

**KINETIC MODELING AND LOWER HYBRID CURRENT DRIVE IN THE REVERSED
FIELD PINCH**

by

David R. Burke

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2010

© Copyright by David R. Burke 2010

All Rights Reserved

ACKNOWLEDGMENTS

This work would not have been accomplished without many contributions from a number of different people, whose support and help made this work possible. I would like to thank all the graduate students and staff on MST, specifically Matt Miller, and Hilary Stephens for friendly conversation, and help on various little scientific matters throughout my graduate career. Likewise, thanks to John Sarff, and Karsten McCollum and Bill Zimmerman for their input and ideas.

Thanks to Charlie Freeman for sparking my interest in Physics at an advanced level. Thanks also to Ed Pogozeleski, Savi Iyer, David Meisel and Charles Fletcher for nurturing that spark. Thanks to my plasma professors, Carl Sovinec, Chris Hegna, and David Anderson. I will proudly say that UW Madison has the best academic plasma curriculum of any school, and credit for that belongs completely to the talented and dedicated professors, who should also be thanked for many fruitful discussions long after their classroom obligation to me had ended.

Many thanks to Josh Reusch for G-S reconstructions, and for listening to and discussing many random crazy ideas, some without merit, some with merit, none dismissed.

Thanks to Bob Harvey for his authorship and help with CQL3D. Thanks to Jay Anderson, for thoughtful conversations and indulging a few flights of fancy. Thanks to Rob O'Connell for his help getting me started on CQL3D. I only wish that I had known then what I know now. Thanks to Abdulgader Almagri, for time, his help and advice, as well as his friendship. All have proven equally invaluable. Thanks to my compatriots Mike Kaufman and Dan Clayton for many discussions both productive and otherwise.

Thanks to John Kulpin, for all his expertise, tinged with exasperation. Your care and precision were important to the success of the lower hybrid project. Look on the bright side!

Thanks to John Goetz, who guided me and watched my transformation from an optimist that saw evidence of current drive in everything to the dreary pessimist I am today. Have a victory snap, buddy.

Thanks to Cary Forest, my advisor, a man whose creativity is matched only by his exuberance and enthusiasm. I have rarely met someone who was simultaneously so well versed in a topic and yet so open to new ideas.

Finally, and most importantly, thanks to my lovely and talented wife Bonita, without whom this thesis would never have been completed. You are the best thing that has ever happened to me, and I can't wait for us to begin our lives together.

TABLE OF CONTENTS

	Page
LIST OF TABLES	vi
LIST OF FIGURES	vii
NOMENCLATURE	xii
ABSTRACT	xiii
1 Introduction	1
1.1 Motivation for Fusion Experiments	1
1.2 The Physics of Fully Ionized Gases	4
1.2.1 Plasmas and Magnetic Confinement	4
1.2.2 Ideal Magnetohydrodynamics	5
1.2.3 Grad-Shafranov Equilibrium	6
1.2.4 Stability in the RFP	7
1.2.5 RFP Profiles and Confinement	9
1.2.6 Plasma rf experiments	12
1.2.7 Electron Bremsstrahlung	13
1.3 Theory of Lower Hybrid Current Drive on MST	14
1.3.1 Landau Damping	15
1.3.2 Wave Propagation	18
1.3.3 Quasilinear absorption	30
1.3.4 Inductive effects	35
1.3.5 System Design	39
1.4 Wave coupling	44
1.5 Non-Local Ohm's Law	46
1.6 Objectives of this work	49
2 Experimental Setup	51
2.1 The Madison Symmetric Torus	51
2.1.1 Operational Modes	53

	Page
2.1.2 Diagnostics	56
2.1.3 Reconstructions	62
2.1.4 X-ray Spectroscopy	62
2.2 The Lower Hybrid System	67
2.2.1 Support Systems	68
2.2.2 The Interdigital Traveling Wave Lower Hybrid Antenna	74
3 Kinetic Simulations and Radial Transport in the RFP discharges	86
3.1 CQL3D	86
3.1.1 Radial Diffusion	88
3.1.2 Electric Field	90
3.2 Fast Electrons in RFP discharges	93
3.3 Methods of Experimental verification of Fokker-Planck predictions	96
3.3.1 Model X-ray flux predictions	97
3.3.2 Data Interpretation	98
3.3.3 Stimulated Bremsstrahlung Emission	102
3.4 Inductively stabilized (PPCD) plasmas	105
3.5 Standard plasmas	107
3.5.1 Diffusive effects in Ohm's law	108
3.5.2 Enhanced Confinement Plasmas	114
3.5.3 Models of standard plasmas	116
4 Lower Hybrid Modeling and Experiments on the MST	122
4.1 Wave Propagation	123
4.1.1 GENRAY	123
4.1.2 The Real LH Antenna	124
4.1.3 Antennas for full stabilization experiments	129
4.2 Wave absorption	130
4.2.1 $Z_{effective}$	134
4.2.2 Radial Diffusion	134
4.2.3 Broad Spectra	136
4.2.4 Electric Field	139
4.2.5 Fast Electron Routing	140
4.3 Evidence of LH current drive	141
4.4 Current Drive for Full Stabilization	151
4.5 Alternative Scenarios	154
4.5.1 The intermediate system	154
4.5.2 The Non-inductive RFP	155

	Page
5 Coupling Power to Improved Confinement Discharges	167
5.1 Coupling to the slow mode in MST plasmas	167
5.1.1 Density Dependence	167
5.1.2 Fast-Slow Cutoff	171
5.2 Coupling to Improved Confinement Plasmas	174
5.3 Localized Gas Puffing Experiments	178
6 In Summary	183
6.1 Review and discussion of original work	183
6.1.1 Introductory Material (Chapters 1 and 2)	183
6.1.2 Kinetic Modeling (Chapter 3)	184
6.1.3 Lower Hybrid Current Drive experiments and modeling (Chapter 4)	185
6.1.4 Gas Puffing to control power coupling (Chapter 5)	187
6.2 Future Work	187
REFERENCES	190
 APPENDICES	
Appendix A: Derivation of the relativistic target probe formula	199
Appendix B: Assumptions about field lines	200
Appendix C: The Cold Plasma Model	202
Appendix D: Electrostatic Dispersion Relation	208
Appendix E: The raytracing equations	212
Appendix F: Lower Hybrid coupling	213
Appendix G: Some key error analysis	216

LIST OF TABLES

Table	Page
1.1 Ideal MHD Equation	6
2.1 MST Parameters	51
2.2 Comparison of waveguide spectra to interdigital spectra	77
2.3 Power Loss fraction, Bench measurements	82
4.1 Power requirements for stabilizing current drive	152
4.2 Experimental efficiency for various machines compared to MST	153
4.3 Parameters for an LH upgrade	154
6.1 Slopes for Linear scaling relationship, LH in PPCD.	186

LIST OF FIGURES

Figure	Page
1.1 Binding Energy per nucleon for common isotopes	3
1.2 Magnetic Shear in the RFP	8
1.3 Model RFP Fields	9
1.4 Stochasticity in RFP Fields	11
1.5 Approximate Scales	16
1.6 Scaling of Landau Damping parameter with thermal velocity	17
1.7 Types of motion which give rise to Cold Plasma waves	19
1.8 Cold Plasma Dispersion Relation regions for MST-like Parameters	23
1.9 Cold Plasma Modes	25
1.10 Stratified Slab results	27
1.11 Poloidal Lensing	29
1.12 Simple 1-D distribution function for Quasilinear current drive	33
1.13 Amperian Loop for Inductance Calculation	36
1.14 400 kA Standard plasma cutoff diagram and outward propagation results	41
1.15 Multiseptum Waveguide	43
1.16 Cartoon of antenna coupling	45
1.17 Impact of the diffusive Ohm's law, Martines model	50

Figure	Page
2.1 Madison Symmetric Torus	52
2.2 Profiles from a standard discharge	54
2.3 Profiles from a PPCD discharge	55
2.4 Profiles from an EC discharge	57
2.5 FIR Chords on MST	58
2.6 Thomson Scattering	60
2.7 Thomson Scattering diagnostic viewing chords	61
2.8 X-ray detector setup	63
2.9 Hard X-ray detectors in place on a boxport flange	65
2.10 Raw Soft X-ray Data	66
2.11 MST LH Power System	69
2.12 Schematic of a Klystron	70
2.13 Klystron Operational Signals	71
2.14 Lower Hybrid Transmission System	73
2.15 Slow Wave antenna structure	75
2.16 Spectra for various antenna types	77
2.17 Schematic sketch of Phase electronics	79
2.18 Vacuum Spectrum	79
2.19 Spectra with damping	80
2.20 Predictions of the 0-D measurement model	85
3.1 Sample CQL3D velocity grid	88
3.2 Flow Chart for CQL3D constant electric field mode	91

Figure	Page
3.3 Flow Chart for CQL3D constant current mode	92
3.4 Consistency between Constant Current Mode and Constant Electric Field Mode	94
3.5 Runaway electron Populations versus field strength	95
3.6 Generation of a Runaway Current	96
3.7 Schematic of X-ray Target Probe	103
3.8 X-ray flux across multiple chords in a PPCD discharge	106
3.9 Parallel Electric field and Current Profile, Standard Plasmas	107
3.10 Diffusion acting on a delta function electric field	109
3.11 Diffusive Ohm's law, CQL3D	111
3.12 Effect of stochastic diffusion on the distribution function	112
3.13 Specific current density showing the impact of stochastic diffusion	113
3.14 Parallel Current profile Evolution during EC Periods	115
3.15 Stochastic model of EC plasmas	117
3.16 Comparison of standard plasma models with measured X-ray Spectra and current profile	119
3.17 X-ray spectrum from Target Probe in standard plasmas at $\rho = 0.9$	121
4.1 Sample GENRAY launch spectrum	124
4.2 Wave Trajectory - Standard Plasma, Real Antenna	125
4.3 Wave Propagation Diagrams - Standard Plasma, Real Antenna	126
4.4 Conversion between the slow and fast mode for Standard plasmas	128
4.5 Wave Propagation diagrams, optimized antennas	131
4.6 Distribution Function Flattening from CQL3D	133
4.7 Current Drive Degradation Due to Z_{eff}	135

Figure	Page
4.8 The Effect of Diffusion on current drive	136
4.9 Impact of Stochastic diffusion on rf driven current	137
4.10 Scaling of current drive peak versus Diffusion	138
4.11 Effects of spectral broadening	139
4.12 Current Amplification due to applied electric field	140
4.13 Fast Electron Routing	141
4.14 Effect of Fast Electron Routing on Current	142
4.15 Toroidal current driven in standard plasmas	144
4.16 Predicted x-ray flux during PPCD discharges	145
4.17 Flux increase due to LH in PPCD plasmas	146
4.18 X-ray spectra for EC plasmas, with and without rf	147
4.19 X-ray flux during EC periods	148
4.20 X-ray spectrum for standard discharges with lower hybrid	158
4.21 Comparison of x-ray spectra with different models of standard plasmas, LH on	158
4.22 Lower Hybrid Associated Target Probe flux versus radial-like coordinate	159
4.23 Estimated Current required for stabilization	160
4.24 Full stabilizing current Predictions	160
4.25 Scaling of peak current with input power	161
4.26 A Possible LH upgrade	162
4.27 Wave propagation diagrams, upgrade system	163
4.28 X-ray flux and current drive, hypothetical upgrade system	164
4.29 Stochastic Transport for non-inductive current drive	164

Figure	Page
4.30 Propagation of a core damping fast wave	165
4.31 Deposition of fast wave current	166
5.1 Scaling of coupled power with electron density	168
5.2 Coupled power fraction, simple model	170
5.3 Coupled power fraction versus vacuum gap length, Golant model	171
5.4 Simple F-S reflection model	172
5.5 The effects of the Fast-Slow cutoff on antenna coupling	173
5.6 Coupling during PPCD periods	175
5.7 Coupling during EC periods	176
5.8 Edge Density Profiles for MST plasmas	177
5.9 Effects of gas puffing on coupling	179
5.10 Mode activity during PPCD with doping	180
5.11 Effects of doping during EC periods	181
D.1 Perpendicular wave number for the slow wave	210

NOMENCLATURE

q_α Electric charge of a single particle of species alpha

n_α Number density of a species alpha

m_α Mass of a single particle of species alpha

B_t Toroidal Magnetic Field

B_p Poloidal Magnetic Field

$v_{th\alpha}$ $\sqrt{nkT/m}$, most probable speed of a particle in a maxwellian. For a magnetized plasma $n = 2$.

KINETIC MODELING AND LOWER HYBRID CURRENT DRIVE IN THE REVERSED FIELD PINCH

David R. Burke

Under the supervision of Professor Cary B. Forest

At the University of Wisconsin-Madison

The reversed-field pinch (RFP) magnetic confinement device has been shown to be susceptible to tearing fluctuations, which lead to the growth of magnetic islands. When these islands overlap, magnetic field lines reconnect, and wander stochastically from the core of the RFP to the edge. This stochastic field line wandering decreases energy confinement as hot particles from the plasma core follow the magnetic field lines to the cold edge. Tearing fluctuations can be stabilized by localized current drive. Lower hybrid wave injection is proposed as a source of this current. Kinetic modeling is employed to make predictions about lower hybrid wave absorption and associated x-ray emission. The strengths and weaknesses of the model used are first explored by applying it to non-rf RFP discharges. The model is shown to provide good results for inductively stabilized plasmas, and a certain subclass of plasmas with stochastic field lines, but it is found to be somewhat deficient for typical discharges. Scaling relationships between current density and injected power are presented, including the effects of a background electric field and radial transport. Stochastic plasmas are shown to be unable to support a sufficiently localized current profile. Inductively stabilized plasmas do provide sufficient confinement, however antenna coupling is not efficient. Low power operations of a lower hybrid antenna are in good agreement with model predictions, and a bremsstrahlung target probe is used to prove that a significant population of electrons is being generated within the plasma. Finally, methods of improving coupling between the antenna and the plasma are developed, studied and described. It is shown that the reduction in coupling associated with improved confinement plasmas can be mitigated by edge fueling.

Cary B. Forest

ABSTRACT

The reversed-field pinch (RFP) magnetic confinement device has been shown to be susceptible to tearing fluctuations, which lead to the growth of magnetic islands. When these islands overlap, magnetic field lines reconnect, and wander stochastically from the core of the RFP to the edge. This stochastic field line wandering decreases energy confinement as hot particles from the plasma core follow the magnetic field lines to the cold edge. Tearing fluctuations can be stabilized by localized current drive. Lower hybrid wave injection is proposed as a source of this current. Kinetic modeling is employed to make predictions about lower hybrid wave absorption and associated x-ray emission. The strengths and weaknesses of the model used are first explored by applying it to non-rf RFP discharges. The model is shown to provide good results for inductively stabilized plasmas, and a certain subclass of plasmas with stochastic field lines, but it is found to be somewhat deficient for typical discharges. Scaling relationships between current density and injected power are presented, including the effects of a background electric field and radial transport. Stochastic plasmas are shown to be unable to support a sufficiently localized current profile. Inductively stabilized plasmas do provide sufficient confinement, however antenna coupling is not efficient. Low power operations of a lower hybrid antenna are in good agreement with model predictions, and a bremsstrahlung target probe is used to prove that a significant population of electrons is being generated within the plasma. Finally, methods of improving coupling between the antenna and the plasma are developed, studied and described. It is shown that the reduction in coupling associated with improved confinement plasmas can be mitigated by edge fueling.

Chapter 1

Introduction

The reversed field pinch (RFP) magnetic confinement device has been proposed as a potential fusion reactor for generating electric power. The RFP has many advantages over other proposed magnetic fusion designs, particularly its lower applied magnetic field. A lower magnetic field can be generated by less expensive and less sophisticated magnets, which are cheaper. It is feasible that copper can be used to generate the required field rather than superconducting magnets, which are vulnerable to damage by the high energy neutrons created by fusion reactions. The RFP is susceptible to large scale tearing fluctuations however, which lead to dramatically increased transport from the hot dense core to the cold plasma edge. These tearing modes are found to be driven by the gradient in the parallel current profile. Therefore, localized current drive is proposed as a means of stabilizing these modes. In the past, rf waves have been used to drive current in tokamak plasmas. The lower hybrid slow mode in particular has been observed to create a fairly localized region of current deposition. This is proposed as a means of stabilizing RFP tearing modes. The current work focuses on the feasibility of this scheme, and compares the predictions of theoretical models to experimental data.

1.1 Motivation for Fusion Experiments

Energy expenditure is directly related to standard of living in modern industrialized nations. Currently, energy production around the world relies primarily on non-renewable resources, particularly hydrocarbon fuels such as gas, coal and oil. These are undesirable because of their limited supply and their potential effect on health and the environment. Many alternative energy sources

have been proposed such as solar and wind power. Unfortunately, both solar panels and wind turbines are energy intensive to construct and subsequently dispose of, and both technologies serve under the caprices of inconsistent fuels. A cloudy, windless day leaves these sources of power silent. Moreover, the most efficient locations for power generation (broad uninterrupted hills and deserts) are physically remote from the primary consumers of electric power, namely dense urban and industrial areas. Nuclear power is suggested to supplement these other technologies and provide base load power consistently and efficiently.

The heart of nuclear fuel cycles is the strong nuclear force and the binding energy of nuclei. The strong force binds protons and neutrons together. The magnitude of the strong force far exceeds that of electric (Coulomb) repulsion. Unlike the Coulomb force, the strong nuclear force is very short ranged, and so does not extend beyond the nucleus of an atom. Just as electrons exist in discrete energy levels around nuclei, nucleons (that is protons and neutrons), bound together by the strong force, exist in distinct energy levels within a nucleus. Because these levels are created by an attraction due to the strong force (as opposed to the electromagnetic force), they are much higher energy than the levels of a bound electron. The energy required to break up a nucleus is called the binding energy. This is identical to the amount of energy liberated when a nucleus is assembled. The key insight of nuclear fuel is gained by examining the amount of binding energy per nucleon for various isotopes of common elements (Figure 1.1). The particles at the peak of this graph are described as tightly bound, because it takes a larger energy input to liberate a single nucleon than for elements away from this peak.

For a large nucleus (like Uranium or Thorium), the total nuclear binding energy is less than the energy of two smaller nuclei. For instance, ^{238}U has a binding energy of about 7.6 MeV/nucleon (a total of about -1.8 GeV). If one were to split the nucleus exactly in half, it would produce two nuclei with 119 nucleons each (for example Palladium 119), at around 8.5 MeV/nucleon, for a total of $2 \times (-1.016 \text{ GeV}) = -2.03 \text{ GeV}$. This extra 200 MeV of energy will be released as either kinetic energy, electrons, neutrons or γ -rays, depending on the specifics of the reaction. On the small nucleus side of the plot, we see that energy can be liberated by creating larger nuclei. Importantly, a good deal more energy is created per nucleon. Progressing from small elements




Figure 1.1: Binding Energy per nucleon for common isotopes

to larger elements is called fusion (as particles are being “fused” together) and progressing from larger to smaller elements is called fission.

While it is a well developed technology and satisfactory in the short term, fission is ultimately unattractive due to weapons proliferation concerns. As the standard of living in the developing world improves, energy demands will increase, and creating a network of fission power plants in less developed countries presents a latent threat to national security. Fusion on the other hand is somewhat resilient against this concern. Although fusion weapons certainly exist, ignition is only achieved by the use of fission based weaponry, thus making possession of highly enriched uranium or plutonium a prerequisite for constructing a fusion based weapon. Controlling the proliferation of these elements is a much easier task if they are not required in power plant operations. Proliferation of tritium and activated materials from fusion plants is somewhat of a concern, but not as serious as the proliferation of uranium and plutonium. A second advantage is that the basic fuel of the most accessible fusion reaction is a combination of deuterium (^2H , usually referred to as D) and tritium (^3H , referred to as T). Deuterium is readily available in the form of heavy water, as it accounts for 0.015% of all naturally occurring hydrogen [1], and tritium can be bred from abundant (albeit

industrially useful) lithium. Additionally, the byproduct of a D + T fusion reaction is helium (this is in fact why helium is naturally abundant in the sun, giving rise to its name), a very stable and nonreactive element, as opposed to cesium, rubidium and a whole host of other toxic, radioactive and/or violently reactive elements produced by fission [2].

In order to harness fusion power on a large scale, the electromagnetic repulsion between nuclear protons must be overcome, either in large quantities, for long amounts of time, or (preferably) both. The sun satisfies both of these requirements by its gravitational force, which is strong enough to hold an ionized hydrogen cloud together and cause a portion of it to heat to temperatures where the thermal (kinetic) energy can overcome electric repulsion. The two major methods of achieving fusion reactions on Earth are called inertial confinement and magnetic confinement. In inertial confinement experiments, deuterium is “shocked” by applying a high powered pulse (usually by firing a laser or directing x-ray radiation) to a solid target, creating a short lived but dense and high temperature plasma. Here, the term confinement refers to being held together long enough to undergo a fusion reaction, as this situation is not long lived due to the intense pressures involved, and the limited application of external force. This scheme is inherently pulsed, and will not be dealt with further in this work. The second and arguably more promising scheme for power generation is to heat a large population of deuterium and tritium nuclei to the point where the thermal energy can overcome the Coulomb repulsion as in the sun (hence it is described as “thermonuclear”).

1.2 The Physics of Fully Ionized Gases

1.2.1 Plasmas and Magnetic Confinement

The temperatures required for thermonuclear energy production are much higher than the ionization energy of hydrogen. Thus, hydrogen atoms in a reactor quickly dissociate into their constituent ions and electrons. When a large population of hydrogen atoms is heated to these temperatures (on the order of $10 \text{ keV} \approx 100 \text{ million Kelvin}^1$), the thermal motion of electrons and ions

¹It is found to be far more convenient to use units of energy to represent the temperatures of plasma species (i.e. ions or electrons). The conversion is accomplished by simply applying the Boltzmann constant, $k = 8.617 \times 10^{-5} \text{ eV/K}$, i.e. $T_{eV} = kT_K$

is so great that recombination into hydrogen atoms is not energetically favorable, precisely in the same way that a sufficiently fast moving asteroid is not bound into orbit by the gravitational pull of the Earth. A soup of freely streaming, dissociated electrons and ions forms, which is called a plasma. Because these particles are both charged and fast moving, a plasma generates and is affected by electric and magnetic fields. Due to the extreme heat, it is necessary to isolate these particles from its surroundings. A natural way to contain such a plasma is by applying a large magnetic field. Charged particle motion in a magnetic field is dictated by the Lorentz force, specifically $F_m = q\mathbf{v} \times \mathbf{B}$, where q is the particle charge, \mathbf{v} is the velocity and \mathbf{B} is the magnetic field. As a consequence of the cross product, this force is perpendicular to the particle motion and the magnetic field. Therefore, motion parallel to the magnetic field is unrestricted, and the particle is forced to undergo a circular orbit around the magnetic field direction. This type of motion is called “cyclotron motion”. In deference to the importance of this effect, plasma dynamics are usually divided into two directions: “parallel”, along the field and “perpendicular”, around the field (this convention will be used throughout this work). Cyclotron motion typically occurs at a very high frequency ($\omega_{c\alpha} = q_\alpha B / m_\alpha$, $\mathcal{O} \sim \text{GHz}$ for electrons, $\mathcal{O} \sim \text{MHz}$ for ions, where q_α is the charge of the particle B is the magnetic field magnitude and m_α is the particle mass) and very small length scales ($r_{cyc\alpha} = v_\perp / \omega_{c\alpha}$, $\mathcal{O}(r_e) \approx 0.1 \text{ mm}$, $\mathcal{O}(r_i) \approx 1 \text{ cm}$, where v_\perp is the velocity perpendicular to the field direction), though these values vary considerably with B and T (the temperature). Because the particles restricted in this way carry their own charge however, redirecting their motion alters the total magnetic field. All of the fundamental challenges in plasma physics arise from the complex interchange of externally applied and internally generated fields. Particles are still free to move parallel to the magnetic field, and so field lines are often formed into closed circles or nested helices. Because of this restriction, fusion reactors are typically bent around to form a torus.

1.2.2 Ideal Magnetohydrodynamics

The most highly developed and simplest effective model of plasma dynamics is ideal magnetohydrodynamics (MHD). The core conceptual basis of the model is derived by considering a plasma as a single perfectly conducting fluid. Whereas a true laboratory plasma consists of upwards of

$d\rho/dt = -\rho\nabla \cdot \mathbf{v}$	Continuity equation
$\rho (d\mathbf{v}/dt)_\perp = \mathbf{J} \times \mathbf{B} - \nabla_\perp p$	Force balance (perpendicular direction)
$\rho (d\mathbf{v}/dt)_\parallel = -\hat{\mathbf{b}} \cdot \nabla p$	Force balance (parallel direction)
$dp/dt = -\gamma p \nabla \cdot \mathbf{v}$	Adiabatic Compression
$\mathbf{E} + \mathbf{v} \times \mathbf{B} = 0$	Ohm's Law
$\nabla \times \mathbf{E} = -\partial \mathbf{B} / \partial t$	Faraday's Law
$\nabla \times \mathbf{B} = \mu_0 \mathbf{J}$	Ampere's Law (non-relativistic)
$\nabla \cdot \mathbf{B} = 0$	Magnetic Divergence Constraint

Table 1.1: The 14 ideal MHD equations listed with the equation from which they are derived.

10^{21} particles of various energy, charge and mass, ideal MHD abstracts these to a single mass density (ρ), net fluid velocity (\mathbf{v}), and pressure (p). Because the various species can have different net velocities and charges, there is a non-zero net current density (\mathbf{J}), which gives rise to and interact with electromagnetic fields \mathbf{E} and \mathbf{B} . Charge density is assumed to be zero. This is easily justified, because if there is charge separation within the plasma, a very large electric field will quickly develop, and relatively light, easily moved electrons will be attracted or repelled very quickly in order to cancel it out (this attribute of plasmas is called “quasineutrality”). This allows us to derive the 14 coupled ideal MHD equations [3] (given in Table 1.1), which are essentially derived from force balance, the continuity equation, momentum and energy conservation and Maxwell's Equations in the non-relativistic limit.

1.2.3 Grad-Shafranov Equilibrium

A simple model used to describe plasma equilibria can now be derived. This is accomplished by taking the MHD equations and setting the time derivatives to zero. Further, it is assumed inertia and stress are negligible. A cylindrical coordinate system is used, such that the axial dimension is aligned with the hole in the center of the torus. Symmetry is assumed in the angular (“toroidal”, ϕ) direction (this is called “axisymmetry”). What remains is a 2-dimensional slice of the total toroidal

plasma. A stream function, ψ , is then derived for the magnetic field in this plane such that

$$B_R = -\frac{1}{R} \frac{\partial \psi}{\partial Z} \quad B_Z = \frac{1}{R} \frac{\partial \psi}{\partial R} \quad (1.1)$$

where R is the major radial dimension, pointing out from the center of the torus, and Z is the vertical dimension. The ideal MHD equations then provide a nonlinear elliptical differential equation for ψ , called the Grad-Shafranov equation.

$$\Delta^* \psi = -\mu_0 R^2 \frac{dp}{d\psi} - F \frac{dF}{d\psi} \quad (1.2)$$

where $F = RB_\phi$, p is pressure, and

$$\Delta^* \psi \equiv R \frac{\partial}{\partial R} \left(\frac{1}{R} \frac{\partial \psi}{\partial R} \right) + \frac{\partial^2 \psi}{\partial Z^2}. \quad (1.3)$$

Because ψ is related to magnetic flux, magnetic field lines in the plasma are constrained to a surface with a single value of ψ , called a flux surface. Various profiles of F and p that satisfy this equation define the different common machine types, with two notable examples being the tokamak and the reversed field pinch. A somewhat heuristic understanding of the difference between these two configurations can be gleaned by analyzing the stability properties of the one dimensional simplification of an MHD equilibrium, called the screw pinch.

1.2.4 Stability in the RFP

A screw pinch is a theoretical model for one dimensional plasma equilibrium. The theoretical screw pinch is a periodic or infinite cylinder, with a magnetic field that has both an angular and an axial component (but no radial component). The current density J is also allowed to have angular and axial components. Sources for instability in a screw pinch are the current gradient and the pressure gradient. A necessary (though not sufficient) stability condition for screw pinches (including tokamaks and RFPs) is given by the Suydam Criterion, [3]

$$\frac{r B_z^2}{\mu_0} \left(\frac{q'}{q} \right)^2 + 8p' > 0 \quad (1.4)$$

Suydam's Criterion is derived by applying an energy principle to a cylindrical plasma column




Figure 1.2: A cartoon of the magnetic field pitch in an RFP. Field pitch is seen to reverse direction at about 80% of the minor radius forming what is called the reversal surface.

with axial and azimuthal symmetry. Perturbations are then assumed to have a small radial extent. If a configuration does not satisfy this criterion, it will not be stable. Examining this restriction, we see that the first term is positive definite (bringing the value farther above zero, and therefore stabilizing), because radius is defined to only be positive, μ_0 is a fixed positive number, and the other terms are squared. The second term on the other hand is typically negative (destabilizing), because the pressure is generally highest in the hot dense core, and low at the edge because the edge temperature must be low enough to not damage the experimental vessel. Therefore a stable screw-pinch type equilibrium will maximize the first term in order to mitigate the second term.

Here q is defined as

$$q = \frac{\mathbf{B} \cdot \nabla \hat{\phi}}{\mathbf{B} \cdot \nabla \hat{\theta}} \approx \frac{r B_t}{R B_p}$$

and is called the MHD safety factor because of its role in stabilizing instabilities. Unlike the tokamak, which achieves stability through its high toroidal field (B_t), the RFP is stabilized by its high relative safety factor gradient q'/q , usually called magnetic shear. This is achieved by inducing a toroidal field on the order of the poloidal field. As the RFP plasma evolves, the toroidal field is observed to pinch inward, reversing direction at the edge, creating a shear profile that is largest near the largest pressure gradient. This is illustrated in Figure 1.2. Thus an RFP is capable




Figure 1.3: A comparison of some models for the magnetic field in an RFP. The parameters for this plot are $F = -0.2$ and $\Theta = 1.69$ which are typical operating parameters in MST. The MPFM model can be shown to be significantly more accurate, particularly in the edge of the plasma, where pressure gradient effects are the largest.

of satisfying Eq. 1.4 for a given pressure gradient with a smaller field than a tokamak. This makes the RFP an attractive candidate as a fusion reactor, as the powerful field coils required to drive the high fields of the tokamak are expensive and (if superconducting) vulnerable to damage induced by the high neutron flux of a fusion reactor.

1.2.5 RFP Profiles and Confinement

RFP research was first motivated by results of the ZETA device, [4] a toroidal pinch located in the United Kingdom. In some discharges of this device, a long, quiet, stable period was observed. Many years later Taylor developed a theory [5] of self stabilization in these plasmas. Taylor's theory predicts that there exists a natural minimum energy state in which the magnetic field can be described completely by

$$\nabla \times \mathbf{B} = \mu \mathbf{B} \quad (1.5)$$

where μ is a constant. This leads to the solution (assuming a cylinder)

$$B_z = B_0 J_0(\mu r) \quad B_\theta = B_0 J_1(\mu r), \quad (1.6)$$

thus the toroidal field will reverse direction if $\mu a > k_0$, where $k_0 \approx 2.404$ is the argument of the first zero of the Bessel function. Several other models for plasma equilibria exist, for instance the Modified Polynomial function model, which attempts to account for pressure gradients (magnetic fields from these two models are plotted in Figure 1.3). The plasma equilibrium is completely characterized by two parameters, called pinch and reversal. The plasma reversal parameter (F) is the ratio of the edge toroidal field to the toroidal field averaged over the circular cross section. The pinch parameter (Θ) is the ratio of the edge poloidal field to the cross section averaged toroidal field. These two parameters constrain the equilibrium magnetic field and the current density profiles, normalized to the core toroidal field. Magnetic profiles accounting for the toroidicity in experiment can be calculated by numerically solving the Grad-Shafranov equation. [3, 6]

The unique magnetic shear of the RFP gives rise to many interesting phenomenon. Ideal MHD theory states that the most unstable surfaces in a toroidal confinement device will be surfaces where the safety factor is a rational number, usually represented as $q = m/n$. Since the safety factor of the RFP goes through zero, there are infinitely many $m = 0$ and $m = 1$ rational surfaces in an RFP discharge, which are susceptible to tearing fluctuations. These fluctuations lead to the growth of magnetic islands which grow until their width is comparable to the distance separating them. Since magnetic field lines are initially confined to the surface of these islands, when the islands expand to the point where they come nearly into contact, the field lines must reform in such a way as to preserve the zero divergence of the magnetic field. The reformed field lines wander radially in a stochastic, random-walk fashion, increasing heat transport from the core to the edge and degrading confinement. [7] This is shown in Figure 1.4. This image was generated by taking magnetic fields from a cylindrical MHD code (DEBS) [8] and following the field lines. Note, this should not be considered hard data because the interpolation is not divergence free, and is being interpolated as though the field geometry were toroidal when in fact DEBS makes a periodic cylinder assumption. This stochasticity limits the ability of the Ohmic field to heat the plasma as hot core particles are rapidly exchanged with cold edge particles. MHD simulations [9] indicate that the tearing fluctuation magnitude can be mitigated by flattening the current profile near the reversal surface, where a large number of $m = 0$ surfaces overlap. This is predicted to reduce stochasticity, thereby



Figure 1.4: These two images show the effects of stochastic field line wandering. The first image is the ideal (fluctuation free field), a set of nested toroidal helices. The second plot shows the distortion caused by tearing modes and reconnection. Line coloration in the second image is based on the origin of the field line, and coincides with the coloration in the first image. Note change in radius, particularly for the orange, blue and green field lines.

lowering particle diffusion and improving particle confinement, leading to increased temperatures in the core. Experimentally, this has been realized by inducing an additional poloidal electric field in the plasma with an ancillary external circuit. This pulsed poloidal current drive (PPCD) technique has demonstrated improved confinement coupled with core heating and increased plasma β . [10] These plasmas are also observed to have a reduced “MHD dynamo”, an MHD phenomenon associated with large magnetic fluctuations. [11] Unfortunately, this current is driven inductively and so is inherently pulsed. A major focus of present RFP research is sustained current drive. Rf waves have a long history of use in other plasma devices, and have been exploited to drive radially localized current, and so they are proposed as a scheme to provide this stabilizing current.

1.2.6 Plasma rf experiments

Wave interactions within plasmas have been of absorbing interest since the first plasma devices [12]. In 1965, T. H. Stix proposed a scheme [13] whereby fast waves would be launched into a plasma, and mode converted at a critical layer; the upper hybrid resonance to heat electrons and the lower hybrid resonance to heat ions. The frequency would be chosen to be an integer multiple of the electron or ion cyclotron frequency, and the wavelength chosen such that once mode converted to a slow wave, the wavelength would be near the cyclotron radius. Therefore, the rf wave would propagate inward quickly without damping much on the way, and then proceed to be mode converted and damp readily. In this way the core can be selectively heated.

Many experiments in lower hybrid (LH) heating were undertaken. The Stix scheme was discovered to be somewhat inefficient however. Karney [14] predicted a different absorption mechanism for heating ions, which was also pursued. Electron Landau damping [15] was also observed.

Realizing a new potential application for rf in fusion plasmas, Fisch [16] proposed exploiting the Landau resonance using LH waves to drive current, potentially reducing the reliance on an ohmically driven field. Because Landau damping preferentially accelerates electrons that are already quite a bit faster than the thermal particles, and the cross section for Coulomb collision is inversely related to speed, electrons accelerated by this method are far less collisional than their slower cousins. This means that a small number of electrons can carry a great deal of current,

and fairly efficiently since they are less likely to give their energy up to the thermal particles. Experiments were carried out in some linear devices to verify this theory, and were quite successful. [17, 18] Shortly thereafter, two major competitive efforts were launched to drive current with LH in a tokamak, one at the Princeton Large Torus [19] and one at Versator at MIT [20]. Both of these experiments observed reduced voltage around the torus. This indicates that the Ohmic field magnitude was dropping, and since the current was maintained, rf current drive could be deduced. Since these two experiments, LH antennas have appeared on many major tokamaks, including JET [21], ASDEX [22], JT-60U [23], Alcator C-Mod [24], and HT-7 [25]. In addition, two notable fusion devices focusing on long pulse lengths specifically exploit LHCD: TRIAM-1M [26] which has achieved 20 kA plasmas for over 3 hours, and Tore Supra [27] which has created hot, dense plasmas sustained above 0.6 MA for over 2 minutes [28]. LHCD is also being used to access so-called Advanced Tokamak regimes, which exploit novel current profiles for performance improvement. [29] Uchimoto *et al.* proposed [30] a similar application, using a lower hybrid scheme to drive the necessary stabilizing current on the Madison Symmetric Torus (MST) RFP. In the advanced tokamak and stabilized RFP schemes, LH is particularly attractive. The low-collisional LH electrons in these scenarios are further accelerated by the Ohmic field, which is not turned off as it is in plasma sustainment experiments. Therefore there is an amplification effect to any current driven.

1.2.7 Electron Bremsstrahlung

Measuring the impact of lower hybrid experiments is fairly simple in a tokamak as the current driven is essentially completely in the toroidal direction. Thus, by monitoring the toroidal loop voltage (a measure of the Ohmic field) and the toroidal plasma current, one can deduce the current created by rf. If the loop voltage drops and the current remains the same, it is a simple matter to claim that the LH system is supplying the necessary current, and in fact this was observed [19]. Since both of these measurements can be made external to the plasma, small changes can be measured accurately and the effects of small scale current drive can be easily seen. This is not the case in the RFP since the targeted direction of current flow is poloidal. Poloidal current cannot be

measured external to the plasma. Therefore, key diagnostics must focus on direct measurements of the fast electron population through bremsstrahlung, [2] either stimulated by probe insertion or observed directly from electron-ion collisions.

As electrons stream freely in a plasma, they occasionally encounter ions or other electrons. The Coulomb force between the colliding particle and the target causes a deflection of the colliding particle. [31] Electrons emit photons when accelerated due to their charge, and this is called bremsstrahlung ("braking") radiation. Predicting the magnitude of this effect typically requires one of a number of approximations, most of which are well studied. Bremsstrahlung cross sections are derived from quantum mechanics, and various permutations have been derived by many authors, especially Bethe and Sommerfeld. [32] These can be integrated over a distribution function in order to predict the total emission from a plasma, assuming the distribution function is known.

1.3 Theory of Lower Hybrid Current Drive on MST

Roughly speaking, the physics of LHCD can be divided into 4 separate time/length scales: the wave phase scale, the wave propagation scale, the wave absorption scale, and the net current drive scale. Mercifully, these different levels do not coincide with the notable exception of the absorption scale and the net current drive scale, which are intrinsically related through induction (discussed in §1.3.4). The wave phase scale refers to what occurs on the order of a single wave oscillation, and so this scale is defined by the wavelength ($\sim 1 \times 10^{-4}$ m) and the period of the wave, $1/(800 \text{ MHz}) = 1.25 \text{ ns}$. Landau damping occurs on this length scale, though it doesn't have a macroscopic effect until kinetic time scales. Arguably there could be turbulence induced stochasticity on this spatial scale, but it is unlikely that the time scales are fast enough to interact with the wave at this level. The next largest scale is the wave propagation scale. This represents the amount of time the wave takes to get to its destination. This distance is roughly 1/3 of the way around the machine poloidally (~ 1 m), and is used to estimate the time scale by dividing by the group velocity. The final two time scales are the absorption time scale, representing the amount of time it takes deposit significant energy from the wave into electrons, and the net current drive time scale. The net current drive time scale is the amount of time between energy being deposited,

and relaxation of the field which arises due to Faraday's law and Ampere's law in response to an injection of current. This ordering is shown in Figure 1.5.

1.3.1 Landau Damping

LHCD experiments exploit Landau damping to drive current. Landau damping is a means of transferring energy from an electrostatic wave to electrons in plasmas. The basic formalism was outlined by Landau in 1946 [33]². The equations governing the evolution of this process can be derived by studying the behavior of a distribution of charged particles in the presence of a sinusoidal electrostatic perturbation. The dynamics of the distribution are constrained by the Boltzmann equation. For simplicity, the effects of collisions are ignored. This equation can be solved analytically by first separating the distribution function into a constant background Maxwellian distribution, and a small perturbed general distribution which is initialized to be zero. After this, a Laplace transform in time and a Fourier transform in space are applied, and the equation is solved. One must be careful to maintain the orientation of the integral curve around the poles when inverting the Laplace transform. Thorough discussion of this solution is given in books by Swanson [35] and Stix [36]. The damping rate predicted by this prescription is

$$\gamma_L = \sqrt{\pi} \frac{\omega_{pe}^2 \omega^2}{k^3 v_{the}^3} e^{-(\omega/kv_{the})^2} \quad (1.7)$$

where γ_L is the damping rate of the wave, ω_{pe} is the electron plasma frequency (discussed in Section 1.3.2.1), v_{the} is the most probable velocity of the Maxwellian distribution, ω is the frequency of the electrostatic wave (in radians/s), and k is the wavenumber of the electrostatic wave. This damping rate has the scaling shown in Figure 1.6. From this we say the condition for strong damping is $v_\phi = (0.2 - 3)v_{the}$.

In the traditional interpretation of Landau damping one considers a single electron encountering an electrostatic wave and being accelerated by that wave's electric field. An electron moving

²This reference can be difficult to locate. It is reproduced in a slightly less difficult to locate collection of Landau's published works. [34]

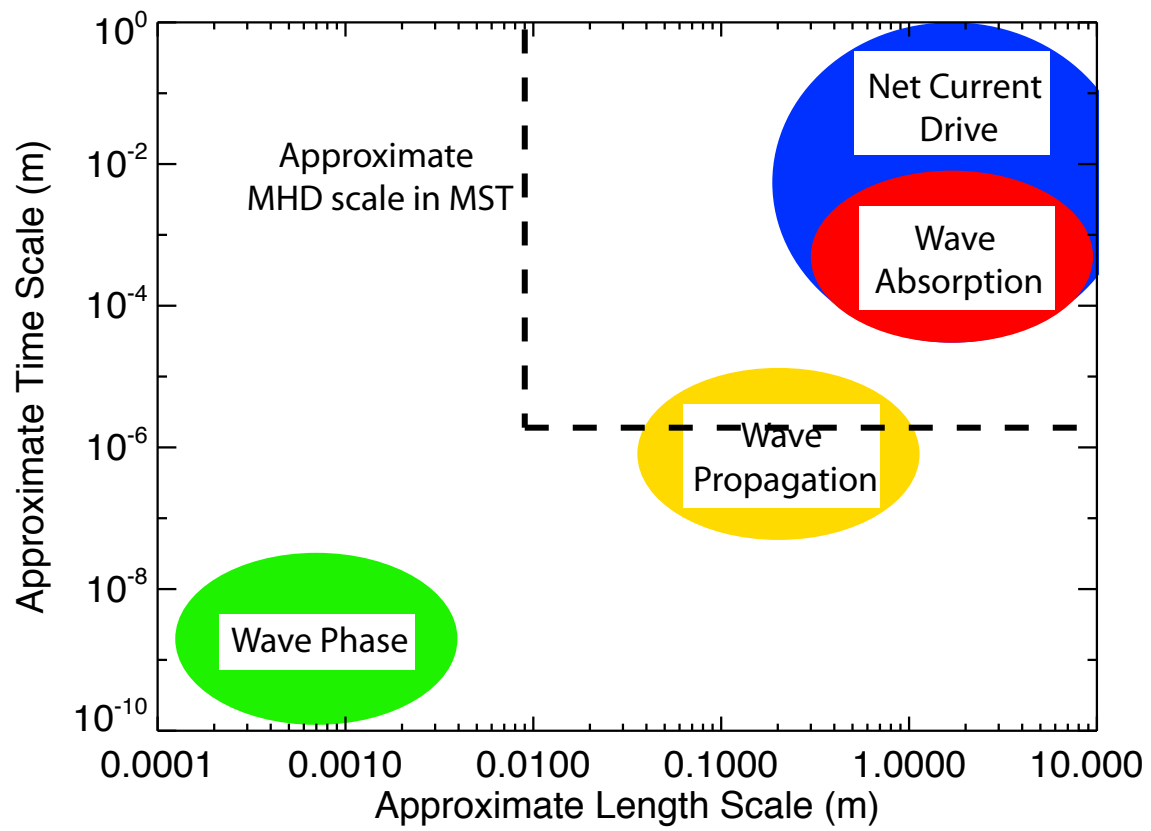


Figure 1.5: Approximate length and time scales for lower hybrid experiments on MST.

