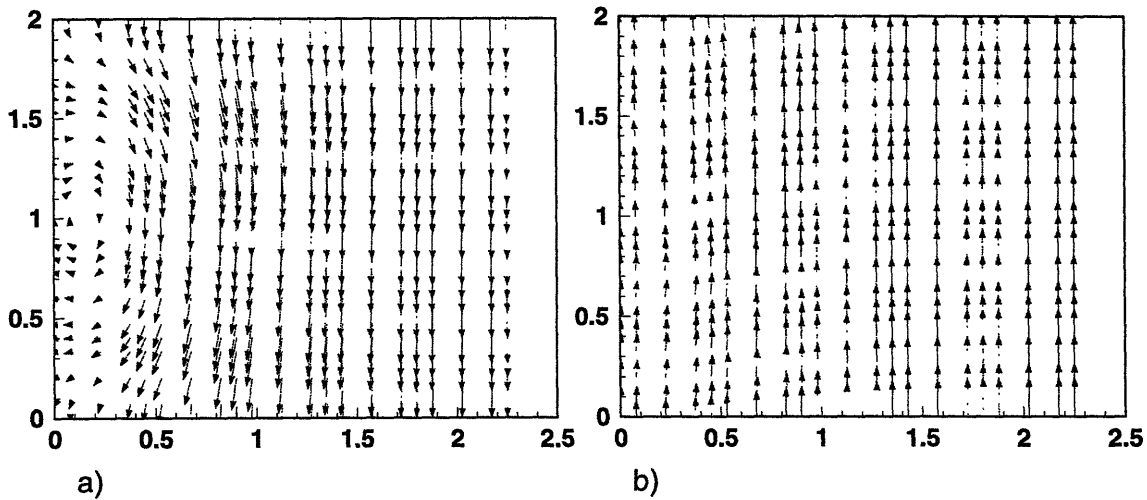


Figure 4-14 Total current for an Ar plasma with 2" gap. a) $t=0.75$, b) $t=1.00$.

It is easy to see how the radial and axial sheaths interact by examining how the electron density varies in both time and space. The electron density is shown in Figure 4-15 for the phases of the RF cycle when the upper electrode has just reached ground potential and for the phase where the upper electrode is the instantaneous anode. In Figure 4-15 a) the density is nearly symmetric and the upper sheath is forming while the lower sheath collapses. In the same way, the electron density indicates that the radial sheath is thickest at the lower electrode, and thinnest at the upper electrode. The rounded corner of the electron density profile indicates that the sheath at the corner of the plasma is a combination of effects from the radial and axial electric fields. As the potential on the upper electrode drops negative, the sheath on the upper electrode builds up, and the sheath on the lower electrode becomes thinner. As the upper sheath builds, it

interacts with the radial sheath, and the electron density at the corner drops precipitously. It appears that the radial sheath is expanding and contracting along with the axial sheath. Another view is that the radial sheath is nearly constant in time, interacting with a time varying axial sheath which varies in thickness along the electrode surface. This view is reasonable since the radial sheath thickness is only a very weak function of space and time along the insulating surface, except at a distance of a few times the axial sheath thickness away from either electrode.




Figure 4-15 Electron density for symmetric Ar plasma. a) $t=0.50$, b) $t=0.75$. The electrons are forced from the axial sheath during the anodic phase in the cycle. The radial sheath varies along the insulating surface, both in space and time.

The interaction of the two sheaths results in an increased local electric field in an area where the electron density is both non-zero and constantly varying in time. For the conditions which are being simulated, the main source of energy input into the plasma is coupling from joule heating near the sheath boundary. One expects that the ionization rate should therefore increase near the corner of the discharge, which is indeed shown in Figure 4-16. This plot also shows that the peak in the electron density near the insulating wall cannot be attributed to a local

generation rate, since the maximum in ionization still occurs near the bulk/sheath interface. In fact, throughout the area where the density is highest, the ionization rate remains nearly zero. The increase in ionization rate near the corner will lead to an increased plasma glow. Since real plasmas are generally viewed through the entire discharge, this effect will be masked by the larger plasma thickness near the discharge center unless the ionization rate is substantially higher in the corner.




Figure 4-16 Time averaged ionization rate for a two-dimensional Ar plasma simulation.

When the gap spacing increases the effect of the radial insulator is expected to increase, as discussed in Section 4.3. The extent of the increase is clear in Figure 4-17, where the gap spacing for the plasma has been doubled. Since the forcing potential has been kept constant, the power, current and plasma density all decrease dramatically, and the sheaths become substantially wider. In fact, the radial sheath has become thick enough that it does not interact very well with the expanding axial sheath and the maximum in the electric field does not result in a maximum in the electron energy. The peak in ionization near the corners has disappeared, and the plasma is now sustained by ionization near the bulk/sheath interface near the radial center of the plasma.

The most useful quantity that the two-dimensional simulation produces for determining the uniformity of plasma etching is ion flux, which is necessary to predict the etching rate, along with the time averaged electric field, which gives an order-of-magnitude of the ion energy. The ion flux must be uniform across the electrode to produce radially uniform etching. The radial insulator has a considerably negative effect on uniformity, as shown in Figure 4-18. At the highest pressure examined, the main effect is a decrease in flux near the reactor wall, due to the

decrease in ion density in the radial sheath. As the pressure drops, the maximum in density near the insulator propagates toward the electrode, and the ion flux peaks near the edge of the electrode. The effect of the insulator becomes more pronounced as the pressure lowers, although the peak ion current decreases due to diffusion of ions. The result is that the edge of the electrode will etch highly non-uniformly. As the pressure decreases, the area where ion enhanced etching is non-uniform will expand, for a constant applied voltage.




Figure 4-17 Time averaged ionization for a 2" gap Ar plasma.




Figure 4-18 Ion Current to the powered electrode as a function of pressure. The current is a weak function of pressure, and is perturbed by the confining insulator. The extent of the perturbation is a weak function of pressure, and is mainly confined to within 1 gap spacing from the insulator.

4.4.2: Experimental setup for comparison to simulation results

As previously stated, the modeling effort is of relatively little use if it is not supported with comparison to a set of experimental data. Very little systematic work has been done to show that various versions of the continuum model correctly predict plasma physics, mainly due to the fact that very few one-dimensional measures of the plasma physics are readily available. The most readily available measures of the plasma physics are the current/voltage characteristic of the discharge, the plasma induced emission (PIE), and sometimes the electron energy distribution through Langmuir probe studies. This information is much more useful for two-

dimensional simulations, since the variations in both radial and axial positions can be examined. Two-dimensional measures are especially useful for the semi-quantitative measures that are generally available, since for PIE the absolute excitation rate is not known, and for Langmuir probe studies the theories to calculate density and plasma potential will give the correct shape to the curves but not necessarily the correct magnitudes.

A reactor with adjustable geometry was built by Liu (1992) to study the scaling laws for power and current as a function of reactor geometry. This reactor is a simple modification of the work of Gogolides et al. (1989), which was used to show that the continuum model is capable of predicting the temporal behavior of the PIE for electropositive and electronegative gases. A brief description of the reactor and the experimental techniques used in the next few sections are the subject of the rest of this section.

The reactor consists of a quartz shell which is sufficiently uniform for optical measurements. The quartz thickness is nominally 2.5 mm, and the outer diameter is 135 mm. Shells are fabricated for varying gap spacing, and one and two inch gaps were used for this work. The electrode is made of aluminum with an anodized coating, and is surrounded by a Teflon spacing ring in an attempt to confine the plasma, Figure . The top electrode has an array of pinholes forming a showerhead to introduce gas into the reactor. There is a gap of 5 to 20 mil between the Teflon and the quartz to allow pumping of the reactor, with the smaller gap allowing higher power but slower pumping speeds. The aluminum electrodes are 4.5 inches in diameter for the symmetric configurations, and the lower (powered) electrode is changed to 3 inches for the asymmetric configurations.




Figure 4-19 Electrode geometry for a) symmetric and b) asymmetric reactor configuration.

A signal of 13.56 MHz is applied to the lower electrode by using a function generator and a 250 kHz to 150 MHz broad-band amplifier. A matching network assures that the reflected power is minimal, and applied power is measured both by a Bird Power meter at the apparatus and at the amplifier. The current and voltage waveforms are also measured very close to the electrodes and collected by a digital oscilloscope, and can be used as a separate measurement of the applied power. The pressure is controlled by setting the gas flow on a mass-flow controller and the conductance to the pump using a throttle valve. Although no specific pressure control was used, the pressure was noted to vary no more than 3% over a 5 hour period.

Plasma induced emission (PIE) was collected using a 0.6 m Jobin-Yvon monochromator set to at 7503 nm imaged onto a silicon photomultiplier. A slit width of 300 μm and a slit height of 1 mm were used to assure proper spatial resolution as well as line resolution on the monochromator. Two lenses of approximately 12 in focal length were used to collimate the light between the reactor and the monochromator, and a series of simple mirrors were used for alignment. The alignment was manually adjusted by passing a He-Ne laser (6328 nm) through the monochromator in place of the photomultiplier, and was adjusted to better than 2 mm both vertically and horizontally. This also allowed for proper registration of the electrode and the glass to ensure that the measurements were at known positions. Spatial position was varied by keeping the optics fixed and using computer controls to move the whole reactor using a set of XYZ stepping motors. By performing realignment at the end of a set of runs, it was determined that the misalignment due to skipping of the stepping motors was less than 1 mm.

The Langmuir probe design used was similar to that described by Surendra (1985). A series of 10 tuned inductors (J.W. Miller 180 μH) were initially used, while later probes used a mixture of 10 of these inductor with 12 other tuned inductors of varying sizes (140,160,200, 220 μH). The impedance of the probes was measured to be 500 k Ω to 660 k Ω , at 13.56 MHz, and remained above 20 k Ω between 3 and 23 MHz. It is unclear whether or not this impedance is high enough to satisfy the criteria described by Godyak et al. (1992) to assure an undistorted probe characteristic. However, there was no significant difference in the results from different probes or probe tips. The inductors were soldered in series, attached to a gold feedthrough on one end and a solid wire on the other, then sealed in a 6 mm Pyrex tube using Torr seal. The reactor was modified by fusing a 10 mm ID quartz tube into the standard configuration to allow

the probe to pass through. A set of Teflon rings were constructed to assure that the space between the probe assembly and the inside of the tube was minimal. Vacuum integrity was maintained by using a Cajon VCO reducing union hand tightened onto both the tube and the probe assembly. This allowed the probe to move into the reactor or rotate (for two dimensional measurements) by loosening one of the seals. The probe tip was formed by spot welding a 3 mil tungsten wire onto a 20 mil soft stainless wire, and encasing these in a 20 mil ID Teflon sheath. A small amount of Torr seal was used to plug the end of the tube to ensure that the plasma was only connected to the exposed tungsten. The feedthrough was also covered by a Teflon sleeve and Torr seal to assure that no plasma electrons leaked in through the connection to the stainless wire. The probe tip was cut to a length of 3 mm, the stainless wire was bent to 90 degrees to allow rotation around the axis, and the tungsten tip was bent parallel to the electrode surface. New probe tips were conditioned by running at high positive voltage, attracting all the plasma electrons and causing the tips to glow, for at least 2 minutes.

The Langmuir probe signal was collected by ramping the voltage on the probe tip and measuring the collection current. The voltage was set by driving a voltage follower with a gain of 50 with a signal of -2.5 to +2.5 volts from the computer using an A/D board. The probe voltage at any given time was then known accurately, and could be ramped at the desired speed within the limits of the A/D board. The current was then measured using a transimpedance amplifier and the same A/D board to read the current signal. The output voltage to the Langmuir probe, and the measured current were examined on an oscilloscope to ensure that no signal pollution was occurring in the electronics. Typical scan rates were 3-10 Hz (for a total scan of 1024 points) and 50 scans averaged. One further modification was necessary to the reactor for Langmuir probe studies: a stainless steel disk was attached to the lower electrode to assure to a good DC ground. It was found initially that the anodization on the electrode prevented the flow of the DC currents necessary for proper probe measurements. With the stainless steel disk, it was also found that the probe could not be run into the electron saturation region, as indicated by a visible perturbation of the plasma. Langmuir probe scans were cut off at a positive voltage which would not perturb the plasma. Run to run stability of the probe was checked, and the plasma parameters were found to vary by less than 5% over the course of an hour.

4.4.3: Comparison to Experimental Measurements

It is necessary to use an approximate model to predict the PIE from the simulation results, since a balance is not included for the metastable states which produce light upon relaxation. The cross-sections for excitation and ionization are known to be very similar, although they have different thresholds. To predict the time dependent ionization rate, one would therefore use the ionization rate convoluted with the exponential relaxation of the excited state. For time averaged results it is sufficient to examine the time-averaged ionization rate. The measurements are made viewing the entire cross-section of the plasma, so comparison of the model results requires deconvolution of the measurement to the emission from a single plane. Since there is unavoidable noise in the experimental measurements, it is generally easier to integrate the predicted emission for the cross-section being viewed in the experimental setup. All the simulated PIE results have this taken into account. It should be noted that the experimental system is not exactly comparable to the simulated geometry, since there is a Teflon spacer between the glass and the electrode to contain the plasma. For comparison between the model and the experimental data, the spacer is treated as equivalent to the beginning of the glass wall.

The comparison between the simulated PIE profiles and the experimental measurements is shown in Figure 4-20 and Figure 4-21. For each pressure, the simulated emission is shown above the experimental PIE, and is displaced to the right to align the insulating boundary with the electrode in the experimental system. The comparison in the features between the experimental data and the simulation results is excellent for the higher pressures as shown in Figure 4-20. At 100 mTorr the continuum model predicts that the peak ionization occurs in the center of the plasma, while distinct sheaths are still visible in the experimental data. The differences at 100 mTorr are due to two effects: the difference between the ionization and excitation thresholds, and the rapid thermalization of the secondary electrons in the continuum assumption. It is most probably that at 100 mTorr the electron energy distribution is no longer close to maxwellian, but has a higher percentage of high-energy electrons. The discharge is known to be sustained by bulk-sheath ionization at higher pressures, but at 100 mTorr the contribution of secondary electrons is probably important.




Figure 4-20 Simulated and experimental plasma induced emission at high pressures. a) 1 Torr, b) 500 mTorr, c) 250 mTorr, d) 100 mTorr. The model predicts the trends in the emission, but overpredicts both the sheath thickness and the bulk glow.

Overall, there will be differences between the model and the experiments because the concentration of metastables is not directly modeled. It is therefore significant that the model is doing a fairly good job at predicting how the sheath thickness varies with pressure, except at 100 mTorr. The model also predicts that the ratio of the intensity of the emission with respect to the emission at 1 Torr should be 1:0.91:0.66:0.66, and the experimental trend is similar, 1:0.69:0.09:0.07, although the magnitude of the change is different.

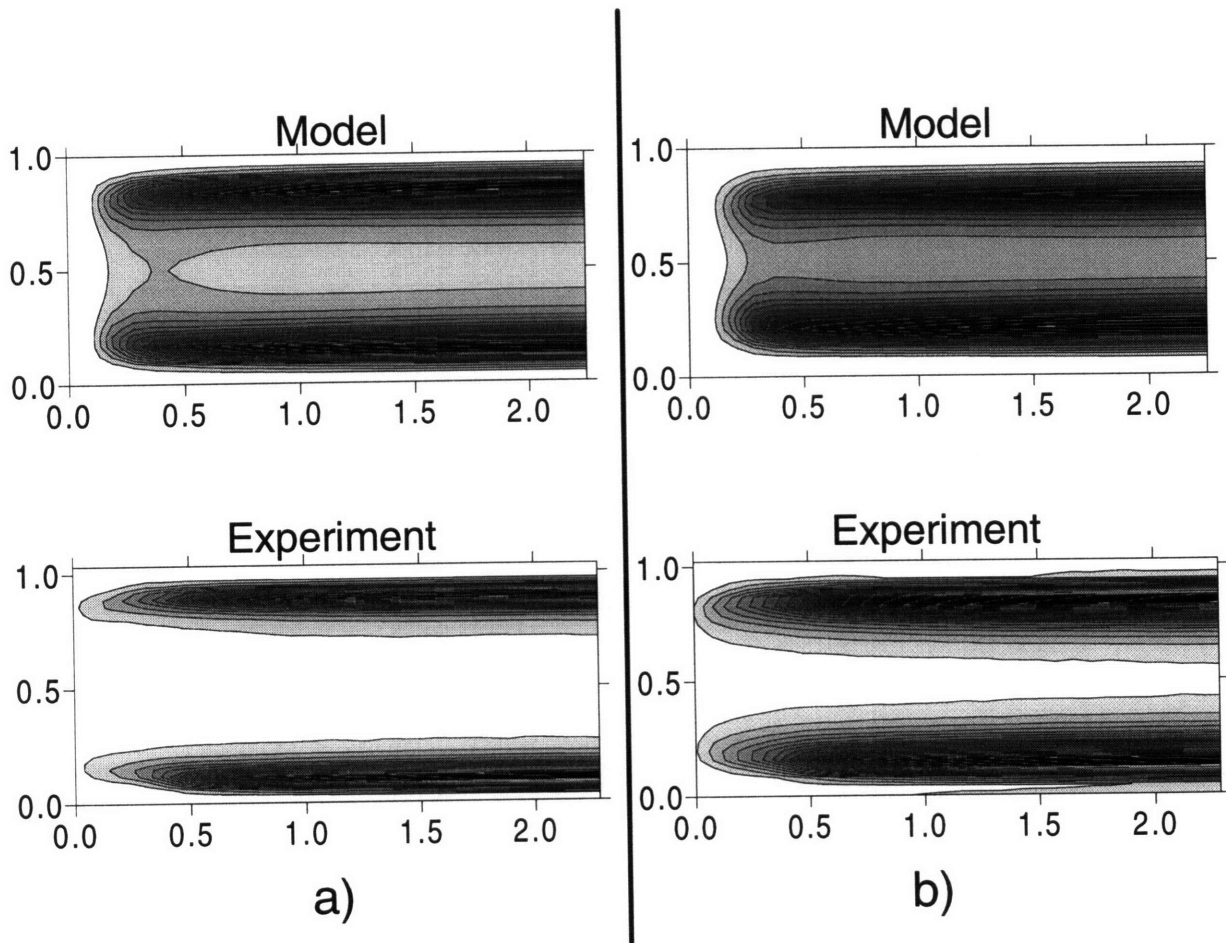


Figure 4-21 Simulated and experimental plasma induced emission at low pressures. a) 250 mTorr, b) 100 mTorr. The continuum assumptions do not hold at 100 mTorr, but the model still does a reasonable job at predicting the plasma glow.

The model predicts that the emission should become more diffuse as the pressure is lowered, and in fact although there is no distinct sheath present in the 100 mTorr simulation results, the experimental results do show that the center of the plasma is glowing quite a bit brighter than at higher pressures: the model results merely over predict (for all pressures) how bright the emission at the center of the discharge is compared to the maximum at the bulk-sheath interface. Simulations also show that the radial confinement of the plasma, except at 100 mTorr, is very weakly dependent on the pressure, and the experiments show this trend as well. At 100 mTorr the simulated plasma appears to have a smaller radius, while the experimental results show no significant change. The difference between the simulations at the lower pressure is again probably due to oversimplified treatment of the secondary electrons.

A simulation was also performed for a 2" gap spacing to compare to the experimental results, shown in Figure 4-22. The simulation predicts that the sheath thickness should increase with respect to an equivalent condition with a 1" gap spacing, but the experimental results show that the sheath thickness is constant. The simulation also predicts that the plasma radius should be smaller than in the 1" gap spacing case, but the experimental data shows that the plasma is contained in pretty much the same way, regardless of the gap spacing. For the real system, the plasma brightness at the sheath edge has increased with increased gap spacing, while the simulations predict that the emission should decrease to a level below the emission at 100 mTorr.

Clearly for this low aspect ratio plasma the sustaining mechanism for the simulations is not the same as in the real system, and as with the low pressure plasma, the reason is probably due to the secondary electrons. To fix this problem, another equation would need to be postulated to account for the non-thermal secondary electrons in the sheath. The other difference which is likely to result in difficulties in the comparison is the difference in geometry between the simulation and the experimental setup. The simulation results suggest that the peak plasma density should decrease by nearly 25%, which will decrease the bulk conductivity of the plasma. This increases the importance of the bulk of the plasma with respect to the plasma sheath. The presence of a sharp corner on the electrode in the experimental geometry, between the confining glass wall and the actual active electrode surface, results in a high field at the edge of the plasma. The enhanced ionization at this point, and the propagation of the discontinuity along the electrode surface, may account for the difficulty in the comparison. The same effects are also present for the 100 mTorr simulation and explain part of the difference in that comparison as well.

The other two-dimensional measures which are available are from Langmuir probe measurements. For these measurements, one places a wire in the plasma and measures the current it draws versus the potential at which it is biased. The only primary information available from the measurement is the floating potential, measured as a zero current when the electron current equals the ion current. Since the Langmuir probe is such a simple measurement to make on plasma systems, it has been studied in great detail. However, there is still quite a bit of disagreement as to how to interpret the probe results (see for example Godyak et al. (1992),

Turkot and Ruzic (1993), and Sudit and Woods (1994)). The results achieved using any one theory should be consistent for a given geometry, yielding at least semi-quantitative results, so the simple theory of LaFramboise is used in this work for interpreting probe results.




Figure 4-22 Plasma induced emission for a wide gap 1 Torr plasma.

If one assumes that the electron energy distribution is Maxwellian, then the electron energy is easy to extract from the Langmuir probe scan by plotting the logarithm of the electron current versus the applied potential. More rigorously, the second derivative of the I/V characteristic will yield the electron energy distribution function. Using the theory of LaFramboise, one can then predict the plasma density from the region where electrons are repelled from the probe. However, the entire theory for extracting information out of the probe scans is derived for lower pressure plasmas, where the mean-free path is much greater than the probe thickness. At 1 Torr with our small probe geometry, the probe theories are at best pushed to their limits, which results in difficulties in quantitative interpretation of the results. The probe design also limits the accuracy of the measurements, and calculation of the full EEDF as suggested by Godyak et al. (1993) is not possible.

The Langmuir probe scans do not span the whole plasma, since the probe cannot function in the plasma sheath where the plasma is not quasineutral, since there are not both positive and negative charge carriers to collect. Data was taken for the full range the probe views: within 2 mm of each electrode, and within 3 mm of the radial wall. The extent of the sheath changes with pressure, and it is not obvious where the probe will function beforehand. Probe traces appear distinctly different in the sheath and bulk regions, as shown in Figure 4-23, and the quality of the fit to probe theory was used as the criterion for accepting or rejecting a specific data point. The electron energy is extracted as a single temperature, assumed to be Maxwellian, and all distributions are assumed to have a single temperature. The fit to a single temperature is not justified, but as shown in Figure 4-24, the typical probe characteristic has the distinct regions for the end of the ion saturation regime, the electron transition regime, and the electron saturation regime. The floating potential is readily measured using a Langmuir probe, and the plasma potential can be predicted to be $V_p = V_f + 5.4 \epsilon_e$ by using the Bohm sheath criterion.

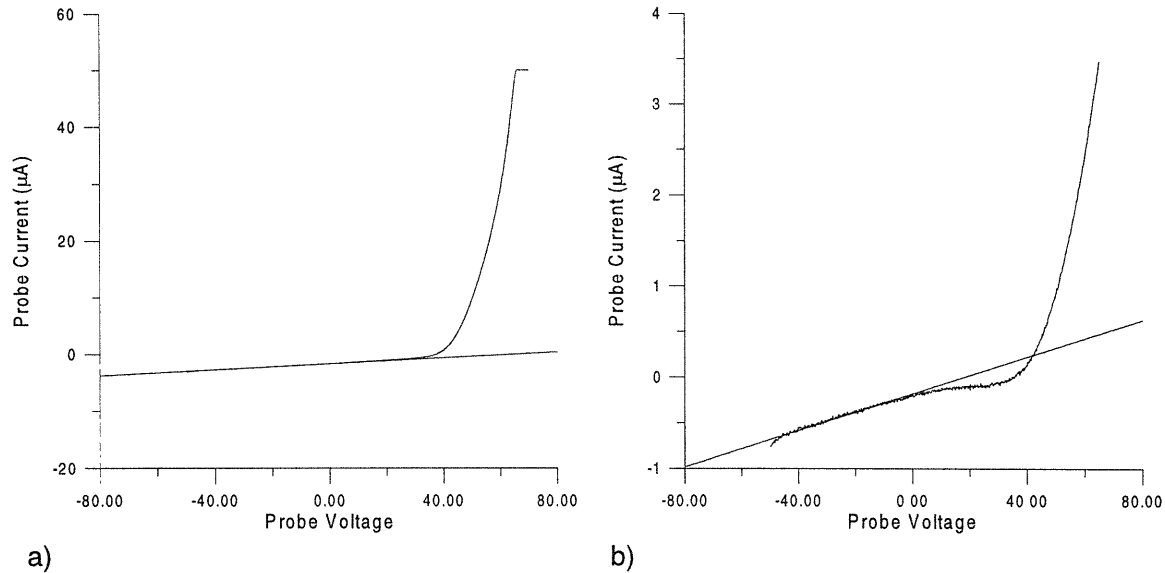


Figure 4-23 Langmuir probe characteristics for Argon Plasmas. a) Normal characteristic b) sheath characteristic. The shape of the probe characteristic is used to determine which data is included in the experimental data set for comparison to the model results.

The comparison of the experimental to simulated plasma potentials for the higher pressure conditions is shown in Figure 4-25. The trend is accurately predicted by the model, with a constant potential throughout the bulk of the plasma at both pressures. The potential does not vary significantly with pressure, as is correctly predicted by the simulations. It is important to note that the experimental data does not extend into the plasma sheaths where the majority of the curvature is seen in the simulation results. The simulation results apparently under predict the plasma potential by almost a factor of 2, but it seems likely that the error lies with the experiments instead of the simulation. The probe tip was designed so that the error should be acceptable at higher pressures. This design leads to some capacitive coupling of the plasma to the probe tip, which probably increases the measured potential by the 20 volt difference seen in Figure 4-25. Unfortunately, the increase in plasma potential will also lead to an erroneously high calculation of the plasma density using LaFramboise theory, although the overall shape of the profile will still be correct.