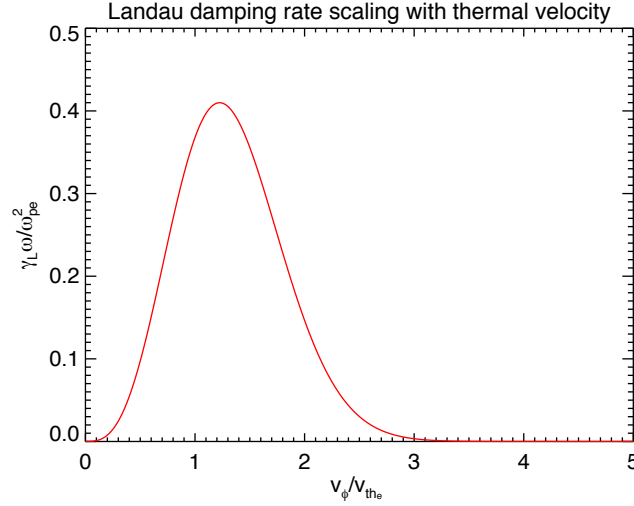


Figure 1.6: Normalized Landau damping rate given by equation 1.7 for constant density and wave frequency. Based on this plot, “strong” Landau damping is interpreted to mean $v_\phi \approx (0.2 - 3)v_{the}$.

slowly relative to the wave’s phase velocity will not be accelerated much before the quickly oscillating wave undergoes a full period, and thus will see virtually no net force. Conversely, a very fast moving electron will see the wave as essentially static, and will have enough momentum to sample an entire period before there is much of an effect on it, thus experiencing a zero average force. Only electrons moving near the speed of the wave will experience acceleration, as they are pushed continually in the same fashion as a surfer on a water wave. Depending on what the phase of the wave is when the electron enters the resonance region it will either be accelerated or decelerated. Thus, our expectation of Landau damping is that it will have a region of effectiveness around $v = v_\phi$, not effecting slow or fast electrons, and it will have the effect of driving resonant electrons both faster and slower, roughly in equal measure. Therefore we expect a flattening of the electron distribution function near $v = v_\phi$. A more thorough investigation of the physics of Landau damping was undertaken by Stubbe and Sukhorukov [37] who enjoin us to recall that the Landau solution requires us to include the electrons which constitute the wave, while the above interpretation ignores those electrons. The distinction is immaterial for the purposes of this work.

1.3.2 Wave Propagation

In order to drive current at a precise radial location, a comprehensive study of wave propagation must be undertaken to determine the correct positioning and launch parameters of the experimental antenna. For lower hybrid experiments in tokamaks, antenna position is typically assumed to be at the outboard midplane due to the spatial constraints on large wave guide antennas. Relaxing this restriction allows for more control over the launched spectrum, and so allows a wider range of frequencies to be considered. Precise predictions about wave spectra and position can be determined by the use of ray tracing codes, but a general description of some of the major phenomena can be derived analytically by employing some fairly simple models. A good derivation of some elements of lower hybrid wave propagation is given by Forest *et al.* , [38], and the elements most relevant to this work are reproduced in Appendices C and D.

1.3.2.1 Cold Plasma waves

Although thermal effects are necessary to understanding wave absorption and are predicted to impact wave accessibility to some degree, [39] a great deal of insight can be derived by the simpler Cold Plasma model, originally developed by Altar and Appleton. [40] This model ignores thermal motion, but has the advantage of tractability, and still captures many of the most important features. The cold plasma model has been observed to predict the propagation of LH waves well. [15]

Before proceeding, I must make a quick digression. In much of the analysis presented in this work, the index of refraction ($n = ck/\omega$) is used. In traditional optics the index of refraction is thought of as a material property. Because plasmas are inhomogeneous and dispersive media, with a dispersion relation related not only to frequency but also polarization, this obfuscates the underlying physics. The index of refraction is favored in many contexts however because it is the figure of merit for many aspects of wave propagation. For instance, the exponential term for Landau damping rate is dependent on k_{\parallel}/ω , not just k_{\parallel} , and so absorption is in some way very dependent on n . Therefore usage of the index of refraction will not be suppressed, but it must be stressed that it is more appropriate to think of it as a wave property, a normalized wave number.

Since we are interested in high frequency, short wavelength (compared to machine size) phenomenon, it seems logical to concentrate on small scale, localized behavior. Bearing this in mind, we consider the forces on individual particles arising from the electric and magnetic fields within the plasma. The primary impacts of the electric field on small scale perturbations are the so-




Figure 1.7: Plasma oscillations are shown on the left. First, the idealized unperturbed electrons.

When bunched together as in the second picture they repel one another, causing bunching follow by depletion. As these newly bunched electrons are forced into one another, they again repel, causing the perturbation to move (in both directions in this case). On the right is an image of cyclotron motion, with ions traveling around the field in one orientation and electrons traveling in the opposite orientation with a smaller orbit size.

called plasma oscillations. Particles of the same charge that move close to one another are repelled away by an ever increasing electric field. Thus, forcing a group of similarly charged particles together gives rise to a repulsive wave which propagates through the plasma, with a frequency of $\omega_{p\alpha} = \sqrt{n_{\alpha}e^2/\epsilon_0 m_{\alpha}}$, where α is the species index (electron, deuteron, carbon, etc.).

The most dramatic effect of the large background magnetic field on plasma particles is cyclotron motion. The collective magnetic field deflects all motion perpendicular to it, redirecting particle trajectories into a helix around field lines. This type of motion is typically very fast, making a full orbit around a field line with a frequency $\omega_{c\alpha} = q_{\alpha}B/m_{\alpha}$, often several GHz for electrons in a fusion reactor. These two types of motion are illustrated in Figure 1.7.

It is the coupling of these two types of motion for all species in a plasma (when combined with drift motion individual particles) that gives rise to the entire host of what are known as cold plasma waves. Cold plasma wave propagation is constrained by a two solution dispersion relation covering a wide range of frequencies. A reasonable derivation of this model is given by Stix [36], and followed by Swanson. [35] The procedure is outlined in appendix C.1. Essentially it arises from solving the single particle equations of motion, in conjunction with Faraday's law and Ampere's law, in the presence of a magnetic field and an electric field perturbation. This analysis leads us to the dispersion relation

$$\mathfrak{D}(\omega, n_{\parallel}, n_{\perp}) \equiv An_{\perp}^4 - Bn_{\perp}^2 + C = 0 \quad (1.8)$$

where

$$A = S$$

$$B = PS + S^2 - D^2 - n_{\parallel}^2(P + S)$$

$$C = Pn_{\parallel}^4 - 2PSn_{\parallel}^2 - D^2P + PS^2$$

with the cold plasma dispersion elements (S , D and P) being

$$S = 1 - \sum_{\alpha} \frac{\omega_{p\alpha}^2}{\omega^2 - \omega_{c\alpha}^2} \quad (1.9)$$

$$D = \sum_{\alpha} \frac{\omega_{c\alpha}\omega_{p\alpha}^2}{\omega(\omega^2 - \omega_{c\alpha}^2)} \quad (1.10)$$

$$P = 1 - \sum_{\alpha} \frac{\omega_{p\alpha}^2}{\omega^2}. \quad (1.11)$$

This equation, quadratic in n_{\perp}^2 , has two solutions, given by

$$n_{\perp}^2 = \frac{B \pm F}{2A} \quad (1.12)$$

where $F = \sqrt{B^2 - 4AC}$. These two solution branches are given many names by many different authors, depending on factors like polarization. For the purposes of this work, these branches will be referred to as the fast wave (the negative branch) and the slow wave (the positive branch). This choice of naming will be justified later (Figure 1.9).

For certain situations we can use a simplified dispersion relation by imposing the electrostatic condition. This approximation is reasonable when the wave polarization is such that the wave is primarily longitudinal. This is derived in Appendix D, and is valid when the wave being studied has a mostly longitudinal polarization. As discussed in Appendix D, this is only really acceptable for the slow branch, far from cutoffs. In our parameter range of interest, the general electrostatic dispersion relation

$$\mathbf{k} \cdot \mathbf{K} \cdot \mathbf{k} = 0, \quad (1.13)$$

simplifies to

$$\frac{k_{\perp}^2}{k_{\parallel}^2} \approx \frac{\omega_{ce}^2}{\omega^2} \approx 50 - 225. \quad (1.14)$$

From this form of the dispersion relation, we can derive an approximate expression for the group and phase velocity of the wave, namely

$$\mathbf{V}_{\phi} = \frac{c}{n_{\parallel}} \frac{\omega}{\omega_{ce}} \left(\frac{\omega}{\omega_{ce}} \hat{\mathbf{b}} + \hat{\mathbf{r}} \right) \quad (1.15)$$

$$\mathbf{V}_{\mathbf{g}} = \frac{c}{n_{\parallel}} \left(\hat{\mathbf{b}} - \frac{\omega}{\omega_{ce}} \hat{\mathbf{r}} \right). \quad (1.16)$$

This derivation is carried out in detail in Appendix D.

1.3.2.2 Cutoffs, Mode Conversion and Evanescence

Since wave analysis assumes a wave of the form $e^{i\mathbf{k} \cdot \mathbf{r} - i\omega t}$, an imaginary value of \mathbf{k} (and therefore n) will lead to a real exponent for our field, i.e. an exponentially growing or decaying field. The exponentially amplifying wave is uninteresting to us, because such a mode does not conserve energy (because it does not conserve power in steady state) without a source in the plasma. Therefore when some component of n becomes imaginary, we say the wave is decaying. This occurs when the right hand side of Equation 1.12 becomes negative, or when the term F becomes imaginary. The first case leads to a purely decaying field and is therefore called a cutoff, whereas the second case leads to a decaying wave, termed an evanescent wave. Moreover, since the fast wave and slow wave solutions converge for $F = 0$, there is a chance for mode conversion. It is important

to bear in mind that when we refer to cutoffs and evanescence, we are talking about exponential decay in the spatial structure of the wave. This means that the wave simply does not penetrate into the region, not that power is being deposited.

The critical point for cutoffs occurs where n_{\perp}^2 crosses zero on its way to becoming negative. In this case, Equation 1.8 must be of the form $C_1 n_{\perp}^2 (n_{\perp}^2 + C_2) = A n_{\perp}^4 + B n_{\perp}^2$, i.e., $C = 0$. This gives rise to the cutoff points

$$\begin{aligned} C &= P n_{\parallel}^4 - 2 P S n_{\parallel}^2 - D^2 P + P S^2 \\ &= P ((n_{\parallel}^2 - S)^2 - D^2) = 0, \end{aligned} \quad (1.17)$$

i.e. $P = 0$ and $n_{\parallel}^2 = S \pm D$. The $P = 0$ condition creates a cutoff barrier at $n = \omega m_e \epsilon_0 / e^2$. The second pair of cutoffs ($n_{\parallel}^2 = S \pm D$) are called the right and left mode cutoffs, because of the impact on certain wave polarizations. The critical point for fast-slow mode conversion and evanescence is the point where $F^2 = 0$, i.e. where $B^2 = 4AC$. While this term is fairly complex, it is perfectly well behaved and can be easily represented on a plot. These results can be combined to create a plot of the propagation region for both the fast and the slow modes. Using approximate MST parameters ($n_e \leq 1.0 \times 10^{19} m^{-3}$, $|B| \approx 0.25$), we can produce Figure 1.8. Qualitatively speaking, increasing the magnetic field causes the cutoff region to become narrower, i.e. the cutoff n_{\parallel} is lower for a given density.

This plot demonstrates several important features. First, it can be seen that while the slow wave propagates in most of the plasma volume (excluding the very edge), the fast wave has a cutoff (the R-mode cutoff, $n_{\parallel}^2 = S + D$) which prevents it from propagating throughout most of the plasma. Thus we need not concern ourselves with unintentionally coupling to the fast wave, as the propagation region is too distant from the antenna to allow a significant amount of power to tunnel through. Conversely, for the slow wave, the non-propagation region is likely on the order of a few millimeters, and easily coupled to by an evanescent wave coming from a very low density or vacuum region.




Figure 1.8: Plotted here are the regions of solution space for n_{\perp} . Magnetic field magnitude is fixed to 0.25 T, typical for MST discharges in the region of interest. Regions I and IV represent purely imaginary solutions (no propagation), region II supports the fast wave and slow wave solution, and region III allows the slow wave only. Detail of the low density section of the slow wave plot is shown.

If the full cold plasma dispersion relation is plotted in the more traditional way (on an ω versus k plot), the relative group velocities can be estimated. Because our profiles increase most dramatically in the radial (perpendicular) direction, we will restrict this discussion to k_{\perp} (though we do not preclude interesting physics in the parallel direction). Plotted in Figure 1.9 are the solutions of the dispersion equation at three different densities, with the experimental frequency marked. Numbers have been intentionally suppressed in this plot because the exact size varies greatly with plasma conditions, and this is only meant to provide a qualitative impression. The two modes correspond to the two places where the marked frequency intersects the dispersion curve (notice sometimes this only corresponds to one mode). Recall that the larger solution branch from Equation 1.12 [$n_{\perp} = (B + F)/2A$] was referred to as the slow wave. We see that the group velocity (the slope of this plot) is smaller for this mode than it is for the fast mode, hence the names. We can also see by examining the slope of the plot that the slow wave and the fast wave have oppositely directed group velocities, and that the group velocity is smoothly varying over the mode conversion region. Thus, one wave is approaching the cutoff point and the other leaving. We can conclude from this that all the power must be transferred from one wave to the other, otherwise there would be nowhere for the power to go. It should be noted that another possibility exists, that of coupling to the opposite frequency branch of the slow wave. However the ingoing slow wave phase velocity matches the outgoing fast wave phase velocity so it is more fitting to think of the slow wave propagating through this region and finding that it has become a fast wave, as opposed to jumping (possibly inefficiently) from one branch to another. This phenomenon has been treated more rigorously elsewhere. [41]

1.3.2.3 Raytracing

In order to determine the motion of a launched wave in the plasma it is necessary to employ the WKB method, also called ray-tracing. It can be shown that for wavelengths and wave periods sufficiently smaller than the length and time scale of changes in the dielectric tensor \mathfrak{D} , the group




Figure 1.9: Here we see the dispersion relation for Cold Plasma waves plotted at three densities. The blue curve only supports a slow wave (gently sloped) at the experimental frequency. The green curve represents the region that supports both fast and slow waves, and the red curve represents the density at which the two modes converge. The exact magnitudes involved depend highly on the plasma conditions.

velocities and wave number rate of change can be determined by the equations

$$\begin{aligned}\frac{dx}{dt} &= -\frac{\partial \mathfrak{D}/\partial k_x}{\partial \mathfrak{D}/\partial \omega} & \frac{dk_x}{dt} &= \frac{\partial \mathfrak{D}/\partial x}{\partial \mathfrak{D}/\partial \omega} \\ \frac{dy}{dt} &= -\frac{\partial \mathfrak{D}/\partial k_y}{\partial \mathfrak{D}/\partial \omega} & \frac{dk_y}{dt} &= \frac{\partial \mathfrak{D}/\partial y}{\partial \mathfrak{D}/\partial \omega} \\ \frac{dz}{dt} &= -\frac{\partial \mathfrak{D}/\partial k_z}{\partial \mathfrak{D}/\partial \omega} & \frac{dk_z}{dt} &= \frac{\partial \mathfrak{D}/\partial z}{\partial \mathfrak{D}/\partial \omega}\end{aligned}\tag{1.18}$$

The requirements for use of the raytracing equations are easily satisfied for our wave. During propagation, the wavelength is a fraction of a millimeter and the period is near one nanosecond, whereas MHD activity takes place at length scales near a centimeter and time scales on the order of hundreds of microseconds.

1.3.2.4 Resounding Chambers

The first model to apply the raytracing equations to is the stratified slab. The stratified slab contains a density profile in the x direction and a uniform magnetic field that points in the z -direction. The x direction is analogous to the radial dimension in a toroidal plasma. Using this model, an expression for the trajectory of a ray can be derived. The chain rule allows us to say

$$\frac{dx}{dz} = \frac{\partial \mathfrak{D}/\partial k_{\perp}}{\partial \mathfrak{D}/\partial k_{\parallel}} = \frac{\partial \mathfrak{D}/\partial n_{\perp}}{\partial \mathfrak{D}/\partial n_{\parallel}},\tag{1.19}$$

which, using 1.8, reduces to

$$\frac{4An_{\perp}^3 - 2Bn_{\perp}}{4Pn_{\parallel}^3 + 2[(P+S)n_{\perp}^2 - PS]n_{\parallel}}.\tag{1.20}$$

We also note that since no profiles are dependent on z , $\partial \mathfrak{D}/\partial z = 0$, so (by our raytracing equations) n_{\parallel} is constant. Because n_{\perp} can be determined at all locations in space by using 1.12, this equation, while analytically insuperable, is quite readily solved by simple numerical methods. Taking a profile from a reconstructed MST discharge (which will be explored in detail later) and using $n_{\parallel} = 5.0$ we get the plot seen in Figure 1.10. The red launched slow wave travels from the $P = 0$ surface, inward to the $F = 0$ mode conversion point. It is then converted into a fast wave. The fast wave propagates outward until it hits the $n_{\parallel}^2 = S + D$ surface, where it is reflected as an inward




Figure 1.10: Numerical solution of Equation 1.20 for $B_0 = 0.25T$ and the given density profile.

The $F = 0$ surface and the R-mode cutoff are overplotted. The red and yellow trajectories are slow waves and the two blue curves are fast waves. The R mode cutoff and $F = 0$ surface were calculated independently of the ray calculations

going fast wave until it returns to the $F = 0$ surface to mode convert back into a slow wave. The slow wave will either reflect off the $P = 0$ boundary, or couple to a surface mode. Thus this wave is incapable of leaving this region of the plasma between the $P = 0$ and $F = 0$ surfaces, trapping the wave in a resounding chamber.

1.3.2.5 Poloidal Upshift

Exploiting the simplified dispersion relation given by Equation 1.14, more complicated geometry can be studied, and more features of LH wave propagation in MST can be explained. Modifying the raytracing equations for toroidal geometry we get

$$\begin{aligned} \frac{dr}{dt} &= -\frac{\partial \mathfrak{D}/\partial k_r}{\partial \mathfrak{D}/\partial \omega} & \frac{dk_r}{dt} &= \frac{\partial \mathfrak{D}/\partial r}{\partial \mathfrak{D}/\partial \omega} \\ \frac{d\theta}{dt} &= -\frac{\partial \mathfrak{D}/\partial \kappa}{\partial \mathfrak{D}/\partial \omega} & \frac{d\kappa}{dt} &= \frac{\partial \mathfrak{D}/\partial \theta}{\partial \mathfrak{D}/\partial \omega} \\ \frac{d\phi}{dt} &= -\frac{\partial \mathfrak{D}/\partial \eta}{\partial \mathfrak{D}/\partial \omega} & \frac{d\eta}{dt} &= \frac{\partial \mathfrak{D}/\partial \phi}{\partial \mathfrak{D}/\partial \omega} \end{aligned}$$

where $\kappa = rk_\theta$ and $\eta = Rk_\phi$. The effect typically considered is that of poloidal lensing, specifically, in the cylindrical approximation, no profiles are dependent on poloidal angle, so it can be said that

$$\frac{d\kappa}{dt} = \frac{\partial \mathfrak{D}/\partial \theta}{\partial \mathfrak{D}/\partial \omega} = 0, \quad (1.21)$$

i.e. κ is a constant. If this is true then $k_\theta = \kappa_0/r$. This says that as radius decreases, k will increase. The physical reason for this can be seen in Figure 1.11. This effect is indeed seen in wave propagation, but this simple analysis obscures another important effect, namely the upshift in the wave number due simply to the wave propagating around poloidally. Naively, we might assume that there is a similar toroidal lensing effect. This is in fact true, but it affects k_ϕ which for our experiment is a small number, and a component of k_\perp in any case. Reexamining the raytracing equations and employing the chain rule we get the equation

$$\frac{d\kappa}{d\theta} = -\frac{\partial \mathfrak{D}/\partial \theta}{\partial \mathfrak{D}/\partial \kappa} \quad (1.22)$$




Figure 1.11: The Poloidal lensing effect. A perturbation with a constant poloidal wave number at two different radii will have two different wave lengths simply due to the reduced arc length.

Further assuming that $k_{\parallel} \approx k_{\theta}$ and $\frac{dr}{d\theta} \approx 0$, this simplifies to

$$\frac{d\kappa}{d\theta} = -\frac{-\frac{\kappa}{r} \frac{\partial}{\partial \theta} \left(\frac{\omega_{ce}}{\omega} \right)}{-\frac{1}{r} \frac{\omega_{ce}}{\omega}} = -\kappa \frac{d}{d\theta} \ln \left(\frac{\omega_{ce}}{\omega} \right) \quad (1.23)$$

i.e.

$$k_{\theta} = \frac{r_0}{r} \frac{\omega_{ce0}}{\omega_{ce}} k_{\theta 0} = \frac{\rho_0}{\rho} \frac{\omega_{ce0}}{\omega_{ce}} k_{\theta 0} \quad (1.24)$$

Notice that this retains the poloidal lensing effect, however it further states that k_{θ} is inversely proportional to B - higher on the outboard side and lower on the inboard side. This still ignores the significant effect of the radial profiles. More complete analysis of wave propagation requires the use of a raytracing code. This will be discussed in more detail in Chapter 4.

1.3.3 Quasilinear absorption

Recall in the earlier discussion of Landau damping that it was predicted that depending on the phase of the electrostatic wave as the electron entered the Landau damping region, it could either be sped up or slowed, with roughly equal probability. When examining a distribution function, this means particles will be driven away from a point in velocity space, either to higher or lower velocities. Collisionality will prevent a divot from forming in the distribution function, so the most heavily distorted distribution function that can be expected is a distribution with a “flat spot”. Furthermore, once the distribution function has been completely flattened in a given region of velocity space, no more flattening can occur, and therefore no additional Landau damping can be induced. It seems reasonable to conclude that when the distribution is slightly flattened, the effects of Landau damping will be somewhat reduced. Therefore, it is critical to consider the absorption process in terms of quasilinear behavior. That is to say, although we still evolve the linearized Boltzmann equation, we allow whatever operator represents Landau damping to be based on the perturbed distribution function as well as the unperturbed distribution function. This seems quite complex for analytic calculations, and indeed it is, but a sufficiently abstract model can still be employed to provide some insight into the effects of Landau damping. A more complicated model is used as the backbone of CQL3D, a quasilinear Fokker-Planck code discussed and used in Chapters 3 and 4. The abstracted model of choice is the one proposed by Fisch [16], using a simplification of the

Lenard-Bernstein Fokker-Planck equation given by (using the notation followed by [36] in Section 16-7)

$$\frac{\partial f}{\partial t} = \frac{\partial}{\partial v} \left[\nu \left(v f + \frac{kT}{m} \frac{\partial f}{\partial v} \right) \right] + \frac{\partial}{\partial v} D_{QL} \frac{\partial f}{\partial v}. \quad (1.25)$$

where ν is the collision frequency (a simplified representation of the Coulomb collision operator) and D_{QL} is the velocity space diffusion coefficient for the damping mechanism, in this case the Landau resonance. A velocity space diffusion coefficient can be interpreted as a force, as it represents the statistical net effect of either making electrons faster or slower. Restricting our discussion to only the equilibrium behavior we can say

$$\frac{\partial f}{\partial v} = -v \left(\frac{kT}{m} + \frac{D_{QL}}{\nu} \right)^{-1} f \quad (1.26)$$

which has the concise solution

$$f = \text{Exp} \left[\int -v \left(\frac{kT}{m} + \frac{D_{QL}}{\nu} \right)^{-1} dv \right] \quad (1.27)$$

Notice that for $D_{QL} = 0$, i.e. no acceleration, the solution is simply

$$f = C \text{Exp} \left[-\frac{mv^2}{2kT} \right], \quad (1.28)$$

which is as we would hope, a Maxwellian (note the C comes from the constant of the integration).

Returning our attention to 1.27, we first substitute $v_n = \sqrt{kT/m}$ (the thermal velocity in one dimension), and then assume that $\nu(v) = \nu_0 v_n^3 / v^3$, which is consistent with velocity dependent collisionality reduction observed in plasmas.

$$f = \text{Exp} \left[\int -\frac{v}{v_n^2} \left(1 + \frac{D_{QL} v^3}{v_n^5 \nu_0} \right)^{-1} dv \right] \quad (1.29)$$

Defining $D_{red} = D_{QL} / v_n^2 \nu_0$ and $w = v / v_n$, this becomes

$$f(w) = C \text{Exp} \left[-\int_0^w \frac{w}{1 + w^3 D_{red}} dw \right] \quad (1.30)$$

For the final step in the analysis, we examine the case

$$D_{red} = \begin{cases} \delta & v_0 \leq v \leq v_1 \\ 0 & \text{otherwise} \end{cases} \quad (1.31)$$

There is a temptation to use a more complicated model, for instance, using the Landau resonance velocity dependence from Figure 1.6, especially when numerical methods are used. It is important to remember however, that the simple Coulomb collision model used above ($\nu \sim v^{-3}$) breaks down as the distribution function flattens. Therefore, in some ways the top hat model employed here is in fact more realistic (albeit more abstract), as the collisionality will eventually come into equilibrium with wave damping. In this way we are presented with the fundamental conflict in analytic lower hybrid theory. A simple model provides some insight but more complicated models require careful accounting of more complicated effects. This often means a model which adds only a slight complication is in practice less accurate than a simple model.

For $\delta \rightarrow \infty$ the distribution is flat in the resonance region. This flat distribution function gives rise to a current density of [16]

$$J \approx 6.5 \times 10^8 n_{20} T_{10}^{1/2} f(w_0) w_0 (w_1 - w_0) \text{ A/m}^2. \quad (1.32)$$

where n_{20} is the density in units of $1 \times 10^{20} \text{ m}^{-3}$ and T_{10} is temperature normalized to 10 keV. Assuming that collisions with the bulk are balancing power deposition, we can calculate the deposition as

$$P_D \approx (5 \times 10^9) n_{20}^2 T_{10}^{-1/2} (2 + Z) f(w_0) w_0 (w_1 - w_0) \text{ W/m}^3 \quad (1.33)$$

Taking the ratio of these two values we derive an efficiency parameter

$$\frac{J}{P_D} = \frac{0.13}{(2 + Z) n_{20}} T_{10} (\text{A/m}^2) / (\text{W/m}^3) \quad (1.34)$$

Interpreting the meaning of this parameter is somewhat complicated. As we are talking about a power density being converted into a current density, the interpretation is largely dependent on the propagation geometry. Moreover, the amount of power absorbed is entirely dependent on the amount of power which has successfully propagated inward (i.e. not already been absorbed). Nevertheless, this is a fairly concise expression which provides some insight into scaling. In particular it shows that there is a linear relationship between the temperature of a plasma and the efficiency of current drive.

Something can be learned by plotting these solutions (this is done in Figure 1.12). Before proceeding to read too much into this plot, note that I have declared the resonance region by




Figure 1.12: A plot of the solutions of Equation 1.30, for varying values of δ . Note that there is nothing particularly special about the length and location of the resonance region, as that is input into the model (I have based it on the Landau resonance). Also note that the normalization is not perfect, and in fact we expect that the extra electrons on the right hand side of the plot will be sourced from the left hand side of the plot.

flat which is why it is in the region we expect it to be. Moreover, since a numerical integral was used to normalize the plot, the normalization is not perfect - we expect the high δ curves to exhibit some depletion left of the resonance region which will act as a source for the additional particles seen right of the resonance region. It can be seen that there is not only a flattening in the resonance region, but also an additional tail which develops on the high speed side. One is entitled to wonder why that might be. The intuitive explanation is that while quasilinear diffusion enforces flatness in the region of resonance, some particles will be forced out by acceleration, or by collisions (remember our model does include collisions, albeit in a simplified form). These will be continuously sourced by electrons coming from the resonance region until the slowing-down of very fast electrons balances out ejection. Thus, LH acceleration does not only directly accelerate a small population of electrons in the region of resonance, it also acts as a sort of electron conveyor, pumping electrons from the low speed end of the distribution function to the high speed end. This

process becomes more important when an electric field is applied, and these electrons are allowed to run away.

A somewhat more complicated version of this type of analysis was undertaken by Uchimoto for MST like parameters. [30] He found that 1 MW of LH power was sufficient to drive a peaked current with a maximum of 100 A/cm^2 . This study was based on the Brambilla ray tracing code, [42] following the trajectory of a single ray until 90% of the power is deposited and approximating the current drive based on an analytic approximation of current drive efficiency. However, in his authoritative work, [43] Fisch asserts that a complete understanding of LH current drive requires a consistent solution of the (quasilinear) Fokker-Planck equation in at least 2 dimensions. Forest *et al.* extended Uchimoto's work to include quasilinear effects, employing the raytracing code GENRAY and the Fokker-Planck code CQL3D. [44] This led to a prediction of 2 MW required to drive stabilizing current. Although this did not include diffusion or electric field effects, the results provided a sound foundation for further experimentation. This last work also separately examined the effects of radial diffusion on driven current. Uchimoto found that the LH driven current had a very narrow radial extent. Forest however showed that while the LH wave can drive current with a great degree of radial localization, radial diffusion will smear this current out, and this effect can be devastating to the peak driven current. Wedding the effects of radial diffusion with quasilinear absorption was left as a tantalizing direction of future research. CQL3D now has a self-consistent radial diffusion operator, and so the approximation of Forest can be improved upon. Moreover, the diffusion coefficient and plasma equilibria in MST are more exactly known. Finally, Forest did not account for the amplifying effects of an electric field working in concert with the electron conveyor effect described above, which is expected to mitigate the effects of radial diffusion somewhat. These complications underscore the need to revisit these studies. [45] It should also be noted that these earlier studies focused on the launch of a single ray at the expected centroid of the launch spectrum, rather than the broader launch spectrum of a real antenna. This affects the efficiency of absorption as the width of the resonance region in velocity space is directly related to this spectral broadening. This is discussed further in §4.2.3.

1.3.4 Inductive effects

In any study of current drive it is important to consider the plasma response. An ideal (perfectly conducting) plasma would not allow current drive, because it would be induced to create an electromotive force exactly canceling any current from fast electrons. In a real plasma, this effect is typically approximated by a simple circuit model. This leads to the simple formula [43]

$$\frac{dJ}{dt} + \frac{J}{\tau_L} = \frac{J_{rf}}{\tau_L}, \quad (1.35)$$

where $\tau_L = L/R$ is the so-called “inductive time”. Before continuing, it is important to note that this model assumes that current is being driven in a filament in free space (i.e. plasma response is not modeled). This is not an accurate representation of the physics, but is used for simplicity. Placing this filament in a conductive medium will increase the inductance, most likely by a fairly large amount. Honestly account for this is best done with a full 3-D MHD model, in order to measure the current associated with the injected current, which is beyond the scope of this work.

For constant J_{rf} , this is solved by $J(t) = J_{rf} (1 - e^{-t/\tau_L}) + J_0$. Notice that as $R \rightarrow 0$, $\tau_L \rightarrow \infty$ and $J(t) \rightarrow J_0$. Now, current filaments created by rf deposition can be considered. These can be divided into two categories. At or very near the reversal surface, the toroidal pitch of the field lines is near zero. As long as these filaments feed back on themselves in one poloidal transit, they can be considered to be isolated loops, dubbed “poloidal filaments”. A poloidal filament wraps once around the magnetic axis and reenters the deposition region immediately. In this case, the inductance of the poloidal filament can be calculated by approximating it as a simple circular current loop. The inductance of a loop of wire is given by Jackson, [46] (page 234)

$$L_p = \mu_0 r \left[\ln\left(\frac{8r}{\delta}\right) - 7/4 \right] \quad (1.36)$$

where r is the minor radius at which the filament sits, and δ is the radius of the filament cross section.

For filaments which *nearly* reconnect with themselves (henceforth “helical filaments”), the situation is more complicated, and seemingly less favorable. Consider a pie-slice shaped Amperian loop (Figure 1.13). The inherent toroidicity of these filaments, while small, is vitally important,




Figure 1.13: Amperian loop used in inductance calculation

as we are now dealing with an induced field that increases *every time the filament crosses the midplane*, which will be quite a lot for a filament that only barely misses itself.

On the first toroidal transit this is somewhat like a toroidal solenoid. Ampere's law tells us

$$\int \mathbf{B} \cdot d\mathbf{l} = \mu_0 I_{enc} \quad (1.37)$$

using B.4, the toroidal angle covered by a field line in a single poloidal transit is approximately

$$\Delta\phi = 2\pi q \quad (1.38)$$

Therefore, for an Amperian loop that subtends a toroidal angle of $\Delta\phi$,

$$B_\phi R \Delta\phi = \mu_0 N I_{rf} = \mu_0 \frac{\Delta\phi}{2\pi q} I_{rf}. \quad (1.39)$$

Note this calculation is isolated to only the induced magnetic field, not the background plasma field. Also notice that we are implicitly using a symmetry argument to cancel out the field in the major radius direction along the two straight legs of the loop.

Therefore

$$B_\phi(R) = \frac{\mu_0 I_{rf}}{2\pi q R} \quad (1.40)$$

Notice, it is not dependent on $\Delta\phi$. The total poloidal flux from this field is

$$\psi_p = \int \mathbf{B} \cdot d\mathbf{a} = 2 \int_0^{r_{rf}} \int_{R_{left}}^{R_{right}} B_\phi dR dz \quad (1.41)$$

$$= \frac{\mu_0 I_{rf}}{\pi q} \int_0^{r_{rf}} \log \frac{R_0 + \sqrt{r_{rf}^2 - z^2}}{R_0 - \sqrt{r_{rf}^2 - z^2}} dz \quad (1.42)$$

$$= \frac{\mu_0 I_{rf}}{\pi q} G \quad (1.43)$$

where G is

$$\begin{aligned} G = & \left[-2\sqrt{R_0^2 - r_{rf}^2} \tan^{-1} \left(\frac{R_0 z}{\sqrt{R_0^2 - r_{rf}^2} \sqrt{r_{rf}^2 - z^2}} \right) \right. \\ & + z \log \left(\frac{\sqrt{r_{rf}^2 - z^2} + R_0}{R_0 - \sqrt{r_{rf}^2 - z^2}} \right) \\ & \left. + 2R_0 \tan^{-1} \left(\frac{z}{\sqrt{r_{rf}^2 - z^2}} \right) \right]_0^{r_{rf}}, \end{aligned} \quad (1.44)$$

an intimidating expression, but there is a great reduction of complexity at the limits of integration, as the arctangent terms reduce to $\pi/2$ at $z = r_{rf}$ and 0 at $z = 0$. Furthermore the log term becomes 0 at $z = r_{rf}$, so G reduces to

$$G = \pi \left(R_0 - \sqrt{R_0^2 - r_{rf}^2} \right) \quad (1.45)$$

This makes the inductive electromotive force

$$\mathcal{E} = -\frac{d\psi}{dt} = -\frac{d}{dt} \left(\frac{\mu_0 G I_{rf}}{\pi q} \right) = -\left(\frac{\mu_0 G}{\pi q} \right) \frac{dI}{dt} \quad (1.46)$$

i. e.

$$L = \frac{\mu_0 G}{\pi q} \quad (1.47)$$

Notice this indicates that the inductance is highest just off of the reversal surface, whereas (as shown earlier) at the reversal surface, the inductance is quite low. This is consistent with intuition, since a nearly closed loop will have many more midplane crossings in a single toroidal transit than a

more toroidally oriented loop. This however is not the end of the story. On any surface which does not satisfy single rotation symmetry in the toroidal direction (meaning $q = m/n, m = 1, n \in \mathbb{Z}$), the filament continues around the machine, amplifying the inductive field leading to an integer amplification factor m . Indeed, on an ergodic surface (of which there are infinite in any given region), the filament will *never* come back to itself. However, and crucially, when calculating the inductive time, the factor of m which the inductance is multiplied by is cancelled out by the factor of m used for the length of the filament in calculating the total resistance. In any event, at some point the finite width of the current channel comes into play and the process is arrested.

Interestingly, the inductive time is not affected much by the distinction between poloidal and helical filaments. Using the simplified formula for field line length (B.15) the inductive time (for a filament of radius δ) is

$$\tau_{L/R} = \frac{L_m \pi \delta^2}{\eta l_m} = \frac{\mu_0 m G}{\pi q} \frac{\pi \delta^2 q}{\eta 2 \pi m r} = \frac{\mu_0 G \delta^2}{2 \pi \eta r} \quad (1.48)$$

We approximate the resistivity using the well known formula derived by Spitzer and Härm [47–49]

$$\eta_{sp} = \frac{\pi e^2 m_e^{1/2}}{(4 \pi \epsilon_0)^2 T_e^{3/2}} \ln \Lambda \quad (1.49)$$

An estimate for the inductive time can be approximated using MST-like parameters of $T_e = 200 - 300 \text{ eV}$, $\ln \Lambda \approx 12 - 14$, $r \approx 0.4 m$, $\delta \approx 0.01$ ($2.241 \mu\Omega m \leq \eta_{sp} \leq 4.803 \mu\Omega m$). The length of helical field lines is derived in Appendix B. This gives us a poloidal filament inductive time at $52.6 - 113 \mu s$, and an off-poloidal inductive time of $0.565 - 1.21 \mu s$. This is 1-3 orders of magnitude lower than that found in similarly sized tokamaks. [20] The inductive load is similar to that experienced in a tokamak but the length of the filament is so much larger that the resistance more than compensates. It must be stressed however that it is by no means clear that this simple 1-D model can capture the complex interplay of the real fields, particularly since the impact of the plasma response to the injected current is ignored. In order to understand dynamics at a deeper level a more complex 3D model should be used and that is beyond the scope of this work, however, the importance of this issue specifically concentrating on the RFP cannot be overstated, and it is not a problem merely for the lower hybrid current drive scheme. Nevertheless, the 1-D circuit model has been observed to be a reasonable approximation in tokamaks. [20]

1.3.5 System Design

With a complete set of tools at our disposal, we can now move on to design a real-world LH system for an MST-like RFP. The first step is to determine the target n_{\parallel} based on Figure 1.6. Our goal is to drive current just inside of the reversal surface. For the plasmas of interest, this corresponds roughly to 250 eV. Since our wave will be approaching the Landau resonance from a colder area of the plasma, v_{ϕ}/v_{the} will go from a large number to a small number. Therefore we expect the wave to begin damping at approximately $v_{\phi} = 3v_{the}$ (based on Figure 1.6). This leads to

$$v_{\phi} = 3v_{the} \quad (1.50)$$

$$\frac{c}{n_{\parallel}} = 3\sqrt{\frac{2kT}{m_e c^2}} c \quad (1.51)$$


$$n_{\parallel} = \sqrt{\frac{m_e c^2}{18kT}} \approx 10.65 \quad (1.52)$$

The next step is to construct a wave propagation diagram for this plasma, that is a diagram like Figure 1.8 but tailored to our specific equilibrium. The best way to construct a meaningful diagram is to create a plot with the flux surface on the y-axis and n_{\parallel} on the x-axis. The salient features to include are the $F = 0$ cutoffs for the high field and low field side, and the Landau resonance, specifically $n_{\parallel} = \sqrt{m_e c^2 / 18kT}$. In order to do this we must choose a frequency. Due to the availability of sources, 800 MHz was chosen for MST, but the procedure would be identical if a different frequency was chosen. If the design space is completely unrestricted, then the process would be carried out several times at different frequencies. Using these results and an equilibrium generated for a typical MST discharge, Figure 1.14a can be constructed. In this plot, the vertical axis is the (normalized) square root of the area of a given flux surface, which would be equivalent to the radius in the cylindrical limit. This parameter is used frequently and is called ρ . Examining this graph, it can be seen that if the wave veers too closely to the left (low n_{\parallel}) side, it will be mode-converted into a fast wave, and be directed out of the plasma, resounding as described above. If on the other hand the wave veers too closely to the right hand side of the diagram before getting to the target region, power will be absorbed before the desired radius is reached. Thus we are attempting


to “thread a needle” with our wave. The importance of this plot cannot be overstated. It contains almost all of the essential elements for LHCD accessibility. Because of this, it will be revisited in a later section, with more elaborate description.

Alternatively, or rather, concurrently, a simulation can be run to discover reasonable launch parameters. The principle of reciprocity states that emitted waves should travel along the same trajectory as injected waves. This is because the plasma does not differentiate the origin of the two waves, and the wave equation has time reversal symmetry. Figure 1.14b was generated by launching waves at the approximate radius where deposition is desired with the n_{\parallel} required to deposit power. These waves are tracked out of the plasma and the value of n_{\parallel} at the edge is plotted versus the poloidal angle that the wave exits the plasma. Multiple reflection and higher order effects are ignored, so this should be regarded as somewhat approximate. Moreover, n_{\perp} is simply assumed to be directly radial when these waves are initialized, which is not necessarily true at this depth within the plasma. This plot should not be taken to imply non-propagation at any point in the spectrum as waves not moving from the core to a particular angle of the edge does not mean they cannot be launched from the edge and subsequently make into the plasma by some complicated path. Figure 1.14 shows the results from this analysis for 400 kA standard plasmas.

From these two plots, $n_{\parallel} \approx 7$ is chosen for inboard launch so that F-S reflection can be avoided, and $n_{\parallel} \gtrsim 12$ is chosen for outboard launch to be consistent with the anticipated n_{\parallel} shift. It is important to note that we can increase n_{\parallel} slightly beyond the number that Figure 1.14b indicates, as a higher n_{\parallel} corresponds to a lower v_{ϕ} , and therefore more efficient absorption at the target region. The key aspect of off-axis current drive is exploiting the sharp gradient in the temperature profile. This allows us to launch a fairly broad collection of rays that will all damp at roughly the same radius due to the sharp decrease in resonant n_{\parallel} right in the region which we are most interested in. This also allows us to drive current in very different conditions. Because the plasma temperature increases rapidly, so long as we choose a wave not resonant in the very edge of our machine, higher temperature plasmas will have a peak absorption radius only slightly larger than the design specification. Since higher temperature plasmas also tend to exhibit either lower densities or higher magnetic fields, the difference in the cutoffs are not a concern.



(a)



(b)

Figure 1.14: The profiles from an equilibrium reconstruction of a 400 kA standard plasma are used to make diagram (a). The temperature profile is used to draw the resonance line (at $v_\phi = 3v_{the}$ and the magnetic field strength on the high field and low field side are used to mark the F=0 Fast-Slow mode conversion point. Launched waves must stay in the region of $n_{||} \sim 7$. The second plot gives us an idea of the launch $n_{||}$ at a given launch angle which should correspond to $n_{||} \sim 10.65$ near the absorption region, generated by launching waves from the absorption layer of the plasma toward the plasma edge

Having established the launch conditions for our wave, attention now turns to the radiative element (antenna). The essential purpose of a radiative element is to create a 2-D perturbation on the electric field at the plasma edge, at high power, and with known wavelength such that it produces the desired n_{\parallel} spectrum. The *au courant* antenna in fusion devices is the multiseptum waveguide grill, shown schematically in Figure 1.15. [50] Essentially, the strategy employed by this design is to launch $TE_{0,1}$ modes down the waveguide, with an interseptum phase staggering which creates a pseudo-sinusoid at the plasma face.

It is interesting to note that it is physically impossible to construct a vacuum filled waveguide capable of launching a wave like the one we are interested in without exploiting a multiseptum design. This can be seen by examining the cutoff frequency for a rectangular wave guide, given in Jackson [46] (equation 8.44) by the expression

$$\omega_{m,n} \geq c\pi \sqrt{\frac{m^2}{a^2} + \frac{n^2}{b^2}} \quad (1.53)$$

which is the frequency at which the wave's phase velocity in the waveguide becomes infinite and below which the phase velocity is imaginary, i.e, the mode is evanescent. In this equation, m and n are the mode numbers in the two waveguide cross section directions, and a and b are the respective waveguide wall lengths in those directions. In order to create a perturbation with a dominant $n_{\parallel} = 7$ Fourier component, the wave must have a wavelength in one direction equal to $\lambda = c/n_{\parallel}f$. Therefore, the length of one side of the wave guide (we will choose a) is determined, $a = m\lambda = mc/n_{\parallel}f$. Equation 1.53 then gives us a restriction on the size of the other dimension, specifically

$$2f \geq c \sqrt{\frac{m^2 n_{\parallel}^2 f^2}{m^2 c^2} + \frac{n^2}{b^2}} \quad (1.54)$$

$$\frac{4f}{c} \geq \frac{n_{\parallel}^2 f^2}{c^2} + \frac{n^2}{b^2} \quad (1.55)$$

$$\frac{n^2}{b^2} \leq (4 - n_{\parallel}^2) \frac{f^2}{c^2}. \quad (1.56)$$

This equation cannot be solved if $n_{\parallel} > 2$ because n , b and f must be pure real numbers (the speed of light c is always a pure real number). Therefore a single waveguide is not acceptable.




Figure 1.15: A multiseptum waveguide can produce different spectra for different phasing. Shown above is $\pi/2$ phasing, meaning every 5th septum has the same phase. This scheme is favored by most fusion experiments in this frequency region because the design is fairly simple and robust, and has been shown to have good performance.