At lower pressures the comparison between the model and the experiments begins to break down. The model predicts, Figure 4-26, that the potential is no longer constant across the bulk of the plasma, but should have significant curvature. The maximum in the plasma potential is predicted to be slightly higher as pressure decreases, and the experimental data show this

clearly. The curvature at the edges of the experimental data is a result of the inaccuracy of the electron energy near the bulk/sheath boundary, before the probe fully enters the sheath and the data is considered invalid. The main difficulty in these comparisons is that as the pressure is reduced, the amount of reliable experimental data decreases as well, and the spatial resolution of the probe becomes a limiting factor.




Figure 4-24 Electron Current characteristic for Langmuir probe showing fit to electron energy as a Maxwellian.




Figure 4-25 Plasma potential in the symmetric Ar plasma at high pressures. a) 1 Torr simulation b) 500 mTorr simulation c) 1 Torr experimental, d) 500 mTorr experimental. The experimental data does not extend into the sheath region, where the simulation results have the greatest drop.

The electron energy should be much less strongly perturbed by the coupling of the probe sheath to the plasma, so the comparison between model and experimental results is quite a bit better, Figure 4-27. The outlying edges for the experimental data are generally exaggerated because the plotting routine extrapolates the data from the nonuniform spacing of data points. For both the simulation results and experiments the electron energy is highest at the bulk/sheath edge, and decreases toward the center of the discharge. The experimental data show that the electron energy in the center of the discharge should be around 3.5 eV, while the simulation is pinned near 5.5 eV because of the ionization cross-sections which were used. The maximum measured energy is 5-7 eV, depending on the conditions, and the simulation shows that the maximum energy at the bulk/sheath interface should be higher than 6 eV. There is no significant

radial structure to the measured electron energy, and the simulations predict that the only variation in energy radially occurs in the sheath, where no experimental data is available.




Figure 4-26 Plasma potential in the symmetric Ar plasma at low pressures. a) 250 mTorr simulation b) 100 mTorr simulation c) 500 mTorr experimental, d) 100 mTorr experimental. The experimental data does not extend into the sheath region, where the simulation results have the greatest drop.

At lower pressures, the match between the experimental and simulated electron energy is still quite good, Figure 4-28. The simulations indicate that the bulk energy is slightly higher, around 6.5 eV, and that the maximum energy in the sheath is increased to around 7.5 eV. At 250 mTorr, there was less valid data for the experimental measurements, so there is no indication of the energy at the bulk sheath interface, but the profile has more curvature in the plasma bulk, as is seen in the simulation results. At 100 mTorr, the energy is not really constant anywhere in the plasma, and is between 6-7 eV at the bulk/sheath interface. The simulation predicts that there is increased curvature for the energy as well as a higher maximum at the bulk/sheath interface.

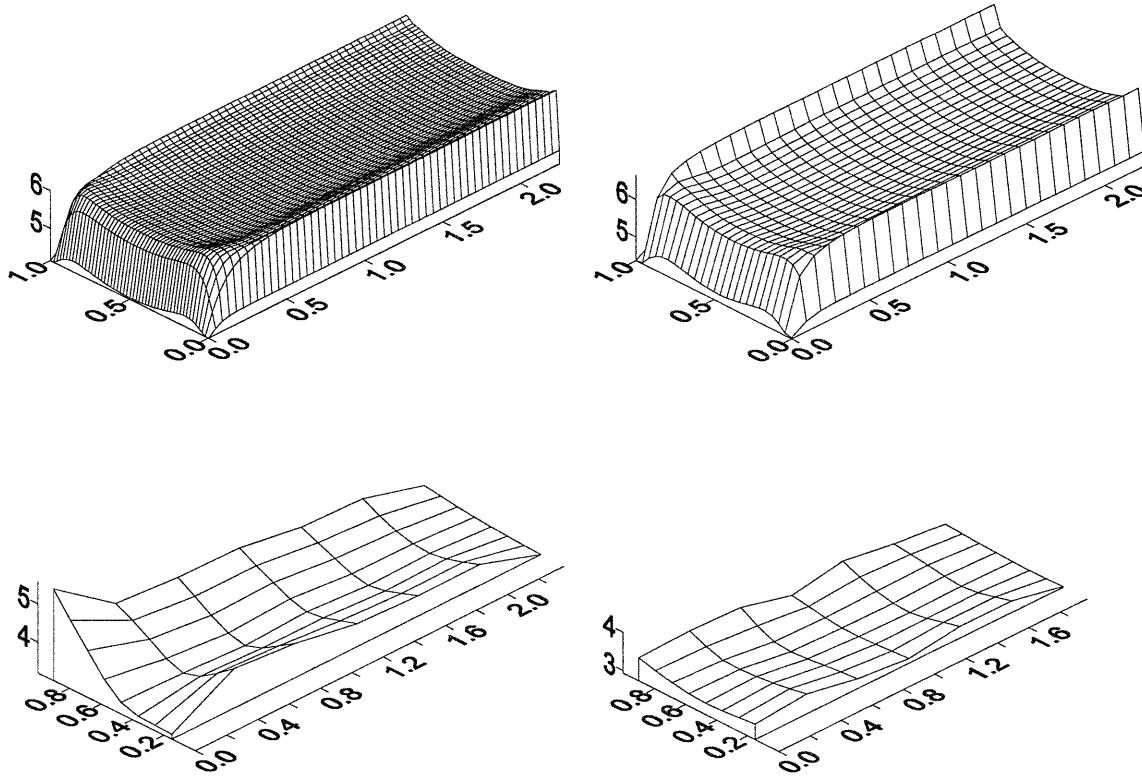


Figure 4-27 Electron energy in the symmetric Ar plasma at high pressures. a) 1 Torr simulation b) 500 mTorr simulation c) 1 Torr experimental, d) 500 mTorr experimental. In all cases, the electron energy is nearly constant in the plasma bulk, probably set by the ionization threshold for argon.

The results for the plasma density are a little more difficult to interpret, because the density is not measured but is calculated using LaFramboise theory. A prerequisite for accurate calculation of the plasma density is an accurate evaluation of the electron energy, which should be reasonable as previously discussed, as well as an accurate calculation of the plasma potential. Since it seems likely that the plasma potential may not be accurate from the experimental measurements, the baseline for the plasma density is expected to be considerably higher than in reality. This effect is evident in Figure 4-29, where the measured densities are more than an order of magnitude higher than the simulations predict. The experimental trend of higher density with lower pressure is probably an artifact because LaFramboise theory was used at the highest pressure. The experimental results show curvature in both the radial and azimuthal directions, in good agreement with the simulation results, and the radial sheath is less than 0.5” in thickness, as

shown in the simulated density profiles. The simulations show that there should be a maximum in density near the insulating wall, which is not apparent in the experimental data. However, at the higher pressures the maximum is a small perturbation on the overall profile, and may be smaller than the experimental error.

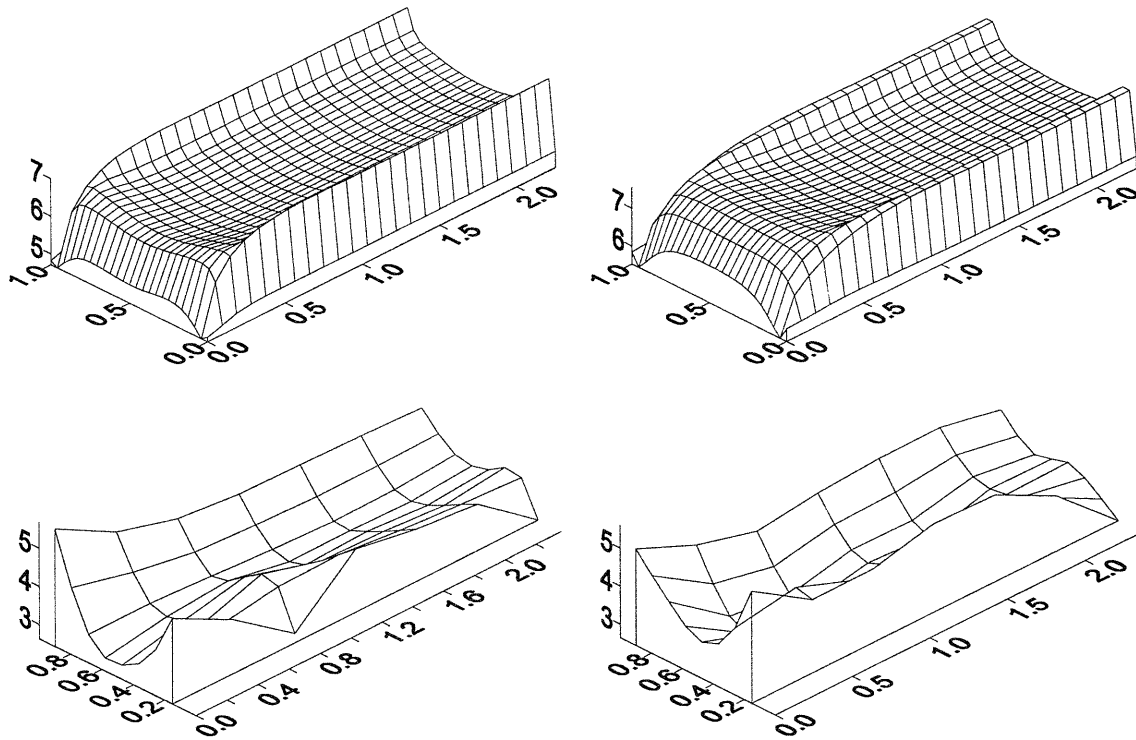


Figure 4-28 Electron Energy in the symmetric Ar plasma at low pressures. a) 250 mTorr simulation b) 100 mTorr simulation c) 250 mTorr experimental, d) 100 mTorr experimental. . In all cases, the electron energy is nearly constant in the plasma bulk, probably set by the ionization threshold for argon.

At the lower pressures the probe theory is expected to perform better in accurately representing the plasma density. At 250 mTorr, Figure 4-30, the measured density profile is more strongly curved, and closely resembles the simulated results, although the maximum near the insulating wall is still absent. The maximum in the experimental density profile has not changed significantly with respect to the higher pressures, although the simulation predicts that the density should nearly double with a doubling of the pressure. For 100 mTorr, the density has decreased by a factor of close to 10, and is more closely approaching the simulation results. In

fact, it appears that the maximum in the measured density is no longer at the center of the discharge. The increase in experimental plasma density very close to the wall is an artifact of the extrapolation from the plotting package, since these points are in the plasma sheath. However, the experimental data definitely show a maximum around 1/2" from the wall. It is likely that only the 100 mTorr density is accurate, due to the combination of the coupling of the plasma to the probe sheath, and the inapplicability of LaFramboise theory at higher pressures. For the experimental data, it appears that the plasma density is less curved as the pressure increases, and is nearly flat at the highest pressure. The simulated plasma density profiles are similar in shape at all pressures, except near the insulating wall. It is not clear whether the differences are due to inaccuracy in the measurements or an inadequate model for the ion diffusivity and mobility.

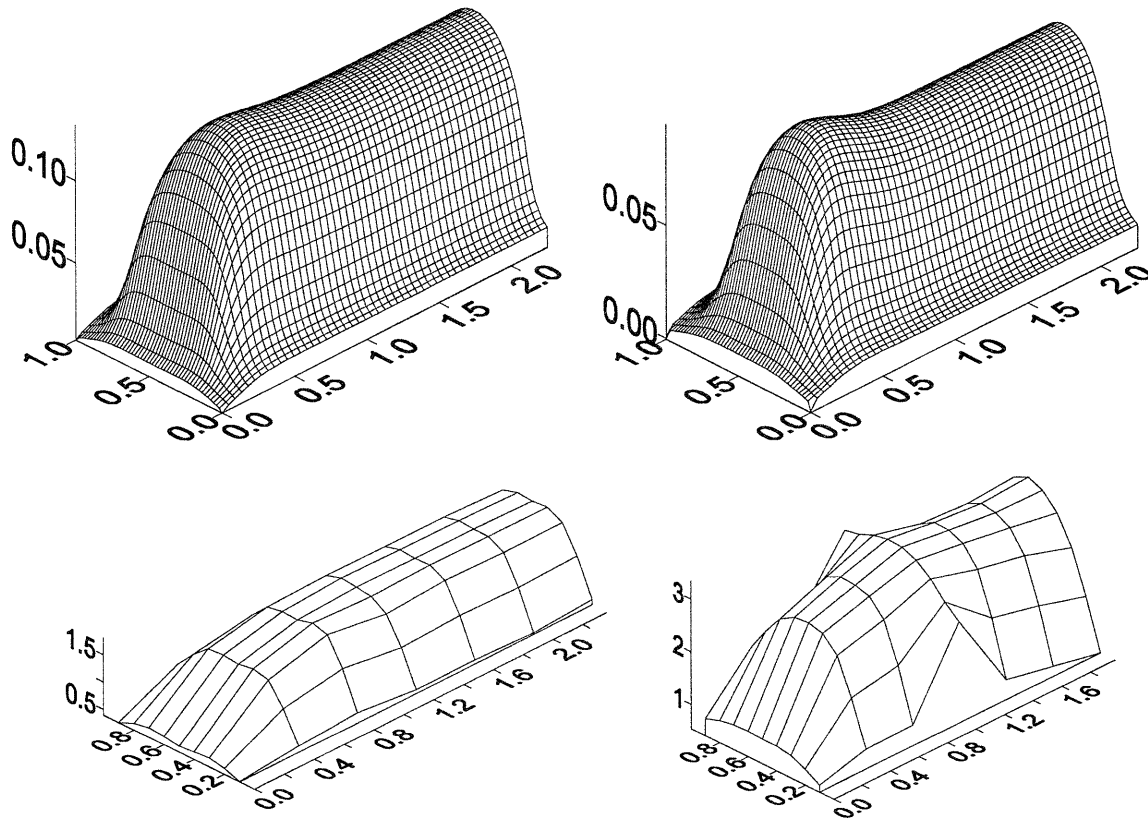


Figure 4-29 Plasma density in the symmetric Ar plasma at high pressures. a) 1 Torr simulation b) 500 mTorr simulation c) 1 Torr experimental, d) 500 mTorr experimental. The overall shapes of the curves are correct, but the theory used to generate the experimental data results in uncertainty in the absolute value of the density measurements.

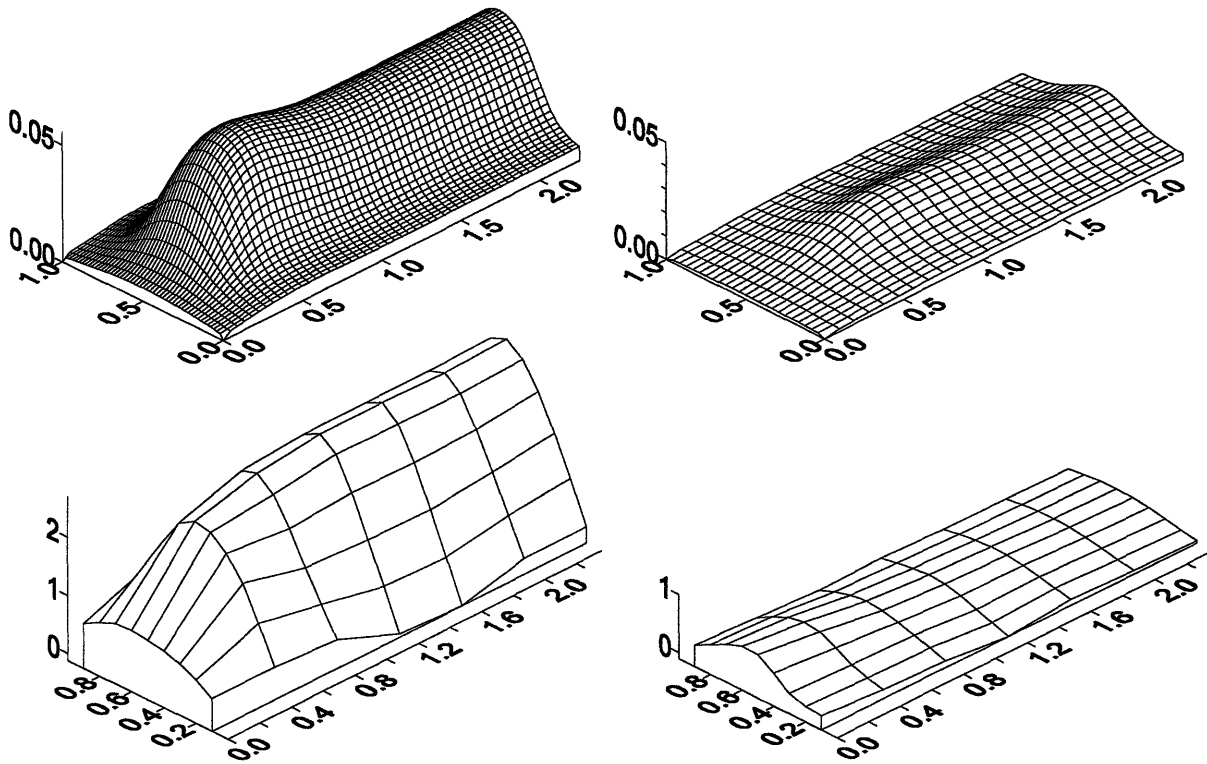


Figure 4-30 Plasma density in the symmetric Ar plasma at low pressures. a) 250 mTorr simulation b) 100 mTorr simulation c) 250 mTorr experimental, d) 100 mTorr experimental. The overall shapes of the curves are correct, but the theory used to generate the experimental data results in uncertainty in the absolute value of the density measurements.

4.5: Asymmetric Plasmas

The following two sections examine the physics of two-dimensional discharge where the grounded electrode is smaller than the powered electrode. Previous sections have dealt with the general plasma physics in one-dimension, and the effect of an insulating bounding wall. The focus of these sections is the effect of asymmetry in the geometry on the simulation results, and how well the simulations compare to experimental data.

It was shown in an earlier section that the change in boundary conditions where the insulator meets the grounded or powered electrode in the symmetric cases results in a solution which is not as smooth. Since the density of both electrons and positive ions is very low in the corner of the discharge, the discontinuity is effectively hidden and does not seriously affect the

convergence of the simulations. For the asymmetric simulations, the electric field is formally singular at the point where the insulator meets the smaller electrode. The ion density is not expected to be small at this boundary, and the overall convergence properties of the simulation are expected to change.

The singularity at the insulator/electrode boundary seems to have the greatest effect on the charging of the insulator adjacent to the electrode. The iterative solver (GMRES) for the Newton shooting method was observed to converge very slowly for the asymmetric cases. It was shown that the major change in convergence rate was due to the charging of the insulator next to the electrode. Since the electron density is very sensitive to changes in the local electric field, and the insulator charge is very sensitive to the electron density, these variables are very strongly coupled. Charge was allowed to diffuse along the boundary to mitigate this problem. The diffusion constant for insulator charge was chosen so that the characteristic diffusion length was less than 1 cm for timescales on the order of the ambipolar diffusion time. Since the overall plasma physics is not very sensitive to the charge on the insulating boundary, this change should not affect the simulations significantly.

Changing the diffusion constant on the boundary charge by two orders-of-magnitude for the symmetric case lead to a change of less than 5% in the boundary charge, and less than 1% change in all variables not directly related to the surface charge (since the voltage on the charged surface is directly proportional to the charge, it changes by 5% as well). It was necessary to use a surface charge diffusivity three orders-of-magnitude higher in the asymmetric cases with respect to the symmetric cases, for stability of the numerical methods, so a greater perturbation to the solution is expected. Changing the surface charge diffusion constant by one order-of-magnitude results in less than 10% difference value of the charge far away from the corners, where the charge reaches its maximum value. However, the charge near the insulator/electrode boundary is reduced by up to two orders-of-magnitude, which does effect the other variables near the charged surface. The differences in the TPSS are minor, with less than a 10% change in the positive ion density everywhere, the electron energy and electron density change by less than 10% outside the sheaths (but can change significantly near the insulator and near the singularity since they are already close to 0), and the effective field changes by less than 10% where it is not close to the insulator and does not vanish. Overall, the changes in the solution are less than 10^{-2}

except very close to the insulating surfaces, and the overall physics is not greatly perturbed by even a large change in the surface charge.

Even with this small change in the plasma physics, converging the simulations to a rigorous TPSS is very difficult. At some point in the RF cycle, the electron density at the boundary between the insulator and the electrode becomes very small, and rises rapidly to either side. The effective field also varies quite strongly near the singularity, and generally is the last variable to converge to a TPSS. Simulations for the asymmetric cases generally have a higher final residual in finding the TPSS ($\|U(T)-U(0)\|_2$ near 10^{-3}) than the symmetric cases. The 0.5 Torr and more asymmetric discharges were the most difficult to converge, and therefore have the highest error in the TPSS. Lower pressure simulations were not possible on any discretization which was tried. At lower pressures the plasma can exhibit longer range coupling of variables because the diffusivities increase. The singularity at the interface between the smaller electrode and the insulator for these cases probably perturbs the plasma physics on a longer length scale, resulting in our difficulty in finding a suitable discretization.

Finally, a blocking capacitor (100 nF) was added between the forcing voltage and the powered electrode to examine the effect of the asymmetry on the self bias potential. The simulations do not appear to be at all sensitive to the DC potential across the blocking capacitor (the DC potential may change by 50% without a 5% change in any plasma parameters not directly related to the potential on the powered electrode), but the potential is very sensitive to all the variables close to the powered electrode. The DC self-bias potential calculated by the simulations is therefore not a reasonable measure of the plasma physics. It was not possible to tell whether or not the DC self-bias had converged to a TPSS, since the TPSS residual for the DC potential always remained small, but even small changes in any of the variables close to the powered electrode could result in a change in the DC self-bias of up to 10%.

4.5.1: Explanation of Plasma Physics

The base case simulation for the asymmetric simulations is an Argon discharge at 1 Torr, 13.56 MHz, with a 1" gap spacing, larger (powered) electrode 4.5" in diameter, and smaller

(grounded) electrode 3" in diameter. This is directly comparable to the symmetric base case simulations: the only significant change is that the lower electrode is smaller in the asymmetric case. The effect of geometry is examined by changing the gap to 2" while maintaining the other plasma and geometric parameters constant, and with a highly asymmetric case at 1 Torr with a 1.5" lower electrode.

The main effect having of asymmetric geometry on the plasma physics is to change where the maximum in density occurs. For symmetric cases, the perturbation from the insulating wall, which is present to contain the plasma, results in a maximum in density approximately 0.5 gap spacings from the wall. For the asymmetric case, Figure 4-31, the density generally increases as the radial center of the plasma is approached, and there is a minor maximum near the smaller electrode: there is not a high enough radial field near the insulating wall to greatly affect the plasma confinement. Higher fields near the smaller electrode lead to a higher generation of species, and a slight asymmetry in the axial density profiles. The plasma is maintained mainly at the bulk/sheath interface, and diffuses out to be collected at the electrodes, resulting in a gently sloping region from the smaller electrode to the insulating wall. Even though the entire large electrode is treated the same way, the plasma is mostly confined in the region shadowed by the smaller electrode.




Figure 4-31 Positive ion density contours for an asymmetric Argon plasma with an area ratio of 2.25. 1 Torr discharge, 13.56 MHz, 1" gap spacing. The density profile is not highly asymmetric, and the radial sheath thickness increases with respect to the symmetric case.

The singularity in the field does effect the density profiles, creating a maximum in ionization at the boundary between the electrode and the insulating surface. The maximum in

density is a minor effect, but since the major source of charged species is moved far away from the insulating wall, the density in the corners decreases dramatically. The singularity in the electric field also propagates into the effective field which is seen by the ions. The sudden drop in the ion convective velocity results in a small oscillation in the ion concentration at the boundary, but this effect is on a length scale which is too small to be of any great importance, see Figure 4-32. Since the singularity is completely a local effect, perturbations in the ion density are limited to at most the two elements closest to the insulator/electrode boundary. The perturbation in electric field is slightly longer range, but still on a relatively small length scale with respect to the overall geometry of the simulation.




Figure 4-32 Positive ion density at the grounded electrode for the same conditions as Figure 4-31.

The singularity in the electric field is not evident in Figure 4-33 because of the coarseness of the mesh used to generate the vector plot. All of the important features of the asymmetric discharge previously discussed are evident in the time averaged electric field. There is a very slight asymmetry which can be noted by finding the point where the electric field vanishes. The

field vanishes at the same point on the small and large electrodes, and begins to disappear at the point where the electrode meets the insulator on the lower surface. The difference between the symmetric and asymmetric discharges is very evident in Figure 4-33 as well, in that the absolute minimum in the electric field is now much further away from the insulating wall, and the radial field is higher in a much larger section of the plasma. This occurs because the plasma is mainly generated at the bulk/sheath interface in the area shadowed by the smaller electrode, and a radial electric field is necessary to maintain a net neutral charge flux to the insulating wall.




Figure 4-33 Time averaged electric field for 1" gap, 1 Torr, asymmetric Argon plasma. The point of minimum field occurs closer to the edge of the smaller electrode, and the radial maximum in plasma density is minimized.

The total current flow also shows quite clearly that the plasma is asymmetric, Figure 4-34. For symmetric cases, the current is a mirror image with respect currents at a phase $1/2$ RF cycle later, but for the asymmetric simulations the current at any phase in the RF cycle is unique. Again for the current, the effect of the insulating wall is reduced, and the currents in the corner region are very small. The current is also always small in the area which is not shadowed by the smaller electrode. In Figure 4-34 a) and c) we see the main effect the geometry has on the total current flow. When the current is flowing toward the smaller electrode, the radial fields cause the total current to be constricted toward the smaller electrode, while the current spreads out when flowing mainly toward the larger electrode. This is, however, for the parts of the RF cycle when the current is close to a minimum, and this represents mainly the flow of electrons. When the current is at its peak value, near the time slices shown in Figure 4-34 b) and d), there is very little current spreading, and the current is very small in the area outside the smaller electrode.