A multiseptum design can be considered however. The multiseptum design has the advantage of maneuverability, by which I mean the launch spectrum can be modified, by changing the phasing of the array. In order to exploit this technique with optimal phasing ($\pi/2$), the MST system would need to have an array of $\lambda/4 \approx 1.33$ cm. Unfortunately, in order for a waveguide to support an 800 MHz $TE_{0,1}$ mode, the cross-field length must be at least half the vacuum wavelength (i.e. 18.75 cm) in order to satisfy the cutoff restriction 1.53. This is unacceptably large for MST, and so a different solution must be pursued. The solution implemented is a traveling wave antenna, and is the subject of Section 2.2.2.

1.4 Wave coupling

The process of moving power from the antenna face to the plasma is called wave coupling. Although a satisfactory analytic interpretation of wave coupling is difficult to develop, a qualitative understanding of the process is fairly straightforward. This is illustrated in Figure 1.16. The area around the the antenna is divided into four regions. The antenna itself (I), the vacuum region (II), a non-propagation region (III), and a propagation region (IV). The vacuum cannot support an electrostatic wave, and so the field in the vacuum decays exponentially. The first region of plasma is very low density, and does not support a cold plasma wave. This corresponds to the region where $P > 0$, from the cold plasma analysis. It should be noted that width of this layer is likely to be quite small as $P = 0$ occurs (at 800 MHz) to a density of $n_{crit} = \omega^2 m_e \epsilon_0 / e^2 \approx 8 \times 10^{15} \text{ m}^{-3}$, and the smallest density measured at MST's edge is about $1 \times 10^{17} \text{ m}^{-3}$. The far edge of this layer, the point where the evanescent plasma mode turns into a propagating mode ($P = 0$) is called the critical layer. What we expect from this simple picture is the farther the critical layer is from the antenna, the more the mode will decay before coupling, and the less power will radiate into the plasma.

Difficulties with the critical layer arise in high confinement plasmas as the plasma column is observed to “pinch” inward, away from the antenna. As this work will show, these plasmas are of the most interest to us, so it is vitally important to mitigate the effect of this pinching, without disturbing the plasma stability. It is observed in these high confinement discharges that fueling




Figure 1.16: Cartoon of antenna coupling. There are four regions in this problem, the slow wave structure (I), the vacuum between the plasma and the antenna (II), the region of the plasma without wave propagation i.e. cold plasma $P > 0$ (III), and the full propagation region (IV).

with the main MST puff valves leads to a confinement degradation, though the reasons behind this are not well understood. Experiments run on JET [51] indicate that edge fueling can improve slow wave coupling when the local plasma density is low. Some similar coupling studies have been done on MST, including some studies looking at plasma edge fueling. [52] Unfortunately, these experiments (while promising) were run with the Mark II lower hybrid antenna which did not have the diagnostic capabilities of the Mark III lower hybrid antenna. Because of these less sophisticated diagnostics this work did not allow for sufficient understanding of coupling dynamics, and therefore could not distinguish between coupling improvement due to doping and coupling improvement due to plasma destabilization. Both the problem and the solution are not well understood.

1.5 Non-Local Ohm's Law

One of the more intriguing effects of the stochastic diffusion mentioned in Section 1.2.5 is the breakdown in the localization of Ohm's law with respect to flux surface averaged fields. In traditional derivations of Ohm's law, the presence of an electric field is observed to create a particle flow, as charged particles are accelerated by the Lorentz force. The acceleration from the electric field is resisted by drag imposed on the electrons from Coulomb collisions with other charged particles in the medium in which the electrons are traveling (a plasma in this case). The net effect of this is represented by a scalar relationship between the electric field and current (this becomes a tensor relation when anisotropy due to magnetic field is taken into account), with resistivity being the relational constant, derived by applying these assumptions to the Boltzmann equation.

It is important to exercise semantic caution when discussing this effect. To clarify this point, the plasma can be modeled with a litz wire. A litz wire is a bundle of insulated wires braided together in a random fashion. This can be thought of as a simple model of field line stochasticity, as (in a stochastic plasma) electrons are bound to follow magnetic field lines to the first order (hence the insulation), and like wires, they cannot cross directly, they can only wrap around one another. A radially dependent potential is applied to one end of the cable bundle and the other end is grounded. In this case, the standard, local conception of Ohm's law (simply $E = \eta J$) holds. What breaks down is our interpretation of the term local. That is to say, the electric field that determines the

current at a given point in a wire is dependent on the potential on the ends of that wire, and any inductive field from the neighboring wires. It is not accurate to claim that an average of the electric field at a given radius represents the local field at any point at that radius. It is likely not even a particularly good approximation. This insight can be applied to a stochastic plasma. Unlike wires however, it is experimentally difficult to determine what measurements occur on connected field lines. Therefore, from a diagnostic standpoint, this picture is abstracted to some extent, and rather than imagining a number a field lines that have a definite trajectory in real space (which is what occurs in the actual system), the model of a random walk of field lines is employed. The transfer of momentum is conceived of (in some average sense, as described below) as a diffusive process rather than a direct consequence of following a somewhat erratic path. It is only in this context that Ohm's law can be described as non-local.

Rechester and Rosenbluth showed [53] that in a plasma with large overlapping magnetic perturbations, field lines are well described by the classic random walk model. Specifically, the radial deviation of the field lines is given by the formula

$$\langle (\Delta r)^2 \rangle = 2LD_m \quad (1.57)$$

Where $\langle (\Delta r)^2 \rangle$ is the average of the square of the radial displacement of a field line given a step size of L along the field line. The coefficient D_m is related to the size of the magnetic perturbations in the plasma. The actual magnitude of this term is dependent on the interpretation of the correlation length in Rechester-Rosenbluth diffusion, and is given by [54]

$$D_m = \Lambda \left(\frac{\langle b_r \rangle}{B} \right)^2 \frac{1}{(1 + \Lambda / \lambda)} \quad (1.58)$$

$$(1.59)$$

Our model assumes that the collisional mean free path λ is much larger than the correlation length Λ , which we assume must be smaller than the minor radius. This means that $D_m = 10^{-4}m$ corresponds to a radial field that is about 4% of the total field magnitude using the somewhat pessimistic value of $\Lambda = a$ the minor radius.

A corrected Ohm's law can be derived from the Boltzmann transport equation to illustrate the difficulties posed by the diffusive assumption and non-localized complication. The full Boltzmann

transport equation is

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f + \frac{q\mathbf{E} + q\mathbf{v} \times \mathbf{B}}{m} \cdot \nabla_v f = \frac{\partial f}{\partial t} \Big|_{coll}, \quad (1.60)$$

and (when the equation is solved for all plasma species) describes without approximation, the state of all the particles in the plasma. This equation can be simplified by dividing it into a Maxwellian distribution and a perturbed distribution, i.e., $f = f^{(0)} + f^{(1)}$ where $f^{(0)}$ is a Maxwellian and $f^{(1)}$ is a modification to the Maxwellian. Recall that the Maxwellian by definition solves the Boltzmann equation without the additional force term. So without loss of generality

$$\frac{\partial f^{(1)}}{\partial t} + \mathbf{v} \cdot \nabla f^{(1)} + \frac{q\mathbf{E} + q\mathbf{v} \times \mathbf{B}}{m} \cdot \nabla_v f = \frac{\partial f^{(1)}}{\partial t} \Big|_{coll}, \quad (1.61)$$

Because electrons are constrained to follow wandering field lines, they will radially transfer momentum at a rate proportional to their velocity. Jacobson and Moses [55] argue that the contributions of this momentum transfer can be accounted for in the Boltzmann equation by averaging the convective term ($\mathbf{v} \cdot \nabla f$) and the magnetic portion of the Lorentz force $[(q/m)\mathbf{v} \times \mathbf{B} \cdot \nabla_v f]$ over the spatial scale of the wandering magnetic field lines. Assuming the collision length is much larger than the correlation length (related to the amount of distortion experienced by field lines), this transport term takes the form of a Fick's law like diffusion with the parallel velocity dependence predicted by Rechester and Rosenbluth. In this simplification of the Boltzmann equation implicitly requires the assumption that the collision length is much larger than the correlation length. This gives a kinetic equation (for large spatial scales) of

$$\frac{\partial f^{(1)}}{\partial t} - \frac{eE_{\parallel}}{m} \frac{\partial f^{(0)}}{\partial v_{\parallel}} = \frac{\partial f^{(1)}}{\partial t} \Big|_{coll} + |v_{\parallel}| \frac{\partial}{\partial x} D_m \frac{\partial f^{(1)}}{\partial x}. \quad (1.62)$$

Taking the Krook approximation (that collisions can be represented by a collision time such that $\partial f / \partial t|_{coll} = -f_1 / \tau$), and solving for a steady state solution for f_1 , we get

$$f^{(1)} = \tau \frac{eE_{\parallel}}{m} \frac{\partial f^{(0)}}{\partial v_{\parallel}} + \tau |v_{\parallel}| \frac{\partial}{\partial x} D_m \frac{\partial f^{(1)}}{\partial x} \quad (1.63)$$

Finally the current density can be calculated by usual formula

$$J_{\parallel} = - \int e v_{\parallel} f d^3 v = -e \int v_{\parallel} f^{(1)} d^3 v + J_{drift} \quad (1.64)$$

where J_{drift} is the current from the drift velocity of the electron species. Plugging in the value of the perturbed distribution (equation 1.63), we get

$$J_{\parallel} = -e \int v_{\parallel} \tau \frac{eE_{\parallel}}{m} \frac{\partial f^{(0)}}{\partial v_{\parallel}} d^3v - e \int \tau |v_{\parallel}| \frac{\partial}{\partial x} D_m \frac{\partial f^{(1)}}{\partial x} d^3v + J_{drift} \quad (1.65)$$

To put this into simpler terms, we define

$$\eta \equiv \left(\int v_{\parallel} \frac{\tau e^2}{m} \frac{\partial f^{(0)}}{\partial v_{\parallel}} d^3v \right)^{-1} \quad (1.66)$$

making our final Ohm's law

$$\eta J_{\parallel} = E_{\parallel} + \eta \int e v_{\parallel} \tau |v_{\parallel}| \frac{\partial}{\partial x} D_m \frac{\partial}{\partial x} f_1 d^3v, \quad (1.67)$$

For non-stochastic discharges, such as high confinement MST discharges and tokamak plasmas, this nonlocal portion of Ohm's law is term is fairly unimportant. This is mainly because the diffusion coefficient is relatively small, and (more importantly) it does not carry with it an extra factor of parallel velocity [56, 57]. This extra weighting means fast electrons are preferentially diffused. The effects of the this diffusive element of Ohm's law have been studied both numerically [58] and analytically [54] with assumed current profiles in cylinders (although the profile in the numerical treatment was allowed to evolve). The effects are seen in Figure 1.17

1.6 Objectives of this work

The main body of this work seeks to address three essential questions for kinetic modeling and lower hybrid current drive on the RFP. Each of these has secondary questions associated with it.

- **Does the Fokker-Planck Model provide a reasonable representation of RFP plasmas?**

Inductively stabilized plasmas are shown to be well described by Fokker-Planck Models. What problems arise when considering Standard RFP discharges? Can Enhanced confinement discharges be described accurately? Can Fokker-Planck Modeling predict the diffusive portion of Ohm's law? How does this compare with theory?

- **How much LH power is needed to drive the required amount of stabilizing current?**

What range/types of discharge is LH good for? Can LH drive sufficient current even when




Figure 1.17: The effect of diffusion on parallel current, using the model from Martines [54] and MST like electric field and density profiles. Notice as diffusion is increased, edge current will eventually run opposite the edge electric field. The black line is the current profile assuming Spitzer resistivity.

radial diffusion is taken into account? Can the required amount of current be driven with a physically realizable system? How is this effected by using a realistic launch spectrum instead of a single ray? What physical effects are important for LH deposition?

- **Can antenna coupling be controlled in the RFP?** What are the primary impediments to antenna coupling? Can these be overcome without compromising the plasmas under consideration? What are the primary constraints on the conditions of the plasma edge?

Chapter 2

Experimental Setup

2.1 The Madison Symmetric Torus

This work was done at the Madison Symmetric Torus (MST), a world class RFP facility located in Madison, Wisconsin. The MST is a long running plasma physics experiment, having operated continuously for more than 20 years. MST is involved in a large number of basic plasma science experiments, with applications to astrophysics and fusion energy science. Figure 2.1 is a photograph of the MST machine area. A rundown of typical parameters is given in Table 2.1.

MST Parameters	
Major Radius	1.5 m
Minor Radius	0.52 m
Plasma Current	200 – 600 kA
Core Electron density	$0.4 - 2 \times 10^{19} \text{ m}^{-3}$
Core Electron Temperature	200 – 2000 eV
Core Ion Temperature	100 – 1000 eV
Magnetic Field on Axis	0.3 - 0.6 T
Edge Magnetic field	0.1 T
Plasma β	5-10 %

Table 2.1: Approximate experimental parameters in MST

Unlike many other plasma experiments, MST's vacuum vessel plays a vital role in MST operations beyond providing a low pressure environment [59]. The close fitting shell stabilizes surface




Figure 2.1: A picture of the Madison Symmetric Torus. The Lower Hybrid antenna is below the boxport on the left side of the picture, and the X-ray array can be seen on the box port on the right side of the picture.

instabilities. Image currents on the inner surface of MST's shell passively provide the vertical field which stabilizes the hoop force. Finally the toroidal field is formed by passing current through the shell of MST. Because of these vital tasks performed by the shell, MST cannot have large holes bored in it because of the field errors that arise. This is an important constraint for RF experiments since it provides a limitation to the usage of waveguides.

2.1.1 Operational Modes

MST has several different operational modes, of which the most interesting for this work are standard RFP mode, Enhanced Confinement mode and Pulsed Poloidal Current Drive mode (henceforth Standard, EC and PPCD respectively). Standard plasmas are traditional RFP plasmas as described in the introduction. Experimental profiles for standard plasmas are shown in Figure 2.2. PPCD and EC are discussed in detail below.

2.1.1.1 Pulsed Poloidal Current Drive

In an attempt to stabilize RFP plasmas, a secondary inductive circuit has been attached to MST. This circuit is used to induce an additional poloidally-directed electric field, after a shot is already underway. This in turn leads to reduced magnetic fluctuations, reducing the effect of the MHD dynamo. [11] These so called Pulsed Poloidal Current Drive (PPCD) discharges are shown to have improved confinement characteristics, hotter electron and ion temperatures and reduced magnetic fluctuations. [10] Additionally, hard x-ray diagnosis has been used to determine that transport is velocity independent in PPCD plasmas, [60] implying reduced stochastic field line wandering, and possibly restored flux surfaces. PPCD plasmas provide good experimental evidence of the theory of RFP tearing mode stabilization using current drive. Unfortunately the current driven in this manner is inherently pulsed as induction is used, and so it is not acceptable for the very long pulse lengths, usually considered required for power plant operations. Profiles from PPCD plasmas are shown in Figure 2.3.




Figure 2.2: Reversal Parameter ($B_\phi(a)/\langle B_\phi \rangle$), $n=4$ mode activity (a mode which does not have an $m=1$ resonant surface), core line-averaged density and core temperature for a typical MST discharge. Core temperature is typically between 200 and 350 eV. A regular sawtooth pattern is seen with a period near 8 ms, though this is not particularly regular. Fast electrons are lost at the peak of the sawtooth. This discharge (the most studied) had a plasma current of 400 kA.

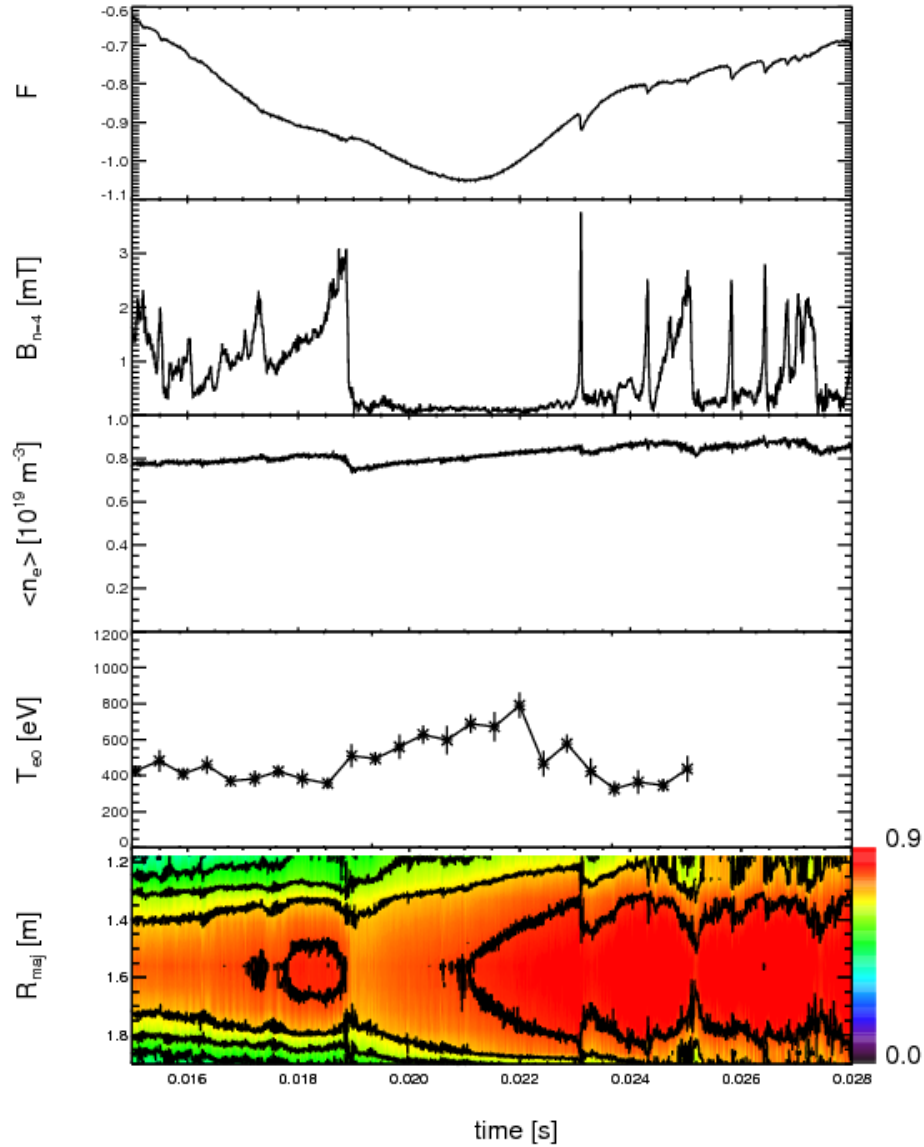


Figure 2.3: Pertinent profiles from a PPCD discharge. Plasma temperature is typically between 800 and 1000 eV (for 400 kA PPCD plasmas). Density typically increases in the core, and the plasma column is observed to pinch inward. Magnetic activity is greatly reduced. The bottom plot is a contour plot of the density across the major radius from interferometry.

2.1.1.2 Enhanced Confinement

Like the tokamak's H-mode, MST has an operational mode that is associated with anomalously improved temperature and reduced magnetic mode activity. This so-called Enhanced Confinement (EC) regime is shown to have strong $E \times B$ flow shear [61] as is observed in tokamak H-mode plasmas. [62] The reduction in mode activity in EC plasmas is strongly correlated with quasiperiodic $m = 0$ MHD activity (denoted by Chapman as "Small Dynamo Events"), separated by quiet periods. Enhanced Confinement plasmas are observed in MST when running low densities and deep reversal. The thresholds for this transition vary with plasma current. Enhanced confinement is observed to onset at densities of around $6 \times 10^{18} \text{ m}^{-3}$ and a reversal parameter (F) of around -0.3 for 0.4 MA plasmas, and are seen to become more dramatic as F or the density are reduced. It has not been previously established whether EC plasmas exhibit stochastic transport in the same way that standard plasmas do. The electron temperature increases observed in these plasmas are (in the best conditions) comparable to those achieved during PPCD discharges. The large scale sawtooth oscillations seen in standard plasmas are seen in EC plasmas however, which is found to limit their temperature. EC plasmas have been observed to have a density pinch in the same way as PPCD plasmas. Plasma profiles for EC plasmas are displayed in Figure 2.4.

2.1.2 Diagnostics

MST has a comprehensive diagnostic suite, only a portion of which are discussed here. The following diagnostics provide crucial constraints for Grad-Shafranov reconstructions which are used extensively in this work.

2.1.2.1 Faraday Rotation Diagnostic

Faraday rotation is used help constrain the magnetic field profile. A laser is fired directly down into the MST plasma. It travels nearly perpendicular to the plasma magnetic field. There are two wave modes available to electromagnetic waves in plasmas, with different polarizations and phase velocities (these were discussed in the last chapter). When a laser of appropriate wavelength hits the plasma, it decomposes into an R-mode wave and an L-mode wave. These waves have different




Figure 2.4: Profiles from a EC discharge. Core plasma temperature is seen to increase to well above the levels seen in standard plasmas, sometimes reaching that seen in PPCD plasmas, although the density is quite a bit lower. Magnetic activity is generally confined to small bursts separated by relative quiet periods. Again, the bottom plot is a density profile from the FIR interferometry diagnostic. Notice the decay in the density.

dispersion relations, and consequently, travel through the plasma at different rates (the R-mode is faster). The discrepancy in the speeds of the waves leads to a shift in polarization when the wave recombines on the other side of the plasma. From this, the magnetic field can be inferred. MST is equipped with an 11 chord far-infrared laser system which can be used for Faraday rotation. The chords of this laser system are vertically oriented and span the torus midplane from major radius of 1.18 m to 1.93 m. This translates into coverage of 75 cm of MST centered near the presumed magnetic axis, as shown in Figure 2.5. [63] This beam can be run as a traditional interferometer to measure high frequency (> 250 kHz) density fluctuations [64], or as a polarimeter [63]. More recently the system has been used to measure current profile evolution and the Hall dynamo effect. [65,66]




Figure 2.5: A poloidal cross section of MST showing the FIR chords.

2.1.2.2 Motional Stark Effect

The motional Stark effect (MSE) is used as a diagnostic for direct measurements of the plasma magnetic field on axis. [67] The Stark effect is a perturbative quantum effect, whereby an applied electric field breaks the spin degeneracy of the excited states of the atom. For the plasma diagnostic, a beam of neutral particles is injected. When considered in the atom's frame of reference, the lab

frame's $\mathbf{v} \times \mathbf{B}$ Lorentz force appears as an electric field. Stark splitting is calculated in this frame (hence the term “motional”) and then boosted into the lab frame, where it is measured. Thus by determining the polarization of the emitted light near the H_α Balmer line, the magnetic field can be determined. In high field devices, this is sufficient on its own, however in a lower field device such as MST, the line separation is not large compared to the profile broadening due to finite temperature effects and large sampling volume. Thus a more complicated model must be used. [68] In practice this means that the MSE spectrum is fit with a model Doppler broadened profile. This diagnostic is unique in the RFP community, and it provides vital information on RFP dynamics due to the pivotal role of magnetic perturbations. This work only concerns itself with the longer time scale behavior of the magnetic field for Grad-Shafranov reconstructions.

2.1.2.3 Thomson Scattering Diagnostic

Plasma temperature is measured using a Thomson scattering system. [69] The process of Thomson scattering can be understood with a fairly simple quasiclassical model. When photons are fired at an electron, they are reflected off at some angle (more properly quantum physics tells us they are absorbed and re-emitted). When this scattering occurs, for a static electron, the emitted photon will have a wavenumber equal to the projection of its original wave number along the scattering chord (see Figure 2.6). Motion causes a Doppler shift in the emitted light. When a laser is applied to a Maxwellian distribution, the reemitted spectrum will have a Gaussian shape, and the width of that Gaussian will be related to the temperature. Optical ingenuity can be used to isolate photons from a small volume. Thus the temperature across the plasma can be determined. The core of MST's Thomson scattering system is a pair of Nd:YAG lasers. Recent advances in the system allow [70] for a total of 30 2 J pulses at a rate of 1-12.5 kHz [71], allowing for multiple well diagnosed reconstructions [72] of MST discharges when combined with FIR data. The data from the Thomson scattering system is vital to this work as both input for reconstructions and as a measurement of absorption regions. The Thomson lasers enter MST vertically from the top, and a lens is used to focus and disperse photons coming from different radii. These photons are then picked up by fibers and fed into polychrometers for spectral analysis (see Figure 2.7).




Figure 2.6: Here an incident wave (on the left) scatters off of a charged particle. if the the particle is stationary then the emitted photon is just a projection of the incident photon. Particle motion causes a doppler shift in the emitted light. The aggregate of this Doppler effect on an entire distribution gives information about the thermal portion of the distribution function. This image is in the public domain.

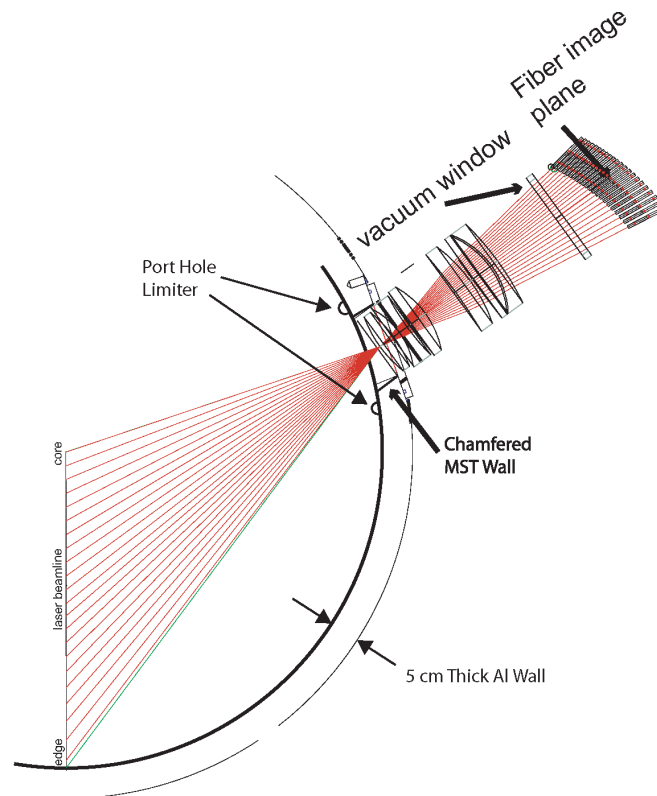


Figure 2.7: The viewing chords of the Thomson scattering diagnostic on MST. The laser scatters at various points along a vertical line and is picked up by fibers and fed to polychrometers where the data is analyzed, and the temperature at different radii is determined.

2.1.3 Reconstructions

The temperature, density and magnetic field measurements from these diagnostics are leveraged to gain a more complete understanding of the configuration of the plasma in an MST discharge. They can be combined to predict F and P profiles, which constrain the current profile in the plasma by the Grad-Shafranov equation. It should be noted that Grad and Shafranov's model ignores the effects of large magnetic fluctuations, which are an important factor in standard RFP discharges. However, it is generally accepted that any oscillations in the plasma must be centered around a Grad-Shafranov equilibrium, otherwise the plasma will collapse fairly rapidly. It must be stressed that reconstructions are based on real, measured data, and therefore, they can be considered as at least an "Educated Guess" splining routine, if not completely accurate. As will be shown later, standard plasmas are not the most important discharges for rf experiments.

MSTFit [73, 74] is a plasma equilibrium code which reconstructs the state of the plasma at a given time using experimental data to constrain the fit. A close fitting conducting shell is assumed, as is required for RFP stability against external mode activity. The grid is initialized (usually with one of the cylindrical analytic models of RFP discharges). This is used to reconstruct values for F (the toroidal field times the major-radial coordinate R) and p , the plasma pressure, as well as the poloidal flux ψ . The grid solution is then adjusted in order to minimize a chi-squared goodness-of-fit statistic. The results are then transposed to a rectangular grid for use in other codes.

2.1.4 X-ray Spectroscopy

X-ray observations can be used to study electron bremsstrahlung. X-ray detectors aim to determine the emission from an infinitely thin column of plasma (traditionally called a pencil beam). In practical experiments, this is accomplished by placing a detector and a series of apertures near the plasma. A realistic detector will not see a pencil beam but will see a cone whose size is determined by the size of the apertures and the distance from the plasma. This cone is characterized by a value called etendue (G), which is a Lagrange invariant for any optical system with two apertures. [75] An example of this setup is given in Figure 2.8


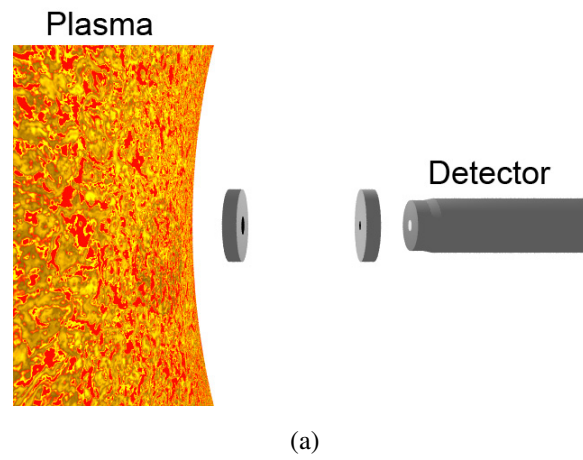


Figure 2.8: A cartoon of the x-ray detectors. Figure (a) shows a typical setup. The aperture nearest the the plasma is called the collimator and the aperture nearest the detector is called the pupil. Figure (b) shows the observation cone of a given aperture setup

MST is equipped with an array of single photon x-ray detection equipment, utilizing either the XR-100CR from Amptek for soft x-ray detection, or the ev-CZT detectors from eV Microelectronics (now EI Detection and Imaging Systems, a division of Endicott Interconnect Technologies, Inc.) for the detection of hard x-rays. These are primarily located on a boxport flange at 135° toroidal. This boxport allows detection of x-rays at many different radial locations (see Figure 2.9). Each detector chain consists of a PIN photodiode operated in avalanche mode, a preamplifier, and an amplifier. The diode material determines the sensitivity of the crystal at different x-ray energies, Silicon (Si) for x-rays from 3-10 keV and Cadmium Zinc Telluride (CdZnTe) for x-rays from 10-150 keV. Although there is some variation in the specifics, the general operation of all avalanche PIN photodiodes is essentially the same and is described in many references, for example [76], and is summarized here.

A diode is placed in the path of photons in the correct energy range. A bias is applied to the diode so that it is nearly in breakdown. An incoming x-ray interacts with the photoactive material, either through the photoelectric effect or Compton scattering, creating a hole-electron pair in the intrinsic semiconductor layer. The hole and electron are drawn in opposite directions by the electric field. During this process, the hole and electron will collide with the atoms of the crystal, releasing electrons or holes in proportion to their energy. These will then be accelerated by the field and the process will repeat, leading to an avalanche of holes and electrons. The total charge released in this avalanche is proportional to the energy of the original electron/hole pair, and therefore to the energy of the incoming photon. The diode is then connected to a charge sensitive preamplifier which conditions the signal into a more useful form. This preamplified signal is then shaped with a Gaussian pulse shaper and digitized. Raw data from the soft x-ray detector are presented in Figure 2.10. A fitting routine determines the time center and energy peak of each individual photon. The results are used for spectroscopy, using the scheme outlined in §3.3.2.

X-ray spectra have traditionally been analyzed in terms of intensity, i.e. total power generated by x-ray emission within a certain energy range. This can be attributed to the limitations of technology. In the past, most commercially available x-ray detectors were integrated, that is they produced an output current proportional to the total amount of power deposited within the diode at a given



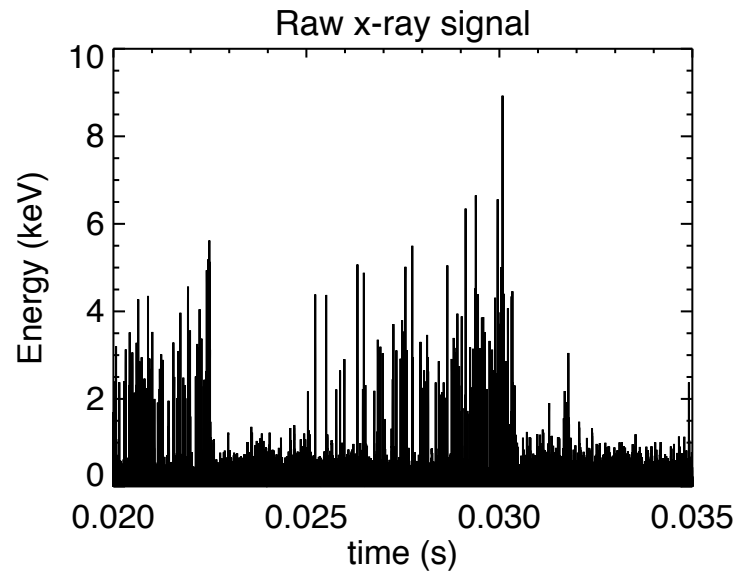



Figure 2.9: Hard x-ray viewing chords shown in a cutaway, and detectors in place. Marked with the viewing chords is the geometric axis (in blue) and the presumed magnetic axis at $R_{maj} = 1.56m$. Detector positions are numbered starting with 1 at the inboard-most chord. Chord 8 looks at the presumed magnetic axis.



(a)



(b)

Figure 2.10: Raw data for the Amptek soft x-ray detector. Figure (a) shows raw data with many pulses. Figure (b) shows an individual x-ray, with the time relative to the time of incidence. Individual points from the digitizer are also plotted. Note the change in scale.

time. With a system of this type, energy discrimination can be achieved by using filters of various materials and thickness, with known transmission properties. Thus, an array of x-ray detectors can detect an entire spectrum, assuming they are looking at the same plasma volume. Alternatively, many shots can be ensembled together with the filters being altered in order to produce a full, energy resolved spectrum. Since there is no way to determine the energy of individual incoming photons within a given energy bin, the only rational way to analyze these data is to consider the total deposition power over a certain period of time in a detector binned for certain energies.

More recently, single photon detectors have become available on the commercial market. Although the data can be readily analyzed using a similar method to the integrated detector (by binning the photons detected and summing their energies), we are free to interpret the data in a variety of other ways. The method of most relevance to this work is calculation of the number flux (or counting flux, henceforth flux), i.e., the number of photons contained in a given bin defined in energy and time. These two methods must be compared for relative accuracy. This point will be returned to in Section 3.3.2.

There is a major caveat that merits further study. Both of the above mentioned interpretations rely on the "pencil-beam" assumption. That is, we assume that the emission coefficient is constant regardless of the size of the detection cone (i.e. doubling the etendue of the optical system will merely double the detected emission). This assumes that our viewing cone is smaller than any perturbations in the plasma. While likely true in some average sense, this can feasibly be a significant source of error which will be amplified for setups with larger etendues.

2.2 The Lower Hybrid System

Although the focus of this thesis is on the interactions of electrostatic waves with plasmas, the challenge of providing hundreds of kilowatts of high frequency rf power with a precisely controlled wave number presents its own difficulties and requires some innovative solutions. The electrical system at the MST facility cannot support the service required to power a full 800 MHz transmission system directly. In order to facilitate experiments therefore, energy must be stored and quickly discharged. This was accomplished by building a pulse forming network (PFN). The PFN

provides a high voltage DC signal to the klystron which generates the rf for the system. The lower hybrid system is therefore naturally divided schematically into the power generation system and the rf transmission system.

2.2.1 Support Systems

2.2.1.1 Power System

The lower hybrid PFN (Figure 2.11) is an LC network used to supply a DC square pulse as a power source to the system klystrons. When charged to 100 kV the capacitors hold a total of 70 kJ which is discharged over the course of 30 ms, leading to a power supply rate of about 2.3 MW.

This power supply is designed to power three klystrons with a 40 kV 20 A pulse and so there is a water resistor in parallel with the current klystron to dissipate the extra energy. The PFN is controlled by a novel switch stack made up of thyristors. [77,78] When a trigger is sent to the PFN, the first (main) thyristor stack is triggered and the stored energy is discharged into the klystron and the water resistor. If the current from the PFN spikes above the expected amount (indicating an internal arc) or the optical arc detector inside the klystron cavity sees an arc then the crowbar stack will discharge in order to prevent further damage.

2.2.1.2 The Klystron

The klystron used primarily in this work is a Varian 7955 S (S stands for super), reconditioned by Freeland Products to provide higher output power. A schematic of a klystron is shown in Figure 2.12a. Before LH operations, the filament is turned on. This provides a small number of hot source electrons, which strike the emission medium and liberate many times more electrons, which are relatively low energy, but exist for a reasonable amount of time because the klystron interior is evacuated. When a shot occurs, the large DC voltage supplied by the PFN is applied between the cathode and the modulating anode, accelerating the electrons into the main cavity section. This high velocity electron beam is held focused by an axial magnetic field. A low power rf input signal is applied to the first cavity. As the AC signal enters the cavity walls, an electric field is induced between the left and right sides of the cavity. This causes bunching of the electron

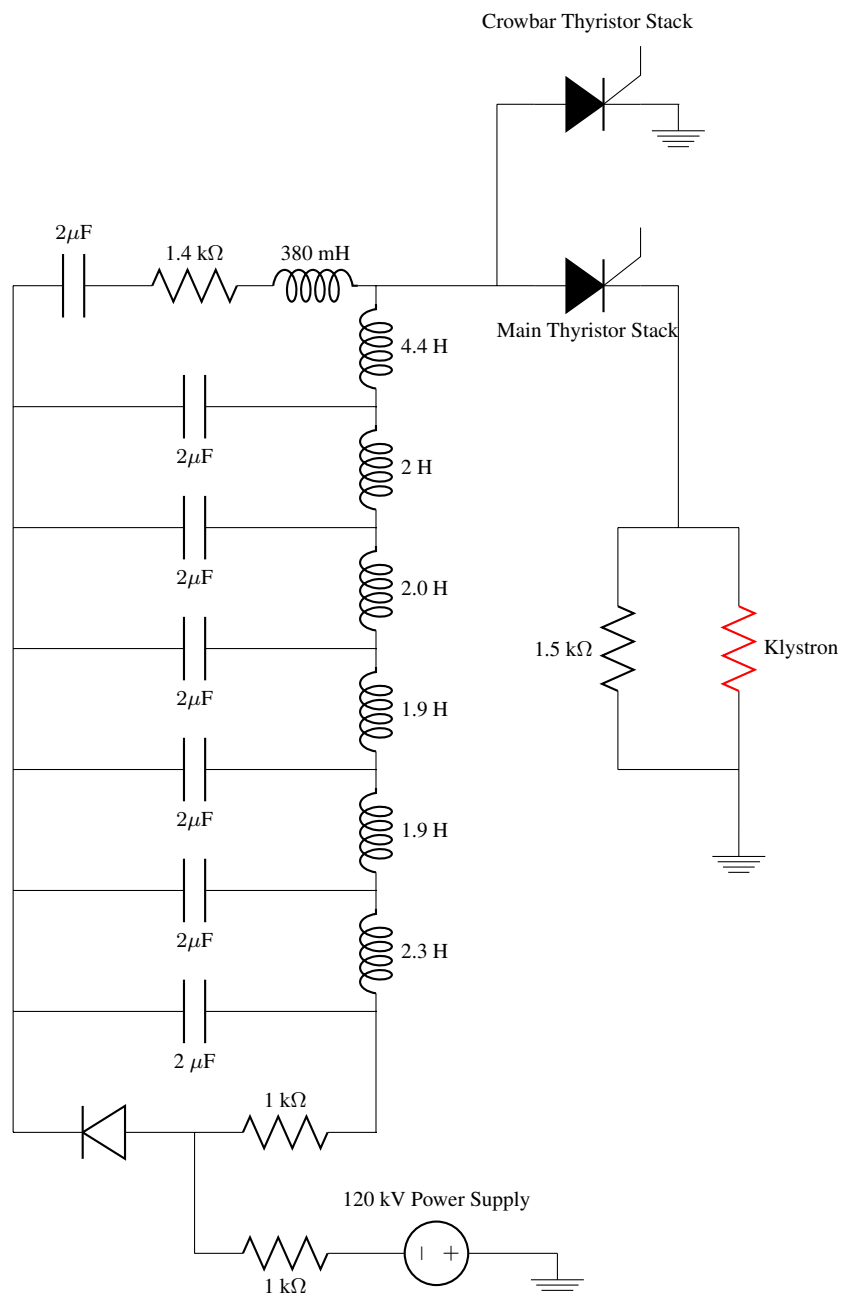


Figure 2.11: The lower hybrid power system


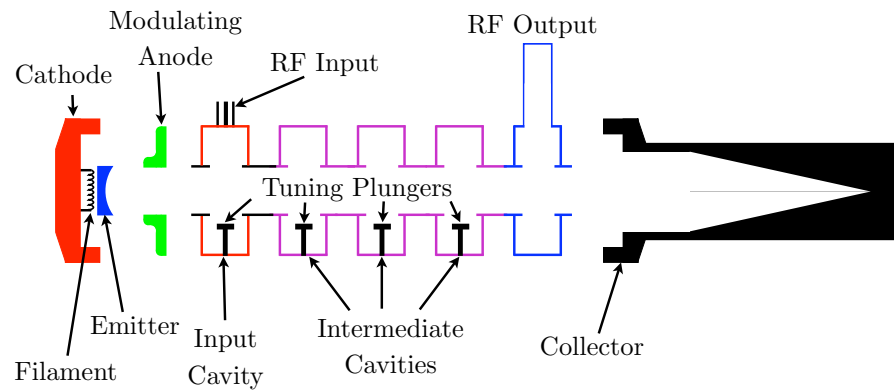


Figure 2.12: (a) shows an operational schematic of a klystron. Notably absent are the magnetic field coils which provide an axial field which focuses the electron beam. Figure (b) shows a cutaway view of a cavity

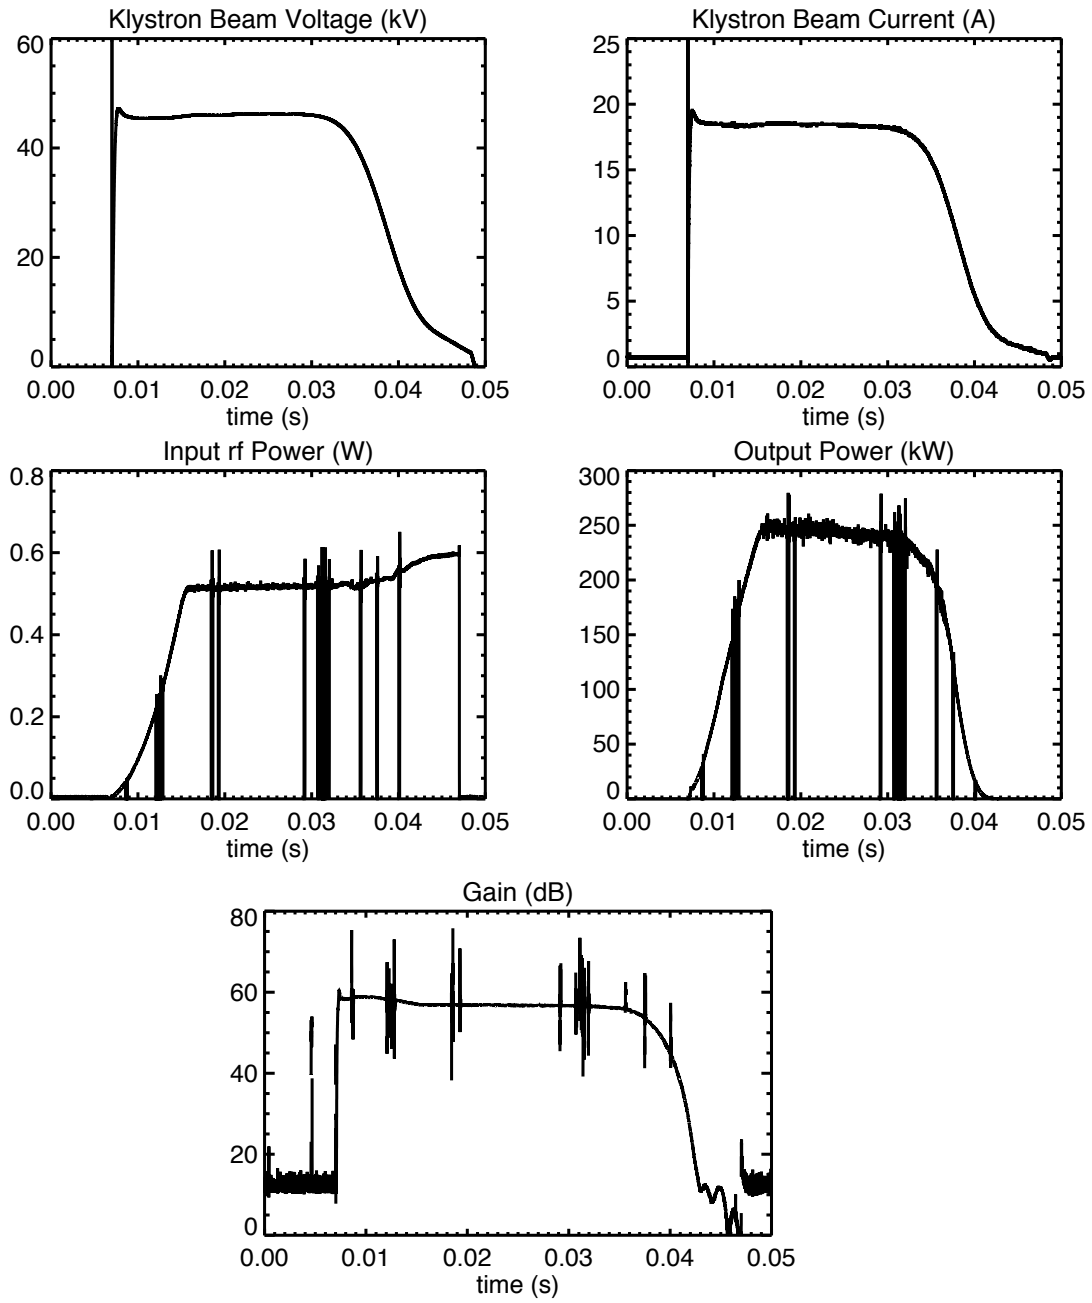


Figure 2.13: Operational signals of the klystron, the Voltage and current of the klystron beam on top, the input power in watts and the output power in kilowatts in the middle, and the gain G in decibels, where $G_{dB} = 10 \log_{10} \left(\frac{P_{out}}{P_{in}} \right)$.

beam which is progressing rapidly down the center of the klystron. When the slightly bunched beam enters the second cavity, the electrons induce a voltage at the cavity entrance. The entrance is capacitively and inductively coupled to the exit. The magnitude of capacitance (and to some extent the inductance) can be altered by moving the tuning plungers inward or outward. If the capacitance and inductance are tuned correctly, then the voltage peak will reach the cavity exit in resonance with the bunched electrons, causing further bunching. This process is repeated in each of the cavities until the electrons are extremely bunched. Because the field induced in the cavity is proportional to the beam current, a higher DC voltage will lead to a larger amount of bunching. The bunched beam is passed through a loop, inducing a high power, quickly oscillating voltage wave. The beam is then dumped into the collector, which is simply a large hunk of metal surrounded with a great deal of copper to aid in cooling. This system is very robust, and can be run rather hot (they are typically “cooled” with high pressure steam at 400° F). Because of this robustness, and the short pulse duration (~ 30 ms), the klystron can be overdriven, meaning the applied DC voltage is substantially higher than the specification, so long as the resonant cavities are tuned for such operation.

2.2.1.3 Lower hybrid transmission system

The signal input to the klystron is a fairly low power (< 10 W) rf signal generated by a bench sized signal generator (an Anritsu MG3641N), which is amplified by a 50 W amplifier. The input signal can be modulated, and in this way the power supplied to our antenna can be varied dramatically during a shot. This functionality is only used in practice to turn the antenna off for periods of time and to allow for a “soft start up”, gently ramping the power up to maximum instead of dumping the full power immediately in order to prevent arcing in the antenna. The rest of the transmission system is sketched out in Figure 2.14. Power from the klystron is directed to a circulator, which efficiently moves the power around once clockwise. This is done so that any reflected power can be redirected away from the klystron (in order to protect it) to a $50\ \Omega$ dummy load, essentially a large, cooled resistor which is capable of dissipating up to 300 kW of power. From the circulator, the power is brought into the MST machine area where the ground is stripped to

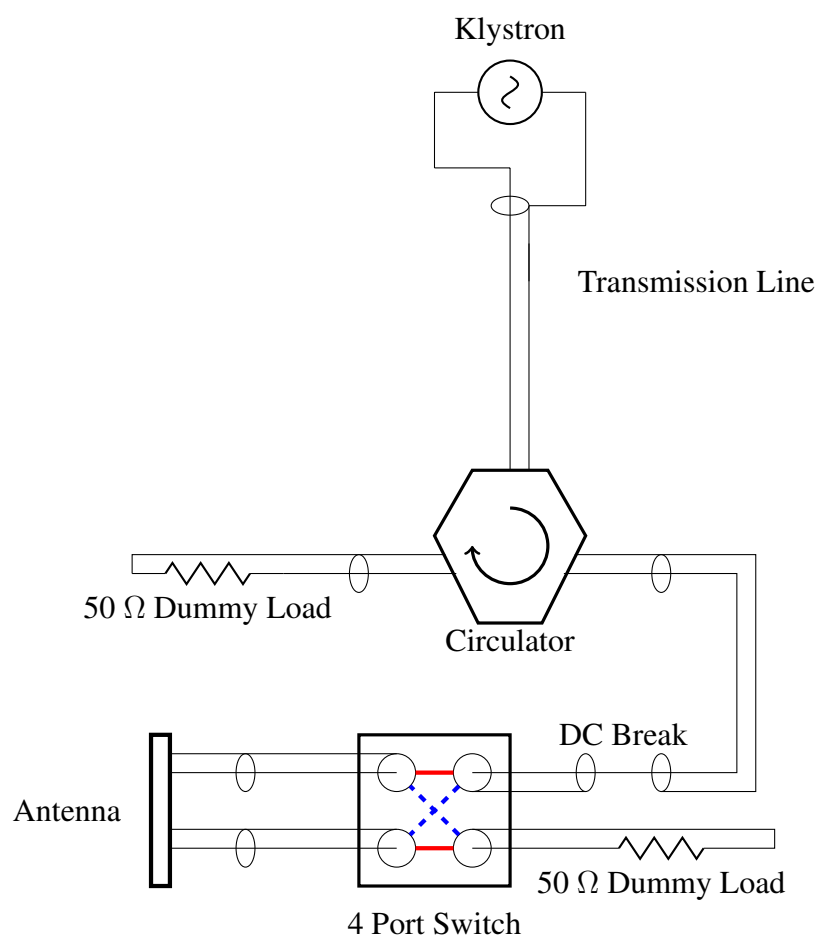


Figure 2.14: The Lower Hybrid transmission system

prevent a ground loop. This is followed by set of bidirectional couplers, power diagnostics which are capable of separately measuring the forward going and backward going waves, allowing us to measure both the input power to the antenna, and the power reflected from the antenna. Power then travels to a 4-port switch, allowing us to control the direction that power is flowing down the antenna face. Whatever fraction of power is not coupled to the plasma or reflected at this point is fed into a dummy load (after being measured). This is done to prevent a reflection, which would cause the antenna to launch a wave in the direction opposite the main launched wave, canceling out a portion of the driven current.

2.2.2 The Interdigital Traveling Wave Lower Hybrid Antenna

Power is coupled to the plasma using an interdigital traveling wave antenna. This antenna is similar in design to the stripline and combline antennas used on several other fusion devices [79, 80], but is somewhat novel in many ways. A closer analogue is the bandpass filter described by Cristal. [81] The structure of the slow wave antenna can be seen in Figure 2.15. It is composed of an array of cylindrical molybdenum rods (“digits”), each attached to opposite sides of the antenna structure (hence “interdigital”). There is a copper back plane which has been coated in silver, and copper side rails. A copper aperture is placed over the antenna face to provide a second ground plane, and has a molybdenum cover to protect it from damage. Each digit has an open side and a grounded side, and thus can support a wave of a fixed wavelength, just as a length of rope can sustain a standing wave when one side is fixed. Power is transmitted down the face of the antenna by capacitive and inductive coupling of neighboring rods, and of rods to the ground plate, which sets up a perturbation on the antenna face. The wavelength down the antenna’s length is fixed by the spacing of the digits. The cross face wavelength is fixed by the length of the digits. The large copper sections at either end are impedance matching elements, meant to ease the transition from the load in the transmission line to the load on the antenna. It is interesting to note that the antenna cannot support a DC voltage, as there is a direct electrical connection between the transmission line’s center conductor and outer conductor. A more detailed explanation of basic antenna operation can be found in Kaufman. [52]




Figure 2.15: Slow wave antenna structure, shown partially completed (a) and in exploded view (b).

2.2.2.1 Theoretical launch spectrum

The spectrum of waves launched from this type of antenna design is easily predicted. A circular grid representing a poloidal slice is constructed. The electric field is assumed to be zero all around the grid, except where the antenna aperture is. For the traveling wave design, as power moves down the face of the antenna, a fraction of it is coupled to the plasma. Thus the electric field magnitude is assumed to decay exponentially from one side of the aperture to the other in a plasma, and is assumed to not decay in a vacuum. In order to compare this scheme with, say the Brambilla waveguide [50] scheme mentioned in the last chapter, we must construct some meaningful statistics which describe the spectrum. Here I introduce Directivity, Unanimity, and Centrality. Each is an important efficiency parameter. A perfect LH system should have a reasonably narrow peak (not infinitely narrow) centered around a single n_{\parallel} that corresponds to the desired launch n_{\parallel} . The broadness of the peak (Centrality) can aid in damping as it distributes the resonance region, but if the peak is too broad then power is lost to less desirable waves. For the same reason, the relative magnitude of the primary peak to the secondary peaks (Unanimity) is also an important quantity. Finally and perhaps most importantly, any power launched counter to the specified direction (e.g. $n_{\parallel} > 0$ for experiments attempting to launch $n_{\parallel} < 0$) will drive current against the desired current, and therefore is specifically deleterious. Therefore, we must also be interested in the relative amount of power on the desired half of n_{\parallel} space (Directivity).

The theoretical curves for our antenna and four and ten septa Brambilla-type wave guides with $\pi/2$ phasing firing in a vacuum are plotted in Figure 2.16, and the various parameters are summarized in Table 2.2.

2.2.2.2 Spectral Measurements

The Mark III MST Lower hybrid antenna (the current incarnation, so named because it replaces two previous versions) is equipped with 20 pickup loops set into the antenna backplane. These loops are aligned to pickup the electric field in the down-face direction. The 800 MHz signal is brought from the antenna to phase electronics which downsample the data to a more manageable frequency. This is accomplished by the use of a set of rf mixers (one for each pickup loop), which




Figure 2.16: The spectra for various antenna types. The interdigital antenna is in black, the 10 element waveguide is in blue and the 4 element waveguide is in red. The waveguide is seen to have broader main peaks and more secondary peaks, leading to a reduced total power in the main peak.

Antenna	Directivity	Unanimity	Centrality	Secondary Peak $n_{ }$	Secondary Peak Power
Interdigital	98.4 %	90.3 %	0.924	8.79	2.3 %
Brambilla 4 $\pi/2$	76.3 %	71.4 %	3.2	-20.9	12.1%
Brambilla 10 $\pi/2$	79.5 %	71.0 %	1.33	-22.3	10.9%
Brambilla 10 $\pi/4$	80.3 %	78.4%	2.42	-47	5.08 %

Table 2.2: A Comparison of various waveguide setups to the interdigital antenna. $n_{||}$ was chosen to be near 7. The interdigital line antenna is seen to have better directivity, unanimity and centrality. Additionally, the secondary peaks are clustered near the primary peaks, which is preferred to the waveguide spectrum.

take two high frequency signals in and output the product of the two signals. One of the input signals is the pickup loop and the other is called the local oscillator, with a known power and frequency. In the case of the MST LH antenna, the local oscillator frequency is 800.1 MHz. The resulting signal (as shown in Figure 2.17) consists of two waves, one at a frequency of 1600.1 MHz and one at a frequency of 100 kHz. A high pass filter is then used to eliminate the high frequency component. The resulting signal is then processed in order to determine the relative power and phase at each loop. Relative phase is used because the phase of the local oscillator is not recorded (and phase is only defined up until an arbitrary constant), and relative power is used because the absolute pickup efficiency of the loops is unknown (though they are calibrated against one another).

The phase and power are used to reconstruct the relative electric field along the face of the antenna. The signal measured from the pickup loops is an AC voltage wave, so the power is proportional to the square of the magnitude of the field at the pickup face. Thus we assume that, if the relative phase and power measured are ϕ_l and P_l , $E_{lf} \propto \sqrt{P_l} e^{i\phi_l}$. The electric field is Fourier transformed so that the parallel launch spectrum can be seen, and the loop power value is fitted with an exponential curve to help determine the coupling efficiency, the parameter $\alpha = (1/\alpha_{ohm} + 1/\alpha_{coup})^{-1}$, as will be discussed in §2.2.2.3. Figure 2.18 shows the the vacuum spectrum measured from this technique compared with the theoretical vacuum spectra. The directivity, unanimity and centrality of the measured data are 95%, 88% and 0.989 respectively compared with the theoretical values of 98%, 90% and 0.925. The difference should not be considered significant.

The extra exponential decay present on the wave in a plasma is expected to distort the spectrum somewhat. This is calculated in terms of damping angle, but is converted into an arc length (λ_{damp}) for spectral calculations. Significant distortion does not occur in the model until the arc length of the damping parameter is comparable to the wavelength. Experimental measurements and the corresponding theoretical curves are shown in Figure 2.19

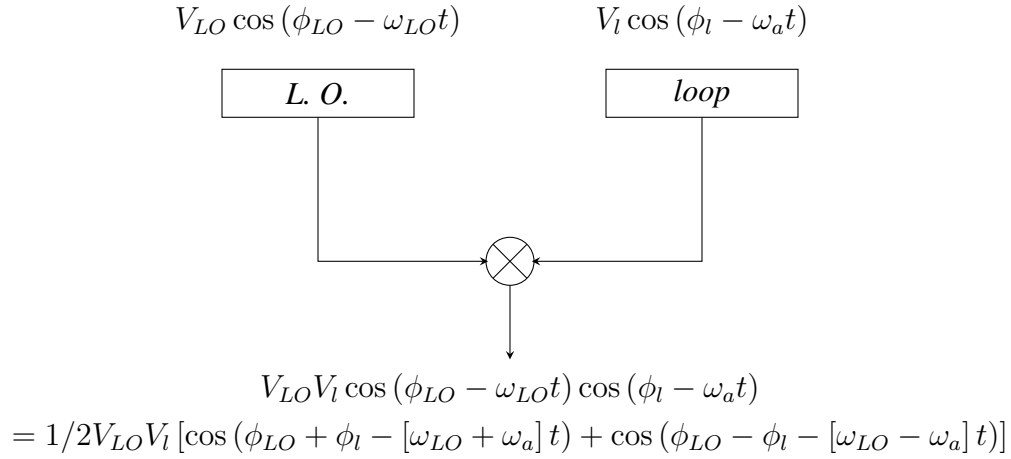


Figure 2.17: The LH Phase electronics take a local oscillator signal (L.O.) and multiply the signal with the 800 MHz signal from the pickup loop in the back of the antenna. The local oscillator is set to a known power and 800.1 MHz. The multiplicative rule of trigonometric functions shows that this creates a high frequency and low frequency wave which are used to calculate the relative phase and power at the loop.




Figure 2.18: The vacuum spectrum of the antenna. The theoretical curve is in black and the measurement is in red.




Figure 2.19: Experimental (a) and theoretical (b) antenna spectra. The black curve is for $\lambda_{damp} \approx 20$ cm, the blue curve is $\lambda_{damp} \approx 6$ cm, the green curve corresponds to $\lambda_{damp} \approx 4$ cm and the orange curve represents $\lambda_{damp} \approx 2$ cm. For comparison, the wavelength at $n_{||} = 7$ is 5.3 cm, and the aperture is 17.8 cm long.

2.2.2.3 Radiated Power Fraction

The important figure of merit for LH experiments is the power radiated from the face of the antenna. Radiated power for the traveling wave antenna is derived from the measurements of forward (F), reflected (R) and through (T) power. The method of doing so has been a matter of considerable debate due to uncertainty in representation of the Ohmic loss fraction. I will first present the previous model and then introduce a more sophisticated model which will be used for the remainder of this work.

The conservative (pessimistic) model assumes that the measured reflected power is reflected from the far side of the antenna. Thus we assume that the amount of power detected has been reduced by a factor of the Ohmic loss fraction as it propagated back through the antenna. Therefore, in order to discover the true amount of reflected power, we must divide it by the Ohmic loss fraction. Furthermore, we assume that all Ohmic losses occur instantaneously, before the radiative portion of the antenna. So the total expression for radiated power becomes (with β being the power radiation fraction)

$$P_{rad} = \beta F - R/\beta - T \quad (2.1)$$

It is further assumed that in a vacuum little or no power is emitted because our wave is electrostatic, so the Ohmic loss fraction can be measured by measuring the power characteristics in a vacuum, and solving equation 2.1 with $P_{rad} = 0$, i.e.

$$\beta = \frac{T + \sqrt{T^2 + 4FR}}{2F} \quad (2.2)$$

where we discard the negative branch since it is less than zero. This leads us to a radiated power fraction of about 35%. This was the approach used by Kaufman. [52]

A more sophisticated model depends on exploiting the antenna geometry. We can safely assume that the two most reflective parts of the antenna are the impedance matching elements at either end, because that is where the wave experiences the most dramatic propagation barrier. When a plasma is present, it seems reasonable to assume that only the impedance matching section at the high power end of the antenna is relevant because, assuming they both reflect roughly

the same power fraction, the power at the far end of the antenna is quite a bit lower due to wave damping. Thus we can safely discard it.

We also reason that the Ohmic power loss is something that we can directly estimate. In order to test this we use the lower hybrid test antenna, which is a replica of the antenna in the machine, except that it is longer and has type-N feedthroughs. First the power attenuation through the test antenna is measured, along with the reflection coefficient. In order to make sure the antenna isn't radiating, a copper cover is placed over the aperture. When this was done, new reflections were setup within the antenna, and the antenna reflection coefficient went up, causing the transmission to go down. The final test was to place a copper mesh in front of the antenna, as a second, hopefully less obtrusive method of preventing reflection. The results are shown in Table 2.3. The mesh aperture appears to reflect less power, but we can deduce by the missing power fraction that either complicated internal reflections are being setup, or that the mesh induces an electromagnetic component, acting as a secondary antenna. We see that the aperture free setup has the least lost power, which is interpreted to mean that the small radiation assumption is reasonable.

cover type	T/F	R/F	P_{loss}/F
none	59.26 %	14.64 %	26.10%
solid	34.71%	35.57%	29.28%
mesh	34.23%	12.22%	53.55%

Table 2.3: The power measurements for the three different aperture setups. P_{loss} is the power that is not accounted for by the two measurements (no radiation is assumed).

Using this result we attempt to build a better model. We assume the power decays exponentially down the face of the antenna due to Ohmic losses at a constant rate. In plasma conditions, this is assumed to be in addition to the losses from wave coupling. An initial reflection is also assumed as discussed above. Therefore the power at the beginning of the antenna structure is $F - R$, and the power at the end of the structure is T . The power along the structure (as a function of angle) is

given by

$$P(\theta) = P_0 \text{Exp}(-\theta [1/\alpha_{ohm} + (\Theta(\theta - \theta_{ap1}) - \Theta(\theta - \theta_{ap2})) / \alpha_{coup}]) \quad (2.3)$$

where α_{coup} is the power coupling angular damping length, α_{ohm} is the Ohmic loss angular damping length, and $\Theta(\theta - \theta_0)$ is the Heaviside function. It is also important to define the angular length of the antenna, $\alpha_{ant} = 37^\circ$, and the aperture $\alpha_{ap} = 20.7^\circ$, as well as the angular positions of the beginning and end of the aperture, $\theta_{ap1} = 7.7^\circ$ and $\theta_{ap2} = 28.3^\circ$. The power at the beginning of the aperture is then given by

$$P(\theta_{ap1}) = (F - R)e^{-\theta_{ap1}/\alpha_{ohm}} \equiv P_1 \quad (2.4)$$

at the end of the aperture by

$$P(\theta_{ap2}) = (F - R)e^{-(\theta_{ap2}/\alpha_{ohm} - \alpha_{ap}/\alpha_{coup})} \equiv P_2 \quad (2.5)$$

and the power that leaves the antenna is given by

$$T = (F - R)e^{-\alpha_{ant}/\alpha_{ohm} - \alpha_{ap}/\alpha_{coup}}. \quad (2.6)$$

Notice that for a vacuum shot $\alpha_{coup} \rightarrow \infty$, so

$$T = (F - R)e^{-\alpha_{ant}/\alpha_{ohm}} \quad (2.7)$$

$$\ln\left(\frac{F - R}{T}\right) = \alpha_{ant}/\alpha_{ohm} \quad (2.8)$$

$$\alpha_{ohm} = \alpha_{ant} / \ln\left(\frac{F_{vac} - R_{vac}}{T_{vac}}\right) \quad (2.9)$$

and for a plasma shot

$$\alpha_{coup} = \frac{\alpha_{ap}}{\ln\left(\frac{F-R}{T}\right) - \alpha_{ant}/\alpha_{ohm}} \quad (2.10)$$

To solve for radiated and Ohmic power, we deploy our damping assumption

$$\frac{dP_{ohm}}{d\theta} = \frac{P}{\alpha_{ohm}} \quad (2.11)$$

and

$$\frac{dP_{rad}}{d\theta} = \frac{P}{\alpha_{damp}} (\Theta(\theta - \theta_{ap1}) - \Theta(\theta - \theta_{ap2})) \quad (2.12)$$

Therefore

$$P_{ohm} = \int_0^{\theta_{ap1}} \frac{P}{\alpha_{ohm}} d\theta + \int_{\theta_{ap1}}^{\theta_{ap2}} \frac{P}{\alpha_{ohm}} d\theta + \int_{\theta_{ap2}}^{\theta_{ant}} \frac{P}{\alpha_{ohm}} d\theta \quad (2.13)$$

$$\begin{aligned} P_{ohm} &= (F - R)(1 - e^{-\theta_{ap1}/\alpha_{ohm}}) \\ &\quad + P_1 \frac{\alpha_{coup}}{\alpha_{coup} + \alpha_{ohm}} (1 - \text{Exp}[-\alpha_{ap}(1/\alpha_{ohm} + 1/\alpha_{coup})]) \\ &\quad + P_2 (1 - e^{-(\theta_{ant} - \theta_{ap2})/\alpha_{ohm}}) \end{aligned} \quad (2.14)$$

For the radiated power it is only necessary to consider the middle integral since power cannot be radiated where there is no aperture, so

$$P_{rad} = P_1 \frac{\alpha_{ohm}}{\alpha_{ohm} + \alpha_{coup}} (1 - \text{Exp}[-\alpha_{ap}(1/\alpha_{ohm} + 1/\alpha_{coup})]) \quad (2.15)$$

Using a high power vacuum shot (1081003017) and Equation 2.9, α_{ohm} is calculated to be 27.8 ± 0.2 degrees. It is possible to determine the angular power loss parameter $\alpha = (1/\alpha_{coup} + 1/\alpha_{ohm})^{-1}$ independently using the power/phase diagnostics in the back of the antenna. The results are presented in Figure 2.20. The agreement between the technique just outlined (the 0-D technique) and the direct power measurements far exceeds expectations. This indicates that the loop diagnostics (1-D) are not necessary for radiation measurements and could be left off of future antennas in favor of robustness, although they provide a lower uncertainty for the radiated power, and are also important for measuring the launch spectrum. This analysis leads us to two important figures of merit for antenna coupling experiments, namely the coupling efficiency and the damping parameter. The coupling efficiency, defined as $\eta = P_{rad}/(F - P_{ohm})$ represents how much of the available power has been launched into the plasma. Ohmic losses are excluded to make interpretation simpler and more intuitive, however before dismissing them it should be noted that they typically represent 30-40% of the applied power. While important for interpreting data for a given experimental setup, it should also be noted that if the coupling efficiency is unacceptably low for a given plasma setup, a longer antenna can be used. Therefore, we also consider the damping parameter (simply α_{coup} as defined above) because it is applicable to general antenna design considerations.

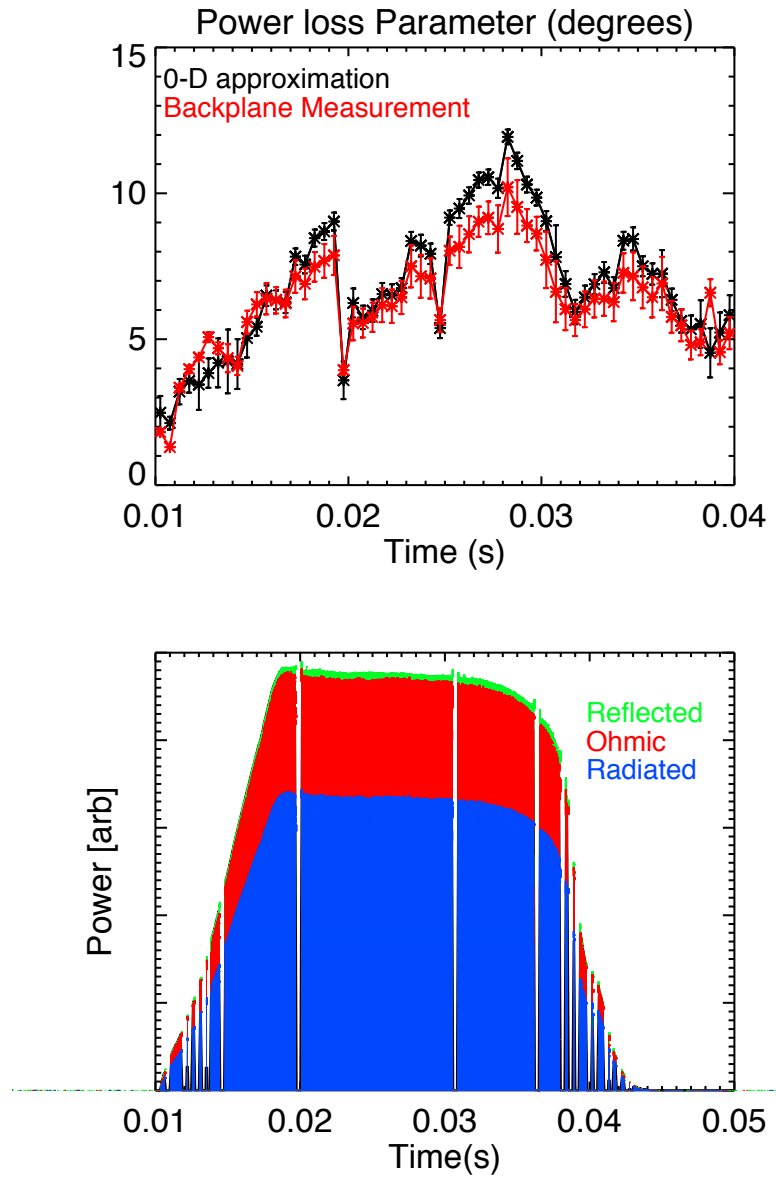


Figure 2.20: Shown in figure (a) is the damping parameter measured with the radiated power estimation (in black) and the antenna backplane pickup loops (in red), with great agreement. The second plot shows the relative proportion of power for this particular shot. The blue area represents the radiated power, the red area represents the Ohmic losses and the green area represents the reflected power.

Chapter 3

Kinetic Simulations and Radial Transport in the RFP discharges

Before examining the effects of lower hybrid excitation on RFP plasmas, it is important to explore some key concepts about transport in RFPs. These will form the basis for studies involving lower hybrid. The results discussed here will be revisited in the next chapter, because lower hybrid excitation can be exploited to help resolve some of the questions posed below. RFP kinetics may be investigated with the Fokker-Planck code CQL3D. Although interesting in its own right, the rf-free distribution functions are vital to interpreting experimental data during rf operations. Moreover, transport plays a vital role in the efficiency of fast electron based current drive techniques (such as rf current drive). This will be discussed in more detail in the next chapter.

3.1 CQL3D

CQL3D is a code which models electron distributions in plasmas. In order to do this, a simplified version of the Fokker-Planck equation (itself a simplification of the Boltzmann equation) is employed. Three Maxwellian species are modeled (ions, electrons and impurities) along with one species of electrons that is allowed to evolve. CQL3D is capable of modeling a non-Maxwellian evolving ion species, but that was not done in this work in favor of numerical stability. The Fokker-Planck equation is averaged over an electron gyroperiod, allowing us to couple the two perpendicular directions, reducing the order of the equation in velocity space from three to two. The equations are then again averaged over a bounce period, assuming zero banana width (however particle trapping effects are retained). Finally the distribution is assumed to be axisymmetric and so can be averaged over a flux surface. This results in a 3 dimensional quasilinear model with

the deceptively simple form

$$v_{\parallel} \frac{\partial f_0}{\partial l_B} = 0 \quad (3.1)$$

$$\frac{\partial f_0}{\partial t} + v_{\parallel} \frac{\partial f_1}{\partial l_B} = -\frac{qE_{\parallel}}{m_e} \frac{\partial f_0}{\partial v_{\parallel}} + C(f) + Q(f) + R(f) \quad (3.2)$$

where l_B is the bounce path, E_{\parallel} is the inductive electric field, and C , Q , and R are the major quasilinear operators acting on the distribution f . f is ordered by the ratio of the bounce time (τ_b) over the collision time (τ_{coll}) such that $\mathcal{O}(f_i) = (\tau_b/\tau_{coll})^i$. Note that the velocity in this equation (v_{\parallel}) is meant to represent the guiding center motion of electrons, as the convective term arising from drift velocity has already been used by *ansatz* in the radial diffusion operator (covered in the next section). C represents the Coulomb collision term, Q is the rf quasilinear operator, and R is the radial diffusion operator. Here f is a function of ψ , a normalized radial-like coordinate, v_{\parallel} and v_{\perp} and implicitly of time. Specifically, CQL3D uses the coordinates ψ , v (speed) and θ_v (pitch angle). A sample speed - pitch angle grid is shown in Figure 3.1. The type of coordinate used for ψ can be selected, and for these experiments the square root of the area of the poloidal cross section of a given flux surface is used. Notice that if Shafranov shift and toroidicity are ignored then this reduces to the normalized radius, r/a . CQL3D assumes up-down symmetry (which is reasonable for MST because of its circular cross section) and axisymmetry. Magnetic field information is determined by reading in a Grad-Shafranov equilibrium (generated by MSTfit). Density and temperature are separately input, as is the effective ion charge, Z_{eff} , given by the formula

$$Z_{eff} = \frac{\sum n_i Z_i^2}{\sum n_i Z_i}, \quad (3.3)$$

where the summation is over all ion species. Although this can be given any sort of profile through a splining algorithm, it is held at a constant value across the plasma because it is not a well known parameter, and many RFP profiles are observed to be fairly flat. CQL3D uses Z_{eff} to determine the concentration of the impurity species. The inductive electric field can either be specified or determined from the MSTFit equilibrium current profile (as discussed below). The electron distribution function is divided into a Maxwellian component which is held fixed, and a general component which is discretized onto the solution grid, and allowed to evolve. The total grid then has some




Figure 3.1: A distribution function from CQL3D with the grid overplotted. The grid is normalized (here it is normalized to the speed of a 150 keV electron). Dreicer runaways can be seen on the right side of the plot, caused by reduced collisionality and the electric field. Also, the trapped-passing boundary can be seen (at approximately 60 and 120 degrees). Notice the pitch angle (θ) grid has extra packing near this area.

radial-like component (which identifies the surface), and a speed and pitch angle of the general distribution function at the outboard midplane on each surface. CQL3D then evolves Equation 3.2 from a Maxwellian until steady state is achieved. Speed is normalized to a specified energy which defines the highest speed tracked by the grid. It is also important to note that CQL3D is a relativistic code, and so rather than traditional Newtonian speed, the momentum over mass is used (this reduces to speed in the non-relativistic limit). For the sake of brevity this will be called speed in this work.

3.1.1 Radial Diffusion

As mentioned above, the convective term connected to the drift velocity in the Boltzmann equation is used by *ansatz* to create the radial diffusion operator for CQL3D. This operator takes a

Fick's law-like form, specifically

$$R(f) = \frac{1}{r} \frac{\partial}{\partial r} r \left[D \frac{\partial}{\partial r} - V_r \right] f \quad (3.4)$$

where $D = D(r, v_{\parallel}, v_{\perp})$ is called the diffusion coefficient, and V_r is a model pinch velocity, which acts to balance the particle losses due to diffusion. This *ansatz* has a long history of usage in kinetic studies. [56, 57, 82–84] It has also been used as part of CQL3D successfully to model cyclotron absorption [85, 86], and runaway current in “killer pellet” experiments in DIII-D. [87] The pinch term is an ad-hoc addition, but it is reasonable to include it because the electron density has been determined by a measurement (and reconstruction) and so (neglecting model beam/pellet injection) it is reasonable to mandate $\partial n / \partial t = 0$. This term also naturally accounts for a radial electric field. Comparing the radial diffusion operator given in 3.4 to our previous derivation of the non-local approximation of Ohm's law (specifically equation 1.63, we postulate another *ansatz*, namely that this diffusion operator can be made to model stochastic field line wandering by replacing D with $|v_{\parallel}| D_m$. Recall this earlier derivation required the assumption that the collision length was much longer than the stochastic correlation length. This is a reasonable assumption for CQL3D to make because it has already implicitly assumed a long collision length when the Fokker-Planck equation was ordered, using the ratio of bounce time to collision time. The final general form of D used in the code is

$$D = D_r D_{rr}(\psi) D_{vv}(v) \quad (3.5)$$

$$D_{rr}(\psi) = (c_1 + c_2 \psi^{c_3})^{c_4} \left(\frac{n_e}{n_{e0}} \right)^{c_5} \left(\frac{T_e}{T_{e0}} \right)^{c_6} \left(\frac{Z_{eff}}{Z_{eff0}} \right)^{c_7} \quad (3.6)$$

$$D_{vv}(v_{\parallel}, v_{\perp}) = \left| \frac{v_{\parallel}}{v_{the0}} \right|^{b_1} (1 + l_{ac}/\lambda_{mfp})^{-b_2} \left(\frac{v_{\perp}}{v_{the0}} \right)^{b_3} \gamma^{b_4} \quad (3.7)$$

The scaling for D_{rr} is given enough free parameters that most any shape can be made. D_{vv} is constructed so that the first two terms (each with a coefficient of 1) give momentum transport as predicted by Rechester and Rosenbluth for a stochastic plasma, [53] and as discussed in the introduction The third term is used to test the capabilities of the code, and the fourth term gives

relativistic retardation of diffusion. Neither the third nor fourth term are relevant to the work at hand. Notice that for stochastic diffusion, $D_r = v_{the0}D_m$.

It should be noted that though I have separated D into velocity and spatial components, this is not a restriction of the code, as D is calculated on the entire domain grid (ψ, v, θ_v) , and in fact D_{vv} has a radial dependence which enters through the thermal velocity, autocorrelation length and mean free path. For non-stochastic simulations, the coefficients used for the radial operator are $c_1 = 1$, $c_2 = 3$, $c_3 = 3$, $c_4 = 1$, $c_5 = -1$, and $c_6 = c_7 = 0$, which are found to be consistent with many experiments, as discussed in O’Connell. [60] In simulations of plasmas with stochastic transport, these are set to zero due to the ambiguity in the true value (though other profiles have been tried). The velocity coefficients ($b_1 - b_4$) are all set to zero, except in stochastic simulations where $b_1 = b_2 = 1.0$, consistent with Rechester-Rosenbluth transport (this is what defines them to be stochastic from CQL3D’s point of view).

3.1.2 Electric Field

CQL3D can be run in one of two modes, constant field mode or constant current mode, which determine how the electric field is calculated. In constant field mode (shown schematically in Figure 3.2), the electric field is specified in the input file and not allowed to evolve. It is then simply a matter of letting the surfaces come to equilibrium with the field, any collisions, and the radial diffusion operator. The magnetic profiles are not adjusted based on this current profile, so the dynamics are not necessarily accurate. In constant current mode (shown schematically in Figure 3.3), the field is initialized using the current profile from the equilibrium reconstruction multiplied by Spitzer resistivity. [49] The distribution function is then advanced one time step. Parallel current is calculated by evaluating the integral

$$J_{\parallel} = -e \int_{-\infty}^{\infty} v_{\parallel} f d^3v$$

This is then compared to the current from the equilibrium reconstruction. If there is a mismatch, CQL3D adjusts the electric field and reruns the time step. This mode takes longer to run but is occasionally preferred because the parallel electric field profile is not well known in MST. Because

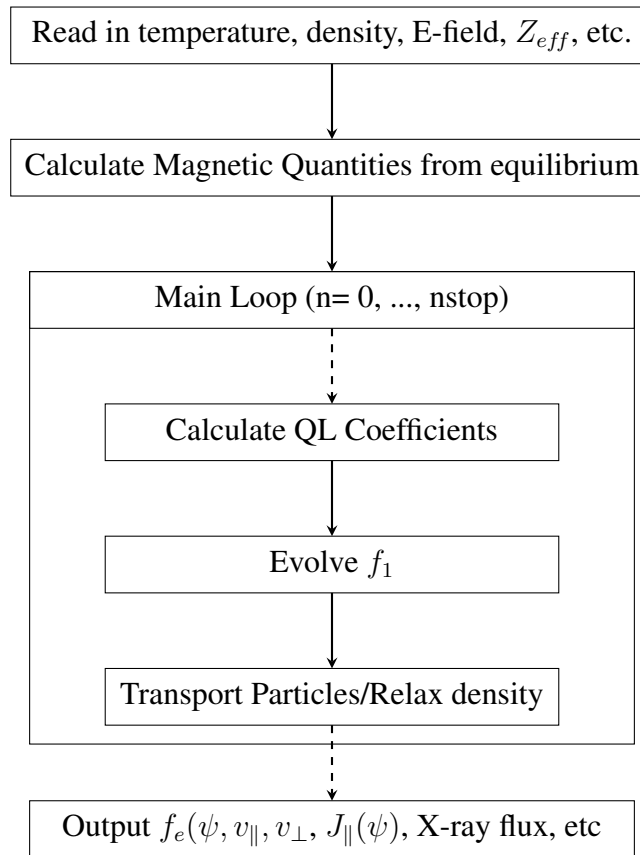


Figure 3.2: Flow chart for CQL3D in constant electric field mode.

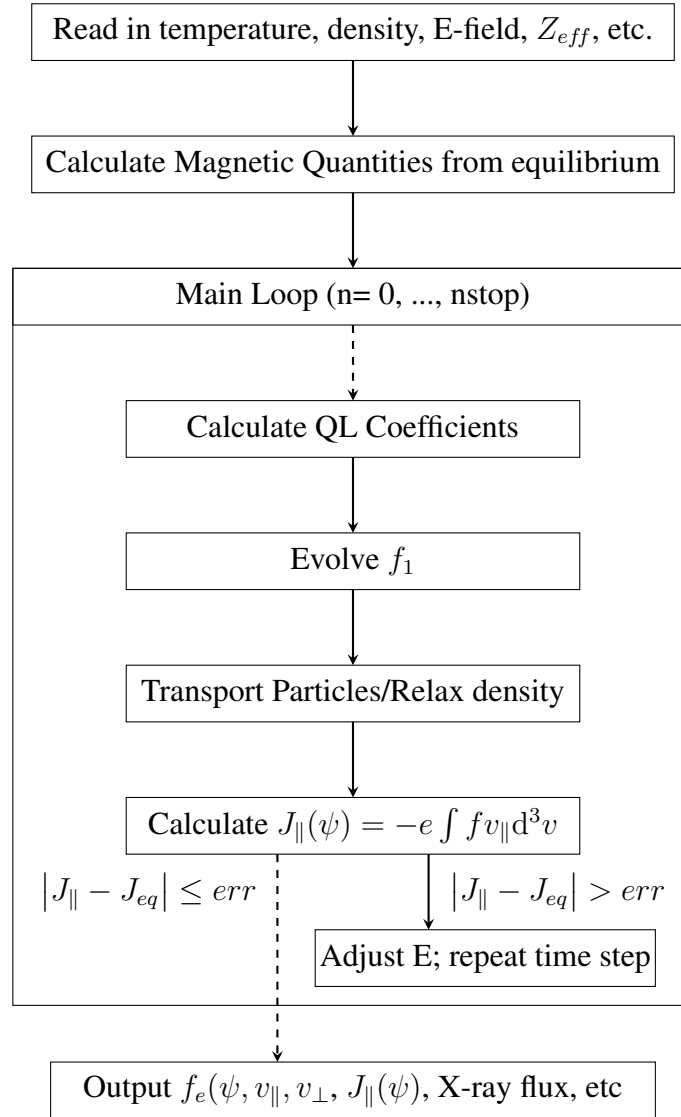


Figure 3.3: Flow chart for CQL3D in constant current mode.

the electric field is adjusted to fit a known profile however, one must be very careful not to ascribe too much of the effects of simulation to the underlying physics. For example, two simulations with different Z_{eff} will settle to two different values of electric field in order to produce the same current profile. Therefore differences between them may be more closely related to the electric field difference than Z_{eff} itself. It should also be noted that the dynamics of this method are somewhat dubious. A serious concern with the model is that the electric field is being evolved without self-consistently evolving Faraday's law and taking induction into account. This means that large electric fields can develop quickly, or that quasi-static states can be reached which are not physical, where the electric field oscillates up and down near the runaway generation threshold. It can be hoped however that if the plasma is in some reasonably steady state that this effect will be minor. This can be tested to some extent by seeding a constant electric field simulation with the final field from the constant current simulation and seeing if the results match, which is observed to be true for PPCD plasmas (see Figure 3.4). In order to account for this completely accurately, the kinetic code would have to be coupled to a nonlinear MHD code. A similar argument also applies to the current density profile in the constant field case, but this is found not to be a serious limit unless there is an exceptionally large runaway population.

It must be stressed that CQL3D does not impose Ohm's law, nor does it employ it in any way except implicitly in creating an initial guess for constant current mode. Ohm's law is an integrated version of the Fokker-Planck equation. In directly solving the Fokker-Planck equation, CQL3D attempts to accomplish the same task as Ohm's law, but it does not directly use it.

3.2 Fast Electrons in RFP discharges

An important component of MST's current is provided by Dreicer runaway electrons. [88] This process is described rigorously in any good basic plasma physics textbook (see Refs. [47, 89]), but the key conceptual elements are fairly intuitive. All electrons on a given flux surface are accelerated equally (regardless of speed) by an applied electric field. The primary source of drag for plasma particles is Coulomb scattering off of other charged particles. Faster particles are less likely to experience significant Coulomb collisions, because they spend less time near each particle




Figure 3.4: This figure shows the parallel current that arises from simulation when the fitted electric field profile from constant current mode is fed into constant electric field mode. Better agreement can be found by slightly adjusting the edge field, but it is clear that agreement is fairly good. This is not true in standard plasmas.

in the plasma. Specifically, the Coulomb collision cross section runs (approximately) like $v^{-1/3}$. Therefore, fast electrons will be accelerated by the electric field, but will experience a smaller drag force. For electrons moving above a certain critical speed, the force of the electric field overpowers the drag force, and they continue to accelerate. These electrons will only be arrested when their population is large enough to increase the drag force, when relativistic effects come into play, or when they diffuse out of the plasma into the vacuum vessel. The amount of runaways created by a given electric field strength is usually characterized by the size of the field relative to a value called the Dreicer field, E_D . If the field strength is above a certain critical fraction of the Dreicer field, a significant population of fast electrons will runaway in this manner. Since they are moving very fast, a fairly small population can carry a great deal of current. These electrons are particularly important for rf experiments because (as mentioned in Chapter 1) the electron conveyor aspect of current drive will pump electrons into this runaway regime. Roughly speaking, the Dreicer field scales with the number of slow-moving particles in a plasma, because they make up the largest

contribution to collisionality. That is to say a dense, cold plasma will require a higher electric field to create a large runaway population than a hotter, more diffuse plasma. CQL3D provides a good model for exploring the effect of this population in detail. It is important to note however, that as the runaway current gets larger, it is no longer sufficient to consider kinetic evolution without also considering the effects of induction. Figure 3.5 shows CQL3D’s predictions for a standard




Figure 3.5: The electron distribution function for various electric field strengths, versus velocity normalized to the thermal velocity. As the field gets closer and closer to the Dreicer field the runaway population increases dramatically.

MST discharge, as the background electric field is increased. The runaway population gets larger and larger until the current carried by it becomes comparable to the background current (at which point the simulations become more difficult to converge). This is further clarified by Figure 3.6. Recall that CQL3D does not impose Ohm’s law. It can be seen however that an Ohm’s law-like relationship exists for $E/E_D \lesssim 0.09$. The blue line on the figure assumes an Ohm’s law with the resistivity calculated from a simulation with a very low electric field ($E = 1 \times 10^{-8} \text{ V/m}$). The red line represents Spitzer resistivity with a Z_{eff} correction. The green shaded region can be described as a “runaway” current, reflecting the anomalous current amplification caused by a large




Figure 3.6: The impact of a large runaway population on the total current density. The region where $E/E_D \lesssim 0.09$ where Ohm’s law (approximately) holds is described as the “Ohmic” regime. The region with large runaway currents is described as the “runaway” regime. The shaded area is called the “runaway” current.

runaway population. Standard MST discharges have a normalized electric field of about 0.25, but when stochastic diffusion is included, the runaway population is significantly diminished.