When the largest current is passing through the discharge, it flows almost directly from the smaller electrode to the larger electrode. In this way, the simulation is very similar to a two-dimensional case with smaller electrodes.

Examining the current to the larger electrode, Figure 4-35, it is obvious that the overall perturbation to the simulation by changing the geometry is not a first order effect. The asymmetric current is identical in form to the symmetric current, with a small ion component, a pulse of electrons once per RF cycle, and the displacement current dominating the current at the electrode. Because the area and the density of the plasma are both reduced, the total current is slightly reduced as well. The asymmetric geometry does not, however, result in any significant higher harmonics in the current through the plasma. Similarly, if one examines the current density to the larger electrode, Figure 4-36, the density drops off away from the center of the discharge and the discharge becomes more capacitive exactly the same trends seen in the symmetric simulations. However, the current density has already dropped off by 15% by the point where the smaller electrode stops. Because the plasma is dominated by the smaller electrode, the current is reduced over a much larger fraction of the large electrode than is seen in the symmetric case.

When the area ratio of the smaller to larger electrode is increased the plasma does not become increasingly nonuniform, but rather tends to become more and more confined in the area defined by the smaller electrode. The positive ion density, Figure 4-37, shows the effect of the lower electrode best. The density is maximum in the area defined by the smaller electrode, and decreases rapidly radially (through ambipolar diffusion). The asymmetry in the plasma is limited to a very minor spreading of the plasma between the two electrodes, so that the plasma effectively sees an upper electrode of similar size to the smaller electrode.

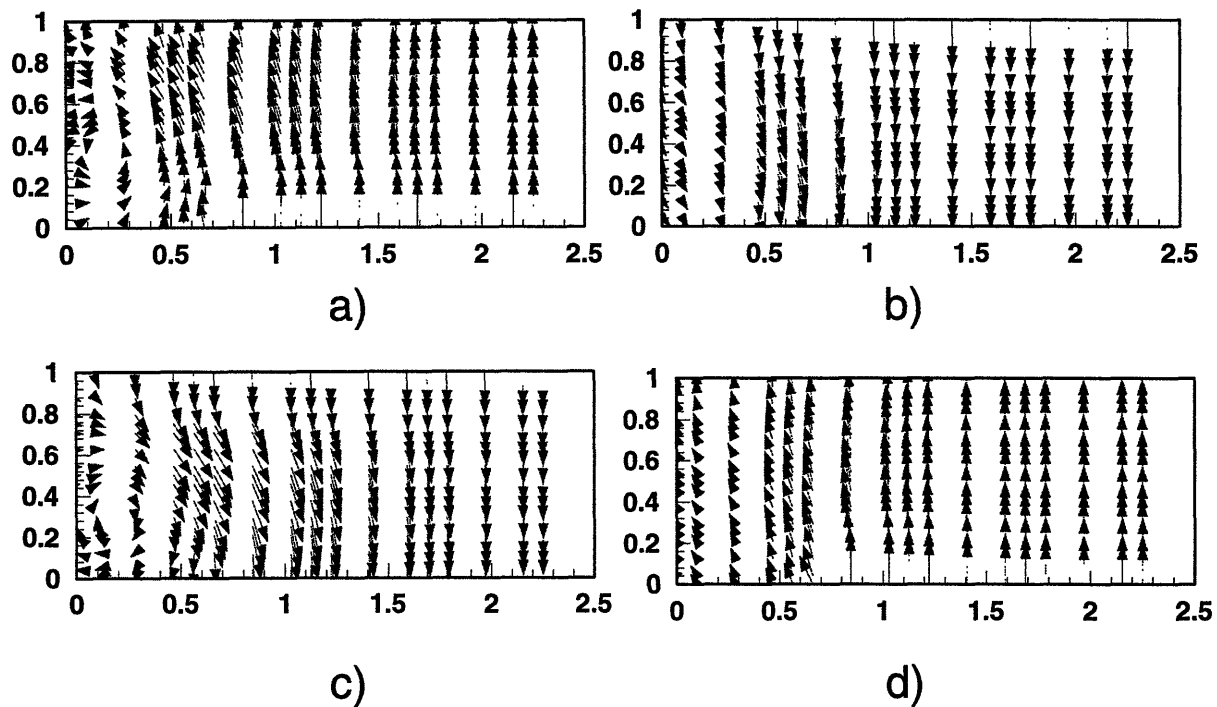


Figure 4-34 Current for asymmetric 1" gap discharge. Simulations at 1 Torr for different phases in the RF cycle: a) 0.25, b) 0.50, c) 0.75, d) 1.00. The asymmetry in the current flow near the edge of the smaller electrode is clear. Most of the discharge between the small and large electrodes is still 1D.

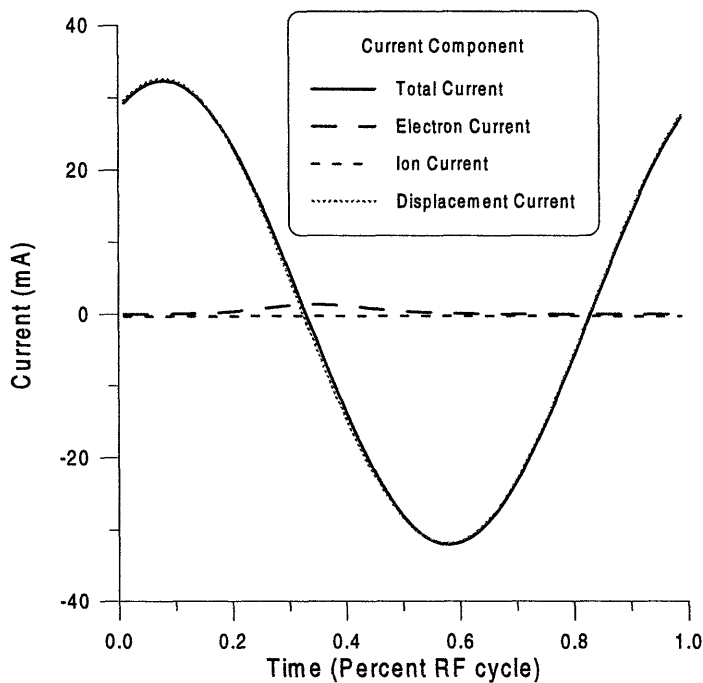


Figure 4-35 Current at the powered (larger) electrode for the asymmetric Argon Discharge at 1 Torr.




Figure 4-36 Total current density on the powered electrode for an asymmetric Argon discharge at 1 Torr.

Examining the time averaged electron energy, potential, and positive ion density in Figure 4-38, it is evident that all of these quantities are nearly constant from a radial distance from the insulating wall of 1.5" and upward. The electron energy plot has very low values (below the ionization threshold) everywhere except at the bulk sheath interfaces, and even on the larger electrode is very low outside of the area defined by the smaller electrode. The potential is nearly constant inside the active plasma volume, and has two well defined sheaths near either electrode. One interesting effect of the highly asymmetric discharge is that the DC potential on the larger electrode has become negative. This occurs because the plasma is sustained mainly at the smaller electrode, the density near the smaller electrode is higher, and the structure of the sheath has changed significantly to confine the plasma.




Figure 4-37 Positive ion density contours for a highly asymmetric Argon discharge. Simulation at 1 Torr with an electrode area ratio of 9. The results are asymmetric, with the highest density near the smaller electrode. The plasma is mostly confined by the smaller electrode.

The easiest way to see the difference in sheath thickness between the small and large electrodes is to examine the predicted PIE profile for the highly asymmetric case, Figure 4-39, which is essentially equivalent to the ionization rate within the plasma. The maximum in the PIE occurs one third closer to the smaller electrode, and the intensity at the smaller electrode is nearly double. If we take any of the lower contours as the visual sheath edge, the upper sheath appears to have a larger radius than the lower sheath, although it still does not even come close to covering the larger electrode.

The asymmetric discharges can be related very closely to the symmetric cases because area ratio of the two electrodes does not have an effect, to the first order, on the simulation results. Since the plasma remains confined by the smaller electrode, and since the relevant quantity for the symmetric discharges was the ratio of the radius to the electrode gap, the relevant ratio for the asymmetric system should be the ratio of the radius of the smaller electrode to the gap spacing. The positive ion density at the centerline of the discharge and at the edge of the smaller electrode is shown in Figure 4-40; the symmetric profiles are shown for comparison. For the low aspect ratio plasma, at the plasma centerline we are well beyond one gap spacing away from the edge of the smaller electrode (which is 1.5" long), so the plasma behaves almost identical to the symmetric and one-dimensional cases. If the aspect ratio of the asymmetric discharge is not close to one, since there is a DC bias on the larger electrode, there is not an exact equivalent in a one-dimensional simulation unless a DC component is added onto the RF forcing potential. For the high aspect ratio plasma, which has a small electrode which is less than

one gap-spacing long, even at the centerline the discharge does not approach the symmetric simulation results.




Figure 4-38 Time averaged quantities for a highly asymmetric Argon discharge. a) Electron Energy, b) Potential, c) Positive ion density.


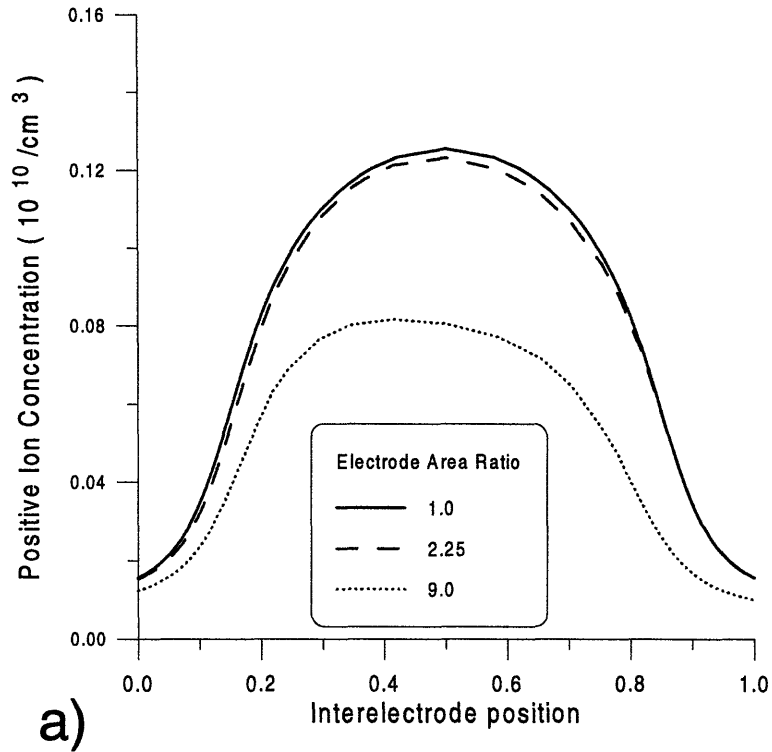



Figure 4-39 Simulated Plasma Emission for a highly asymmetric Argon plasma. The plasma is clearly active in the area defined by the smaller electrode.



a)



b)

Figure 4-40 Effect of plasma asymmetry on the positive ion density a) at the radial centerline, b) at the edge of the smaller electrode. In b) the density for the area ratio 1.0 case is plotted at the radial centerline instead of at the edge of the electrode, which would be on the insulating wall surface.

If the plasma were totally confined by the smaller electrode, we would expect to see the positive ion density approach zero at the edge of the smaller electrode. This is clearly not the case in Figure 4-40 b), although the density has decreased by 20% from the value it attains at the radial centerline of the discharge for both aspect ratios. Since the distance between the edge of the smaller electrode and the insulating wall is less than a gap spacing, one does not expect the lower aspect ratio density to reach zero for any radius. The distance between the edge of the small electrode and the insulating wall is greater than a gap spacing for the high aspect ratio simulation, and indeed the positive ion density is essentially zero beyond one gap spacing from the electrode edge. These observations indicate that for asymmetric discharges of the type examined here, the aspect ratio defined as the ratio of the smaller electrode radius to the discharge gap spacing is the important spatial scaling.

The aspect ratio of the plasma determines how much the asymmetry perturbs the plasma physics in the bulk of the discharge, but unlike in the symmetric case, the radial boundary layer is quite different for different aspect ratios, Figure 4-42. The radial boundary layer changes significantly because the distance ions can diffuse to reach the insulating wall has changed significantly, and the discharge is being sustained differently at very low aspect ratios. Note that the density for the lowest aspect ratio case does not approach the symmetric solution even at the centerline. Since the effective aspect ratio for this simulation is less than 1.0, we can expect that the simulation will not resemble the symmetric case at any spatial position. For a simulation with the same electrode area ratio, and a gap spacing of 1/2", the results should be nearly one-dimensional beyond for a distance of greater than 2" from the insulator, and comparable to the 1/2" gap symmetric case at the radial centerline, since the effective aspect ratio would be 1.5. This is indeed the case as demonstrated in Figure 4-41.




Figure 4-41 Positive ion density contours for a 1/2" gap, highly asymmetric argon discharge. The plasma is clearly contained by the smaller electrode, and has minimal radial variation more than 1 gap-spacing away from the edge of the smaller electrode.




Figure 4-42 Positive ion density at the radial centerline of the discharge for various sized smaller electrodes. The effective aspect ratio is the ratio of the small electrode length to the gap spacing.

Although the majority of the plasma physics is relatively unchanged by changing the area ratio of the electrodes, for plasma etching the most critical parameters are what the wafer actually experiences. The absence of large changes in the bulk plasma physics implies that the plasma should be relatively unchanged above the smaller electrode regardless of the ratio or the areas of the small and large electrodes. Examining the ion flux to the lower electrode (smaller electrode

and adjacent insulator), Figure 4-43, it is evident that there are some changes in the ion flux which may have a significant effect on the etching of a wafer placed on the lower electrode. The ion flux in all cases increases away from the center of the plasma until a maximum 15% higher is reached about 1/4 of a gap spacing away from the insulator/electrode boundary. The singularity in the field causes a large rise in the ion flux at the boundary between the insulator and the electrode, along with a large radial ion flux (not shown).




Figure 4-43 Ion current to the grounded (smaller) electrode at various pressures and aspect ratios.

The effect of pressure on the ion flux is relatively small, as was the case for the symmetric simulations. Lowering the pressure from 1 Torr to 500 mTorr only slightly raises the ion current to the electrode, and results in a very similar shaped profile. Even changing the size of the smaller electrode to the point where the plasma is not expected to be radially uniform at any point (a change in the aspect ratio between the smaller electrode radius and the gap spacing from 1.50 to 0.75) only results in a 25% decrease in the total ion flux. The ion flux profile across the grounded surface is still similar to the case with a larger grounded electrode, although since

the electrode is now nearly the same size as the gap spacing the minimum in ion flux is never reached, and the ion flux across the electrode appears more uniform.

When the gap spacing is increased, the effect of the insulator bounding the plasma has already shown to be more significant. Since the plasma is already quite two-dimensional for an aspect ratio close to one, reducing the radius of the lower electrode by one third does not have a very significant effect. The main difference between the symmetric and asymmetric simulations in this case is that the density is asymmetric due to an increase in the ionization rate near the lower electrode. The peak in the positive ion density is also nearly 25% less for the asymmetric case than for the symmetric case. The maximum in the ion density is lower because the effective aspect ratio for the geometry is now significantly less than unity so the plasma is not even close to one-dimensional at the radial centerline.




Figure 4-44 Positive ion density for 1 Torr, 2" gap asymmetric Argon discharge.

4.5.2: Comparison to Experimental Measurements

The comparison between the asymmetric simulations and the experimental results is done the same way as the symmetric simulations. Since there is still a Teflon spacer between the insulating glass wall and the upper electrode, a direct comparison is not possible. Instead, all

simulation results are compared to the experimental results assuming that the bounding insulator is effectively the same as the insulating wall. The comparison of the predicted PIE to the experimental data is facilitated by correctly juxtaposing the two data sets, and assuring consistency of the spatial scales.

A direct comparison of the simulation results and experimental measurements is shown in Figure 4-45. The simulation results match the experimental measurements at least as well as in the symmetric case, and all of the qualitative features are correctly captured: the sheath at the lower electrode is glowing about two times brighter than at the upper electrode, the sheath at the small electrode is slightly thinner than at the upper electrode, the glow covers more of the large electrode than on the small electrode, the glow diffuses out more at lower pressures and becomes overall more uniform, and the intensity of the maximum glow decreases with pressure. For the experimental PIE profiles the brightest glow on the lower electrode does not extend beyond the electrode edge, while there appears to be a small but significant glow beyond the edge of the electrode. This is the one point where the comparison between the simulation and the experimental data seems to fail. The most likely reason that the simulation result is different is that the a perfect insulator with a small charge leakage term is not a good approximation under plasma bombardment. Besides the absence of the confining ring between the insulating wall and the electrode, the simulation also has one other difference in geometry which will affect the comparison. The ground plane on the lower electrode in the experimental geometry is very complex, and may be as close as one gap spacing away from the edge of the smaller electrode, while in the simulations the approximation is that the ground plane is nearly infinitely far away (2.5 gap spacings, which should be beyond the range of perturbation of the plasma). Certainly having a ground plane much closer to the plasma should result in a less asymmetric discharge, and will change the simulation results, but simulation of the exact experimental geometry is beyond the scope of this thesis.




Figure 4-45 Comparison of simulated plasma induced emission for asymmetric discharges. The top plot is model results, and the experimentally measured emission is presented in the bottom. a) 1 Torr and b) 0.5 Torr.

For larger gap spacings the asymmetric simulations have about the same short-comings as in both the symmetric simulations with wide gaps, and in the high-pressure low-gap spacing asymmetric cases. The simulation correctly predicts that the sheath at the smaller electrode is both brighter and thinner, it appears that there should be a smaller sheath on the upper electrode, while the experimental observation shows that the sheath on the larger electrode is the same size for 1" and 2" gap spacings. There is obviously, as with the 1" gap spacing, a significant difference in the way the plasma is being sustained in the simulation versus the experimental setup. There is also the possibility that the measurement of the applied power is in error, so that the power being applied is significantly different between the experiments and simulations.




Figure 4-46 Plasma induced emission for wide gap, asymmetric discharges. Comparison of model (top) to experimentally measured plasma emission (bottom) for a 1 Torr, 2" gap spacing asymmetric Argon plasma. The simulation results are dominated by edge effects since the aspect ratio is close to one.

The comparison of the time averaged potential for the simulations and the measured plasma potential using a Langmuir probe is shown in Figure 4-47. For both the simulations and the experimental data the potential is essentially flat throughout the plasma except in the sheath regions, and does not change very significantly with pressure. The experimental plasma potential is always higher than the simulated potential, as was seen in the symmetric cases. In fact, there is virtually no difference in the experimental results for the symmetric and asymmetric cases, while the simulations require that there is a positive DC potential on the larger electrode. The exact value of the simulated DC floating potential is not very well determined as discussed earlier due to problems in convergence of the TPSS shooting algorithm, but it does appear to decrease by at least a couple of volts when the pressure is decreased from 1 Torr to 0.5 Torr.




Figure 4-47 Plasma potential for asymmetric plasmas. a) 1 Torr and b) 0.5 Torr; simulation results c) 1 Torr and d) 0.5 Torr; calculated from Langmuir Probe measurements. There is no experimental data in the plasma sheath, where the simulations contain the largest variations.

The electron energy for the simulations and as measured by Langmuir probe is shown in Figure 4-48. For the pressures investigated both the experimental and simulated electron

energies are between 4 and 6 eV in the bulk of the discharge, and are essentially as flat as the plasma potential. The only difference between the simulated electron energy for the symmetric and the asymmetric cases is that the energy is higher next to the smaller electrode. For the experimental measurements, this effect is not detectable because predicted the increase in electron energy occurs in the sheath where the Langmuir probe measurements are not possible.




Figure 4-48 Time averaged electron energy for asymmetric discharges. a) 1 Torr and b) 0.5 Torr; simulation results c) 1 Torr and d) 0.5 Torr; calculated from Langmuir Probe measurements. All results show a nearly constant energy in the plasma bulk.

The plasma density in the quasi-neutral region of the plasma should be directly comparable to the density profile as predicted by Langmuir probe measurements. As was seen in the symmetric comparison, however it is very difficult to interpret the density calculated from LaFramboise theory at higher pressures. With the measured densities for the asymmetric cases, Figure 4-49, we again see a decrease as the pressure increases, and a discrepancy of an order of magnitude between the simulation results and experimental measurements. However, for these cases there is more radial variation, both for the experimental and simulated density profiles. The experimental plasma density is relatively constant radially next to the smaller electrode, and

slowly decreases as the radial insulator is approached. The simulated results show a very similar trend. The main difference is that the slope of the density seems to be gentler for the experimental results, but this is partly because of the difference in the geometry between the simulation and the experiments. The simulations also show that the density does not drop off as quickly at lower pressures, but the experimental error makes it difficult to tell how different the slopes are at different pressures.

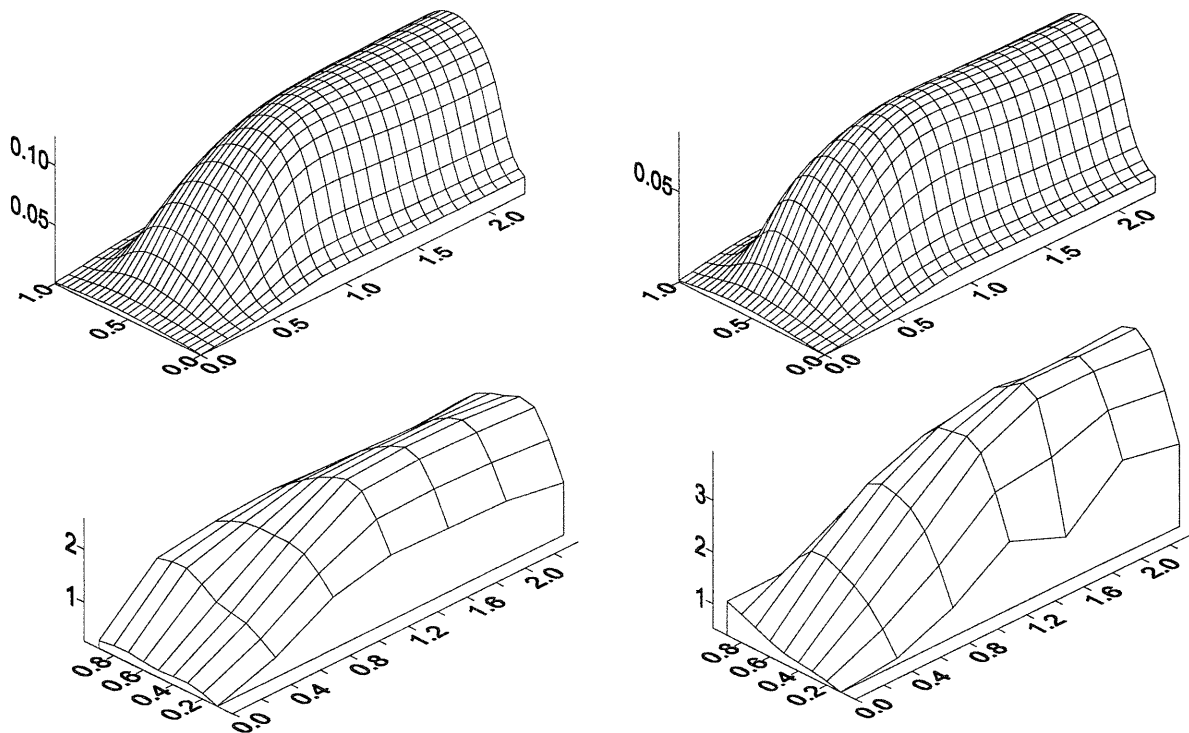


Figure 4-49 Time averaged positive ion density for asymmetric discharges. a) 1 Torr and b) 0.5 Torr; simulation results c) 1 Torr and d) 0.5 Torr; calculated from Langmuir Probe measurements.

Chapter 5 : Argon Plasmas In Complex Geometries

The continuum model results predict one well defined characteristic feature that should be readily observed experimentally: a strong enhancement of the ionization rate where the insulator meets the electrode surface. However, when one accounts for the averaging effect that is produced since the PIE is viewed through an entire cross-section of the plasma, the enhancement of ionization in the corner is nearly wiped out for the simulated PIE profiles. This occurs because the enhanced ionization is taking place at a point where the smallest cross-section is being viewed, and all other cross-sections also contain a view of the enhanced ionization. The difference in the cross-sections is shown by the lines on the upper electrode in Figure 5-1 a): as the scan one proceeds from the left to the right, the optics average over a larger area of plasma.

In order to demonstrate the enhanced ionization, it is necessary to produce a geometry where the enhanced ionization will occur at a point where the plasma cross-section is near its maximum. Experiments were performed with rings and bars, shown in Figure 5-1, to capture this effect. The following sections deal with the experimental observation of a few such systems, and an explanation of the physics based on simplified arguments from the simulations. The units of distance shown on the plots are in centimeters.




Figure 5-1 Geometry for perturbation studies with a) ring b) angled bar on lower electrode.

5.1: Symmetric Plasmas with Focusing Rings

Clamping rings are a common feature in industrial plasma etching tools. The purpose of these features is merely to hold the wafer in place, in which case very little effort is usually spent on designing the geometry of the clamp. Focusing rings are used to increase the plasma uniformity across the wafer, and are designed by trial and error for this purpose. Since one already would predict from the plasma modeling that a ring with a right angle to the electrode surface should produce an enhancement in the ionization rate, we examined four separate cases at two different pressures. The difference in the 4 cases were the thickness of the ring, which should affect the magnitude of the perturbation, and the material of construction, since a conducting or insulating ring should have significantly different effects on the plasma.