3.3 Methods of Experimental verification of Fokker-Planck predictions

Beyond direct comparison with the current density profile from the equilibrium reconstruction (which is moot in constant current mode), the predictions of CQL3D are best tested by x-ray emission. Whereas current density reflects electron motion in some average sense, x-ray spectroscopy is directly related to the balance of thermal and fast electrons within the plasma. As shall be seen, this will challenge some of the long standing interpretations of various plasma phenomena. Furthermore, x-ray spectra have a long history as a diagnostic for LH experiments. X-ray emission was mentioned in the first work to demonstrate LH current drive on a tokamak. [19] Not

long after these first experiments, x-ray emission was used as a constraint to predict the distribution function on this same device. [90] LH associated x-ray flux was also used to confirm current drive on JFT-2M and HT-7. [25, 91] In this work, the distribution functions generated by CQL3D are used to predict x-ray emission by applying the Born-Sommerfeld approximation [32] of the bremsstrahlung cross section and integrating over the distribution function on each surface along a viewing chord. This technique also has a history on MST, however, it has only at present been applied to non-rf plasmas, and has never used real detector locations. [60, 92] Beyond previous works, a method of interpreting Fokker-Planck simulation results for stimulated bremsstrahlung emission is also outlined. Before moving on to data interpretation, a brief discussion of the model is presented.

3.3.1 Model X-ray flux predictions

A discussion of model x-ray flux can get quite complicated and bogged down in details, and so what follows should be regarded as somewhat simplified. The aim of this Section is to provide a qualitative interpretation of these data so some of the details are suppressed. The probability of x-ray emission at a given energy arising from collisions between an electron and another charged particle can be determined by the emission cross-section (that is, the likelihood of emission as a function of energy). This prescription can be generalized to a distribution of electrons and ions by integrating over the respective distribution functions, i.e.

$$\phi = \int \int f_e f_i \left(\frac{d\sigma}{df_i} \right)_i + \int \int f_e f_e \left(\frac{d\sigma}{df_e} \right)_e \quad (3.8)$$

Where df_α is meant to imply integration over all of the velocity dimensions of a given distribution function. Further splitting the electron distribution function into a general distribution and a Maxwellian distribution,

$$\phi = \int \int f_{e0} f_i \left(\frac{d\sigma}{df_i} \right)_i + \int \int (f_{e0} + f_{e1}) (f_{e0} + f_{e1}) \left(\frac{d\sigma}{df_e} \right)_e + \int \int f_{e1} f_i \left(\frac{d\sigma}{df_i} \right)_i \quad (3.9)$$

CQL3D directly calculates this function, but for the purposes of interpretation, a little manipulation is in order. The first integral is a constant during the evolution of the simulation because the

Maxwellian distributions are fixed by measurements. For the middle integral, the $f_{e0}f_{e0}$ term is constant. Taking only the time varying components, it is seen that

$$\phi(t) - \phi(0) = 2 \int \int f_{e0}f_{e1} \left(\frac{d\sigma}{df_e} \right)_e + \int \int (f_{e1}f_{e1}) \left(\frac{d\sigma}{df_e} \right)_e + \int \int f_{e1}f_i \left(\frac{d\sigma}{df_i} \right)_i \quad (3.10)$$

Finally, the perturbed distribution function is split into a perturbation due to rf excitation (f_{rf} and a perturbation which exists before the rf is applied (f_{e11}). Subtracting the no rf flux from the rf flux we are left with

$$\begin{aligned} \Delta\phi_{rf} = & 2 \int \int f_{e0}f_{rf} \left(\frac{d\sigma}{df_0} \right)_e + 2 \int \int f_{e11}f_{rf} \left(\frac{d\sigma}{df_e} \right)_e \\ & + 2 \int \int f_{rf}f_{rf} \left(\frac{d\sigma}{df_e} \right)_e + \int \int f_{rf}f_i \left(\frac{d\sigma}{df_i} \right)_i \end{aligned} \quad (3.11)$$

The important point here is that when the background spectrum is subtracted from the rf spectrum, there is only one term that relates to the initial perturbed electron density. The largest emission is expected to come from the ion collision term - due to the large mass of ions, they are not perturbed much and they do perturb the electrons significantly. Koch and Motz [32] predict a much smaller cross-section for electron-electron x-ray emission. Due to the dependence of collisionality on velocity, the largest electron-electron term is the the first one, the scattering of rf generated electrons off of the Maxwellian distribution. Therefore, it is anticipated that the direct impact of the non-rf perturbed (NRP) distribution on x-ray emission will be quite small. This justifies examining the flux difference between rf on and rf off shots even if the NRP distribution function is not well known (so long as the temperature and density *are* well known). Note, this is not to say that the background distribution does not effect rf absorption. There are still quasilinear effects by which the efficiency of rf absorption is reduced due to the distorted background distribution, though this is not a significant effect at low power levels.

3.3.2 Data Interpretation

Before presenting results, the method of analysis is presented. As alluded to in the introduction, there is a finer point often overlooked in this type of analysis, that is the statistical bias in traditional data interpretation. It is an important first step to clarify what is actually being measured by

counting individual x-rays. We seek to estimate either the intensity ϕ , or the flux θ as a function of energy. We do this by dividing the energy spectrum into a number of bins and calculating the mean of the quantity of interest in each bin. If a comparison is to be undertaken, then it is equally important to quantify the statistical and experimental uncertainty in each number.

3.3.2.1 Intensity Method

The intensity method begins by considering the total power absorbed by the detector. Cast in terms of an emission probability this is

$$P = \iint \frac{d^2\lambda}{dGdE} E \, dE \, dG, \quad (3.12)$$

where P is the total power, $\lambda = \lambda(E, t, G)$ is the Poisson emission rate, E is energy and G is the etendue of the optical system. Inverting these integrals we get

$$\phi \equiv \frac{d^2P}{dGdE} = E \frac{d^2\lambda}{dGdE} = E \frac{d^3n}{dGdEdt} \quad (3.13)$$

Due to the discrete nature of x-ray emission, ϕ cannot be directly measured. Instead, we can only estimate the mean of ϕ over an energy and time bin. Therefore, the quantity we are trying to measure is given by the expression

$$\phi_{mean}(E, t) = \frac{1}{\Delta E \Delta t} \int_{t_0}^{t_1} \int_{E_0}^{E_1} E \frac{d^3n}{dGdEdt} dE dt \quad (3.14)$$

We simplify this expression by applying the pencil beam assumption, and assuming that λ is not a function of time (or at least not strongly so). This simplifies our expression to

$$\phi_{mean}(E, t) = \frac{1}{\Delta G \Delta E} \int_{E_0}^{E_1} E \frac{d\lambda}{dE} dE \quad (3.15)$$

Now we construct a statistical estimator for ϕ , $\hat{\phi}$. The estimator commonly used is defined as

$$\hat{\phi} = \frac{\sum E_i}{\Delta G \Delta E \Delta t} \quad (3.16)$$

where the summation is over all photons in a bin.

Before accepting this estimator, it must be checked for statistical bias. Bias is defined as

$$\text{Bias} \left[\hat{\phi} \right] = \text{E} \left[\hat{\phi} \right] - \phi_{mean} \quad (3.17)$$

$$= \frac{\text{E} [n] \text{E} [E_0]}{\Delta G \Delta t \Delta E} - \frac{1}{\Delta G \Delta E} \int_{E_0}^{E_1} E \frac{d\lambda}{dE} dE \quad (3.18)$$

$$(3.19)$$

where (assuming emission to be a Poisson process)

$$\text{E} [n] = (\text{rate}) \times (\text{time}) \quad (3.20)$$

$$= \Delta t \int_{E_0}^{E_1} \frac{d\lambda}{dE} dE \quad (3.21)$$

and

$$\text{E} [E_0] = \int_{E_0}^{E_1} E dE \equiv \mu_E \quad (3.22)$$

So our total bias equation becomes

$$\text{Bias} \left[\hat{\phi} \right] = \frac{\mu_E}{\Delta G \Delta E} \int_{E_0}^{E_1} \frac{d\lambda}{dE} dE - \frac{1}{\Delta G \Delta E} \int_{E_0}^{E_1} E \frac{d\lambda}{dE} dE \quad (3.23)$$

$$= \frac{1}{\Delta G \Delta E} \left(\mu_E \int_{E_0}^{E_1} \frac{d\lambda}{dE} dE - \int_{E_0}^{E_1} E \frac{d\lambda}{dE} dE \right) \quad (3.24)$$

$$= \frac{1}{\Delta G \Delta E} \left(\mu_E (\lambda_1 - \lambda_0) - \left[E_1 \lambda_1 - E_0 \lambda_0 - \int_{E_0}^{E_1} \lambda dE \right] \right) \quad (3.25)$$

So, in order for this to be a bias free indicator, $\lambda(E)$ must be constant over our energy bin, which it certainly is not (the form expected is an exponential decay). Nevertheless, we can calculate the variance in order to estimate the uncertainty, making use of the law of total variance.

$$\text{Var} \left[\sum_n E \right] = \text{E} \left[\text{Var} \left[\sum E | n \right] \right] + \text{Var} \left[\text{E} \left[\sum E | n \right] \right] \quad (3.26)$$

$$= \text{E} \left[n s_E^2 \right] + \text{Var} \left[n \bar{E} \right] \quad (3.27)$$

$$= n s_E^2 + n \bar{E}^2 \quad (3.28)$$

$$= n \left(s_E^2 + \bar{E}^2 \right) \quad (3.29)$$

where \bar{E} is the mean of the detected photon energies and s_E is the standard deviation of the incoming photon energies. so we can say the error in ϕ is

$$\delta_\phi = \phi \sqrt{\left(\frac{\delta_G}{G}\right)^2 + \frac{\text{Var}[\sum_n E]}{(\sum_n E)^2}} \quad (3.30)$$

$$= \phi \sqrt{\left(\frac{\delta_G}{G}\right)^2 + \frac{n(s_E^2 + \bar{E}^2)}{(\sum_n E)^2}} \quad (3.31)$$

3.3.2.2 Flux method

In order to construct a bias free statistic, a different formulation must be considered. This can be accomplished by using the counting flux instead of the energy-biased intensity measurement. Start by examining the total number of x-rays

$$N = \iint \frac{d^2\lambda}{dGdE} dE dG \quad (3.32)$$

Inverting these integrals we can define θ ,

$$\theta = \frac{d^2\lambda}{dGdE} = \frac{d^3n}{dGdEdt} \quad (3.33)$$

It is helpful to note here that $\theta = \phi/E$.

Following the derivation above, we calculate θ_{mean} over some range of time and energy

$$\theta_{mean}(E, t) = \frac{1}{\Delta E \Delta t} \int_{t_0}^{t_1} \int_{E_0}^{E_1} \frac{d^3n}{dGdEdt} dE dt \quad (3.34)$$

$$= \frac{1}{\Delta E \Delta t} \int_{t_0}^{t_1} \int_{E_0}^{E_1} \frac{d^3\lambda}{dGdEdt} dE dt \quad (3.35)$$

$$\theta_{mean} = \frac{1}{\Delta G} \left\langle \frac{d\lambda}{dE} \right\rangle_E \quad (3.36)$$

where

$$\langle x \rangle_E = \frac{1}{\Delta E} \int_{E_0}^{E_1} x dE$$

We declare our estimator to be

$$\hat{\theta} = \frac{\sum 1}{\Delta G \Delta E \Delta t} \quad (3.37)$$

so we can calculate the bias as

$$\text{Bias} [\hat{\theta}] = E [\hat{\theta}] - \theta_{mean} = \frac{E [n]}{\Delta G \Delta E \Delta t} - \frac{1}{\Delta G} \left\langle \frac{d\lambda}{dE} \right\rangle_E \quad (3.38)$$

$$= \frac{1}{\Delta G \Delta E \Delta t} \int_{E_0}^{E_1} \frac{d\lambda}{dE} \Delta t dE - \frac{1}{\Delta G} \left\langle \frac{d\lambda}{dE} \right\rangle_E \quad (3.39)$$

$$= \frac{1}{\Delta G} \left\langle \frac{d\lambda}{dE} \right\rangle_E - \frac{1}{\Delta G} \left\langle \frac{d\lambda}{dE} \right\rangle_E = 0, \quad (3.40)$$

which is to say that number flux is unbiased regardless of the distribution of λ . Next the error must be calculated, which is easily accomplished using the well known Poisson counting error (\sqrt{n}).

$$\delta_\theta = \theta \sqrt{\left(\frac{\delta_G}{G} \right)^2 + \frac{\text{Var} [n]}{n^2}} \quad (3.41)$$

$$= \theta \sqrt{\left(\frac{\delta_G}{G} \right)^2 + \frac{1}{n}} \quad (3.42)$$

3.3.3 Stimulated Bremsstrahlung Emission

Fast electrons can also be observed by stimulating bremsstrahlung with a target inserted into the plasma. A probe was developed consisting essentially of a tube with a hole in it pointed at a target. The target (tungsten in these experiments) is struck by an electron beam coming from the plasma. This is shown schematically in Figure 3.7. Emission can be predicted using the theory of Kramers. Specifically, the intensity from a thick target struck by an electron of energy E_e is given by [93]

$$I(E_e, E_\gamma) = CZ (E_e - E_\gamma)^\alpha, \quad (3.43)$$

or in terms of number flux

$$N(E_e, E_\gamma) = \frac{CZ}{E_\gamma} (E_e - E_\gamma)^\alpha, \quad (3.44)$$

where E_e is the energy of the incident electron, E_γ is the energy of the emitted photon, Z is the atomic number of the target atom, and C and α are $1.68 \times 10^{-6} \text{keV}^{-1}$ and $\alpha = 1$ in theory, but are in practice found to be material dependent constants.




Figure 3.7: A schematic of the x-ray target probe. The tungsten target is in green. Electrons enter the probe through the hole and strike the tungsten target, which is grounded to prevent charge accumulation. Ions are prevented from entering the probe because the entrance hole is designed to be small enough that their cyclotron motion will cause them to be absorbed in the boron nitride sheath.

The emission then is simply the emission probability given above, integrated over the incoming electron flux. Assuming no electrons are lost in the probe and that the hole is directed along the magnetic field, the flux on the target is equal to the flux at the entrance hole. This is given by

$$\frac{dN_{el}}{dv^3 d^3x} = f_e \mathbf{v} \cdot \mathbf{a} = av_{\parallel} f_e, \quad (3.45)$$

where a is the area of the entrance hole. It should also be noted that this should only be integrated over the forward going portion of the distribution function. Also, there is a zero by energy conservation for all values $E_{\gamma} > E_e$, so we must restrict our velocity magnitude integral to run from $[v_{\gamma}, \infty]$ where $v_{\gamma} = 2E_{\gamma}/m_e c^2$. Finally, we assume that emission is isotropic, and since we are looking at a point source, we normalize to the solid angle. This gives rise to quantity ψ , namely

$$\psi(E_{\gamma}) = \frac{a}{4\pi} \int_{v_{\gamma}}^{\infty} \int_0^{\pi/2} \int_0^{2\pi} (v_{\parallel}) (f_e) \left(\frac{CZ}{E_{\gamma}} (E_e - E_{\gamma})^{\alpha} \right) v^2 \sin \theta_v d\phi_v d\theta_v dv \quad (3.46)$$

$$= \frac{CZa}{4\pi E_{\gamma}} \int_{v_{\gamma}}^{\infty} \int_0^{\pi/2} \int_0^{2\pi} v^3 f (E_e - E_{\gamma})^{\alpha} \cos \theta_v \sin \theta_v d\phi_v d\theta_v dv. \quad (3.47)$$

Assuming no preference in cyclotron phase, the integral over ϕ_v becomes 2π , so the final formula is

$$\psi(E_{\gamma}) = \frac{CZa}{2E_{\gamma}} \int_{v_{\gamma}}^{\infty} \int_0^{\pi/2} v^3 f (m_e v^2/2 - E_{\gamma})^{\alpha} \cos \theta_v \sin \theta_v d\theta_v dv. \quad (3.48)$$

It is important to note that this analysis assumes non-relativistic motion. In practice, this integral is typically calculated numerically, and so the relativistic formula (derived in appendix A) is used.

Consider two datasets with slightly different distribution functions, for instance one set containing lower hybrid current drive and the other not, or two sets from different radial locations. We represent this difference by defining the distribution in the first case as $f_e = f_0$ (not necessarily Maxwellian) and in the second case as $f_e = f_0 + f_1$. This is still completely general, as f_1 is not restricted in anyway, except that $f_0 + f_1 \geq 0$ everywhere. The difference in the measured flux is given by

$$\Delta\psi = \frac{a}{4\pi} \int_{v_{\gamma}}^{\infty} \int_0^{\pi/2} \int_0^{2\pi} v_{\parallel} f_1 \varepsilon v^2 \sin \theta_v d\phi_v d\theta_v dv \quad (3.49)$$

Notice that this value does not depend in anyway on the distribution of the unperturbed electrons, unlike the plasma bremsstrahlung formula. This makes it a particularly good value for comparing

perturbative effects in areas of different temperature (for example different flux surfaces), because it provides insight into the size and speed of the perturbed distribution and does not require an identical unperturbed distribution. That is to say, the flux change on two x-ray chords looking at different radii in the plasma (not through a target probe) will be different even if the perturbed distribution is the same because the one of the scattering targets (the unperturbed electrons) will have a different temperature. If the flux change from the target probe is different, then this is a definitive statement that the perturbed population is different.

3.4 Inductively stabilized (PPCD) plasmas

PPCD plasmas have been successfully modeled by several previous authors [60, 92], nevertheless, as this is the only successful application of Fokker-Planck modeling of MST to date, it bears revisiting (furthermore, LH absorption predictions focusing on these discharges have not been presented). CQL3D was first applied to PPCD discharges for comparison with x-ray observations by O’Connell. [60] It was run in constant current mode and used to predict both Z_{eff} and diffusion coefficient. This work also presented a simulation of a 400 kA standard plasma, but the results are not accurate, as the hard x-ray emission of standard plasmas is very close to zero, and so the uncertainty is very large. PPCD is amenable to a solution of this type, as the relationship between electric field and current is shown to be direct and simple. PPCD plasmas have very good reproducibility and exhibit consistent x-ray spectra. Therefore modeling them with CQL3D is a fairly straight forward task. Although both of the previous major works studying PPCD plasmas focused on a single central x-ray chord, a similar analysis can be performed for a multi-chord setup like that on MST. The results are in very good agreement, as can be seen in Figure 3.8. This is particularly important when considering LH deposition, because the process of wave damping and driving current is highly dependent on both the electric field and the background distribution function, and observing the x-ray spectrum across the plasma builds confidence in these values.

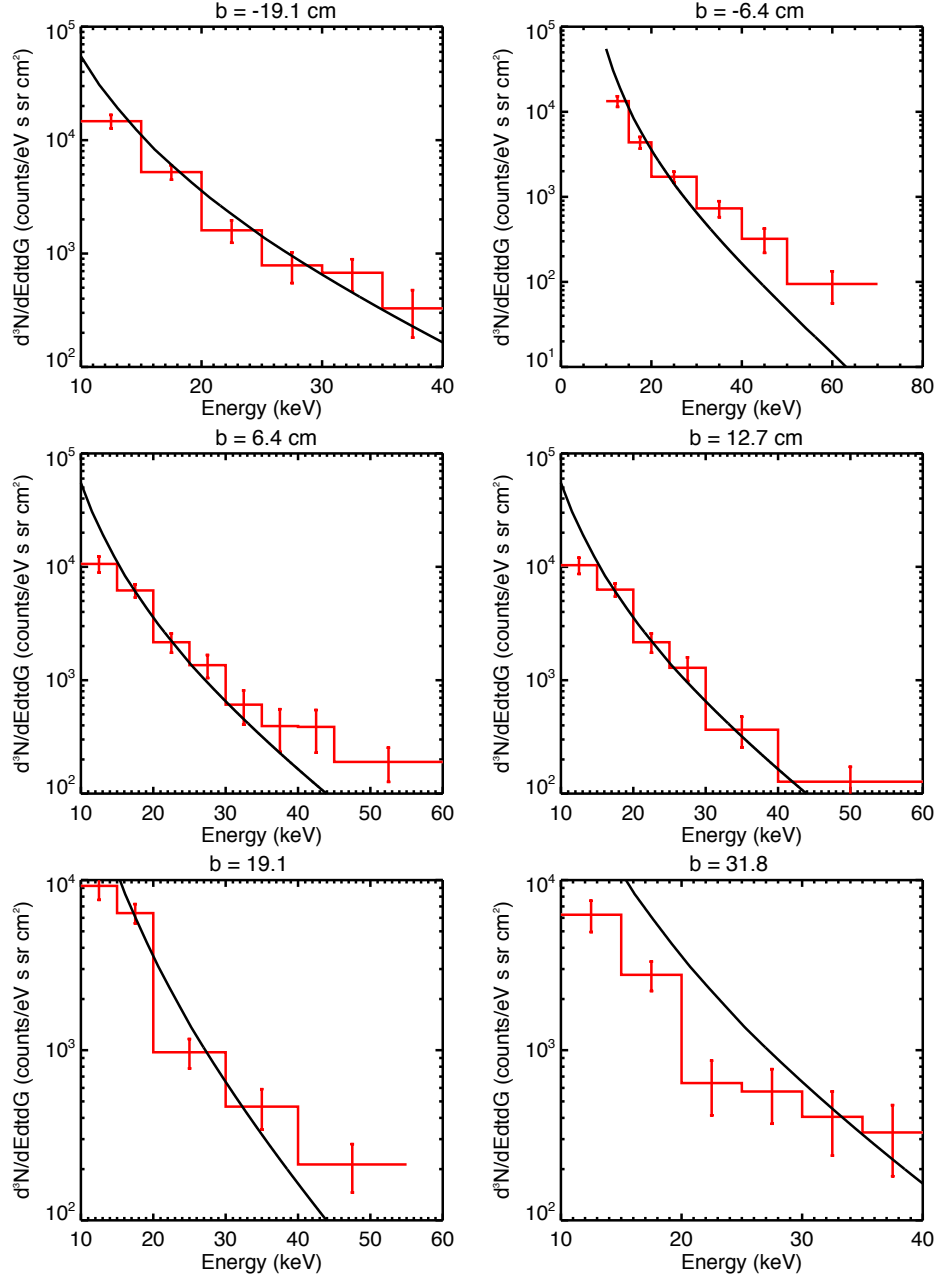


Figure 3.8: X-ray flux during PPCD discharges, with diffusion coefficient of $5 \text{ m}^2/\text{s}$ and a flat $Z_{\text{eff}} = 4$. Each plot represents a different detector, characterized by the impact parameter b , which is the shortest distance between the viewing chord and the presumed magnetic axis at $R = 1.56 \text{ m}$. It should be noted that these results are qualitatively similar for a range of diffusion coefficients.

3.5 Standard plasmas

By far the most complex plasmas studied on MST are standard RFP plasmas. These are not attractive targets for LH experiments, as will be shown, however they exhibit a lot of unique behavior and present many open questions, due to field line stochasticity. The current profile in MST plasmas is observed to be antiparallel with the magnetic field, across the entire plasma volume. This presents a serious problem for a localized interpretation of Ohm's law using the mean electric field (see Figure 3.9). The mean inductive electric field is directed toroidally, and so (because the




Figure 3.9: Electric field and current profile from standard plasmas (this data set comes from Jay Anderson). [11] The measurement of inductive electric field comes from a time derivative of the magnetic field from reconstructions. The electric field reverses direction relative to the magnetic field at the edge, but the current profile does not.

toroidal B field reverses direction), the projection of the electric field into the magnetic field direction reverses sign at the edge (that is the projection becomes parallel instead of antiparallel). The current is observed to not reverse sign however. A moment should be dedicated to consider the magnitude of this incongruity. These electrons are moving directly opposite to the direction of the force being applied to them. This is analogous to observing a waterfall with a net fluid flow that

is directed upward. Needless to say, some element has not been considered. Recall the nonlocal form of Ohm's law for large spatial scales in the RFP from Chapter 1:


$$\eta J_{\parallel} = E_{\parallel} + \eta \int ev_{\parallel} \tau |v_{\parallel}| \frac{\partial}{\partial x} D_m \frac{\partial}{\partial x} f_1 d^3v. \quad (1.67)$$

It is proposed that the nonlocality accounts for at least some of this deviation. Several models are going to be considered here. The two main constraints on our fit are the soft x-ray spectrum and the current profile.


3.5.1 Diffusive effects in Ohm's law

Diffusive effects become important in Ohm's law (and more so for understanding fast electrons) when the momentum of current-carrying electrons accelerated in a high field area allows them to maintain their speed even as they diffuse radially to areas with a different field strength. The most compelling instance of this effect is when these electrons move against the force imposed on them by the local electric field. This of course can only be maintained if electrons are constantly being supplied, and if the local electric field is low enough to not immediately redirect diffused electrons, or the electrons do not remain at a given radius for long. In order for this effect to be significant, there must be an electron population that carries a significant current and is subject to high diffusion. This will either be fast electrons in the presence of a stochastic field, where they are preferentially diffused, or any electrons in the presence of very high radial diffusion.

The first step in understanding the impact of diffusion is to study the impact of a delta function electric field. With no transport, this is expected to create a delta function current profile. As stochastic diffusion is added however, the diffusive term in Equation 1.67 will become more and more prominent, leading to a radially diffused profile as anticipated qualitatively. Examining Figure 3.10, it can be seen that the peak acts essentially as we expect for diffusion in a cylinder. Note that any radial fields created by particle diffusion (as theorized by Harvey [94]) are accounted for by CQL3D's *ad hoc* pinch velocity balancing particle density. It is important to note that inductive fields associated with the motion of these currents are not accounted for in this model. This is less of a concern when real profiles are being used because the density is held fixed to a measured



(a)



(b)

Figure 3.10: Diffusion acting on a delta function electric field (on a linear plot, and then a log plot) derived by CQL3D. In the highest diffusion case, the location of the original field is scarcely visible.

value, acting as a constraint for long term evolution. Nevertheless, a lot of very interesting studies could be done if Ampere's law and Faraday's law were being self-consistently evolved along with the kinetic evolution. This has proposed as the next major upgrade to CQL3D.

Now that the general action of diffusive effects has been demonstrated, the next step is to reproduce the work of Martines, [54] whose analytic formulation of the high diffusion RFPs was used to create Figure 1.17. The simulations (results of which are shown in Figure 3.11) use the same Z_{eff} profile as Martines, but rather than approximate cylindrical fields, measured fields and densities are taken from a reconstruction of an MST discharge. Notice that the best match to the experimental current density profile is observed for the case with $D_m = 10^{-4}$ m. The exact meaning of this magnetic diffusion coefficient is subject to interpretation, but to put it in context, a commonly used model states that [54]

$$D_m^0 = D_m (1 + \Lambda / \lambda) \quad (3.50)$$

$$D_m^0 = a (b_r / B)^2 \quad (3.51)$$

Where Λ is the auto-correlation length, λ is the mean free path of particles, and a is the minor radius. A $D_m = 10^{-4}$ m corresponds to a radial magnetic field of a few percent. In the context of previous experiments, a magnetic diffusion coefficient of 10^{-4} m corresponds to a diffusion coefficient of $1800 \text{ m}^2/\text{s}$ for a 1 keV electron. This is a factor of 4 higher than the observations made by Stoneking, [95] who measured a diffusion rate for fast electrons of $500 \text{ m}^2/\text{s}$.

The plots in Figure 3.11 demonstrate a lot of general information about the impact of diffusive effects. The largest parallel current in the RFP is generated in the core, because the magnetic field lines align directly with the inductive field there, i.e. E_{\parallel} is highest. The effect of diffusion is to reduce this peaked core current (since electron momentum is being drawn away from there), and increase the current in the midradius. Current is seen to be antiparallel to B across most of the plasma radius for the high diffusion cases, even beyond the reversal surface, where it is moving against the electric field. This counter field current is fairly substantial. Figure 3.12 shows contours of the distribution function on the first surface past the reversal surface for the case of




Figure 3.11: The effects of the stochastic diffusion with a real MST equilibrium. The simulation is run with a constant electric field, purely toroidal, with a constant magnitude of 1.5 V/m. This shows strong qualitative agreement with Figure 1.17. In the lower plot, detail of the edge of the plasma is shown (note the change in scale). The reversal surface and the zero point of current are marked. Notice that in the cases where $D_m \neq 0$, there is positive current even though the electric field reverses direction.




Figure 3.12: The impact of the stochastic diffusion on the distribution function. This plot shows contours of the distribution function with electrons whose momentum was transferred from another surface highlighted in blue. This data is from the $\rho = 0.85$ flux surface in the $D_m = 10^{-3}$ m.




Figure 3.13: This plot shows the specific current density $dJ_{||}/dv$ versus velocity for the surface just outside the reversal surface, comparing the $D_m = 10^{-4}$ m case to the diffusion free case. This specific current density should be thought of as the contribution to the total current density from electrons of a given speed. The no diffusion case has a Gaussian like shape, whereas in the high diffusion case there is a very large contribution to the net current from very fast electrons, and also a contribution from slow electrons, and both populations providing these currents originate from other flux surfaces.

$D_m = 10^{-3}$ m. The area highlighted in blue represents momentum brought in from other flux surfaces.

In stochastic plasmas, the impact of diffusion is heavily dependent on fast electron populations. This is demonstrated in Figure 3.13. For this reason the impact of stochastic diffusion is very much dependent on the same factors that determine the size of runaway populations, e.g. Z_{eff} and the Dreicer field. The data from Figure 3.11 were generated with a Z_{eff} of 1, flat across the entire plasma, which is not very realistic.

As a side note, the non-locality of Ohm's law interferes with CQL3D's constant current mode. In this mode, CQL3D adjusts the electric field at a given radius by determining the difference in current between the equilibrium being studied and the prediction from the distribution function

evolution. If (for the real experiment) a substantial amount of current is being created in the core and is diffusing outward, then CQL3D will erroneously suppress the core field and increase the edge field to try to match this current by local adjustments. This leads to strange behavior, which is often unstable and not to be considered accurate.

3.5.2 Enhanced Confinement Plasmas

A special subset of standard RFP discharges are so-called enhanced confinement plasmas, which have not yet been subjected to a Fokker-Planck treatment. Operationally speaking, Enhanced Confinement plasmas are distinguished from typical RFP plasmas by their low density and deep reversal. This is observed to give rise to periods of spontaneous stabilization of magnetic mode activity (although they are less stabilized than PPCD plasmas). These discharges are about twice as hot (700-1000 eV) and half as dense ($\langle n_e \rangle \approx 4 - 5 \times 10^{18} \text{ m}^{-3}$) as standard discharges, so the Dreicer field is substantially lower. On the other hand, the applied inductive field is the same as in standard discharges. Thus Enhanced Confinement plasmas will have a large fast electron and runaway electron population. Indeed, they are observed to have a very substantial hard x-ray flux compared to the unobservably low flux in their higher density cousins (the previously discussed standard RFP plasmas). If transport is stochastic or diffusion is high, momentum from these fast moving electrons should be diffused rapidly across the plasma.

There is some evidence for this. The current profile evolution during an EC period from reconstructions is shown in Figure 3.14. It is seen that the core current increases significantly, and this is associated with a small decrease in the current near the reversal surface. This relationship between core current and midradius current is characteristic of the diffusive Ohm's law displayed above (though it could be explained by other effects as well). As the time elapses from the sawtooth event, the magnetic fluctuations are reduced, and so the stochastic wandering of field lines should gradually relax. This would mitigate stochastic diffusion by reducing the flow of electrons from the core to the midradius. Since mode activity still exists however, it is certainly possible that it is not eliminated.






Figure 3.14: The evolution of the parallel current profile during EC periods. The core current can be seen to increase as the EC period continues. This is correlated with a reduction at $\rho \approx 0.8$. Detail is shown in Figure 3.14b. Time is relative to a sawtooth event.

This scenario of a strongly non-local Ohm's plasma can be tested. Two simulations were run with a constant toroidally directed electric field with magnitude of 1.5 V/m (approximately correct for MST discharges) and slightly different values of D_m . A flat $Z_{eff} = 4$ was used. The results are shown in Figure 3.15. It is satisfying that this model provides a meaningful application of the diffusive term in Ohm's law with a realistic Z_{eff} , unlike the very low Z_{eff} required to generate the fast electrons necessary for standard plasmas. Moreover the diffusion coefficient used is not unreasonable, again it corresponds to a radial magnetic field on the order of a few percent.


3.5.3 Models of standard plasmas

Standard plasmas are not well modeled by CQL3D. Several models are advanced to demonstrate how the code fails. These models are primarily differentiated by the method of initializing the electric field. It is important to note that these stand in for collections of runs, as explained for each model. Other models were also considered (for instance the Giruzzi high stochastic diffusion Model [58] and the Martines Model, presented in Figure 3.11), but are not included for concision (the two mentioned models also depend on a Z_{eff} of 1, and so are unlikely to reflect the actual plasma conditions). The models considered are (results plotted in Figure 3.16):

- **Diffusion free Model:** This model assumes that radial momentum transport is not important in relation to the current profile or the fast electron population. Therefore diffusion is turned off. Because there is diffusion, this is run in constant current mode to guarantee a good fit for current profile. This assumes (without justification) that whatever effects are not directly included in the kinetic model employed by CQL3D can be approximated by amplifying the additional electric field. This is equivalent to assuming there is some force in the that is not velocity dependent. The presented results employ a $Z_{eff} = 4$, although a scan of many Z_{eff} values was run.
- **Diffusion and Inductive field model:** The field here is initialized with the inductive field determined from the time derivative of a succession of reconstructions, take from the data of Anderson [11]. This model assumes that the plasma electric field is completely represented



(a)



(b)

Figure 3.15: Simulations are run with a flat toroidal electric field and $Z_{eff} = 4$, and two different values of D_m . Considering the approximate nature of the field, the agreement is very good.

by the applied inductive field, and that no other effects other than diffusion affect the motion of electrons. Many simulations of this setup were undertaken, varying Z_{eff} from 3 to 8, and also varying the central diffusion coefficient. The results plotted were the best fit, $Z_{eff} = 3$, flat $D_m = 10^{-5}m$.

- $\mathbf{E}_{\parallel} = \eta_{neo}\mathbf{J}_{\parallel}$: This model assumes neoclassical resistivity and uses this value to initialize the field. Due to the precipitous drop in temperature at MST's edge, the farthest radial values for resistivity are very high. This leads to a peak in the electric field, which must be reduced a little bit in order to produce reasonable results (otherwise it acts as a pumping mechanism for momentum via the diffusion and throws off the resulting x-ray spectrum). This is a source of ambiguity for this model.
- **\mathbf{J}_{\parallel} -Fit**: The most complicated simulation to run, this method begins by seeding a simulation with a guessed electric field. The current profile after the simulation comes to equilibrium is compared to the profile from the reconstruction, the electric field is adjusted in an attempt to reduce the difference between these two profiles, and the simulation is rerun. Electric field adjustments are done manually and in small increments in order to avoid over adjustments, and this process usually takes about 5-10 runs to complete. This is very similar to constant current mode, but it does not have step to step adjustments and so is able to come to equilibrium.

Unfortunately, none of these models produce results that match both the background current distribution and the x-ray spectrum. This can be seen in Figure 3.16. It is important to point out that both the equilibrium reconstructed current profile and the x-ray spectrum come (of necessity) from several shots and time periods ensembled together, and so an exact match should not be expected. The two ensembles do not include the same shots.

Standard RFP plasmas are very complicated and include important dynamics at many levels. None of these simple models provides reasonable agreement with both constraints (the current profile and the x-ray spectrum), in fact other than the \mathbf{J}_{\parallel} -fit model and the diffusion free model (both of which mandate an accurate current profile) none of the models demonstrate much agreement

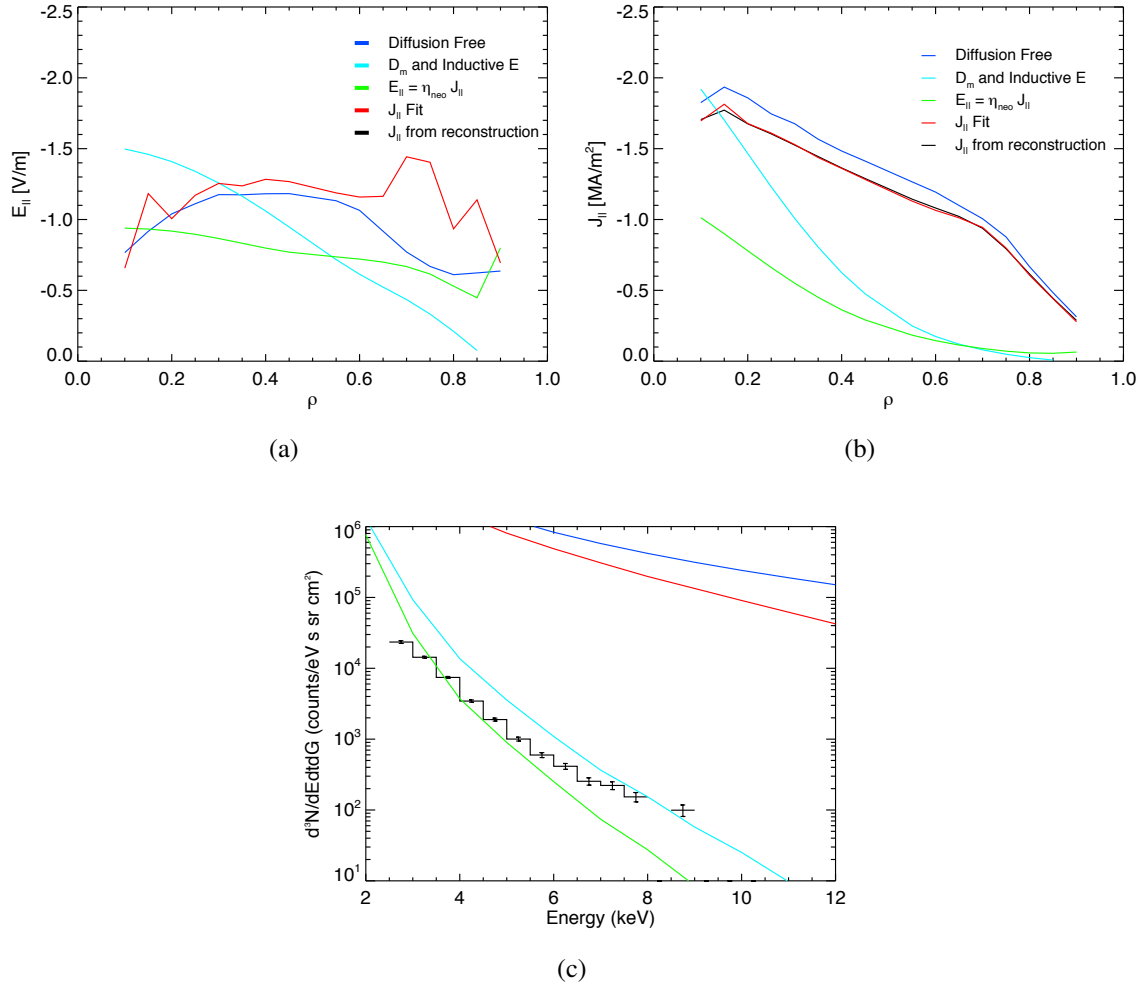


Figure 3.16: Comparison of electric field (a), predicted current profile (b) and x-ray spectra (c) predictions with the equilibrium current profile and the observed x-ray spectrum for standard non-EC plasmas. The no KDE model is in blue, the no MHD model is in cyan, the $E_{||} = \eta_{neo} J_{||}$ model is in green and the $J_{||}$ -fit model is in red.

with the equilibrium current. The J_{\parallel} -fit model and the diffusion free model predict extremely high x-ray fluxes however, and therefore are not acceptable. This high spectrum implies that this model over predicts the fast electron population, i.e. it assumes more current is being carried by fast electrons than is observed in experiment. The best x-ray spectrum fits are the Diffusion and Inductive field model and the neoclassical resistivity model, but their under prediction of the parallel current makes them no more correct than the first two models. This current must be being carried by an electron population that is not observed in the x-ray spectrum. Therefore we conclude that the simplified CQL3D Fokker-Planck model is deficient when applied to standard plasmas, although it is seen to have good agreement in Enhanced Confinement plasmas.

Despite not having a completely satisfactory model for the core x-ray flux from standard plasmas that produces a current profile consistent with reconstructions, the edge distribution function can be determined using the x-ray target probe. The most important factors in determining stimulated emission are Z_{eff} (held at 4), diffusion coefficient and electric field. Z_{eff} and D_m are fixed by estimates from the core x-ray spectrum, and so electric field is the only unknown parameter, and is determined by the model chosen. The neoclassical resistivity ($E_{\parallel} = \eta_{neo} J_{\parallel}$) model is used here. The data are plotted in Figure 3.17. It is found that a simulation using the values that produce the best fitting core x-ray spectrum also produces a target flux that agrees impressively. This is important for establishing the effect of lower hybrid current drive experiments.




Figure 3.17: The x-ray spectrum from the target probe at $\rho \approx 0.875$, with the CQL3D prediction also plotted. The agreement is observed to be much better than the agreement of the core flux.

Chapter 4

Lower Hybrid Modeling and Experiments on the MST

As discussed in §1.3, understanding the process of lower hybrid current drive requires studying the interplay of diverse physics at multiple length and time scales. In order to make accurate predictions, sophisticated numerical tools must be employed. Since much of the physics is separable to some degree, individual codes can be used to tackle the different elements in isolation, and the results can be communicated to the next code forming a simulation chain. Notably absent from this chain is the coupling of electromagnetic and global MHD effects with the kinetic simulation. These two physical processes occur on the same time scales, and understanding their interaction is vital. This is the focus of a major SciDAC effort called the SWIM project. [96] As such, this work seeks to provide a refinement to current understanding and not the definitive statement of current drive in the RFP. Methods of exploiting the results of this work will be discussed in the final chapter.

The simplified simulation chain used in this work consists of an equilibrium solver, a ray tracer and a kinetic code. First, the non-perturbed background conditions are determined from experimental data with MSTFit. Rf wave propagation physics is then studied with GENRAY, a general 3-D plasma ray tracing code. Finally, wave absorption and x-ray emission are studied using the Fokker-Planck kinetic code CQL3D, the subject of the preceding chapter. Ideally, the current deposition profile from CQL3D would be coupled to a nonlinear MHD code such as NIMROD, in order to determine the stability of the resulting configuration, and account for the electric and magnetic field response to driven current, but this level of complexity is beyond the scope of this work.

4.1 Wave Propagation

The first task of simulation is to predict wave propagation. Distortion of the background electron and ion distributions is not accounted for in the cold plasma wave propagation model, and is found to have minor effects for the slow wave in an RFP. [39] Therefore we apply the cold plasma model to an equilibrium reconstruction in order to determine the trajectory of the wave. This is accomplished by use of the code GENRAY.

4.1.1 GENRAY

GENRAY [97] is an general, 3-D, plasma raytracing code. It determines wave trajectories by taking the WKB (geometric optics) approximation in cylindrical coordinates, leading to the following set of equations

$$\begin{aligned} \frac{dR}{dt} &= -\frac{c}{\omega} \frac{\partial \mathfrak{D} / \partial n_R}{\partial \mathfrak{D} / \partial \omega} & \frac{dn_R}{dt} &= \frac{c}{\omega} \frac{\partial \mathfrak{D} / \partial R}{\partial \mathfrak{D} / \partial \omega} \\ \frac{d\phi}{dt} &= -\frac{c}{\omega} \frac{\partial \mathfrak{D} / \partial M}{\partial \mathfrak{D} / \partial \omega} & \frac{dM}{dt} &= \frac{c}{\omega} \frac{\partial \mathfrak{D} / \partial \phi}{\partial \mathfrak{D} / \partial \omega} \\ \frac{dz}{dt} &= -\frac{c}{\omega} \frac{\partial \mathfrak{D} / \partial n_z}{\partial \mathfrak{D} / \partial \omega} & \frac{dn_z}{dt} &= \frac{c}{\omega} \frac{\partial \mathfrak{D} / \partial z}{\partial \mathfrak{D} / \partial \omega} \end{aligned}$$

(where $M = Rn_\phi$) as discussed in §1.3, but cast in terms of index of refraction instead of wavenumber. Recall the left hand side of the left equations are the spatial coordinates, \mathfrak{D} is the dispersion relation, and n is the index of refraction. Note that the cylindrical system used is a system with the axial direction (z) pointing up and the radial direction corresponding to the major radius (R). The fields and profiles are represented in a fully toroidal fashion.

Although these equations are too complicated to solve analytically, they form a set of coupled first order ordinary differential equations which are amenable to a Runge-Kutta numerical solution. GENRAY takes as inputs the antenna location (in radial, toroidal and poloidal coordinates). It then traces a specified number of discrete rays from the antenna to the point where power has mostly dissipated (based on an approximation of quasilinear absorption) or until a specified number of steps have elapsed. The power in each ray in the launch spectrum is set so that when the initial points of all the rays are taken together, they form a Gaussian when plotted against n_\parallel , the




Figure 4.1: Sample of the n_{\parallel} spectrum launched in GENRAY. Ray power has a Gaussian shape. These initial powers are fed into CQL3D for the wave absorption calculation. The experimental spectrum is over plotted in black.

characteristics of which are provided as an input, as shown in Figure 4.1. Ray trajectories are then fed into CQL3D for a 2-dimensional (in velocity space) kinetic treatment of power absorption.

4.1.2 The Real LH Antenna

The lower hybrid antenna in place on MST was designed around the propagation and absorption of a single ray with $n_{\parallel} \approx 7.5 - 8$ (after construction it was discovered that the true n_{\parallel} centroid was closer to 7). Because the antenna is not infinitely long however, the real spectrum has some breadth in Fourier space, which has an impact on the propagation of power from the plasma edge to the absorption layer. This provides an important advantage for absorption as we shall see, but also carries with it a cost in lost power. In simulations of the real LH antenna, all of the important propagation effects discussed in §1.3.2 are observed. The poloidal transit of these waves as determined by simulation for 400 kA standard plasmas is shown in Figure 4.2. The waves can be seen to move mostly poloidally (essentially parallel to the magnetic field in the region of interest), and slowly wend their way inward. The spectrum is also seen to disperse as rays with different n_{\parallel}




Figure 4.2: Wave trajectories projected into the poloidal plane. The color in these plot corresponds to initial $n_{||}$. Plot (a) shows positive $n_{||}$ launch and plot (b) shows negative launch. Note that even though these rays are traced into the core of the plasma, most of the power is expected to be dissipated in the edge. Waves are traced beyond this point merely to ensure that the entire trajectory through the absorption region is calculated.

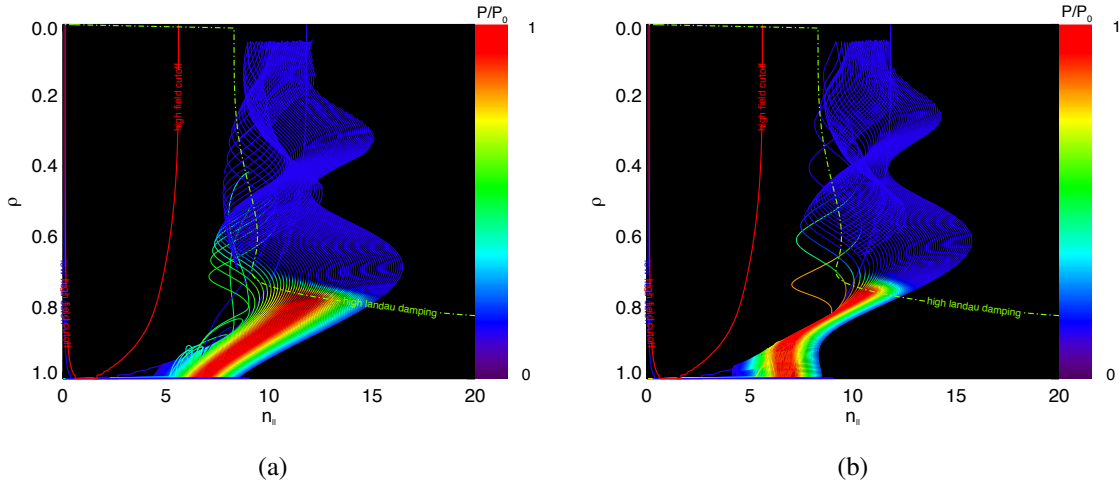


Figure 4.3: Wave propagation diagrams for the standard plasma equilibrium, with the launch spectrum of the current antenna. These plots show the wave's motion in n_{\parallel} space verses radial like coordinate. The color represents the relative power in the ray at a given point. The Fast-Slow Mode conversion point and the region of high Landau Damping are also plotted. Plot (a) shows the positive n_{\parallel} launch and plot (b) show negative launch.

propagate with different radial velocities as predicted in §1.3. A great deal more can be learned by examining the wave-space propagation diagram (based on Figures 1.8 and 1.14a). Raytracing results thus analyzed are presented in Figure 4.3.

In these plots, the horizontal axis represents the parallel index of refraction, on which Landau damping efficiency and Fast-Slow (F-S) mode conversion are based. The vertical axis represents the radial-like parameter ρ (the normalized area of a poloidal cross section of a given flux surface), thus the top of the plot is the center of the plasma and the bottom is the edge. The red line indicates the Fast-Slow mode conversion point on the high field side and the blue line (mostly obscured) indicates the F-S conversion point on the low field side. The region of efficient Landau damping (where $3v_{the} > v_{\phi}||$) is demarcated by the green dashed line. The other lines are the wave trajectories, with the coloration indicating the relative amount of power in that particular wave at that particular point (using an approximation of the quasilinear absorption rate). Note that the

$n_{\parallel} < 0$ direction is expected to drive current opposite the B field, and since J is anti-parallel to B , this is the net current drive direction.

There are several key aspects to note in these plots. First, note that the wave n_{\parallel} trends generally upwards and also oscillates up and down. The general upshift is due to the poloidal lensing effect, and the oscillations are due to the difference in magnetic field on the outboard and inboard sides as discussed in §1.3.2.5. This section predicted a correction to the poloidal lensing effect, stating that n_{\parallel} that is inversely proportional to $r|B|$ rather than just r . Similarly, notice that although the $n_{\parallel} > 0$ waves immediately increase in n_{\parallel} magnitude, the $n_{\parallel} < 0$ waves first down shift in n_{\parallel} magnitude. This is because the parallel wave group velocity is dependent on the sign of n_{\parallel} as predicted in Equation 1.16. Waves with a positive n_{\parallel} move parallel to the magnetic field, downward at the inboard midplane and upward at the outboard midplane. Thus, due to the antenna location ($\sim 45^\circ$ below the midplane on the inboard side), the negative spectrum travels inboard toward a higher magnetic field for a short while, reducing $|n_{\parallel}|$, until it crosses the midplane and the trend is reversed. This is an important effect because it brings the edges of the spectrum into collision with the F-S mode conversion layer. Since the antenna as constructed has a slightly lower n_{\parallel} than expected, this effects a larger portion of the spectrum than originally intended.

The impacts of this fast-slow mode conversion can be seen clearly if Figure 4.3 is examined near the region where power is deposited. The positive launch direction appears to have a broad spectrum of waves which damp away rapidly, whereas the negative launch direction is missing almost half of the initial spectrum. These waves are lost to fast-slow mode conversion. One of these mode converted rays is examined in detail in Figure 4.4. This figure shows the wave transferring solution branches from the slow mode to the fast mode at $F = 0$ (from Equation 1.12). The bottom plot indicates that this causes the formerly inward propagating wave to travel outward, reflecting off of the R-mode cutoff, then mode converting again, never reaching the target layer. This is identical to the process shown in Figure 1.10.

This mode conversion has profound negative consequences, as the power in that part of the spectrum is unavailable for current drive. Therefore, when studying full stabilization scenarios,




Figure 4.4: Profiles of one of the rays ($n_{\parallel 0} \approx -5$) plotted versus length along the ray. The top plot is F from the cold plasma solution (Equation 1.12). The second plot is n_{\perp} , with the fast and slow mode solutions also plotted. The third plot shows n_{\parallel} (which does not change much), with the R-mode cutoff (recall this is only a cutoff for the fast wave). The bottom plot shows the radial like coordinate, and reflection can be seen. As predicted, when $F \rightarrow 0$ the wave transfers from one branch to the other, resulting in a resounding chamber. Also, reflection from the R-mode cutoff can be seen.

more optimal antennas are considered. Because a multi-megawatt system will require many antennas, it stands to reason that such a project could include a refashioning of the current antenna with a more optimal launch spectrum. Full stabilization analysis is focused around PPCD plasmas, because, as will be shown, standard plasmas are not an attractive target for LH stabilization experiments.

4.1.3 Antennas for full stabilization experiments

Although a full understanding of LH deposition requires a 2-D Fokker-Planck simulation, ray-tracing results do provide sufficient information to optimize an antenna for propagation (avoiding the F-S mode conversion layer). The goal is to drive current near the tip of the temperature gradient (typically just inside of the reversal surface). n_{\parallel} can be expected to have a general upshift due to poloidal lensing. There will also be further oscillation associated with magnetic field magnitude as per the modification to poloidal lensing. Increasing n_{\parallel} too much will cause our ray to be absorbed farther out radially in the plasma than is desired (though this is not nearly as serious as trapping power in a resounding chamber). Therefore a succinct statement of the objective of antenna prototyping is to design an antenna with the lowest n_{\parallel} that will avoid the F-S mode conversion layer (perhaps with a little bit of wiggle room), and still be absorbed as near to the end of the temperature gradient region of an MST discharge as feasible.


Given a constraint on antenna size (roughly 20 cm) and an approximate damping parameter ($\sim 8^\circ$, based on operations of the current antenna), an educated estimate of the antenna launch spectrum can be made, using the formalism developed in §2.2.2.2. Equilibria reconstructed from two PPCD discharges are studied here (with plasma currents of 380 kA and 560 kA) and two reconstructions of EC periods are also studied (with nominal currents of 400 kA and 470 kA), because they are interesting targets for transport studies. These studies all focus on the $n_{\parallel} < 0$ case because it is the current drive direction. The results of these simulations are presented in Figure 4.5. Notice that for the EC cases and high current PPCD the Fast-Slow mode conversion layer is less obtrusive. This is because for the EC cases, the density is lower than for PPCD or standard plasmas by approximately a factor of 2, and the high current PPCD case has a high magnetic field.

Comparing these cases, it can be seen that the effect of a fairly modest increase magnetic field (The high current PPCD case has a core magnetic field of about 0.56 T compared with the 380 kA PPCD case's 0.37 T) is as dramatic as a substantial decrease in density. Moreover, it is seen that a system designed for a 380 kA PPCD plasma will provide a spectrum that is at least sufficient to study transport in EC plasmas, and demonstrate the feasibility of current drive in 560 kA plasmas. A separate system is considered for 560 kA PPCD plasmas however, because the 380 kA PPCD optimized system will drive current near or just outside of the reversal surface in a 560 kA PPCD system, making it less useful for stabilization.


4.2 Wave absorption

In order to study wave absorption, the trajectories predicted by GENRAY are fed into CQL3D, which calculates a quasilinear diffusion coefficient representing Landau damping, and includes it in the evolution of the Boltzmann equation. Before considering the full system, several elements of the model are examined individually to allow us to get a better feeling for their effect. These basic simulations include a full quasilinear Coulomb collision operator, and (obviously) the rf quasilinear operator. Diffusion and background electric field are not included except where explicitly stated.


A quick aside on the relative importance of the electric field. In analyzing rf driven current specifically for RFPs, it is important to separately consider the directly driven current and the electric field current amplification. Although there is certainly a large electric field in a normal RFP discharge, it is not clear what an LHCD stabilized RFP would be like. Because the inductive field is toroidal, the “parallel” inductive field near the reversal surface (where the magnetic field is mostly poloidal) is expected to be quite small. In a standard RFP discharge, this is supplemented by the so-called MHD dynamo, which is a nonlinear coupling between velocity perturbations and magnetic field perturbations). It may or may not be reasonable to consider this as equivalent to an electric field in some sense, but regardless of whether or not it will aid in the creation of a fast electron population, the explicit goal of running LHCD experiments is to eliminate this effect. PPCD experiments rely on applying a poloidal electric field, but again, the entire purpose of LHCD is to drive current in a non-inductive fashion, and so relying on this field seems counter to the overall




(a) 380 kA PPCD



(b) 560 kA PPCD



(c) 400 kA EC



(d) 470 kA EC

Figure 4.5: Wave propagation diagrams for optimized antenna in each high confinement case. The $n_{||}$ launch is centered around 8 for 380 kA PPCD, 6.5 for 560 kA PPCD, and 6 for the two EC cases.

mission. Therefore it is important to know the current drive prediction in the absence of an electric field as well as the current drive prediction with the electric field observed in experiments. These should provide a best and worst case scenario, with the actual setup lying somewhere in between. One can imagine a successful experiment pulsing PPCD at intervals and using the residual electric field for an LHCD experiment only to fill the gaps between PPCD periods. This would not necessarily be an optimal system, but would be an interesting demonstration of LHCD which could (in theory) require less power and be used as a stepping stone to higher power, full LHCD discharges. It must be stressed that the case with no electric field is overly pessimistic, because most of the region where current is being driven has some toroidal pitch, and therefore will have some field even in the absence of fluctuations and additionally applied poloidal field.

Attention is first turned to the simplest case, rf excitation with no diffusion and no electric field. Specifically of interest is whether or not the predictions of the Fisch model of distribution flattening discussed in Section 1.3.3 are reproduced in the full 2-D distribution function evolution using real quasilinear coefficients. The results can be seen in Figure 4.6. The resonance region can be clearly seen in both figures; it is the region where the distribution function is approaching a plateau. Figure 4.6a shows the effects of ever increasing rf power on the distribution function. It can be seen that for this flux surface, 1 MW of power seems to make the distribution function nearly flat. Figure 4.6b shows the full 2-D distribution function for the 1 MW case. Notice the increase in the population of electrons faster than the resonant velocity in the parallel direction. This is the action of the electron conveyor. It can be seen that there are more high speed electrons generally, in the counter-current direction as well as the co-current direction. This is due to pitch angle scattering, that is, fast electrons that undergo Coulomb collisions, but are primarily redirected rather than slowing down substantially. For the purposes of the no-electric field case, these are merely a curiosity since they do not represent a significant current, but as will be shown, this situation becomes more complicated when an electric field is applied.

Two parameters are not discussed below, temperature and magnetic field. Because we are driving current in a temperature gradient region, increasing the temperature of the plasma has the effect of moving the resonance region farther out radially, but current is deposited at a location with

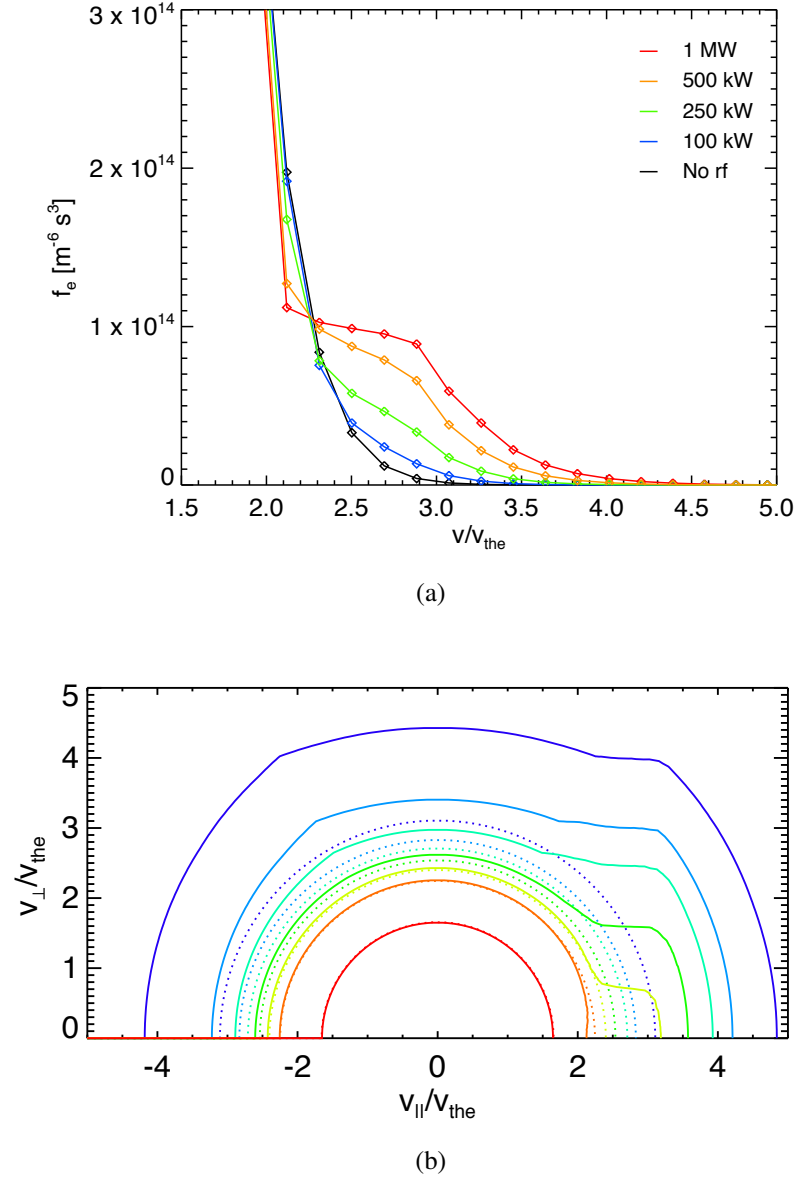


Figure 4.6: Figure (a) shows a slice of the distribution function along the $v = v_{\parallel}$ axis for various power levels. As anticipated, the distribution function flattens, and this flattening occurs for $v \approx 2 - 3v_{the}$, perfectly within expectation for Landau damping. Figure (b) shows a contour plot for the 1 MW case. The contours are essentially logarithmic. The dashed lines are the unperturbed (Maxwellian) distribution function. The effects of the electron conveyor can be seen.

essentially the same temperature. Magnetic field primarily impacts wave propagation as discussed above - lowering the field makes the antenna vulnerable to fast-slow mode conversion, but the absorption calculation is not dependent on magnetic field.

It is tempting to then declare that 1 MW of power should then be enough to excite the maximum amount of rf current drive achievable, but it should be recalled that there are other flux surfaces in the plasma, and as this distribution flattens, more power is allowed through to the deeper surfaces, which will (in general) have a larger resonant region. Further, one must recall that the inclusion of both diffusion and electric field will dramatically alter this situation. This is not to say that 1 MW is *not* enough power, just that the problem is more nuanced.

4.2.1 $Z_{effective}$

$Z_{effective}$ also plays a role in moderating LHCD. Just as a high Z_{eff} can mitigate the growth of a Dreicer runaway population, it also slows electrons expelled from the electron conveyor. In a high Z_{eff} discharge, electrons which would be ejected and come into equilibrium with collisions from ions (and each other) on the high speed side of the resonance region will have to contend with a more potent Coulomb collision force and thus will be prevented from achieving the same speeds seen with lower Z_{eff} . This relationship is demonstrated in Figure 4.7. Z_{eff} is not well known in MST but has been estimated at 4-5. This is consistent with the $1/(2 + Z)$ dependency predicted by Fisch as in Equation 1.34. An interesting side note is that this has little effect on the x-ray flux. The x-ray flux increases linearly with Z_{eff} , so even though fewer electrons are accelerated, the flux does not change substantially.

4.2.2 Radial Diffusion

Radial diffusion is found to have a significant impact on rf driven current. With diffusion, some electrons will certainly be lost to the walls, but the primary negative consequence is profile broadening. As rf accelerated electrons diffuse away from the deposition layer, they bring their current with them. Moreover, when these electrons move to surfaces nearer the core, they encounter a hotter, denser background plasma. This plasma will have a higher collision cross section with fast




Figure 4.7: Current drive degradation caused by increasing Z_{eff} . Higher Z_{eff} imposes a stronger Coulomb collision force, increasing the drag on electrons in the same way as Z_{eff} mitigates Dreicer runaway electrons.

electrons than the previous colder plasma, forcing the reintegration of electrons formerly decoupled from the main distribution. Because they are far from the deposition layer, there is no force driving them back to high speed. The net effect of diffusion is not only to broaden the deposition profile, but also to reduce the driven current. This effect is shown in Figure 4.8.

In stochastic plasmas, this effect is amplified by the stochastic diffusion. As discussed in the previous chapter, stochastic diffusion is intrinsically tied to the fast electrons. When applied to core electrons, where there is a large parallel electric field and therefore a large runaway pump, this amplifies the current in the midradius. That is to say, diffusion is very useful for broadening a centralized current profile. As one might have guessed by now, this is deleterious when attempting to drive localized current. Because lower hybrid current drive relies solely on faster-than-thermal electrons, it is particularly vulnerable to diffusion. The profiles which result are similar to the highest diffusion case in Figure 4.8, or even worse, illustrated in Figure 4.9.

The net impact of diffusion is summarized in Figure 4.10. It can be seen that about 2/3 of the current peak is dissipated in high diffusion plasmas. PPCD plasmas are expected to have a diffusion coefficient between $1 \text{ m}^2/\text{s}$ and $10 \text{ m}^2/\text{s}$ [60], where an acceptable fraction of the current is still driven. Stochastic plasmas have too broad of a profile to be considered good targets for




Figure 4.8: The impact of diffusion on current drive. The qualitative behavior is very similar to that seen by Forest [44], but the magnitude of diffusion required to broaden the peak is seen to be a factor of 10 higher. The diffusion coefficient in PPCD plasmas is between $1 - 5 \text{ m}^2/\text{s}$ [60].

stabilization. This will be countered to some extent by the electric field. Even if they were an attractive target, it would be a mistake to optimize a system for stochastic plasmas. Assuming a successful current drive experiment, temperatures would rapidly begin to increase as confinement improved, thus distorting the current drive profile.

4.2.3 Broad Spectra

Although broader launch spectra are more at risk for Fast-Slow mode conversion, they are absorbed much more efficiently than single ray spectra. This is because they have a broader resonance region (see Figure 4.11). Because the single ray is only resonant very near a single phase velocity, only so much current is driven by distribution function flattening. This means that few electrons are pushed out to higher speeds by the electron conveyor effect. Moreover, because the resonance region is smaller, it takes less power to completely flatten the distribution function. Once it is completely flat, more power is allowed to progress past the target region, broadening the deposition




Figure 4.9: The impact of stochastic diffusion on LH driven current. These plots simulate 60 kW of absorbed power. Because LH current drive focuses exclusively on fast electrons, it is particularly vulnerable to Rechester-Rosenbluth diffusion.




Figure 4.10: The effect of diffusion on current drive. The top plot shows the magnitude of the peak of the driven $J_{||}$ curve. It drops dramatically as diffusion is increased. This drop is even more dramatic for stochastic diffusion. The bottom plot shows distortion of the profile shape as diffusion is increase (note that radial like coordinate is the vertical axis). The points represent the peak of the Gaussian driven current curve, and the bars represent the width. It can be seen that as diffusion increases, the profile broadens and penetrates deeper into the plasma. The profile broadening is much worse for the stochastic case. Note for stochastic diffusion, $D_r = v_{the}(\rho = 0)D_m$, so that $D_r = 100 \text{ m}^2/\text{s}$ corresponds to $D_m \approx 1 \times 10^{-5} \text{ m}$.




Figure 4.11: The effects of a broad spectrum can be seen here. The left plot shows an individual flux surface distribution function. The single ray case (in blue) has a much smaller resonance region (in velocity space) than the full spectrum. This means that less current can be generated directly by distribution flattening and also that the electron conveyor is less effective. Moreover, the deposited current is broader as seen in the right plot. This is due to the drop in absorption efficiency.

profile. This process is so deleterious that Figure 4.11b shows two absorption peaks in the single ray case. Enough power is passing through the resonance region due to the reduction in absorption efficiency, that a significant portion of the ray's power is able to travel through the area of high absorption (the outboard side), having its $n_{||}$ reduced again by modified poloidal lensing. This wave will circle around poloidally, continuing to propagate inward until it reaches the outboard side again where Landau damping is efficient, however it is much deeper in the plasma.

4.2.4 Electric Field

The background electric field in an RFP discharge can be exploited to amplify applied current. Once electrons are pumped into the Dreicer runaway region they are free to accelerate with a minimal amount of drag. This amplifies the applied current considerably. That is, since all of the driven electrons are heading the same direction, a number of them will have a low relative velocity.




Figure 4.12: The current amplification caused by ever increasing electric field. This amplification will be mitigated considerably by diffusion (see Figure 4.10), however the overall trend persists. As mentioned, it is not clear what the field will be like in a stabilized RFP.

This will increase the drag, though still not nearly to the level experience by slow moving electrons in the thermal portion of the distribution. The effects of this can be seen in Figure 4.12.

4.2.5 Fast Electron Routing

The other major impact of the electric field is fast electron routing. When current is driven against the electric field, it is expected that the electric field will resist the fast electrons and drive them back into the thermal region. However some of the fast electrons will undergo pitch angle scattering, as seen in the no electric field case above (Figure 4.6b). Because these electrons retain their velocity, they still have a low collision cross section, only now the electric field is at their back. They are then accelerated just as normal runaway electrons. If the electric field is sufficiently high, the net effect of this is to drive current in the direction of the electric field. This process is illustrated in Figure 4.13. The picture is only slightly more complicated when trapped particle effects are included. In this case, the fast electrons which run parallel to the electric field arise from increased scattering of particles out of the trapped region due to increased collisionality, and electrons that have undergone a larger deflection and miss the trapped region entirely. In this




Figure 4.13: An illustration of the process of fast electron routing. First electrons are accelerated by lower hybrid. Some of them will scatter and be deflected. The electric field then accelerates them without forcing them through the thermal distribution.

way, a strong background electric field prevents anti-current drive experiments. This should not be considered an efficient means of driving current - for most plasmas of interest the field is high enough to drive electrons in this manner, but the overwhelming majority of the electrons are merely forced back into the thermal distribution. The key insight is that rather than there being a current drive and an anti-current drive direction, it is more appropriate to consider these an efficient current drive and an inefficient current drive direction. This effect can be seen in Figure 4.14.

4.3 Evidence of LH current drive

Although it has a long history of success on tokamaks, lower hybrid current drive has not been demonstrated on an RFP. Before moving on to more sophisticated uses of LH, it is important to demonstrate that current is actually being driven. Diagnosing the effects of lower hybrid at low power is particularly difficult. This complication does not exist for tokamaks - since they are replacing the inductive field of their plasma by driving toroidal current, either the inductive field will decay, or the toroidal current will increase (depending on how the inductive field is applied). Thus, current drive should be evident in either the loop voltage (a measure of the strength of




Figure 4.14: This plot shows current amplification verses applied field, where current is driven against the background (negatively directed) current. Note that the field causes the current to reverse direction when it is sufficiently high, though this process is not nearly as efficient as in the current drive direction, Figure 4.12.

the inductive field), or the toroidal plasma current. These are easily measured externally to the plasma. This is important because it means that they can be measured in a much lower noise environment, and small changes can be seen easily. The situation is somewhat more complicated for the RFP. There may be no net change in the applied electric field in the RFP because, as studies of PPCD suggest, [11] current drive might simply exactly cancel the dynamo effect (not impacting the inductive field at all). There certainly will be a modification to the measured loop voltage, but given the fluctuations in this value in standard plasmas, at low powers it is unlikely that this can be measured. More complex diagnostics must be considered.

Succinctly stated, lower hybrid current drive is used in the RFP to stabilize magnetic mode activity by exploiting fast electrons to drive current. Three main quantities of interest are suggested by this statement of objective: current (or current density), magnetic mode activity, and x-ray flux (a measure of fast electrons). These diagnostics will be more or less relevant in different regimes of applied power, and it is important to decide which diagnostics apply to which regime. Magnetic mode activity, while easily measured, is difficult to use as a diagnostic of lower hybrid current drive efficiency because it is unclear what the impact of driving small amounts of current will be. Along these lines, it should be noted that occasionally PPCD plasmas are seen to not achieve stabilization, for various reasons. During these failed PPCD periods, the poloidal inductive field is still applied, but magnetic mode activity is not observed to be significantly reduced. This implies that magnetic mode activity will only be mitigated (at an observable level) very near to full stabilization, and lends itself more to being a binary diagnostic of stabilization than a quantitative diagnostic of driven current. Direct measurements of driven current are the next most promising diagnostic. Unfortunately, this too presents some challenges. Because the primary direction of current drive is poloidal, most of the anticipated driven current cannot be measured externally. Probes can be used to directly measure current density, but in order to measure the small amount of current being driven, fairly delicate probes must be used. Owing to the high fields (and corresponding high temperatures) required to avoid F-S mode conversion (and thus lost power), these probes cannot be inserted deeply into plasmas of interest (specifically 400 kA PPCD plasmas, or even 400 kA standard plasmas). There does remain one possibility for directly measuring current however. Because

of the kinetic leaching effect for standard RFP discharges which diffuses LH driven current across the entire plasma, a significant number of fast electrons will be drawn into the core and accelerated by the inductive field, leading to a toroidal current which could feasibly be measured externally. This will still be a small amount of current, but if the main plasma current can be held constant enough, it might be measurable at high powers. Figure 4.15 shows CQL3D's predicted integrated toroidal current as a percentage of the background current for a 400 kA standard plasma at various power levels. Unfortunately, the change in current is small and the power required is large. This might prove to be a very good diagnostic in the future however, and might help shed light on the effects of induction.




Figure 4.15: The toroidally directed current driven in standard plasmas as a fraction of the background current (400 kA). Although fairly small, this might be measurable at higher powers, especially given MST's expanded capabilities with a programmable power supply.

The only candidate diagnostic which remains is x-ray flux. This diagnostic also presents some difficulties, though they are not insuperable. Examining PPCD plasmas first, difficulty arises due to the fairly low flux increase relative to the background flux. Figure 4.16 shows CQL3D's predictions of integrated x-ray flux from 10-30 keV versus absorbed power. Notice the lowest power case in this plot has a flux of about 1.6 times the background flux, and this point is about twice the present launched power level. This situation is further complicated by the fact that gas must be puffed in




Figure 4.16: The predicted hard x-ray flux (10-30 keV) during PPCD discharges. Notice that at present power levels ($P_{rad} < 60kW$) the predicted flux is less than a factor of 1.5 above the background flux, shown in blue.




Figure 4.17: The difference in flux for PPCD plasmas with 100 kW of lower hybrid power at the antenna face (~ 60 kW radiated). Predictions from CQL3D are also plotted. This flux difference is statistically significant.

order to couple significant power to the plasma (the subject of the next chapter). This is observed to reduce the number of high energy photons, implying a mitigation of fast electrons. The x-ray detectors cannot be operated in such a way as to view the maximum number of photons during the puffing-on time without saturating before the puffing is turned on, thus eliminating the data for the rest of the shot. It is not clear if this effect would impede stabilization experiments. Puffing can be avoided by using a longer antenna however, so in either case this would not be an insurmountable obstacle. Nevertheless, x-ray observations were made and a flux increase was observed, and is reproducible, albeit small. These data are plotted in Figure 4.17. Examining this plot, it can be seen that the flux increase observed during LH is consistent with the predicted flux increase, though the difference is not overwhelming. EC plasmas also exhibit an increase in x-ray flux. This is presented in Figure 4.18.

This provides an interesting constraint on transport in EC plasmas. Recall from the discussion of enhanced confinement periods that due to a low Dreicer field, there is a risk of explosive runaway growth. This would certainly be observed during experiments as the plasma current would rocket




Figure 4.18: The x-ray spectra with and without rf during EC discharges. The rf case is in red and the no rf case is in blue. This is a fairly consistent factor of 2 increase.

up to some very high level (which it does not), and so it is assumed to not occur in reality. Explosive runaway generation can only be mitigated in simulations by reducing the electric field or by adding massive amounts of diffusion. This is natural for stochastic plasmas, where fast electrons (and hence runways) are preferentially lost, but as mentioned in the last chapter, it is also possible that these plasmas exhibit very high but not stochastic diffusion.

Lower hybrid waves, injected into these discharges can determine if the transport in these plasmas is stochastic or non-stochastic. If transport is not stochastic, fast electrons are expected to be long lived in these plasmas even with fairly high diffusion. More precisely, all electrons are expected to be equally short lived, whereas with the stochastic diffusion case fast electrons are expected to disappear more rapidly than slow electrons. The background subtracted x-ray spectra from LH experiments is given in Figure 4.19. It is seen that even at low powers (60 kW of radiated power in this case) there is a significant difference in the predictions of the kinetic model for stochastic and non stochastic transport, and the data are more consistent with the stochastic transport case. Increasing the injected lower hybrid power would make this relationship even clearer.




Figure 4.19: X-ray flux during EC periods with CQL3D predictions for both the stochastic and non-stochastic diffusion case.

This provides strong evidence that EC plasmas exhibit stochastic transport, and by extension (vis-à-vis the previous chapter) are strongly impacted stochastic diffusion.

Finally, attention is turned to standard plasmas. Standard plasmas are known to exhibit stochastic transport, and so it is difficult to build up a substantially localized current. At the power levels currently used on MST, the hard x-ray flux predicted for these plasmas is very low, although it is observed (see Figure 4.20). Non-rf standard plasmas have an unobservably low x-ray flux however, so quantifying the flux increase is not possible. Moreover there is a dispute as to the origin of these x-rays. It has been previously conjectured by Kaufman [52] that this flux was entirely due to the so-called near field interaction, where electrons are accelerated because of cyclotron motion in the presence of an electric field gradient near the antenna face. A very large x-ray flux is indeed observed directly above the antenna itself, and can likely be attributed to this effect. Near field electrons are expected to be very energetic, having a very high perpendicular velocity (i.e. a large cyclotron orbit). Since they are created at the very edge of the plasma, it is expected that they will be lost very rapidly, striking the vessel wall and emitting a significant hard x-ray flux. They do not represent a sizable loss of power however.

In order to observe the Landau associated electrons, observations were made toroidally distant from the antenna (the Landau electron population should be axisymmetric). Furthermore, diagnostic capabilities were extended into the soft x-ray spectrum (2.5-10 keV). The flux in this part of the spectrum is high enough that emission is always observed, whether lower hybrid is on or not. This means that the difference in magnitude between background emission and LH-associated emission can be directly quantified. Moreover, because stochastic plasmas have velocity dependent transport, it is anticipated that the longest-lived electrons will be those at lower speed. While the difference in flux is not overwhelming, if many shots are ensembled together, a statistically significant flux difference is observed. The spectrum for these discharges (with a radiated power of approximately 100 kW) along with predictions from CQL3D are shown in Figure 4.20. Notice the strong qualitative agreement between the lower hybrid spectrum prediction and observation. The high energy soft x-rays (5-10 keV) and the hard x-rays are not observed without lower hybrid, as predicted by the kinetic simulation. Also notice the clear agreement in where the high energy soft

x-ray data are trending and where the low energy hard x-ray data begins. This is important because these detectors are at different toroidal angles, and so their relative agreement provides support of the hypothesis that the population being observed is not toroidally localized, as predicted for near field electrons. Although this agreement is certainly not perfect, it is impressively accurate given the uncertainties in the model even for non-rf plasmas. Bear in mind this plot shows rough agreement across 6 orders of magnitude. It is interesting to note that the simulation run here used the neoclassical resistivity model. Despite the uncertainty surrounding the background electric field, we are justified in using this model, because for stochastic diffusion, the electric field is unable to accelerate particles to very high energies without them being rapidly lost. Therefore, the exact nature of the electric field is of little importance. This is illustrated by comparing the predicted x-ray flux for a few of the closely agreeing models, shown in Figure 4.21.

The most important piece of evidence supporting the claim that LH x-ray emission far from the antenna is due to the Landau interaction and not the near field interaction comes from the x-ray target probe. Since the x-ray target probe makes a radially localized measurement, it cannot be argued that edge electrons are distorting the measured flux. The flux difference between LH on and LH off shots from target probe measurements is presented in Figure 4.22. Recall from the previous chapter that target probe flux and target probe flux difference are both directly related to the size and speed of the fast electron population. The enhanced x-ray flux is observed to be consistent shot to shot, and increasing as the probe is inserted deeper into the plasma. The circles in this plot are taken 180 degrees toroidally distant from the LH antenna and the triangles are taken 90 degrees away from the antenna. The agreement at these two locations provides strong evidence that this population of electrons is axisymmetric. The CQL3D prediction for 60 kW of absorbed power is also plotted. This appears to be roughly consistent with the observed population, although some extrapolation is required. There is an interesting plateau like feature in these plots. Notice that it can be seen in both the LH on, LH off, and Δ data points. This feature is resilient, as these data points represent a very large number of shots, and demands further scrutiny.

4.4 Current Drive for Full Stabilization

Having established that an LH associated, axisymmetric fast electron population is generated and increases with plasma depth, a full stabilization scenario can be considered. The most important computational work dealing with auxiliary current drive for RFP stabilization not directly focused on PPCD is the work of Sovinec and Prager. [9] In this work, many different current drive scenarios are simulated, exploring a wide parameter space, the most important of which for the purposes of this study are those denoted as case B, and I. These two scenarios, (properly unnormalized) are presented in Figure 4.23, with the reversal surface predicted by MPFM model also plotted. This work goes on to say that an RFP with a particularly large Ohmic field will not experience stabilization of core resonant modes with edge localized current drive. Unnormalizing the fields quoted, we find that MST operational parameters are very near the region shown to resist stabilization. It is also stated in this work that drive that is too localized near the reversal layer cannot stabilize core modes. Because the reversal surface in experiments is observed to be somewhat deeper than that presented in these simulations, the deposition layer based on experimental profiles is closer to the reversal surface. Due to the differences in geometry (the inherent toroidicity) and magnetic topology however, it is difficult to quantify the restriction on current drive location, and so it is unclear what the impact of this effect will be. However, in the absence of a better approximation, these two cases are assumed to be sufficient. It is the hope of this author to spur interest in this topic by presenting several realistic current drive profiles based on actual MST discharges. Finally, the different parts of wave propagation can be assembled and a meaningful prediction of current drive for real MST plasmas can be presented. The results of these simulations are presented in Figure 4.24 and Table 4.1. It can be seen in these plots that the case with electric field requires less power and is broader. This is due to runaway amplification across the entire profile. It can be seen that between 1.6 and 2 MW of power are needed for full stabilization experiments.

Just as or perhaps more important than the predicted amount of power required for current drive is establishing a scaling relationship. This relationship is observed to be quite linear for stabilized PPCD plasmas and is shown in Figure 4.25. It is also consistent with the scaling relationship of

case	P [MW]	J_0 [A/cm ²]	B_0 [T]	ρ_0	σ_a	ρ_{rev}
380 kA PPCD No E field	2.0	55.5	0.37	0.70	0.07	0.76
380 kA PPCD E field	1.6	55	0.37	0.72	0.11	0.76
560 kA PPCD No E field	2.4	85	0.56	0.74	0.11	0.8
560 kA PPCD E field	1.8	82	0.56	0.73	0.13	0.8

Table 4.1: Simulations determining the power required for stabilization based on the prediction for case B in Figure 4.23. J_0 is the peak current, B_0 is the magnitude of the core field, ρ_0 is the centroid of the peak, σ_a is the width of the peak and ρ_{rev} is the location of the reversal surface for reference. For comparison, case B requires a current of $\sim 150B_0[A/cm^2/T]$, and case I requires $\sim 110B_0[A/cm^2/T]$, but recall there are a lot of inherent assumptions in this approximation.

Experiment [ref]	η ($\times 10^{19}$ A/m ² W)
ASDEX [98]	0.168
HT-7 [99]	0.2-0.4
380 kA PPCD	0.275/0.297
560 kA PPCD	0.534/0.78
Alcator C [100]	1.2
JT-60U [101]	1.5-2.3
JET [21]	2.6

Table 4.2: The experimental efficiency of various tokamaks compared to the MST simulation results. The efficiency parameter is related to the temperature and it can be seen that the larger, hotter tokamaks have higher efficiencies. The double values listed for MST are with and without electric field.

Fisch discussed in Chapter 1, whereby rf is more efficient at driving current in higher temperature plasmas. The density profile for these two cases is quite similar. The exact prediction of the Fisch model depends on assumptions about the absorption volume, but for absorption layers on the order of 1 cm the efficiency is approximately 12 A/cm²/MW. This is about a factor of 3 lower than the CQL3D prediction, as Fisch noted for tokamaks. [43]

Accurately measuring radially localized J_{\parallel} and absorbed power (P_{abs}) in a hot plasma experiment is quite difficult, particularly when trying to measure a narrow band of increased current density. For this reason, the efficiency parameter usually cited by experiments is $\eta = \langle n_e \rangle I R_0 / P$ (the product of line-averaged density, driven current, and major radius divided by input power), which can be approximately related to J_{\parallel} / P_{abs} by the formula $\eta \approx 0.0002 T_e (keV) J_{\parallel} / P_{abs}$. [16] All of this is based on the geometry and scheme of the tokamak however, and it is not clear how directly relevant it is. Nevertheless, there is good agreement between the predictions of the MST simulations and some modern tokamaks as shown in Table 4.2. MST lies very near to tokamaks with similar temperatures.

Poloidal Angle	Peak n_{\parallel}	Spectral width
-134°	8.4	4.8
45.6°	8.5	4
86.9°	8	4

Table 4.3: Parameters for an upgraded LH antenna system. Note, this includes a redesign of the current lower hybrid antenna located at $\sim -134^{\circ}$ poloidal.

4.5 Alternative Scenarios

4.5.1 The intermediate system

The above sections assumed antennas placed in the same, fairly optimized location on MST. This would be a fine idea if one were building a brand new RFP and wanted to run full power stabilization experiments, however it is impractical in MST as it stands. Port access is limited, and there are precious few ports of appropriate size with a poloidal angle of -135° (the drilling of new holes at this poloidal angle would be both technically complex and politically infeasible). The MST LH group currently possesses three (potentially) 300 kW klystrons, and so, assuming this is near the limit of antenna power handling (a limit that has not been established), a different, more realizable system is considered, pictured in Figure 4.26.

The design presented here is optimized for a 380 kA PPCD discharge. Higher current discharges should be fine targets as well (though the deposition layer might be unacceptably far out radially), as would low density discharges of many different currents (EC plasmas for instance). This system calls for two additional antennas, located at the same toroidal angle as the current system. There is already a boxport at this angle, currently used (infrequently) for probes. This positioning is advantageous, as localized gas puffing is required for coupling to PPCD plasmas (as will be discussed in the next chapter) and the amount of gas puffed must be limited in order to prevent interruption of stabilization. Optimal parameters for the three antennas are presented in Table 4.3. Notice the new antennas are a little longer than the current antenna, leading to a narrower n_{\parallel} spectrum. Also note that the current antenna is redesigned, to better accommodate the

temperature profile of a full PPCD plasma (though it should still operate satisfactorily at lower or higher temperatures). The propagation diagrams of these antennas are presented in Figure 4.27. Although some of the features in these diagrams are slightly different than for the previous setup, they are completely within our expectations based on the theory outlined in §1.3.

A system at this power level should produce very significant current drive. Whereas the current system is not expected to produce much observable x-ray flux during PPCD discharges, an upgrade system of this sort should provide unequivocally increased x-ray spectra. X-ray emission and current drive predicted for this system are shown in Figure 4.28. Due to damage to the antenna, the current system is observed to have a limit of about 50 kW of radiated power. An upgraded system, operating at modest power levels previously observed with the present antenna when it was in better health, should be able to provide 450 kW of radiated power, and if the antenna design specification is reached, could launch 600 kW of rf power, a factor of 10 increase over the present system. As can be seen, the amount of current driven with the amplification caused by the electric field is almost half of the required current, in line with the predictions of Figure 4.25. A system at this power could be coupled with PPCD in order to see global stabilization. Specifically, PPCD could be initiated and after stabilization is achieved, the applied inductive field could be reduced. If this is done while rf is fired, the efficacy of current drive will be visible in the magnetic mode activity as compared to the same scenario without rf. This scenario is called a PPCD-rf handoff scenario.

4.5.2 The Non-inductive RFP

It has often been speculated [102] that rf current drive is not efficient enough to provide full current sustainment for an RFP (as opposed to profile control). Of course the term “efficient” is subject to interpretation, but it is a question worth addressing. In order to facilitate this, two approaches are considered. In both approaches, the inductive field is set to zero. Although this is not particularly physical, it is assumed that if an RFP was run for a long time after the inductive field had shut off, the field would decay away in some fashion, and it is not clear what the magnitude of

the field which resulted would be. Therefore, these should be interpreted as a worst case scenarios rather than an optimized scenarios.

First, motivated by the profile shape observed for stochastic transport (Figures 1.17, 3.11 and 4.9), it is speculated that core current could be driven by intentionally driving massive amounts of edge current and exploiting stochastic diffusion to draw the current into the plasma. Because lower hybrid current drive exploits fast electrons, these will be preferentially distributed across the plasma by stochastic transport. This is expected to be a fairly inefficient means of current drive, but would certainly be interesting scientifically. Stix and Ono proposed a similar scheme for non-inductive current drive using DC current injection. [103] This could be achieved with a slow wave antenna array, identical to the current system, albeit much larger in order to accommodate more power. The results simulations of this scenario are seen in Figure 4.29. This simulation focused on an Enhanced Confinement plasma because fast electrons are more easily generated due to the low Dreicer field. Note that even though this simulation is run without a background inductive field, the Dreicer field is still a critical parameter because it represents the ease of generating runaway currents. More power would be required for a standard plasma, where a large runaway population is more difficult to develop. This is shown to require at least 50 MW of power, a large but not unattainable amount of power. Although this is a lot more power than is currently used by the Ohmic system on MST (by a factor of 2.5), the advantage of arbitrary pulse length alone would justify a system of this size. Of course, based on the current antenna, 50 MW of radiated power would correspond to about 75 MW of provided power, and this is probably unacceptably large. This would also require 250 antennas operating at the design spec, and with a width of about 5 cm, the array would have to wrap around the machine twice. This is not as impossible as it sounds, as two arrays could be placed at different poloidal angles, but it still seems like a Herculean effort to keep an admittedly poorly confined plasma going.