The PIE scans for all four geometries at 1 Torr are shown in Figure 5-2. The dashed square in the figure is equivalent to the cross-section of the ring in Figure 5-1 a) at the center of the plasma. The first obvious feature is that the emission intensity is greatly enhanced near the corner where the ring meets the electrode, at a distance of about one sheath thickness above the electrode surface. The position of the enhanced emission is not sensitive to the size of the focusing ring or the material of construction. Since the ring blocks most of the rest of the electrode, the rest of the lower sheath is not visible on the plots, and a similar enhancement occurs inside the ring which is also not visible in the PIE scans. The enhancement seems to be significantly greater for a thinner focusing ring, although it still is not as intense as the sheath is on the powered electrode. There is also an enhancement at the top edge of the focusing ring, which is due to the exterior corner of the ring. For the thinner rings, these two effects apparently interact to further enhance the plasma emission.




Figure 5-2 Plasma emission for 1 Torr Ar plasmas with focusing rings on the electrode. a) 1/2" Al ring b) 1/2" Teflon ring c) 1/4" Al ring d) 1/4" Teflon ring.

The second interesting effect noticeable in the PIE in Figure 5-2 is that the aluminum ring does seem to focus the plasma glow from the center of the gap toward the electrode, and is more effective for a thinner focusing ring. The Teflon ring, on the other hand, seems to have the opposite effect, enhancing the plasma emission through most of the center of the discharge. The conductive ring apparently results in stronger fields near the electrode, while the insulating ring bends the electric fields and allows for more of the bulk of the discharge to effectively see the ground plane.

The model predicts that the sheath thickness should greatly increase with decreasing pressure, and that the effect of the corner should become less important (at 100 mTorr, for the insulating wall only one peak in ionization occurs, on the centerline of the plasma, with a characteristic length scale of 1/2"). Therefore one expects that as the pressure decreases and the focusing ring becomes contained in the sheath region, the perturbation to the plasma should

diminish. The effect of the insulating rings does indeed diminish as the pressure decreases, as shown in Figure 5-3. The critical length scale for plasma perturbations has been shown to be the aspect ratio of the plasma. For the rings, the pertinent dimension is the ring height. For all ring geometries, the plasma does not appear to be perturbed more than one ring-thickness away. For the 1/2" rings, the enhancement in ionization is still noticeable since the ring is more than double the thickness of plasma sheath. When the ring is only 1/4" thick, which is on the order of the sheath thickness, it sits mainly within the sheath and no longer greatly focuses the plasma toward the electrode. In this case, the plasma sheath simply follows the contour of the ring, and the geometry and material of construction are of secondary importance. For the 1/4" rings the lower sheath reappears, but since the position of the sheath is different at the ring surface and at the electrode surface, the intensity at the bulk/sheath interface is greatly reduced at the lower electrode.




Figure 5-3 Plasma emission for 100 mTorr Ar plasmas with focusing rings on the electrode. a) 1/2" Al ring b) 1/2" Teflon ring c) 1/4" Al ring d) 1/4" Teflon ring.

Two-dimensional scans with a Langmuir probe were performed to see how the rings perturbed the electron energy profiles within the plasma. Since the enhanced glow occurs in a sheath region, we do not expect to observe a large peak in the electron energy using the Langmuir probe. Since the rings should perturb the electric fields throughout the plasma, the electron energy should also change globally. Scans were performed near the powered electrode, near the center of the plasma, and close to the electrode with the ring. The scans near the electrodes are not shown because they contain too many points which fell within the sheaths, and no significant trends were observed. The perturbation of the electron energy at the center of the plasma by the rings is shown in Figure 5-4 for high and low pressures.

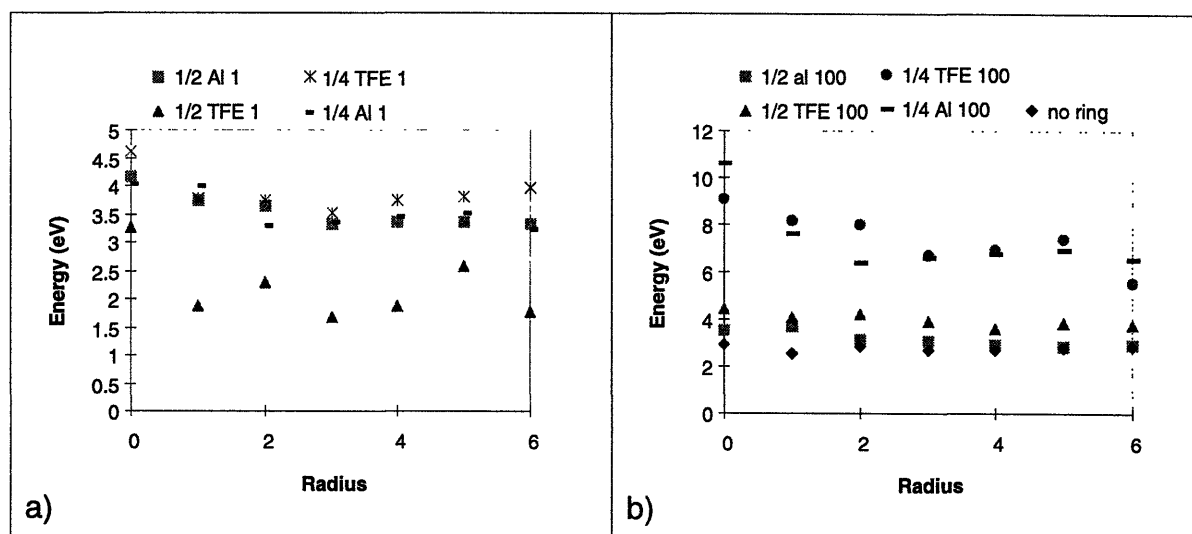


Figure 5-4 Electron energy as determined by Langmuir probe for the center of a perturbed Ar discharge. a) 1 Torr b) 100 mTorr.

At 1 Torr, the electron energy is not affected by the rings in any significant way, and is almost constant across the plasma. Although the 1/2" Teflon ring appears to have lower energy, the differences are within the experimental error. For the 100 mTorr scans, the energy is again mostly unchanged, with no variation across the radius of the plasma. The 1/2" rings have electron energies nearly identical to the case with no ring. For the 1/4" rings, the electron energy is still nearly constant across the discharge, but appears to be around double the energy for the other cases. The PIE shows no enhanced glow for these discharges to indicate that there should be higher energy. The most likely explanation for the difference is that these runs were the last 2 performed, and either the contact to ground was not sufficient or the runs which had lower energy

had a small air leak. As little as 0.5% air in Argon would produce this change, as discussed by Surendra (1985).

5.2: Symmetric Plasmas with Bars

The PIE studies for Argon plasmas with focusing rings showed that an enhancement occurs in the plasma emission due to the ring, but it is still possible that this apparent enhancement was due to the way the plasma traverses the ring. It is not reasonable to suggest that the enhanced glow is merely an experimental artifact due to the cross-section of the plasma that is viewed at the edge of the ring, since the glow at that point is already at least as strong as at the center in the upper sheath. Since it is possible that the ring geometry obscured some of the physics, we looked at bars which completely traversed the plasma as a simple limiting case. Since the bar traverses the whole plasma, the effect of a limited cross-section for viewing the perturbation is not important.

The bars used for this study were made of either Teflon or Aluminum, and were 1" wide and 3/8" high. A bar was fabricated with a 45° angle to see how important the right angle is to the perturbation of the plasma. The angled bar is used in two different configurations: up, where the bar forms an angled step of 135° with respect to the electrode, and down, where the bar forms a retrograde angled step of 45° with respect to the electrode. For reference, the geometry with the bar angled up is shown in Figure 5-1 b). The plasma emission for aluminum bars with all three geometries is shown in Figure 5-5. The striking feature is that the corners do greatly enhance the plasma emission well above the base level which is seen just from the glow of the bulk/sheath interface. In Figure 5-5 a), the bar is symmetric with respect to the centerline of the plasma, and the plasma glow is symmetric except for a slight mis-alignment of the bar with the collection optics. The plasma sheath follows the surface of the bar pretty consistently, and is not significantly perturbed by the upper edge of the bar. The corners represent an increase of more than a factor of two in the plasma emission.

The enhanced emission is not surprising in light of the model results showing that the ionization rate should be enhanced where the two surfaces meet, since the sheaths along both surfaces will interact complexly. In the case where both surfaces are conductors it is reasonable

to view the interaction of the sheaths as being additive at a small enough length scale, since in the absence of a corner there would be separate non-interacting sheaths on each surface. The result of the doubling of the sheath at the corner is that the electrons are pushed out with a field which is much greater, and energy coupling at the corner is much higher.




Figure 5-5 Plasma perturbations by Al bars placed on grounded electrode at 1 Torr.

In light of this explanation, the results in Figure 5-5 b) and c) are easily explained. Since the effect of the corner is a vector summation of the field which couples energy into the electrons, as the angle decreases the effect becomes higher. Even though the field should be excluded at a small enough gap, the electrons will still be pushed out more strongly from any point they can reach. With greater block angles, the effect lessens until at 180° the effect disappears. Note that the intensity from the corner furthest to the right is nearly identical in all 3 configurations, as is the sheath brightness where it is not perturbed. This indicates that there is no long range coupling, and over a length scale of 1" in the plasma the geometry does not interact.

Equivalent emission contours are shown for Teflon bars using the same configurations in Figure 5-6. Even though the sheath on the insulating surface has a very different structure, as is apparent from the emission profiles above the Teflon blocks, the effect in the corner is very similar. The plasma is still concentrated in the corner, and the intensity of the perturbation still is inversely proportional to the angle between the block and the electrode. The biggest difference for the Teflon block is that the plasma no longer follows the contour of the block, but is weaker along the insulating surface. The Teflon bars also have a smaller effect on the emission, and the intensity is less perturbed for every case.




Figure 5-6 Plasma perturbations by Teflon bars placed on grounded electrode at 1 Torr.

In fact there is a very interesting effect along the upper edge of the insulating surface. For the aluminum bars the plasma followed the surface contour, and decreased monotonically toward the center of the discharge. There appears to be a weak glow directly above the Teflon bars,

which is followed by a region which has decreased glow with respect to the bulk plasma. The most likely explanation for this phenomenon is that the insulator surface is charging up, and the result is a weaker field directly above.

The perturbation is significantly lessened at lower pressure, as shown in Figure 5-7 and Figure 5-8. All of the features persist at the lower pressures, but the enhancement of the ionization is diminished, and the length scale over which the perturbation exists increases. The effect of angle is the same as at higher pressures, although the upper angle now apparently decreases the plasma emission slightly. These effects are mostly due to the fact that the block is now approaching the length-scale of the sheath which it is perturbing, as was seen with the very thin rings. However, it is important to note that the plasma is still perturbed by the small angle even though it is excluded from most of the area below the angle. This indicates that the perturbation in the electric field from the corner reaches out into the plasma at least one sheath thickness.




Figure 5-7 Plasma perturbations by Al bars placed on grounded electrode at 100 mTorr.

The Teflon bars still perturb the plasma differently than the aluminum bars, and there is still a hole in the plasma above the Teflon bars. At 100 mTorr, the inverted angle is perturbing the plasma much less than at 1 Torr, and the increase in intensity over the right angle bar is almost within the experimental error. For this low pressure, since the field can penetrate the Teflon, the inverted corner has very little additional effect.




Figure 5-8 Plasma perturbations by Teflon bars placed on grounded electrode at 100 mTorr.

Chapter 6 : Conclusions and Future Work

6.1: Conclusions

A suite of efficient numerical techniques were developed to solve the continuum model equations for both one- and two- dimensional simulations. The spectral element technique allows use of relatively coarse discretizations while maintaining very highly accurate solutions. Spatial discretizations designed to work well for hyperbolic equations contain artificial diffusion which leads to significant differences in the resultant solutions. Application of the spectral element technique allows calculation without this limitation, and allows rigorous computation of the exact plasma solutions. Convergence of the spectral method with grid refinement is possible even for the two-dimensional case because a solution approaching the exact solution can be readily calculated by merely doubling the order of the interpolating polynomial. Two-dimensional simulations show spectral convergence in the spatial grid except at the corner, where there is a discontinuity in the boundary conditions. The effect of the discontinuity is limited to points very close to the corner node, and overall convergence of the method is still spectral.

Direct calculation of the time-periodic steady-state is at least an order-of-magnitude more efficient than following the full transients in the simulations. The combination of direct calculation of the TPSS and the spectral discretization results in a more efficient and accurate simulation effort at the GEC benchmark. For the two-dimensional simulation, the cost of the TPSS Newton shooting algorithm is further reduced by iteratively solving the Newton problem instead of computing the full TPSS Jacobian. The cost of the two-dimensional simulations is minimized by using a split explicit/implicit integrator instead of the fully implicit integrator used in the one-dimensional simulations. The final result is a nearly equal split between the cost of calculating the change in the solution over one cycle, and the cost of calculating one TPSS Newton step.