A second, more viable technique is to focus on a broad deposition profile, restricting the plasmas studied to high current (560 kA) PPCD plasmas. The LH slow wave is not a suitable means of driving current across the plasma core in a stabilized discharge. This is due to the poloidal upshift, which causes the wave to oscillate back and forth in $n_{||}$ space. Because the wave motion

is directed mostly parallel to the magnetic field (therefore not radial), it will oscillate many times before reaching the core (or even the midradius) and will either strike the F-S mode conversion layer, or damp away in the region of high Landau damping. Rather, the LH fast wave is proposed. The wave propagation diagram for a broad deposition fast wave is shown in Figure 4.30.

The fast wave will not oscillate as much as the slow wave. The dispersion relation is significantly more complicated, but it is sufficient to point out that any poloidal shifting is not as important as the general poloidal lensing effect as the wave propagates inward radially. Moreover, because the fast wave is *not* nearly electrostatic, the picture of absorption is more complicated, and the polarization must be taken into account. [104] This reduced efficiency makes the fast wave unattractive when current localization is desired, but is actually a benefit for a non-inductive scenario because the wave is able to propagate in the region of strong Landau damping for a time without being completely absorbed, penetrating deeper into the plasma. In order to move the R-mode cutoff closer to the plasma edge (to facilitate launching fast waves) a much lower frequency (100 MHz) is used. This wave is launched from the outboard midplane, as it is assumed that a lot of real estate would be necessary in order to accommodate the amount of power that will be required. The launch spectrum covers the range $n_{\parallel} = [-5.5, -9.5]$. The results of these simulations are presented in Figure 4.31.

Based on Figure 4.31 it appears that 20 MW of rf power could replace the inductively driven current in a high current PPCD shot. If that seems like a lot, it is important to bear in mind that MST expends approximately 20 MW of power creating discharges anyway. Moreover, a fast wave driven RFP could be run for arbitrary periods of time, assuming long term MHD stability of an RFP driven in this way. Perhaps an ancillary system (either slow waves or OFCD) could provide some smoothing of this profile. Nevertheless in the context of the general efficiency of MST and the potential benefits of arbitrary pulse lengths, it's fair to describe this as a possible scenario to consider, or more likely part of a broad portfolio of current drive devices which could provide noninductive current. Even the 5 and 10 MW cases demonstrate significant current drive.




Figure 4.20: The x-ray spectrum for a lower hybrid discharge (co-current direction). Notice the good qualitative agreement over many orders of magnitude. The current profile in these simulations does not match the observations in reconstructs, as discussed in Chapter 3.






Figure 4.21: Comparison of the predicted x-ray flux during standard discharges with lower hybrid on. Notice the no electric field case under predicts the x-ray flux, as would be expected, but the two other models predict roughly the same flux, both fairly close to the experimental observation, and best agreement is achieved with $E_{||} = \eta_{neo} J_{||}$ model.



(a)



(b)

Figure 4.22: The flux increase associated with LH excitation versus radial-like coordinate. The dashed line indicates the CQL3D prediction for 60 kW absorbed power (the approximate amount of power radiated). The triangles are data 90 degrees away from the antenna and the circles are data 180 degrees away from the antenna. This perturbation is seen to be axisymmetric and increasing at deeper radii.




Figure 4.23: Current drive used for Cases B and I in Ref. [9], normalized to core field magnitude. The reversal surface predicted by the MPFM model is also plotted as well.

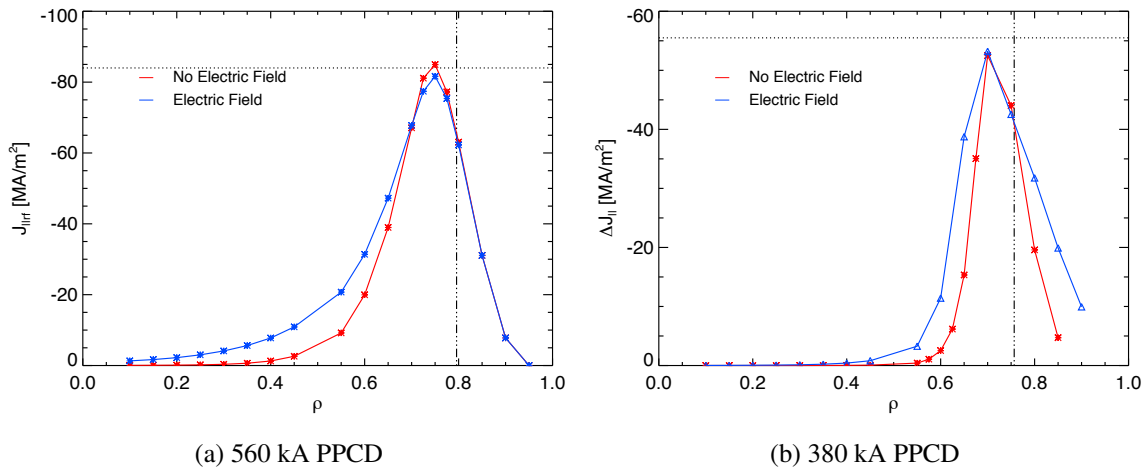


Figure 4.24: Current profiles for full stabilization, with parameters given in Table 4.1. The reversal surface and the current required for case B are overplotted.




Figure 4.25: The relationship between peak current and input power, for the optimized antennas, in both 560 kA PPCD and 380 kA PPCD. The slope of the lines is also provided. This should allow for a detailed analysis when coupled with an MHD code.




Figure 4.26: A possible lower hybrid upgrade system. This system would be placed on a boxport located at 90T, the same location as the current antenna. The total system would be able to deliver at least 450 kW, and possibly as much as 600 kW.




Figure 4.27: Wave propagation diagrams for the hypothetical upgrade system. Antennas at poloidal angles of (a) 45° , (b) 90° and (c) -135° .




Figure 4.28: Predicted x-ray flux and current drive for the hypothetical upgrade system. Because the radiated power is a factor of 10 higher than the current system and the power levels are still in the near linear regime, both driven current and x-ray flux are about a factor of 10 higher than currently observed. Plot (a) shows the driven current with and without an applied electric field. Plot (b) shows the x-ray flux with the electric field enhancement.




Figure 4.29: A plot of the current driven by 50 MW of LH power launched from the same poloidal angle as the current system, with no electric field, in red, the equilibrium profile is also plotted in black.




Figure 4.30: The propagation diagram for a 100 MHz fast wave in a 560 kA PPCD plasma in MST. Notice the topology is quite different at this different frequency, the F-S cutoff point is safely tucked away at the edge of $n_{||}$ space. Also notice the less localized deposition. The oscillating $n_{||}$ behavior is somewhat less pronounced.




Figure 4.31: Fast wave current deposition for very high power launch. The current from the equilibrium reconstruction is also plotted (the dashed line). From this plot it appears that 20 MW of power should be sufficient, assuming the profile is stable. More control can be gained by adding slow waves to the edge.

Chapter 5

Coupling Power to Improved Confinement Discharges

5.1 Coupling to the slow mode in MST plasmas

Having determined the current drive target and demonstrated the efficacy of the lower hybrid system, a detailed study of the process of coupling to the lower hybrid slow mode from the plasma edge is undertaken. Recall the mode cutoff diagram, Figure 1.8. The process of coupling is essentially the process of creating a perturbation in a vacuum (region IV) with n_{\parallel} near 7, and forcing that perturbation through the vacuum and low density region where the electric field decays exponentially, into the propagating zone, where a wave solution is supported. Because of this exponential decay, the more distant the propagating zone is from the source of the perturbation (the antenna), the less power will be available to be coupled to the plasma. The borders between the non propagating regions and the propagating region are the low density cutoff ($P=0$) and the Fast-Slow Cutoff, $F=0$ (which in the cutoff diagram is distant from the plasma edge). Recall that reducing the magnetic field brings the $F=0$ surface to lower densities. Therefore antenna coupling will be less efficient if density is decreased, or if the magnetic field is decreased severely (bringing the F-S mode conversion zone in front of the antenna face).

5.1.1 Density Dependence

The most comprehensive analytic work dealing with the coupling of this type of slow wave structure to a plasma was undertaken by Golant, [105] who focused solely on the effects of edge density (this derivation is briefly summarized in appendix F in a simplified form). Golant attempts to describe the electric field penetration for the four region problem presented in Figure 1.16,




Figure 5.1: An ensemble of many shots showing the relationship between line averaged density and the coupled power fraction. The biggest source of error is most likely the implicit assumption that the edge density directly tracks the line averaged density.

where the four regions are the slow wave structure, a vacuum region, a low density non propagating region, and the propagating region. The model derived from this analysis contains many potentially interrelated free parameters however, and is subject to much interpretation. The prediction of damping parameter seems to provide the wrong scaling. This can be seen by examining the derived value of the damping parameter (α_{coup}), calculated from Golant's equation 44. Recall that the power at the antenna face as a function of angle is given by $P(\theta) \propto \text{Exp}(-\theta/\alpha_{coup})$ (neglecting Ohmic losses), i.e a large α_{coup} means the wave at the antenna face damps away slowly and little power is coupled to the plasma, and a small α_{coup} means the wave at the antenna face damps very quickly so coupling is very efficient.

$$\alpha_{coup} = \frac{cn_{z0}}{4s\omega r_{ap}(n_{z0}^2 - 1)^{7/6} F} \left(\frac{\omega l_c}{c} \right)^{-1/3}. \quad (5.1)$$

In this equation, F represents the reduction in wave intensity due to exponential decay of field magnitude in the vacuum region (and has a somewhat complicated form), n_{z0} is the unperturbed (vacuum) launch spectrum peak, r_{ap} is the radius at which the antenna aperture is placed, and l_c is the distance of the $P = 0$ surface from the plasma edge. This scaling can be immediately

dismissed because it implies that when the non-propagating layer is enlarged (increasing l_c), the damping parameter will decrease, meaning the power at the antenna face will decay faster, leading to more efficient coupling. Indeed if the critical layer is moved to infinity, this formula predicts that all the power at the antenna face will be instantly coupled to the plasma. This is in direct opposition with our physical intuition and experimental observations, shown in Figure 5.1. This is not necessarily surprising as arriving at this solution is difficult work and requires many expansions and assumptions about scale and symmetry.

Something can be learned by studying a fairly simple model. Assume for a moment that the non-propagation region of the plasma acts just as the vacuum. Neglecting the reflected portion of the wave then, the electric field within the non-propagation region is given by

$$E \sim \text{Exp}[-\sqrt{n_{\parallel}^2 - 1} \frac{\omega}{c} x] \quad (5.2)$$

where x is the radial distance from the antenna (derivation of this is covered in Appendix F). The power in the launched wave (being electrostatic) is proportional to the square of the electric field magnitude. The total field decay before a propagating mode is excited will then be related to the total distance of the propagation region from the aperture face, specifically $l_{vac} + l_c$, the sum of the length of the vacuum region and the length of the non propagating region. The width of the vacuum region must be greater than the height of the antenna limiter (approximately 0.8 mm from the aperture face). As previously stated, it is expected that the non propagation region will be fairly small, so we estimate its size by Taylor expanding the plasma density around the vacuum - plasma interface (where the density is zero), i.e.

$$n_e(x) = 0 + \left. \frac{dn}{dx} \right|_{x=0} x \quad (5.3)$$

Holding this to be true, l_c can be calculated as

$$l_c = n_{crit} / (dn/dx)_{x=0}. \quad (5.4)$$

Finally, the density gradient is approximated by assuming that at some fixed distance Δ , the plasma has achieved some density n_{e0} (i.e. we say $(dn/dx)_{x=0} = n_{e0}/\Delta$). This leads to Figure 5.2.




Figure 5.2: The coupled power fraction predicted by the simple coupling model. These different curves represent different Δ parameters in our density profile assumption. For red, $\Delta = 1$ cm, for blue $\Delta = 5$ mm and for green $\Delta = 1$ mm. Notice the density here is edge density not line averaged density, which is why the magnitude differs from Figure 5.1

While these curves do not lie exactly atop the curves of Figure 5.1, they do give the same qualitative scaling, approximately the same limiting value of η , and given the number of uncertain terms (specifically l_{vac} , Δ , and the relation between line-averaged n_e and n_{e0}), agree impressively enough.

It seems prudent to revisit the Golant model, applying the fundamental assumption that the non-propagation region is not important. Rather than assume (as Golant does) that the plasma density begins at zero and increases continuously (necessary for the Taylor series expansion for l_c), the plasma is instead considered to have a minimum density, which is roughly constant over the (very small) region of interest. Thus, our density is zero for a distance of l_{vac} , and then steps up to a large value, presumably near $1 \times 10^{18} \text{ m}^{-3}$, the nominal edge density for these plasmas. The basic assumption of this interpretation is that the plasma ends abruptly at some value well above the critical density for non-propagation ($n_{crit} \approx 8 \times 10^{15} \text{ m}^{-3}$). This requires that a plasma with a high core density will be physically closer to the antenna, which is not altogether unreasonable since the pressure for a plasma such as this will be larger. Plots exploiting this assumption are




Figure 5.3: The coupled power fraction predicted from the Golant model with the plasma step function interpretation. The original width of the boron nitride limiter is marked. This is found to be generally quite reasonable, as the vacuum gap is expected to be on the order of millimeters.

shown in Figure 5.3. These results do indicate that as the plasma region is moved away from the antenna face, coupling is reduced, which is consistent with intuition.

5.1.2 Fast-Slow Cutoff

Determining the possible impact of the fast-slow cutoff is slightly more complicated. Because the antenna is located in the evanescent region for the density cutoff, moving the $P = 0$ critical layer directly inhibits antenna-plasma coupling. For the Fast-Slow equivalence point however, the wave can exist before encountering the critical cutoff layer, and then be reflected back the way it came. Again, a simple model is deployed to interpret how this might impact coupling. The underlying assumption of the model is that if the Fast-Slow cutoff is brought close enough to the antenna face, a launched slow wave will either undergo the resounding chamber phenomenon on a small enough spatial scale that (after converting to a fast wave and back) it will re-enter the antenna, or perhaps more appropriately there is not a large enough propagating region to support a full wave and so less power is able to couple into the plasma. This situation is sketched in Figure 5.4a. While




Figure 5.4: Details of the simple model of F-S reflections. Figure (a) shows a sketch of the process being modeled. Figure (b) shows the portion of the spectrum assumed to be cutoff (the shaded region). The radiated power is then calculated by integrating over the rest of the spectrum (assumed to be a Gaussian for simplicity).

the physical process might be fairly complex, the overall impact can be approximated easily. The plasma density is estimated at a point near the antenna face (near enough that a reflection will reenter the antenna). The nominal edge density of MST $\sim 1 \times 10^{18} \text{ m}^{-3}$ is a good choice. The n_{\parallel} that will be cutoff at that density for a given magnetic field strength is then calculated. Finally, it is assumed that the entire spectrum below that n_{\parallel} is not coupled to the plasma. The integral of the power in the remaining portion of the spectrum should be what appears to be radiated. This is illustrated in Figure 5.4b.

An experiment is devised to test this hypothesis. Although low current (and therefore low field) plasmas can be run to test this effect, it is more efficient to run plasmas with a slow current ramp down (with an associated field ramp down). The coupling efficiency could then be calculated fairly easily for a wide range of magnetic fields. It can be argued that the plasma density might be somewhat dynamic during this period, but the data used in this experiment were selected on the




Figure 5.5: The results of coupling experiments varying the magnetic field strength. The curves represent the predictions of the simple model, with three different densities used as the critical density (green is $n_e = 8 \times 10^{17} \text{ m}^{-3}$, blue is $n_e = 1 \times 10^{18} \text{ m}^{-3}$ and red is $n_e = 1.2 \times 10^{18} \text{ m}^{-3}$). The antenna begins to arc when the coupled power fraction gets below about 30%. 800 MHz is the cyclotron frequency when $B = 28.5 \text{ mT}$

basis of constant core density, and the results exhibit good agreement, given the relative simplicity of the model employed (see Figure 5.5).

5.2 Coupling to Improved Confinement Plasmas

This discussion of coupling to a plasma edge gains a new urgency when examining improved confinement plasmas. It is observed that during periods of high confinement, there is a degradation in the coupled power fraction. In §4.2.2 it was demonstrated that standard, stochastic RFP discharges do not provide a good target for stabilization experiments, due to the velocity dependence of the diffusion coefficient, and so coupling power to high confinement plasmas is of the utmost importance.

The coupling degradation in both PPCD and EC plasmas can be seen in Figures 5.6 and 5.7. As these plots demonstrate, although the average coupling fraction is certainly reduced, considering just the mean coupled power fraction over the entire high confinement period obscures a lot of the physics that is occurring.

In order to mitigate this loss of power coupling, the source must be identified. The discussion of coupling presented in §5.1 suggests that the two most important parameters are magnetic field and edge density. It can be seen in Figure 5.6 that during PPCD the field magnitude increases slightly. This is part of the normal ramping process and also occurs when PPCD does not succeed. In either case, this should improve coupling, although it should be noted that this field strength is well above the cutoff field. The pitch of the field line at the antenna face does change dramatically during PPCD. It is not expected that field *direction* will have much impact on coupling however, and indeed, it can be seen that after PPCD ends, the field pitch angle increases at roughly the same rate that it decreased when PPCD was underway, and yet the coupled power fraction instantly returns to standard (high) levels. Moreover, as can be seen in Figure 5.7, the pitch angle in EC plasmas is not all that different than it is during standard periods, like the one occurring right after EC ends.

In light of this, it is proposed that this coupling degradation is due to an edge particle transport effect. That is to say, the plasma density at the edge of the plasma column is significantly reduced




Figure 5.6: Coupled power fraction during PPCD periods. Notice the steady decrease in coupled power. This cannot be attributed to change in field pitch because after PPCD ends, the field pitch gradually returns to zero, but the coupled power fraction is restored instantly. Rather, a density pinch is indicated to be responsible.




Figure 5.7: Coupled power fraction during Enhanced Confinement periods. Notice this period is much shorter than the PPCD period. The coupled power fraction is still seen to be significantly reduced during the high confinement period, and seems to be steadily decreasing, however (perhaps because the high confinement period is shorter), it is not reduced quite as much as in PPCD. Notice that the coupled power fraction is somewhat erratic, seeing increases clustered around the small dynamo events.




Figure 5.8: The edge density profiles for standard, EC and PPCD plasmas. These data are ensembles of many many shots. The LH antenna limiter radius is marked. Note that these data are for discharges with significantly lower current than normal LH targets, in order to allow for Langmuir probe insertion. These data are courtesy of B. E. Chapman. [10, 106]

in PPCD plasmas. This was observed by Chapman, [10] who also observed a very similar effect in EC plasmas. [106] These data are reproduced in Figure 5.8 (the EC data set is unpublished). This is believed to be caused directly by the improved confinement, as particles are not driven to the edge region of the plasma as rapidly. This also means that fewer particles will strike the walls, leading to a reduction in neutral gas recycling. It is therefore proposed that poor coupling during these discharges can be mitigated by the puffing of neutral gas. Unfortunately it is generally observed that fueling both enhanced confinement and PPCD plasmas from the edge tends to degrade confinement improvement. Therefore it is not only necessary to study the dynamics of coupling during these discharges, it is also important to study the mode structure to see if stability is maintained.

5.3 Localized Gas Puffing Experiments

Additional edge fueling gas is puffed into MST by means of a piezoelectric valve. This valve is connected to a tube which releases the gas very near to the LH antenna, and at the same toroidal angle. The system has a maximum flow rate of 33 Torr L/s. Because the field is mostly poloidal at the edge of MST, the tube outlet is magnetically connected to the antenna face, meaning there is a field line that directly links the tube and the antenna aperture. In order to preserve the purity of ion species within the plasma, deuterium gas is used.

When localized gas is puffed during PPCD, it is observed to have a dramatic impact on coupling behavior, as shown in Figure 5.9. This figure shows that 15.4 Torr L/s of deuterium flow improves coupling essentially to standard plasma levels, maintaining a coupled power fraction near 90%, as compared with the undoped minima near 50%. It is interesting to note that the coupled power fraction is still reduced at the very end of PPCD even with puffing. This is not enough of a reduction to seriously hinder LH experiments, but enough to indicate a density depletion. This feature is observed on all shots with doping.

It is important to concurrently consider mode activity. Doping is seen to improve coupling, but if this is accomplished by derailing the improved confinement state, then it is without value. To this end, the mode activity for the $n = 4$ mode is examined. The mode activity is observed to be low over the entire PPCD period even with puffing.

The improvement of coupling and reduction of mode activity are somewhat more difficult to discern in Enhanced Confinement plasmas. EC periods arise spontaneously and sporadically. They differ from one another and vary fairly erratically as they progress. Moreover they occur at different times in the discharge, with different properties depending on where they occur. Because of this, the dynamics are even more important to consider than they are in PPCD plasmas, and the effect is less pronounced. Nevertheless, improvement in coupling associated with doping is observed (Figure 5.11). While the coupled power fraction in the non-doped EC period varies wildly, the doped EC period is able to maintain consistently high coupled power.




Figure 5.9: Coupled power for three different discharges which display large scale PPCD-driven stabilization, with three different values of puffing. The lines indicate uncertainty. The blue curve is from a shot with no puffing. The red curve is from a shot with puffing of 6.2 Torr L/s. The green curve is from a shot with 15.4 Torr L/s of gas flow. It can be seen that the undoped plasma exhibits a coupling fraction as low as 50% whereas the fully doped shot exhibits almost no coupled fraction reduction at all.




Figure 5.10: The $n = 4$ mode is displayed. The top plot is a shot with 15.4 Torr L/s localized gas flow. This is the very same shot as the green case in Figure 5.9. The second plot is a no doping shot where PPCD was fired, but failed to achieve stability, for the purposes of comparison. It can be seen that the shot with doping remains stable despite the additional edge fueling.




Figure 5.11: Coupled power during EC periods with and without gas puffing. The red data represent the gas puffing case and the blue data represent the no gas puffing case. Looking at the mode activity, it can be seen that the both plasmas achieve EC at about 26 ms. This is seen to effect the coupled power fraction, although it remains above 80% throughout the discharge. In the no gas case however, η is observed to fluctuate as low as 60%. These two discharges were chosen because they have roughly equivalent EC periods based on onset time, soft x-ray emission, and SDE burst severity.

It is clear that edge density is the most important factor affecting edge coupling in improved confinement plasmas, and it is equally clear from these data that localized gas fueling can mitigate these effects. It is not clear however what impact if any this has on current drive.

Chapter 6

In Summary

6.1 Review and discussion of original work

Before discussion of the future directions in which this work can be taken, a simple review and discussion of present original results is given, separated by chapter.

6.1.1 Introductory Material (Chapters 1 and 2)

The theory of plasma wave propagation and some elements of absorption are laid out here. This chapter contains all the tools necessary to gain a good qualitative understanding of any RFP LH slow wave system. In particular, the study of accessibility vis-à-vis the F-S mode conversion layer and the $P = 0$ cutoff by exploiting wave-propagation diagrams with MST like parameters provides a means of estimating a required launch n_{\parallel} . These diagrams also provide a simple method for instantaneously determining the possibilities of reaching a given radial target, once frequency is fixed. The modified poloidal lensing effect was also derived and explained. This is an important propagation effect (as demonstrated in later chapters) as it limits the penetration of lower hybrid slow waves. Because they oscillate back and forth in n_{\parallel} space, they cannot be expected to penetrate to the core of the plasma without encountering the F-S mode conversion point or damping away in the region of high Landau damping. This can be mitigated by choosing a lower frequency or increasing the magnetic field, which would move the F-S mode conversion layer closer to low n_{\parallel} . However, one must then contend with the problem of encouraging the wave to reach the region of high Landau damping at a precise radial location without relying on the n_{\parallel} oscillations because are largest at the plasma edge.

Inductive effects in the RFP were also studied. It was found that the inductive time scale shouldn't be a major impediment if one accepts the simple circuit model. It is not clear if it is right to do this however, and this problem should be studied more thoroughly with a 3D MHD code, and (experimentally) with high power rf before any firm conclusions can be drawn.

Finally, a more complicated model of power radiation is derived, and found to provide a good, dynamic measurement of power radiated to the plasma. This is vitally important to interpreting the results of Fokker-planck modeling, and more important for the coupling studies of Chapter 5.

6.1.2 Kinetic Modeling (Chapter 3)

This chapter was directed at the first objective question from the introduction, specifically “Does the Fokker-Planck Model provide a reasonable representation of RFP plasmas?”. Some time was spent developing and verifying a bias free statistic for interpretations of x-ray data which is not the traditional interpretation used. The theory of an x-ray bremsstrahlung stimulation target probe was developed in order to present predictions from modeling in a meaningful way, and the qualitative difference between these two measurements was discussed. PPCD plasmas had previously been studied by other authors, but it is important to point out that studies done in support of this work indicate that the Fokker-Planck technique can produce reasonable results for a range of diffusion coefficients, and so a certain amount of ambiguity exists when using this technique as a constraint on this coefficient.

The Rechester-Rosenbluth model of stochastic diffusion was employed as a model for the net effect of field line wandering. The impact of this was studied in great detail, starting with a Green's function approach (that is assuming the electric field is a delta function). It should be noted that unlike previous studies, this study used full toroidal fields, reconstructed from actual discharges. The results of Martines [54] were then reproduced for MST using Martines's Z_{eff} profile, to demonstrate some general effects of stochastic diffusion (Giruzzi's work [58] was also reproduced, but not presented as it is very similar to Martines'). Enhanced Confinement plasmas were predicted to have a large fast electron population (because of the low Dreicer field), and also observed to have a high x-ray flux, which is indicative of this population. CQL3D was shown

to have excellent agreement with both the reconstructed current profile and the x-ray spectrum, exploiting the Jacobson and Moses diffusive term to account for stochastic field line wandering.

Standard plasmas were then studied with several models. CQL3D was shown to be unable to reproduce both the observed current profile and the observed x-ray spectrum. The general result of these simulations is that when the model is coaxed to produce the correct current profile, the fast electron population is overestimated. Correspondingly, when the fast electron population is appropriately represented, the current profile predicted is too low. I am forced to conclude that there is a break down in the model. In the experiment there must exist a means of carrying current that is not observable in the x-ray spectrum. Presumably this arises from a population of electrons which are not correctly captured by the CQL3D model. .

6.1.3 Lower Hybrid Current Drive experiments and modeling (Chapter 4)

Evidence for an LH associated fast electron population was presented in this chapter, along with computational studies of wave propagation, absorption, driving stabilizing current and alternative scenarios. X-ray flux during standard PPCD and EC plasmas is shown to agree well with modeling, although this agreement is only qualitative for standard plasmas. This chapter contains the most important plot for establishing the efficacy of the present lower hybrid wave system on MST, that is Figure 4.22, the results of target probe experiments. These results demonstrate the existence of a radially localized, axisymmetric population of fast electrons, associated with lower hybrid power, that is consistent with the absorption of 10s of kW of power. This is not the same as demonstrating current drive, as fast electrons are directly injected by Landau damping, and any inductive field will affect all electrons equally. Therefore, fast electrons can be excited, and the inductive field might slow electrons nearer the thermal velocity, and net current might not yet be driven. Nevertheless, as the inductive field decays, some net current is expected. The present power level is not high enough to establish if current is being driven.

Full stabilization scenarios were studied, and it was determined that (depending on the plasma and applied field scenario), 1.5-2 MW of power would be required for full stabilization of PPCD

plasmas, allowing for the exploitation of electric field current enhancement. The scaling relationships for peak current density and power were established, and found to be linear with the slopes given in Table 6.1. This is enough information to begin studying this stabilization scheme with 3D

380 kA PPCD with Electric field	32.8 A/cm ² /MW
380 kA PPCD without Electric field	29.1 A/cm ² /MW
560 kA PPCD with Electric field	43.2 A/cm ² /MW
560 kA PPCD without Electric field	35.4 A/cm ² /MW

Table 6.1: Slopes for Linear scaling relationship, LH in PPCD.

nonlinear MHD codes, at least in an *ad-hoc* fashion.

Wave propagation was studied and shown to have strong qualitative agreement with the analytic studies of Chapter 1. Electric field, radial diffusion, spectral broadening, and Z effective were studied individually to isolate their physical impact. Effects like fast electron routing were discussed.

A realizable upgrade system was designed employing the tools developed in Chapter 1, using three antennas and the present rf system. This upgrade would bring MST's system to 450 - 600 kW and potentially allow puffing as described in Chapter 5 in one location for three antennas (thus reducing the total amount of required gas). Such a system should provide a large enough fast electron population that its effects would be easily discernible even in PPCD plasmas. A significant amount of current is predicted to be driven, enough that a PPCD assisted LH handoff scenario could be considered. (perhaps this would arise as a sustainment of the PPCD period beyond inductive field application).

Finally, non-inductive RFP scenarios are considered and found to be plausible, though not necessarily efficient. As mentioned in the chapter, it is for others to decide what constitutes an efficient system, however, it should be noted that the Landau interaction (as discussed in Chapter 1) becomes more efficient at higher temperatures such as those anticipated for a reactor. This at least provides a motivation to explore these scenarios in more detail with a reactor-like equilibrium.

6.1.4 Gas Puffing to control power coupling (Chapter 5)

This chapter leveraged the radiated power framework developed in Chapter 2 to study the dynamics of antenna coupling. Radiated power is found to be strongly dependent on the density in front of the antenna, though there is some ambiguity as to the exact nature of this relationship since the local density is not measured. Nevertheless, it is theorized that bad coupling during PPCD and EC periods is due to this density dependence. It is shown that not only is coupling improved generally by localized gas puffing, but that coupling efficiency, which decays constantly in PPCD and EC plasmas is maintained as the period develops. That is to say not only is the average coupling improved, but the dynamic degradation of coupling is mitigated. This provides strong evidence that the antenna can be made to couple power efficiently regardless of plasma dynamics - in a reactor for instance, one could imagine a dynamic system which puffed gas whenever the coupling efficiency dropped below a certain level. This is an important point, as it is unclear if a fully LH stabilized discharge would exhibit pinching or not, and if it did, it would be self defeating without some sort of ancillary means of improving coupling (a puffing system).

6.2 Future Work

There has been much discussion about upgrading CQL3D to self-consistently solve Ampere's law and Faraday's law (assuming static background fields) concurrent with distribution function evolution. This would do much to assuage concerns about the effects of induction during rf excitation and would provide a good model to study the importance of these factors and any possible time dependence of driven current. Of course a complete study would have to couple time advance of the kinetic equation with time advance of the MHD equilibrium, but the scope of such a study would likely demand a substantial team, and it seems unlikely to be pursued in the short term.

Experimentally, there is a strong impetus to follow the upgrade plan laid out in this thesis. A *circa* 500 kW launched power rf system should provide almost an order of magnitude increase in x-ray flux over PPCD background, an impressive amount that should be incontestably observable. This would also provide a serious amount of current, perhaps enough to stabilize a low current

plasma (assuming reasonable wave propagation). Perhaps sustainment of stabilization could be observed. Finally, this large increase in flux could be exploited to further constrain the diffusion coefficient, as the fast electron population is very sensitive to changes in this number.

REFERENCES

REFERENCES

- [1] D. R. Lide, editor *CRC Handbook of Chemistry and Physics* 81st edition CRC Press (2000)
- [2] K. S. Krane *Introductory Nuclear Physics* John Wiley and Sons 2nd edition (1988)
- [3] J. P. Freidberg *Ideal Magnetohydrodynamics* Plenum Press (1987)
- [4] D. C. Robinson and R. E. King, *Factors Influencing the Period of Improved Stability in ZETA*, Plasma Physics and Controlled Nuclear Fusion Research **1** (1969)
- [5] J. B. Taylor, *Relaxation of Toroidal Plasma and Generation of Reverse Magnetic Fields*, Phys. Rev. Lett. **33** (19) pages 1139–1141 (1974)
- [6] D. Biskamp *Nonlinear Magnetohydrodynamics* Cambridge Monographs on Plasma Physics. Cambridge University Press (1993)
- [7] G. Fiksel, S. C. Prager, W. Shen and M. Stoneking, *Measurement of Magnetic Fluctuation Induced Energy Transport*, Phys. Rev. Lett **72** (7) pages 1028–1031 (1994)
- [8] D. D. Schnack, D. C. Barnes and Z. Mikic, *Semi-implicit Magnetohydrodynamic Calculations*, Journal of Computational Physics **70** pages 330–354 (1987)
- [9] C. R. Sovinec and S. C. Prager, *Magnetohydrodynamic effects of current profile control in reversed field pinches*, Nucl. Fusion **39** (6) pages 777–790 (1999)
- [10] B. E. Chapman, A. F. Almagri, J. K. Anderson, T. M. Biewer, P. K. Chattopadhyay, C.-S. Chiang, D. Craig, D. J. D. Hartog, G. Fiksel, C. B. Forest, A. K. Hansen, D. Holly, N. E. Lanier, R. O’Connell, S. C. Prager, J. C. Reardon, J. S. Sarff, M. D. Wyman, D. L. Brower, W. X. Ding, Y. Jiang, S. D. Terry, P. Franz, L. Marrelli and P. Martin, *High confinement plasmas in the Madison Symmetric Torus reversed-field pinch*, Phys. Plasmas **9** (5) pages 2061–2068 (2002)
- [11] J. K. Anderson, T. M. Biewer, C. B. Forest, R. O’Connell, S. C. Prager and J. S. Sarff, *Dynamo-free plasma in the reversed field pinch*, Phys. Plasmas **11** (5) page L9 (2004)
- [12] P. Thonemann, W. T. Cowhig and P. A. Davenport, *Interaction of Travelling Magnetic Fields with Ionized Gases*, Nature **169** pages 34–35 (1952)

- [13] T. H. Stix, *Radiation and Absorption Via Mode Conversion in an Inhomogeneous Collision-Free Plasma*, Phys. Rev. Lett. **15** (23) pages 878–882 (1965)
- [14] C. F. F. Karney, *Stochastic ion heating by a lower hybrid wave*, Phys. Fluids **21** (9) page 1584 (1978)
- [15] P. M. Bellan and M. Porkolab, *Experimental Studies of Lower Hybrid Wave Propagation*, The Physics of Fluids **19** (7) page 995 (1976)
- [16] N. J. Fisch, *Confining a Tokamak Plasma with rf-Driven Currents*, Phys. Rev. Lett. **41** (13) page 873 (1978)
- [17] R. McWilliams, E. J. Valeo, R. W. Motley, W. M. Hooke and L. Olson, *Steady-State Currents Driven by Collisionally Damped Lower-Hybrid Waves*, Phys. Rev. Lett. **44** (4) pages 245–248 (1980)
- [18] K. L. Wong, *Experimental Observation of Current Generation by Unidirectional Electron Plasma Waves*, Phys. Rev. Lett. **43** (6) pages 438–441 (1979)
- [19] S. Bernabei, C. Daugney, P. Efthimion, W. Hooke, J. Hosea, F. Jobes, A. Martin, E. Mazzucato, E. Meservey, R. Motley, J. Stevens, S. V. Goeler and R. Wilson, *Lower-Hybrid Current Drive in the PLT Tokamak*, Phys. Rev. Lett. **49** (17) pages 1255–1258 (1982)
- [20] S. Luckhardt, M. Porkolab, S. F. Knowlton, K. I. Chen, A. S. Fisher, F. S. McDermott and M. Mayberry, *Generation of rf-Driven Currents by Lower-Hybrid-Wave Injection in the Versator II Tokamak*, Phys. Rev. Lett. **48** (3) pages 152–155 (1982)
- [21] A. Ekedahl, Y. F. Baranov, J. A. Dobbing, B. Fischer, C. Gormezano, T. T. C. Jones, M. Lennholm, V. V. Parail, F. G. Rimini, J. A. Romero, P. Schild, A. C. C. Sips, F. X. Söldner and B. J. D. Tubbin, *Profile Control Experiments in JET Using Off-Axis Lower Hybrid Current Drive*, Nucl. Fusion **38** (9) page 1397 (1998)
- [22] F. X. Söldner, F. Leuterer, R. Bartiromo, S. Bernabei, R. Büchse, O. Gehre, R. W. Harvey, M. Kornherr, K. McCormick, H. D. Murmann and G. V. Pereverzev, *Profile control with Lower Hybrid Waves on ASDEX*, Nucl. Fusion **34** (7) page 985 (1994)
- [23] S. Ide, O. Naito, T. Oikawa, T. Fujita, T. Kondoh, M. Seki, K. Ushigusa and J.-. Team, *LHCD Current Profile Control Experiments Towards Steady State Improved Confinement on JT-60U*, Nucl. Fusion **40** (3Y) page 445 (2000)
- [24] J. R. Wilson, R. Parker, M. Bitter, P. T. Bonoli, C. Fiore, R. W. Harvey, K. Hill, A. E. Hubbard, J. W. Hughes, A. Ince-Cushman, C. Kessel, J. S. Ko, O. Meneghini, C. K. Phillips, M. Porkolab, J. Rice, A. E. Schmidt, S. Scott, S. Shiraiwa, E. Valeo, G. Wallace, J. C. Wright and t. Alcator C-Mod Team, *Lower Hybrid Heating and Current Drive on the Alcator C-Mod Tokamak*, Nucl. Fusion **49** (11) page 115015 (2009)

- [25] G. Kuang, Y. Liu, J. Shan, W. Xu, X. Zhang, D. Liu, F. Liu, Y. Zhu, C. Zhang, G. Zheng, J. Wu, J. Lin, B. Ding, H. Xu, Y. Fang, J. Li, J. Luo, X. Zhang, B. Wan, Q. Zhao, J. Mao, X. Gao, S. Zhang, C. Li, X. Gu, P. Qing, H. Fan, S. Liu, B. Ling, B. Ding, Y. Li, Z. Wu, Y. Jie, S. Liu, J. Xie and Y. Wan, *Lower hybrid Current Drive Experiments and Improved Performance on the HT-7 Superconducting Tokamak*, Nucl. Fusion **39** (11Y) page 1769 (1999)
- [26] K. Hanada, S. Itoh, K. Nakamura, H. Zushi, M. Sakamoto, E. Jotaki, Y. D. Pan, M. Hasegawa, S. Kawasaki and H. Nakashima, *Enhanced Current Drive Efficiency in a Long Discharge on TRIAM-1M*, Nucl. Fusion **41** (11) page 1539 (2001)
- [27] G. T. Hoang, C. Bourdelle, B. Pégourié, B. Schunke, J. F. Artaud, J. Bucalossi, F. Clairet, C. Fenzi-Bonizet, X. Garbet, C. Gil, R. Guirlet, F. Imbeaux, J. Lasalle, T. Loarer, C. Lowry, J. M. Travère and E. Tsitrone, *Particle Pinch with Fully Noninductive Lower Hybrid Current Drive in Tore Supra*, Phys. Rev. Lett. **90** (15) page 155002 (2003)
- [28] B. Saoutic, *Status of Long Pulse Experiments in Magnetic Fusion Devices*, Plasma Physics and Controlled Fusion **44** pages B11–B22 (2002)
- [29] P. T. Bonoli, R. R. Parker, M. Porkolab, J. J. Ramos, S. J. Wukitch, Y. Takase, S. Bernabei, J. C. Hosea, G. Schilling and J. R. Wilson, *Modelling of Advanced Tokamak Scenarios with LHCD in Alcator C-Mod*, Nucl. Fusion **40** (6) page 1251 (2000)
- [30] E. Uchimoto, M. Cekic, R. W. Harvey, C. Litwin, S. C. Prager, J. S. Sarff and C. R. Sovinec, *Lower-Hybrid Poloidal Current Drive for Fluctuation Reduction in a Reversed Field Pinch*, Phys. Plasmas **1** (11) pages 3517–3519 (1994)
- [31] I. H. Hutchinson *Principles of Plasma Diagnostics* Cambridge University Press (1987)
- [32] H. W. Koch and J. W. Motz, *Bremsstrahlung Cross-Section Formulas and Related Data*, Rev. Mod. Phys. **31** (4) pages 920–956 (1959)
- [33] L. Landau, *On the Vibration of the Electronic Plasma*, J. Phys. (USSR) **10** page 25 (1946)
- [34] L. Landau *Collected Papers of L.D. Landau* Pergamon Press (1965)
- [35] D. G. Swanson *Plasma Waves* Series in Plasma Physics. Institute of Physics Publishing 2nd edition (2003)
- [36] T. H. Stix *Waves in Plasmas* Springer 2nd edition (1992)
- [37] P. Stubbe and A. I. Sukhorukov, *On the Physics of Landau Damping*, Phys. Plasmas **6** (8) pages 2976–2988 (1999)
- [38] C. B. Forest and D. R. Burke, *Propagation of LH Wave in the edge of MST*, MST Internal Report (2006)

- [39] R. A. Dziubek, R. W. Harvey, S. A. Hokin and E. Uchimoto, *Lower Hybrid Accessibility in a Large, Hot Reversed Pinch*, Nucl. Fusion **36** (4) pages 453–464 (1996)
- [40] E. V. Appleton, *The influence of the earth's magnetic field on wireless transmission*, Papers of the general assembly held in Washington D.C. in October 1927 (1928)
- [41] A. M. Woods, R. A. Cairns and C. N. Lashmore-Davies, *A Full Wave description of the accessibility of the lower hybrid resonance to the slow wave in tokamaks*, Phys. Fluids **29** (11) pages 3719–3729 (1986)
- [42] M. Brambilla, *Ray Tracing of Lower Hybrid and Ion Cyclotron Waves*, Comp. Phys. Rep. **4** (3-4) page 71 (1986)
- [43] N. J. Fisch, *Theory of Current Drive in Plasmas*, Rev. Mod. Phys. **59** (1) page 175 (1987)
- [44] C. B. Forest, P. K. Chattopadhyay, M. D. Nornberg, S. C. Prager and M. A. Thomas Radio Frequency Wave Experiments on the MST Reversed Field Pinch In *Proceedings of the 13th Topical Conference on Radio Frequency Power in Plasmas*(1999)
- [45] D. R. Burke, J. A. Goetz, M. C. Kaufman, A. F. Almagri, J. K. Anderson, C. B. Forest and S. C. Prager Diagnosis of Lower Hybrid on MST In P. M. Ryan and D. A. Rasmussen, editors, *Proceedings of the 17th Topical Conference on Radio Frequency Power in Plasmas*(2007)
- [46] J. D. Jackson *Classical Electrodynamics* John Wiley and Sons 3rd edition (1998)
- [47] F. F. Chen *Plasma Physics and Controlled Fusion 1: Plasma Physics* Plenum Press 2nd edition (1984)
- [48] L. Spitzer and R. Härm, *Transport Phenomena in a Completely Ionized Gas*, Phys. Rev. **89** (5) page 977 (1953)
- [49] L. Spitzer *Physics of Fully Ionized Gasses* Interscience Publishers 2nd edition (1962)
- [50] M. Brambilla, *Slow-Wave Launching at the Lower Hybrid Frequency Using a Phased Waveguide Array*, Nucl. Fusion **16** (1) page 47 (1976)
- [51] A. Ekedahl, G. Granucci, J. Mailloux, Y. Baranov, S. Erents, E. Joffrin, X. Litaudon, A. Loarte, P. Lomas, D. McDonald, V. Petržilka, K. Rantamäki, F. Rimini, C. Silva, M. Stamp, A. Tuccillo and J. E. Contributors, *Long distance coupling of lower hybrid waves in JET plasmas with edge and core transport barriers*, Nucl. Fusion **45** (5) page 351 (2005)
- [52] M. C. Kaufman *Lower Hybrid Experiments Using an Interdigital Line Antenna on the Reversed Field Pinch* PhD thesis University of Wisconsin - Madison(2009)
- [53] A. B. Rechester and M. N. Rosenbluth, *Electron Heat Transport in a Tokamak with Destroyed Magnetic Surfaces*, Phys. Rev. Lett. **40** (1) pages 38–41 (1978)