The aspect ratio of the plasma is the most important quantity in determining the overall uniformity of the discharge, both for symmetric and asymmetric discharges. The perturbations from the insulating wall bounding the plasma become more important at lower pressures, and are less severe as the area ratio of the electrodes increases. The presence of the insulating wall can

account for nonuniformity in the ion flux up to one gap-spacing away, which would directly affect the radial uniformity during etching processes. For the asymmetric geometries, which include an insulator between the smaller electrode and the insulating wall confining the plasma, radial nonuniformity decreases, but is most significant next the electrode/insulator interface. Plasmas with aspect ratios greater than one approach the one dimensional solution within a few percent a gap spacing away from the insulating wall, and can be modeled as one dimensional further away. Plasmas with aspect ratios near one will not approach the one dimensional solution at any point, and exhibit a decreased density at the radial centerline.

The asymmetric geometry used for this study results in a plasma which is defined by the smaller electrode, and therefore does not look highly asymmetric. The singularity at the boundary between the smaller electrode and the insulator results in a more challenging simulation, but does not appear to greatly affect the local physics. However, this limited the simulations to higher pressures where a stable spatial grid was easily obtained. There is a strong coupling between the potential near the coupling, the electron properties, and the surface charging on the insulator adjacent to the small electrode which results in a very stiff coupling for the time periodic shooting problem. A similar difficulty exists in defining the DC floating potential on the blocking capacitor between the voltage source and the larger electrode. However, the simulation is not sensitive to the exact value of any of these variables.

Comparison of the two-dimensional model to experimental results shows good agreement for the pressure range where the continuum model is expected to work best. All of the major trends in the PIE are correctly predicted for high aspect ratio plasmas. Edge effects may become important enough in the lower aspect ratio plasmas to force the comparison to break down. Comparison to Langmuir probe studies indicates that the model is predicting the right trends, but experimental uncertainty limits direct quantitative comparison. Experimental examination of some more complex geometries shows interesting effects in the PIE at corners due to interaction of the radial and axial sheaths, which should be readily evident in model results for the more complex geometries.

6.2: Suggestions for Future Work

The greatest difficulty for the simulations is finding a grid which is not overly refined but still contains a stable solution. The simulations are very sensitive to the exact grid chosen for the plasma sheath region, and it is very easy to “waste” degrees-of-freedom assuring that the method remains stable. This is especially true in asymmetric geometries, and would be much worse on a complex, non-rectangular simulation. The electron energy presents the most problem in all cases, and might be simulated on a separate, refined grid. It would be possible to use either the same degree polynomial for all variables, or the energy could be represented with a higher order approximation. A locally refined, adaptive meshing scheme would be most useful to select and maintain a stable spatial grid, and is a necessity in order to simulate more complex geometries. This would allow a much more complete investigation of the effect of geometry on the plasma physics. The spectral element method is very useful for proving that the simulations converge with mesh refinement, but the results are generally more accurate than is justified by the complexity of the model or the accuracy of available experimental data. The accuracy of cubic polynomials is generally sufficient for plasma modeling, and the increase in computational efficiency for higher order spatial grids is small enough to be unimportant. Future efforts for two-dimensional simulation should therefore use either cubic or quadratic basis functions to represent the spatial discretization.

It was shown that the physics of the one-dimensional electronegative discharge is very complex, with the negative ions responding in the sheath on a very small length scale analogously to electrons in the electropositive case. This structure is on a very small length scale with respect to the overall plasma dimension, and is very difficult to simulate in two-dimensions, especially since the bulk of the plasma is not as simple as for electropositive cases. It would be of great interest to find a substitute set of transport parameters or boundary conditions which would eliminate the negative ion sheath behavior in the simulation. Since the discretizations for electronegative discharges are already much finer than for electropositive discharges, a nonconforming mesh scheme would prove most effective in limiting the computation cost. This would then allow simulation of many other complex and reactive gases which are generally used for industrial applications. The paradigm for modeling electronegative gases has been SF_6 , but most gases of industrial interest are generally not as strongly attaching.

The should be a significant difference in the plasma physics when the concentration of the two negative species are of the same order of magnitude, and there are multiple references in which such plasma are shown to have bifurcations to two or more spatial configurations. If the model, as it is posed in this work, is capable of predicting these bifurcations, it would be excellent proof that sufficient physics have been included. A full set of accurate rate and transport parameters for CF_4 or C_2F_6 would also be of great benefit for this purpose.

The operator splitting for the time integrator worked quite well, but could not be implemented for the TPSS Newton shooting algorithm during this project. It is not clear whether this was due to a programming error or is related to the sensitivity of the insulator charge to the local variables. In the final form, the cost of the simulations has nearly equal contributions from the time integration and from the Newton shooting step. Since the Newton shooting step is based on the integrator used to calculate the change of the nonlinear system over an RF cycle, any improvement in the time integrator for the nonlinear system translates into an equivalent improvement in the Newton shooting computation. Proper implementation of the integrator splitting for the TPSS Newton shooting algorithm would reduce storage and computational requirements by approximately a factor of 4, making the simulation much more useful for equipment design and to further investigate the two-dimensional plasma physics. A further result would be that the simulations should then be readily performed on a high end workstation.

Since both the calculation of the transient simulations for a single RF cycle and the GMRES Newton shooting step have similar costs, totally eliminating the GMRES calculation will at best result in a factor of three enhancement of the simulation speed. Preconditioning the GMRES iterations is therefore useful, but will not have a huge effect on the overall simulation cost. The most efficient way to reduce the cost of the simulations would be to find an alternative time integrator which would either reduce the bandwidth or the total degrees-of-freedom used in the matrix algebra. This is the driving idea behind both the semi-implicit scheme which was not used in this work, and the integrator splitting which was used for the two-dimensional simulations. Similarly, any technique which allows the use of a coarser spatial discretization will greatly enhance the simulation speed.

References

Barnes, M.S., Cotler, T.J. and Elta, M.E., 1988, A staggered-mesh finite-difference numerical method for solving the transport equations in low pressure RF glow discharges. *J. Comp. Phys.* **77**, 53-72.

Barnes, Michael S., Cotler, Tina J. and Elta, Michael E., 1987, Large-signal time-domain modeling of low-pressure rf glow discharges, *J. Appl. Phys.* **61** (1), 1 January, 1987, 81-89.

Boeuf, J.-P., 1987, Numerical model of rf glow discharges, *Phys. Rev. A.*, **36**(6), 2782-2792.

Boeuf, J.-P., 1988, A Two-Dimensional model of dc glow discharges, *J. Appl. Phys.*, **63**(5), 1342-1349.

Brenan, K. E., Campbell, S. L. and Petzold, C. R. 1989. *Numerical Solution of Initial-Value Problems in Differential Algebraic Equations*, North Holland, New York.

Dalvie, M., Surendra, M., and Selwyn, G.S., 1993, Self-consistent Fluid Modeling of R.F. Discharges in Two Dimensions, *Appl. Phys. Let.*, **62**(24), 3207-3209.

Doedel, E.J. and Heinemann, R., 1983, Numerical Computation of Periodic Solution Branches and Oscillatory Behavior of the Stirred Tank Reactor with A->B->C-> Reaction. *Chem. Eng. Sci.*, **38**, 1493-1499.

Finlayson, Bruce, A., 1980, *Nonlinear Analysis in Chemical Engineering*, McGraw Hill, New York.

Godyak, V.A., Piejak, R.B., and Alexandrovich, B.M., 1992, Measurements of electron energy distribution in low-pressure RF discharges, *Plasma Sources Sci. Technol.*, **1**, 36-58.

Godyak, V.A., Piejak, R.B., and Alexandrovich, B.M., 1993, Probe diagnostics of non-Maxwellian plasmas, *J. Appl. Phys.*, **73**(8) 15 April 1993, 3657-3663.

Gogolides, E. and Sawin, H.H., 1992, Continuum modeling of radio-frequency glow discharges. I: Theory and results for electropositive and electronegative gases. *J. Appl. Phys.* **72**, 3971-3987.

Gogolides, E., Sawin, H.H and Brown, R.A., 1992, Direct calculation of time-periodic states of continuum models of radio-frequency plasma, *Chem. Eng. Sci.*, **47** (15/16), 3839-3855.

Gogolides, Evangelos, Nicolai, Jean-Philippe and Sawin, Herbert H., 1989, Comparison of experimental and model predictions for radio-frequency Ar and SF₆ discharges, *J. Vac. Sci. Technol. A* **7** (3), May/Jun 1989, 1001-1006.

Gottlieb, D., and Orszag, S. A., Numerical Analysis o Spectral Methods. *SIAM*, 1977.

Graves, D.B., 1989, Plasma processing in microelectronics manufacturing: a chemical engineering perspective. *J. Amer. Inst. Chem. Eng.* **35**, 1-29.

Graves, David B. and Jensen, Klavs F., 1986, A continuum Model of DC and RF Discharges, *IEEE Trans. on Plasma Sci.*, **PS-14** No. 2, April 1986, 78-90.

Huppert, G. L., Brown, R. A. and Sawin, H. H., 1992, RF Discharge benchmark Modeling, presentation at the 45th annual Gaseous Electronics Conference, Boston MA.

Huppert, G. L., Sawin, H. H., and Brown, R. A., 1993, Spectral Element Analysis of Radio-Frequency Glow Discharges, *Chem. Eng. Sci.* **49** (10), 1601-1611.

Irons, B.M., 1970, A Frontal Solution Program for Finite Element Analysis, *Int. J. Num. Meth. Engng*, **2** 5-32.

Keller, H.B., 1976, Numerical solution of two-point boundary value problems. *SIAM*, Regional Conference Series in Applied Mathematics, Philadelphia.

Keller, H.B., 1987, *Lectures on Numerical Methods in Bifurcation Problems*, Tata Institute of Fundamental Research, Bombay.

Kevrekidis, I.G., Schmidt, L.D. and Aris, R., 1986, Some common features of periodically forced reacting systems. *Chem. Engng. Sci.* **41**, 1263-1276.

Liu, J., 1992, *Scaling Relationships for Power Deposition and Ion Bombardment in Radio-Frequency Plasmas*, Ph.D. Thesis, MIT.

Maday, Y. and Patera, A.T., 1989, Spectral element methods for the incompressible Navier-Stokes equations. in *State-of-the-Art Survey on Computational Mechanics*. *ASME Book No. H00410*, 71-143.

Maday, Y., Patera, Anthony T., and Ronquist, Einar M., 1990, An operator-Integration-Factor Splitting Method for Time-Dependent Problems: Application to Incompressible Fluid Flow, *Journal of Scientific Computing*, **5**(4), 263-292.

Okazaki, Katsuji, Makabe, Toshiaki and Yamaguchi, Yukio, 1989, Modeling of a rf glow discharge plasma, *Appl. Phys. Lett.* **54** (18), 1 May 1989, 1742-1744.

Park, Sang-Kyu and Economu, Demetre J., 1990, Analysis of low pressure rf glow discharges using a continuum model, *J. Appl Phys.* **68** (8), 15 October 1990, 3904-3915.

Paranjpe, Ajit P., McVittie, James P., and Self, Sidney A. 1990, Algorithms for numerical simulation of radio-frequency glow discharges, *Phys. Rev. A*, **41**(12), 6949-6962.

Passchier, J. D. P. and Goedheer, W. J., 1993, A two-dimensional fluid model for an argon rf discharge, *J. Appl. Phys.* **74** (6) 15 September 1993, 3744-3751.

Richards, Albert D., Thompson, Brian E., and Sawin, Herbert H., 1987, Continuum modeling of argon radio frequency glow discharges, *Appl. Phys. Lett.* **50** (9) 2 March 1987, 492-494.

Strang, G., and Fix, G. J., *An Analysis of the Finite Element Method.* Prentice-Hall, 1973.

Sudit, Isaac D., and Woods, R. Claude, 1994, A study of the accuracy of various Langmuir probe theories, *J. Appl. Phys.*, **76**(8), 15 October 1994, 4488-4498.

Surendra, M., 1985, Langmuir Probe Theory and its Application in Radio Frequency Excited Plasmas, MIT undergraduate thesis.

Surendra, M., 1995 , Radiofrequency discharge benchmark model comparison, *Plasma Sources Sci. Technol.* **4** , 56-73.

Thomasset, F., 1981, *Implementation of Finite-Element Methods for Navier-Stokes Equations.* Spriner-Verlag, New York.

Thompson, Brian E. , and Sawin, Herbert H., 1986, Comparison of measured and calculated SF₆ breakdown in rf electric fields, *J. Appl. Phys.* **60** (1), 1 July 1986, 89-94.

Tsai, J. H., and Wu, C., 1990, Two-dimensional simulations of rf glow discharges in N₂ and SF₆, *Phys. Rev. A* **41** (10), 15 May 1990, 5626-5644.

Turkot, R.B., and Ruzic, D.N., 1993, Time- and space-resolved radio-frequency plasma electron energy distributions from a displacement-current-based electric probe diagnostic, *J. Appl. Phys.*, **73**(5), 1 March 1993, 2173-2179.

Young, Fongray F. and Wu, Chwan-Hwa, 1993, Radial flow effects in a multidimensional, three-moment fluid model of radio frequency glow discharges, *Appl. Phys. Lett.* **62** (5), 1 February 1993, 473-475.

Young, Fongray Frank, and Wu, Chwan-Hwa "John", 1993a, Two-Dimensional, Self-Consistent, Three-Moment Simulation of RF Glow Discharges, *IEEE Trans. Plasma Sci.*, **21**(3), 312-321.