- [54] E. Martines and F. Vallone, *Ohm's Law for Plasmas in Reversed Field Pinch Configuration*, Phys. Rev. E **56** (1) page 957 (1997)
- [55] A. R. Jacobson and R. W. Moses, *Nonlocal DC Electrical Conductivity of a Lorentz Plasma in a Stochastic Magnetic Field*, Phys. Rev. A **29** (6) pages 3335–3342 (1984)
- [56] P. J. Catto, J. R. Myra, P. W. Wang, A. J. Wootton and R. D. Bengtson, *Estimating the runaway diffusion coefficient in the TEXT tokamak from shift and externally applied resonant magnetic field experiments*, Phys. Fluids B **3** (8) page 2038 (1991)
- [57] P. Helander, L.-G. Eriksson and F. Andersson, *Suppression of runaway electron avalanches by radial diffusion*, Phys. Plasmas **7** (10) page 4106 (2000)
- [58] G. Giruzzi and E. Martines, *Kinetic Modeling of Fast Electron Dynamics and Self-Consistent Magnetic Fields in a Reversed Field Pinch*, Phys. Plasmas **1** (8) page 2653 (1994)
- [59] J. S. Sarff, *Tutorial on the Magnetic Design of MST*, MST Internal Report (2006)
- [60] R. O'Connell, D. J. D. Hartog, C. B. Forest, J. K. Anderson, T. M. Biewer, B. E. Chapman, D. Craig, G. Fiksel, S. C. Prager, J. S. Sarff, S. D. Terry and R. W. Harvey, *Observation of Velocity-Independent Electron Transport in the Reversed Field Pinch*, Phys. Rev. Lett. **91** (4) page 045002 (2003)
- [61] B. E. Chapman, C.-S. Chiang, S. C. Prager, J. S. Sarff and M. R. Stoneking, *Strong $E \times B$ Flow Shear and Reduced Fluctuations in a Reversed Field Pinch*, Phys. Rev. Lett. **80** (10) page 2137 (1998)
- [62] H. Xia, M. G. Shats and H. Punzmann, *Strong ExB Shear Flows in the Transport-Barrier Region in H-Mode Plasma*, Phys. Rev. Lett. **97** (25) page 255003 (2006)
- [63] N. E. Lanier, J. K. Anderson, C. B. Forest, D. Holly, Y. Jiang and D. L. Brower, *First results from the far-infrared polarimeter system on the Madison Symmetric Torus reversed field pinch*, Rev. Sci. Instrum. **70** (1) page 718 (1999)
- [64] Y. Jiang, D. L. Brower and N. E. Lanier, *Interferometric measurement of high-frequency density fluctuations in Madison Symmetric Torus*, Rev. Sci. Instrum. **70** (1) page 703 (1999)
- [65] S. D. Terry, D. L. Brower, W. X. Ding, J. K. Anderson, T. M. Biewer, B. E. Chapman, D. Craig, C. B. Forest, R. O'Connell, S. C. Prager and J. S. Sarff, *Measurement of current profile dynamics in the Madison Symmetric Torus*, Phys. Plasmas **11** (3) page 1079 (2004)
- [66] W. X. Ding, D. L. Brower, D. Craig, B. H. Deng, G. Fiksel, V. Mirnov, S. C. Prager, J. S. Sarff and V. Svidzinski, *Measurement of the Hall Dynamo Effect during Magnetic Reconnection in a High-Temperature Plasma*, Phys. Rev. Lett. **93** (4) pages 045002–1 (2004)

- [67] F. M. Levinton, R. J. Fonck, G. M. Gammel, R. Kaita, H. W. Kugel, E. T. Powell and D. W. Roberts, *Magnetic field pitch-angle measurements in the PBX-M tokamak using the motional Stark effect*, Phys. Rev. Lett. **63** (19) pages 2060–2063 (1989)
- [68] D. Craig, D. J. D. Hartog, G. Fiksel, V. I. Davydenko and A. A. Ivanov, *First Charge Exchanged Recombination Spectroscopy and Motional Stark Effect Results from the Madison Symmetric Torus Reversed Field Pinch*, Rev. Sci. Instrum. **72** (1) pages 1008–1011 (2001)
- [69] J. A. Reusch, M. T. Borchardt, D. J. D. Hartog, A. F. Falkowski, D. J. Holly, R. O’Connell and H. D. Stephens, *Multipoint Thomson Scattering Diagnostic for the Madison Symmetric Torus Reversed-Field Pinch*, Rev. Sci. Instrum. **79** (10) page 10E733 (2008)
- [70] R. O’Connell, D. J. D. Hartog, M. T. Borchardt, D. J. Holly, J. A. Reusch and H. D. Stephens, *Optimizing a Thomson Scattering Diagnostic for Fast Dynamics and High Background*, Rev. Sci. Instrum. **79** (10) page 10E735 (2008)
- [71] D. J. D. Hartog, J. R. Ambuel, M. T. Borchardt, J. A. Reusch, P. E. Robi and Y. M. Yang, *Pulse-Burst Operation of Standard Nd:YAG Lasers*, Journal of Physics: Conference Series **227** page 012023 (2010)
- [72] D. J. D. Hartog, N. Jiang and W. R. Lempert, *A Pulse-burst Laser System for a High-Repetition-Rate Thomson Scattering Diagnostic*, Rev. Sci. Instrum. **79** (10) page 10E736 (2008)
- [73] J. K. Anderson, C. B. Forest, T. M. Biewer, J. S. Sarff and J. C. Wright, *Equilibrium Reconstructions in the Madison Symmetric Torus Reversed Field Pinch*, Nucl. Fusion **44** (1) page 162 (2004)
- [74] J. K. Anderson *Measurement of the Electrical Resistivity Profile in the Madison Symmetric Torus* PhD thesis University of Wisconsin - Madison(2001)
- [75] W. T. Welford and R. Winston *The Optics of Nonimaging Concentrators* Academic Press (1978)
- [76] S. Kishimoto, *High Time Resolution X-ray Measurements with an Avalanche Photodiode Detector*, Rev. Sci. Instrum. **63** (1) page 824 (1992)
- [77] J. G. Kulpin, *A 100 Kilovolt Phase Control Thyristor Switch Stack*, Proceedings of the 2008 IEEE International Power Modulators and High Voltage Conference page 378 (2008)
- [78] J. G. Kulpin, *High Voltage Switching Device*, US Patent Application 12/031,207 (2009)
- [79] Y. Takase, A. Ejiri, N. Kasuya, T. Mashiko, S. Shiraiwa, L. Tozawa, T. Akiduki, H. Kasahara, Y. Nagashima, H. Nozato, H. Wada, H. Yamada, T. Yamada and K. Yamagishi, *Initial results from the TST-2 spherical tokamak*, Nucl. Fusion **41** (11) page 1543 (2001)

- [80] R. J. L. Haye, C. J. Armentrout, P. I. Peterson and R. D. Stambaugh, *Propagation of very slow lower hybrid waves in the Octopole Tokamak*, Phys. Fluids **23** (9) page 1862 (1980)
- [81] E. G. Cristal, *Coupled Circular Cylindrical Rods Between Parallel Ground Plates*, IEEE Transactions on Microwave Theory and Techniques **12** (4) page 428 (1964)
- [82] I. Shkarofsky and M. Shoucri, *Modelling of lower hybrid current drive in the presence of spatial radial diffusion*, Nuclear Fusion **37** (4) page 539 (1997)
- [83] A. A. Mirin, J. Killeen, K. D. Marx and M. E. Rensink, *A radial transport/Fokker-Planck model for a Tokamak plasma*, Journal of Computational Physics **23** (1) pages 23 – 41 (1977)
- [84] G. Giruzzi, *Modelling of RF current drive in the presence of radial diffusion*, Plasma Physics and Controlled Fusion **35** (SA) page A123 (1993)
- [85] C. C. P. *et al.* In *Proceedings of the 14th Topical Conference on Radio Frequency Power in Plasmas*(2001)
- [86] R. W. Harvey, O. Sauter, R. Prater and P. Nikkola, *Radial Transport and Electron-Cyclotron-Current Drive in the TCV and DIII-D Tokamaks*, Phys. Rev. Lett. **88** (20) page 205001 (2002)
- [87] R. W. Harvey, V. S. Chan, S. C. Chiu, T. E. Evans, M. N. Rosenbluth and D. G. Whyte, *Runaway electron production in DIII-D killer pellet experiments, calculated with the CQL3D/KPRAD model*, Phys. Plasmas **7** (11) page 4590 (2000)
- [88] H. Dreicer, *Electron and Ion Runaway in a Fully Ionized Gas*, Phys. Rev. **117** page 329 (1960)
- [89] J. Wesson *Tokamaks* International Series of Monographs on Physics v. 118. Oxford Science Publications 3rd edition (2004)
- [90] J. Stevens, S. V. Goeler, S. Bernabei, M. Bitter, T. K. Chu, P. Efthimion, N. Fisch, W. Hooke, J. Hosea, F. Jobes, C. Karney, E. Meservey, R. Motley and G. Taylor, *Modelling of the Electron Distribution Based on Bremsstrahlung Emission during Lower-Hybrid Current Drive on PLT*, Nucl. Fusion **25** (11) page 1529 (1985)
- [91] H. Kawashima, T. Matova, Y. Uesugi, K. Hoshino and T. Yamamoto, *Radial Distribution of the Plasma Current Driven by Lower-Hybrid Waves in the JFT-2M Tokamak*, Journal of the Physical Society of Japan **56** (4) page 1348 (1987)
- [92] D. J. Clayton *Fast Electron Transport in Improved-Confinement RFP Plasmas* PhD thesis University of Wisconsin - Madison(2010)
- [93] G. H. McCall, *Calculation of X-ray Bremsstrahlung and Characteristic Line Emission Produced by a Maxwellian Electron Distribution*, J. Phys. D. **15** pages 823–831 (1982)

- [94] R. W. Harvey, M. G. McCoy, J. Y. Hsu and A. A. Mirin, *Electron Dynamics Associated with Stochastic Magnetic and Ambipolar Electric Fields*, Phys. Rev. Lett. **47** (2) pages 102–105 (1981)
- [95] M. R. Stoneking, S. A. Hokin, S. C. Prager, G. Fiksel, H. Ji and D. J. D. Hartog, *Particle Transport Due to Magnetic Fluctuations*, Phys. Rev. Lett. **73** (4) pages 549–552 (1994)
- [96] <http://cswim.org>, last checked September 27, 2010
- [97] A. P. Smirnov and R. W. Harvey, *The GENRAY Ray Tracing Code*, CompX Internal report (CompX-2000-01) (2003)
- [98] F. Leuterer, D. Eckhardt, F. Söldner, G. Becker, K. Bernhardt, M. Brambilla, H. Brinkschulte, H. Derfler, U. Ditte, A. Eberhagen, G. Fussman, O. Gehre, J. Gernhardt, G. v. Gierke, E. Glock, O. Gruber, G. Haas, M. Hesse, G. Janeschitz, F. Karger, M. Keilhacker, S. Kissel, O. Klüber, M. Kornherr, G. Lisitano, R. Magne, H. M. Mayer, K. McCormick, D. Meisel, V. Mertens, E. R. Müller, M. Münich, H. Murmann, W. Poschenrieder, H. Rapp, F. Ryter, K. H. Schmitter, F. Schneider, G. Siller, P. Smeulders, K. H. Steuer, T. Vien, F. Wagner, F. v. Woyna and M. Zouhar, *Recharging of the ohmic-heating transformer by means of lower-hybrid current drive in the ASDEX tokamak*, Phys. Rev. Lett. **55** (1) pages 75–78 (1985)
- [99] D. Bo-Jiang, K. Guang-Li, L. Yue-Xiu, L. Deng-Cheng, Z. Guang-Hua, W. Jun-Shuan, L. Fu-Kun, S. Wei-Ci, L. Jian-An, Y. Chun-Sheng, X. Han-Dong, Y. Jia-Wen, H. Yi-Yun, S. Lian-Quan, S. Yue-Jiang, W. Zhen-Wei, Z. Jin-Song, Y. Fu-Xian, Z. Xiao-Dong, L. Xiao-Ning, X. Ji-Kang and W. Yuan-Xi, *Effects of Plasma Density and Toroidal Magnetic Field on Lower Hybrid Current Drive Efficiency on HT-7 Tokamak*, Chinese Physics Letters **17** (7) page 519 (2000)
- [100] M. Porkolab, J. J. Schuss, B. Lloyd, Y. Takase, S. Texter, P. Bonoli, C. Fiore and R. Gandy, *Observation of Lower-Hybrid Current Drive at High Densities in the Alcator C Tokamak*, Phys. Rev. Lett. **53** (5) pages 450–453 (1984)
- [101] Y. Ikeda, O. Naito, M. Seki, T. Kondoh, S. Ide, K. Anno, H. Fukuda, Y. Ikeda, T. Kitai, K. Kiyono, M. Sawahata, S. Shinozaki, K. Suganuma, N. Suzuki and K. Ushigusa, *Simple multijunction launcher with oversized waveguides for lower hybrid current drive on JT-60U*, Fusion Engineering and Design **24** (3) pages 287 – 298 (1994)
- [102] Report of the FESAC Toroidal Alternates Panel(2008)
- [103] T. H. Stix and M. Ono, *Viscous Current Drive*, PPPL Internal Report **PPPL-2211** (1985)
- [104] S. C. Chiu, V. S. Chan, R. W. Harvey and M. Porkolab, *Fast Wave Current Drive in Low Aspect Ratio Tokamaks*, Nucl. Fusion **29** (10) page 2175 (1989)
- [105] V. E. Golant, *Plasma Penetration Near the Lower Hybrid Frequency*, Soviet Physics - Technical Physics **16** (12) pages 1980–1988 (1972)

[106] B. E. Chapman Private Communication

Appendix A: Derivation of the relativistic target probe formula

Electrons in the unit to tens of keV range are expected to be non-relativistic or mildly relativistic so (with care) we assume that assumption 3.45 is reasonable so long as we replace v_{\parallel} with $u_{\parallel} = p_e/m$. Recall that the emission formula has an implied zero due to energy conservation wherever the photon energy exceeds the maximum possible emission energy. As we anticipate that the electron will survive impact with the target, we take the maximum final energy to be the rest mass of the electron. We use this to derive u_{γ} , the cutoff momentum per rest mass

$$E_{\gamma} = \Delta E = \sqrt{(cp)^2 + E_0^2} - E_0 \quad (\text{A.1})$$

$$u_{\gamma} = c \sqrt{\frac{E_{\gamma}^2}{E_0^2} + \frac{2E_{\gamma}}{E_0}} \quad (\text{A.2})$$

Notice that for $E_{\gamma} \ll E_0$ (the non-relativistic limit) this reduces to

$$u_{\gamma} = c \sqrt{\frac{2E_{\gamma}}{E_0}} = \sqrt{\frac{2E_{\gamma}}{m_e}} \quad (\text{A.3})$$

which is our non-relativistic formula.

The other main consideration is that the electron energy referred to in Equation 3.44 is the total energy the electron can give up, $\Delta E = \sqrt{(cp)^2 + E_0^2} - E_0$. Finally, plugging this in and integrating over cyclotron phase, we get the fully (special) relativistic formula

$$\psi(E_{\gamma}) = \frac{CZa}{2E_{\gamma}} \int_{u_{\gamma}}^{\infty} \int_0^{\pi/2} u^3 f \left(\left[\sqrt{\left(\frac{u}{c}\right)^2 + 1} - 1 \right] E_0 - E_{\gamma} \right)^{\alpha} \cos \theta_u \sin \theta_u d\theta_u du. \quad (\text{A.4})$$

but since we are evaluating our integral numerically, this form is favored. Notice that for $u \ll c$,

$$\left(\sqrt{\left(\frac{u}{c}\right)^2 + 1} - 1 \right) E_0 \approx \left(1 + 1/2 \left(\frac{u}{c}\right)^2 - 1 \right) (mc^2) = \frac{mu^2}{2} \quad (\text{A.5})$$

which matches our non-relativistic formula, (3.48).

Appendix B: Assumptions about field lines

Understanding parallel current drive requires a few facts about magnetic field lines in a toroidal system. First we must find a relation between the poloidal angle and the toroidal angle of a field line.

Recall the field line equations [3]

$$\frac{dx_i}{dB_i} = \frac{dx_j}{dB_j} = \frac{dx_k}{dB_k} \quad (\text{B.1})$$

In toroidal coordinates this gives us

$$\frac{rd\theta}{dB_\theta} = \frac{Rd\phi}{B_\phi} \Rightarrow \frac{d\phi}{d\theta} = \frac{rB_\phi}{RB_\theta} \quad (\text{B.2})$$

We make use of a cylindrical approximation where we define the safety factor

$$q \equiv \frac{\mathbf{B} \cdot \nabla \hat{\theta}}{\mathbf{B} \cdot \nabla \hat{\phi}} \approx \frac{rB_\phi}{RB_\theta} \quad (\text{B.3})$$

which is a constant for a given flux surface (field line). Therefore we can say

$$\phi = q(\theta - \theta_0) + \phi_0 \quad (\text{B.4})$$

Next we use this relation to estimate the length of a field line

$$dl = \sqrt{(rd\theta)^2 + (Rd\phi)^2} \quad (\text{B.5})$$

$$\frac{dl}{d\theta} = \sqrt{r^2 + R^2 \left(\frac{d\phi}{d\theta} \right)^2} = \sqrt{r^2 + q^2 R^2} = \sqrt{r^2 + q^2 (R_0 + r \cos \theta)^2} \quad (\text{B.6})$$

$$= r \sqrt{1 + \frac{q^2 R_0^2}{r^2} \left(1 + \frac{r}{R_0} \cos \theta \right)^2} \quad (\text{B.7})$$

Near the reversal surface ($q \approx 0$), we can expand this to

$$\frac{dl}{d\theta} \approx r \left[1 + \frac{q^2 R_0^2}{2r^2} \left(1 + \frac{r}{R_0} \cos \theta \right)^2 \right] \quad (\text{B.8})$$

$$= r \left[1 + \frac{q^2 R_0^2}{2r^2} \left(1 + 2 \frac{r}{R_0} \cos \theta + \frac{r^2}{R_0^2} \cos^2 \theta \right) \right] \quad (\text{B.9})$$

$$= r \left[1 + \frac{q^2 R_0^2}{2r^2} + \frac{q^2 R_0}{r} \cos \theta + \frac{q^2}{2} \cos^2 \theta \right] \quad (\text{B.10})$$

$$l = \left(r + \frac{q^2 R_0^2}{2r} \right) (\theta - \theta_0) + q^2 R_0 (\cos \theta - \cos \theta_0) + q^2 r (\cos^2 \theta - \cos^2 \theta_0) \quad (\text{B.11})$$

We are most interested in field lines that come back to essentially the same point in m full rotations. More specifically, $\theta = \Delta\phi/q + \theta_0 = 2\pi m/q + \theta_0$, or (subtracting out full rotations for trigonometric functions $\theta \approx \theta_0 + \delta$ where δ is a small parameter. Notice, that this condition is often not met by field lines, but it is met by either 0 or an infinite number of them in any given region.

$$\begin{aligned} l_m &= \left(r + \frac{q^2 R_0^2}{2r} \right) \left(\frac{2\pi m}{q} \right) + q^2 R_0 (\cos(\theta_0 + \delta) - \cos \theta_0) + q^2 r (\cos^2(\theta_0 + \delta) - \cos^2 \theta_0) \\ &\approx \left(r + \frac{q^2 R_0^2}{2r} \right) \left(\frac{2\pi m}{q} \right) - \delta q^2 R_0 \sin \theta_0 - 2\delta q^2 r \sin^2 \theta_0 \end{aligned} \quad (\text{B.12})$$

$$(\text{B.13})$$

Assuming that we are looking at field lines at the midplane, this reduces to

$$l_m = \pi m \left(\frac{2r}{q} + \frac{q R_0^2}{r} \right) \quad (\text{B.14})$$

Finally, for small q this is simply

$$l_m = \frac{2\pi m r}{q} \quad (\text{B.15})$$

Appendix C: The Cold Plasma Model

C.1 Derivation

The cold plasma model begins by analyzing Maxwell's equations. Specifically Faraday's law and Ampere's law, which in differential form are

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad (\text{C.1})$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \frac{1}{c^2} \frac{\partial \mathbf{E}}{\partial t}. \quad (\text{C.2})$$

In our model we assume that our plasma consists of some number of species of charged particles, moving at constant velocity, with no thermal speed variation. Thus, we can say

$$\mathbf{J} = \sum_{\alpha} n_{\alpha} q_{\alpha} \mathbf{v}_{\alpha}, \quad (\text{C.3})$$

where alpha ranges over all species. Since our model is collisionless, \mathbf{v}_{α} can be easily derived from Newton's Second Law for an individual particle. We also introduce the following basic assumptions about the wave nature of our solution.

$$\mathbf{E}(\mathbf{r}, t) = \mathbf{E}_1(\mathbf{r}) e^{i(\mathbf{k} \cdot \mathbf{r} - \omega t)} \quad (\text{C.4})$$

$$\mathbf{B}(\mathbf{r}, t) = \mathbf{B}_0 + \mathbf{B}_1(\mathbf{r}) e^{i(\mathbf{k} \cdot \mathbf{r} - \omega t)} \quad (\text{C.5})$$

$$\mathbf{v}_{\alpha}(\mathbf{r}, t) = \mathbf{v}_{\alpha 1}(\mathbf{r}) e^{i(\mathbf{k} \cdot \mathbf{r} - \omega t)} \quad (\text{C.6})$$

Thus, Newton's Second Law for each species α is reduced to (neglecting second order terms)

$$q_{\alpha} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) = m_{\alpha} \frac{d\mathbf{v}}{dt} \quad (\text{C.7})$$

$$q_{\alpha} (\mathbf{E}_1 + \mathbf{v}_1 \times \mathbf{B}_0) = m_{\alpha} (-i\omega) \mathbf{v}_1. \quad (\text{C.8})$$

To simplify this, we take the background magnetic field to point in the $\hat{\mathbf{z}}$ direction. Splitting this into the component equations (and suppressing the fluctuation subscript) we have

$$\omega v_{x\alpha} = \frac{iq_\alpha}{m_\alpha} (E_x + v_{y\alpha} B_0) \quad (\text{C.9})$$

$$\omega v_{y\alpha} = \frac{iq_\alpha}{m_\alpha} (E_y - v_{x\alpha} B_0) \quad (\text{C.10})$$

$$v_z = \frac{iq}{m\omega} E_z \quad (\text{C.11})$$

Solving this for the different components of \mathbf{v} we get

$$v_{x\alpha} = \frac{iq_\alpha}{m_\alpha (\omega^2 - \omega_{c\alpha}^2)} (\omega E_x + i\epsilon_\alpha \omega_{c\alpha} E_y) \quad (\text{C.12})$$

$$v_{y\alpha} = \frac{iq_\alpha}{m_\alpha (\omega^2 - \omega_{c\alpha}^2)} (-i\epsilon_\alpha \omega_{c\alpha} E_x + \omega E_y) \quad (\text{C.13})$$

$$v_z = \frac{iq}{m\omega} E_z \quad (\text{C.14})$$

Where $\omega_{c\alpha} = |q_\alpha|B_0/m_\alpha$ and $\epsilon_\alpha = q_\alpha/|q_\alpha|$. Having determined \mathbf{v}_α , we can easily determine J_α by C.3

$$J_{x\alpha} = \sum_\alpha \frac{i\omega_{p\alpha}^2}{(\omega^2 - \omega_{c\alpha}^2)} (\omega E_x + i\epsilon_\alpha \omega_{c\alpha} E_y) \quad (\text{C.15})$$

$$J_{y\alpha} = \sum_\alpha \frac{i\omega_{p\alpha}^2}{(\omega^2 - \omega_{c\alpha}^2)} (-i\epsilon_\alpha \omega_{c\alpha} E_x + \omega E_y) \quad (\text{C.16})$$

$$J_z = i\epsilon_0 \sum_\alpha \frac{\omega_{p\alpha}^2}{\omega} E_z \quad (\text{C.17})$$

Plugging this into C.2 and applying the assumptions C.4-C.6, we get

$$\begin{aligned} i\mathbf{k} \times \mathbf{B} &= \mu_0 \mathbf{J} - \frac{i\omega}{c^2} \mathbf{E} \\ &= -i\frac{\omega}{c^2} \mathbf{K} \cdot \mathbf{E} \end{aligned} \quad (\text{C.18})$$

where

$$\mathbf{K} = \begin{pmatrix} S & -iD & 0 \\ iD & S & 0 \\ 0 & 0 & P \end{pmatrix} \quad (\text{C.19})$$

and

$$S = 1 - \sum_{\alpha} \frac{\omega_{p\alpha}^2}{\omega^2 - \omega_{c\alpha}^2} \quad (\text{C.20})$$

$$D = \sum_{\alpha} \frac{\omega_{c\alpha} \omega_{p\alpha}^2}{\omega (\omega^2 - \omega_{c\alpha}^2)} \quad (\text{C.21})$$

$$P = 1 - \sum_{\alpha} \frac{\omega_{p\alpha}^2}{\omega^2} \quad (\text{C.22})$$

Taking the time derivative of C.18 and plugging it into the curl of C.1 we get

$$\frac{c^2}{\omega^2} \mathbf{k} \times \mathbf{k} \times \mathbf{E} = \mathbf{K} \cdot \mathbf{E} \quad (\text{C.23})$$

$$\mathbf{n} \times \mathbf{n} \times \mathbf{E} = \mathbf{K} \cdot \mathbf{E}, \quad (\text{C.24})$$

where $\mathbf{n} = c\mathbf{k}/\omega$. We can then rewrite C.24 as the matrix equation

$$\begin{pmatrix} S - n_{\parallel}^2 & -\imath D & n_{\parallel} n_{\perp} \\ \imath D & S - n_{\parallel}^2 - n_{\perp}^2 & 0 \\ n_{\parallel} n_{\perp} & 0 & P - n_{\perp}^2 \end{pmatrix} \begin{pmatrix} E_x \\ E_y \\ E_z \end{pmatrix} = 0 \quad (\text{C.25})$$

where n_{\parallel} and n_{\perp} refer to the orientation relative to the magnetic field.

In order for this set of equations to be consistent, the determinant of the coefficient matrix must equal zero, i.e.

$$(S - n_{\parallel}^2)(S - n_{\parallel}^2 - n_{\perp}^2)(P - n_{\perp}^2) - (n_{\parallel} n_{\perp})^2(S - n_{\parallel}^2 - n_{\perp}^2) - D^2(P - n_{\perp}^2) = 0, \quad (\text{C.26})$$

and this is our cold plasma dispersion relation.

Note, C.26 is quadratic in n^2 , therefore we can represent it as

$$An_{\perp}^4 - Bn_{\perp}^2 + C = 0 \quad (\text{C.27})$$

where

$$A = S$$

$$B = PS + S^2 - D^2 - n_{\parallel}^2(P + S)$$

$$C = Pn_{\parallel}^4 - 2PSn_{\parallel}^2 - D^2P + PS^2$$

This has solutions of

$$\boxed{n_{\perp}^2 = \frac{B \pm F}{2A}} \quad (\text{C.28})$$

where $F = \sqrt{B^2 - 4AC}$.

C.2 Simplifying the Cold plasma dispersion Elements

For further analytic analysis, it is helpful to consider simplifications of the Cold Plasma Elements for MST parameters. We define MST parameters to be

<i>Quantity</i>	<i>order</i>
$\frac{\omega}{\omega_{pe}}$	ϵ
$\frac{\omega}{\omega_{ce}}$	ϵ
$\frac{m_e}{m_i}$	ϵ^5

This is roughly true throughout MST discharges, except in the very edge where $\omega_{pe} \rightarrow 0$. This means near the $P = 0$ cutoff our approximations breakdown. In practice however, the width of this region is very small and unimportant except when considering antenna coupling.

Under these approximations, our cold plasma elements reduce as follows:

$$\begin{aligned}
S &= 1 - \frac{\omega_{pe}^2}{\omega^2 - \omega_{ce}^2} - \frac{\omega_{pi}^2}{\omega^2 - \omega_{ce}^2} \\
&= 1 - \frac{\omega_{pe}^2}{\omega_{ce}^2 \left(\frac{\omega^2}{\omega_{ce}^2} - 1 \right)} - \frac{\omega_{pi}^2}{\omega^2 \left(1 - \frac{m_e \omega_{ce}^2}{m_i \omega^2} \right)} \\
&= 1 + \frac{\omega_{pe}^2}{\omega_{ce}^2 \left(1 - \frac{\omega^2}{\omega_{ce}^2} \right)} - \frac{\omega_{pi}^2}{\omega^2 \left(1 - \frac{m_e \omega_{ce}^2}{m_i \omega^2} \right)} \\
&\approx 1 + \frac{\omega_{pe}^2}{\omega_{ce}^2} \left(1 + \frac{\omega^2}{\omega_{ce}^2} + \mathcal{O}(\epsilon^4) \right) - \frac{\omega_{pi}^2}{\omega^2} \left(1 + \frac{m_e \omega_{ce}^2}{m_i \omega^2} + \mathcal{O}(\epsilon^6) \right) \\
&= 1 + \frac{\omega_{pe}^2}{\omega_{ce}^2} + \frac{\omega_{pe}^2 \omega^2}{\omega_{ce}^4} - \frac{\omega_{pi}^2}{\omega^2} - \frac{m_e \omega_{pi}^2 \omega_{ce}^2}{m_i \omega^4} + \mathcal{O}(\epsilon^5) \\
&= 1 + \underbrace{\frac{\omega_{pe}^2}{\omega_{ce}^2}}_{\mathcal{O}(\epsilon)} + \underbrace{\frac{\omega_{pe}^2 \omega^4}{\omega^2 \omega_{ce}^4}}_{\mathcal{O}(\epsilon^2)} - \underbrace{\frac{m_e \omega_{pe}^2}{m_i \omega^2}}_{\mathcal{O}(\epsilon^3)} - \underbrace{\frac{m_e^2 \omega_{pe}^2 \omega_{ce}^2}{m_i^2 \omega^4}}_{\mathcal{O}(\epsilon^6)} \\
S &\approx 1 + \frac{\omega_{pe}^2}{\omega_{ce}^2}
\end{aligned}$$

$$\begin{aligned}
D &= \frac{\omega_{ci} \omega_{pi}^2}{\omega (\omega^2 - \omega_{ci}^2)} - \frac{\omega_{ce} \omega_{pe}^2}{\omega (\omega^2 - \omega_{ce}^2)} \\
&= \frac{\omega_{ci} \omega_{pi}^2}{\omega^3 \left(1 - \frac{\omega_{ci}^2}{\omega^2} \right)} - \frac{\omega_{ce} \omega_{pe}^2}{\omega \omega_{ce}^2 \left(\frac{\omega^2}{\omega_{ce}^2} - 1 \right)} \\
&\approx \frac{\omega_{ci} \omega_{pi}^2}{\omega} \left(1 + \frac{\omega_{ci}^2}{\omega^2} + \mathcal{O}(\epsilon^3) \right) + \frac{\omega_{ce} \omega_{pe}^2}{\omega \omega_{ce}} \left(1 + \frac{\omega}{\omega_{ce}} + \mathcal{O}(\epsilon^4) \right) \\
&= \underbrace{\left(\frac{m_e}{m_i} \right)^2 \frac{\omega_{ce}}{\omega} \frac{\omega_{pe}^2}{\omega^2}}_{\mathcal{O}(\epsilon^7)} + \underbrace{\left(\frac{m_e}{m_i} \right)^4 \frac{\omega_{ce}^3}{\omega^3} \frac{\omega_{pe}^2}{\omega^2}}_{\mathcal{O}(\epsilon^{15})} + \underbrace{\frac{\omega_{pe}^2}{\omega^2} \frac{\omega}{\omega_{ce}}}_{\mathcal{O}(\frac{1}{\epsilon})} + \underbrace{\frac{\omega_{pe}^2}{\omega^2} \frac{\omega^3}{\omega_{ce}^3}}_{\mathcal{O}(\epsilon)} \\
D &\approx \frac{\omega_{pe}^2}{\omega \omega_{ce}}
\end{aligned}$$

$$\begin{aligned}
P &= 1 - \frac{\omega_{pe}^2}{\omega^2} - \frac{\omega_{pi}^2}{\omega^2} \\
&= 1 - \left(1 + \frac{m_e}{m_i}\right) \frac{\omega_{pe}^2}{\omega^2} \\
P &\approx -\frac{\omega_{pe}^2}{\omega^2}
\end{aligned}$$

So we end up with these approximations:

Cold Plasma Term	Order	
$S \approx 1 + \frac{\omega_{pe}^2}{\omega_{ce}^2}$	1	(C.29)
$D \approx \frac{\omega_{pe}^2}{\omega\omega_{ce}}$	$1/\epsilon$	(C.30)
$P \approx -\frac{\omega_{pe}^2}{\omega^2}$	$1/\epsilon^2$	(C.31)

Appendix D: Electrostatic Dispersion Relation

A simpler dispersion function can be derived by imposing the Electrostatic condition, ($\mathbf{E} = -\nabla\phi = -i\mathbf{k}\phi$). Combining this restriction with Gauss' Law produces

$$\nabla \cdot \mathbf{K} \cdot \mathbf{E} = 0 \quad (\text{D.1})$$

$$\nabla \cdot \mathbf{K} \cdot \mathbf{k}\phi = 0 \quad (\text{D.2})$$

$$\mathbf{k} \cdot \mathbf{K} \cdot \mathbf{k}\phi = 0 \quad (\text{D.3})$$

Of course in general $\phi \neq 0$ so we can simply say $\mathbf{k} \cdot \mathbf{K} \cdot \mathbf{k} = 0$. In order to examine the region of validity for this approximation, we must take a closer look at the Electric field. First we break the electric field up into longitudinal and transverse components

$$\mathbf{E} = \mathbf{E}_l + \mathbf{E}_t \quad (\text{D.4})$$

$$= (\mathbf{n} \cdot \mathbf{E}) \frac{\mathbf{n}}{n^2} - \frac{\mathbf{n} \times (\mathbf{n} \times \mathbf{E})}{n^2} \quad (\text{D.5})$$

Next, take the dot product of \mathbf{n} and C.24.

$$\mathbf{n} \cdot \mathbf{K} \cdot (\mathbf{E}_l + \mathbf{E}_t) = 0 \quad (\text{D.6})$$

$$\mathbf{n} \cdot \mathbf{K} \cdot \left(\frac{(\mathbf{n} \cdot \mathbf{E}) \mathbf{n}}{n^2} + \mathbf{E}_t \right) = 0 \quad (\text{D.7})$$

Notice for $|\mathbf{E}_t| \ll |\mathbf{E}_l|$ this simplifies to

$$\mathbf{n} \cdot \mathbf{K} \cdot \frac{(\mathbf{n} \cdot \mathbf{E}) \mathbf{n}}{n^2} = 0 \quad (\text{D.8})$$

and since $\mathbf{n} \cdot \mathbf{E}$ is simply a number, this is identical to our earlier dispersion relation, $\mathbf{k} \cdot \mathbf{K} \cdot \mathbf{k} = 0$.

Plugging our full expressions for \mathbf{E}_l and \mathbf{E}_t into C.24 we get

$$-n^2 \mathbf{E}_t + \mathbf{K} \cdot \mathbf{E}_t + \mathbf{K} \cdot \mathbf{E}_l = 0 \quad (\text{D.9})$$

$$\mathbf{K} \cdot \mathbf{E}_l = n^2 \mathbf{E}_t - \mathbf{K} \cdot \mathbf{E}_t \quad (\text{D.10})$$

This implies for a given \mathbf{E}_1 , a very large n^2 corresponds to a very small \mathbf{E}_t . This then becomes our applicability condition for the Electrostatic dispersion relation. As shown Figure D.1, n_\perp (and therefore n) is very large for the slow wave when it is not near a cutoff or the mode conversion layer.

Using this dispersion relation and our approximations, we get a simple form for wave propagation in the bulk of the plasma, namely

$$\mathbf{k} \cdot \mathbf{K} \cdot \mathbf{k} = k_\perp^2 S + k_\parallel^2 P = 0 \quad (\text{D.11})$$

i.e.

$$\frac{k_\perp^2}{k_\parallel^2} = -\frac{P}{S} \quad (\text{D.12})$$

$$= \frac{\omega_{pe}^2}{\omega^2 \left(1 + \frac{\omega_{pe}^2}{\omega_{ce}^2}\right)} \quad (\text{D.13})$$

Taking $\omega_{ce}^2/\omega_{pe}^2 \ll 1$, we can further simplify this

$$\frac{k_\perp^2}{k_\parallel^2} = \frac{\omega_{pe}^2}{\omega^2 \left(1 + \frac{\omega_{pe}^2}{\omega_{ce}^2}\right)} \quad (\text{D.14})$$

$$= \frac{\omega_{pe}^2}{\omega^2 \frac{\omega_{pe}^2}{\omega_{ce}^2} \left(1 + \frac{\omega_{ce}^2}{\omega_{pe}^2}\right)} \quad (\text{D.15})$$

$$\approx \frac{\omega_{ce}^2}{\omega^2} \left(1 - \frac{\omega_{ce}^2}{\omega_{pe}^2} + \mathcal{O}\left(\frac{\omega_{ce}^4}{\omega_{pe}^4}\right)\right) \quad (\text{D.16})$$

$$= \frac{\omega_{ce}^2}{\omega^2} - \mathcal{O}\left(\frac{\omega_{ce}^4}{\omega_{pe}^4}\right) \quad (\text{D.17})$$

$$\frac{k_\perp^2}{k_\parallel^2} \approx \frac{\omega_{ce}^2}{\omega^2}. \quad (\text{D.18})$$




Figure D.1: Perpendicular wave number versus density for a fixed magnetic field of 0.25 T and a parallel wave number of 10. This value is quite large, implying that the slow wave is mostly electrostatic.

For MST like parameters, this gives $k_{\perp} \approx (7 - 15)k_{\parallel}$. This simplified dispersion relation can be used to derive the phase and group velocity.

$$\mathbf{V}_{\phi} = \frac{\omega}{k} \hat{k} = \frac{\omega}{k} \left(\frac{k_{\parallel}}{k} \hat{\mathbf{b}} + \frac{k_{\perp}}{k} \hat{\mathbf{r}} \right) \quad (\text{D.19})$$

$$= \frac{\omega}{k_{\parallel}^2 + k_{\perp}^2} \left(k_{\parallel} \hat{\mathbf{b}} + k_{\perp} \hat{\mathbf{r}} \right) \quad (\text{D.20})$$

$$= \frac{\omega}{k_{\parallel}^2 + \frac{\omega_{ce}^2}{\omega^2} k_{\parallel}^2} \left(k_{\parallel} \hat{\mathbf{b}} + \frac{\omega_{ce}}{\omega} k_{\parallel} \hat{\mathbf{r}} \right) \quad (\text{D.21})$$

$$= \frac{\omega}{k_{\parallel}} \frac{1}{\left(1 + \frac{\omega_{ce}^2}{\omega^2} \right)} \left(\hat{\mathbf{b}} + \frac{\omega_{ce}}{\omega} \hat{\mathbf{r}} \right) \quad (\text{D.22})$$

$$= \frac{\omega^3}{k_{\parallel} \omega_{ce}^2} \frac{1}{\left(1 + \frac{\omega^2}{\omega_{ce}^2} \right)} \left(\hat{\mathbf{b}} + \frac{\omega_{ce}}{\omega} \hat{\mathbf{r}} \right) \quad (\text{D.23})$$

$$= \frac{\omega^2}{k_{\parallel} \omega_{ce}} (1 - \mathcal{O}(\epsilon^2)) \left(\frac{\omega}{\omega_{ce}} \hat{\mathbf{b}} + \hat{\mathbf{r}} \right) \quad (\text{D.24})$$

$$\approx \frac{c}{n_{\parallel}} \frac{\omega}{\omega_{ce}} \left(\frac{\omega}{\omega_{ce}} \hat{\mathbf{b}} + \hat{\mathbf{r}} \right), \quad (\text{D.25})$$

Notice, when deriving the group velocity, care must be take to ensure that the relative signs remain correct in so far as our dispersion relation can reveal them. This is best done by considering

$$\frac{\partial \omega}{\partial k_x} = \frac{1}{2\omega} \frac{\partial (\omega^2)}{\partial k_x} \quad (\text{D.26})$$

so that

$$\mathbf{V}_{\mathbf{g}} \equiv \nabla_{k\omega} = \frac{1}{2\omega} \left(\frac{\partial \omega^2}{\partial k_{\parallel}} \hat{\mathbf{b}} + \frac{\partial \omega^2}{\partial k_{\perp}} \hat{\mathbf{r}} \right) \quad (\text{D.27})$$

$$= \frac{1}{\omega} \left(\frac{k_{\parallel} \omega_{ce}^2}{k_{\perp}^2} \hat{\mathbf{b}} - \frac{k_{\parallel}^2 \omega_{ce}^2}{k_{\perp}^3} \hat{\mathbf{r}} \right) \quad (\text{D.28})$$

$$= \frac{1}{\omega} \left(\frac{k_{\parallel} \omega_{ce}^2}{k_{\perp}^2} \hat{\mathbf{b}} - \frac{k_{\parallel}^2 \omega_{ce}^2}{k_{\perp}^3} \hat{\mathbf{r}} \right) \quad (\text{D.29})$$

$$= \frac{c}{n_{\parallel}} \left(\hat{\mathbf{b}} \pm \frac{\omega}{\omega_{ce}} \hat{\mathbf{r}} \right). \quad (\text{D.30})$$

We accept only the solution which has a negative radial component since the initial wave launch will be at a cutoff (the positive $\hat{\mathbf{r}}$ direction would be outward going, but it cannot go outward and therefore does not exist)

Appendix E: The raytracing equations

This work has made extensive use of the raytracing equations. These equations appear somewhat mysterious but can be derived quite easily. Recall the Dispersion relation \mathfrak{D} (which every one is used) is a function of $(\omega, \mathbf{k}, \mathbf{x})$ and is equal to zero everywhere. Here, the discussion is restricted to one dimension for concision, but the derivation is identical for three dimensions.

Looking at the differential of \mathfrak{D} ,

$$d\mathfrak{D} = \frac{\partial \mathfrak{D}}{\partial \omega} d\omega + \frac{\partial \mathfrak{D}}{\partial k} dk + \frac{\partial \mathfrak{D}}{\partial x} dx \quad (\text{E.1})$$

This means that the full derivative of \mathfrak{D} with relation to k is

$$\frac{d\mathfrak{D}}{dk} = \frac{\partial \mathfrak{D}}{\partial \omega} \frac{d\omega}{dk} + \frac{\partial \mathfrak{D}}{\partial k} + \frac{\partial \mathfrak{D}}{\partial x} \frac{dx}{dk} \quad (\text{E.2})$$

Because \mathfrak{D} is equal to zero everywhere, any and all *full* derivatives are also equal to zero everywhere. Solving this equation for $d\omega/dk$ (recall this is identical to group velocity, dx/dt)

$$\frac{d\omega}{dk} = \frac{dx}{dt} = -\frac{\partial \mathfrak{D}/\partial k}{\partial \mathfrak{D}/\partial \omega} - \frac{\partial \mathfrak{D}/\partial x}{\partial \mathfrak{D}/\partial \omega} \frac{dx}{dk} \quad (\text{E.3})$$

Next, apply the geometric optics assumption, that the wavelength varies slowly in space, specifically $dx/dk \approx 0$. This leaves us with the group velocity equation,

$$\frac{d\omega}{dk} = \frac{dx}{dt} = -\frac{\partial \mathfrak{D}/\partial k}{\partial \mathfrak{D}/\partial \omega} \quad (\text{E.4})$$

In order to study the evolution of the wavenumber, consider the full derivative of \mathfrak{D} with respect to time. This gives

$$\frac{d\mathfrak{D}}{dt} = \frac{\partial \mathfrak{D}}{\partial \omega} \frac{d\omega}{dt} + \frac{\partial \mathfrak{D}}{\partial k} \frac{dk}{dt} + \frac{\partial \mathfrak{D}}{\partial x} \frac{dx}{dt} = 0 \quad (\text{E.5})$$

Finally, we assume the frequency does not change during propagation. Plugging in the previous derived value of group velocity, we get

$$\frac{dk}{dt} = \frac{\partial \mathfrak{D}/\partial x}{\partial \mathfrak{D}/\partial \omega} \quad (\text{E.6})$$

Appendix F: Lower Hybrid coupling

This derivation follows Golant. [105] Starting with the cold plasma dielectric tensor (C.19) and the Helmholtz equation (derived by taking the time derivative of Ampere's law and plugging it into Faraday's law), we see that the total expression of the fluctuation fields in a plasma is given by

$$\nabla \times \nabla \times \mathbf{E} + \frac{1}{c^2} \frac{\partial^2}{\partial t^2} (\underline{\epsilon} \cdot \mathbf{E}) = 0 \quad (\text{F.1})$$

Now we assume $\mathbf{E} = \mathbf{E}_0(\mathbf{x})e^{ikx - i\omega t}$. We also assume the dielectric tensor is not time dependent - this means we cannot account for pondermotive effects. The Helmholtz equation (F.1) then simplifies to

$$x : -\frac{c^2 k^2}{\omega^2} E_x - i \frac{c^2 k}{\omega^2} \frac{\partial}{\partial x} E_z + S E_x - i D E_y = 0 \quad (\text{F.2})$$

$$y : -\frac{c^2 k^2}{\omega^2} E_y + \frac{c^2}{\omega^2} \frac{\partial^2}{\partial x^2} E_y + S E_y + i D E_x = 0 \quad (\text{F.3})$$

$$z : \frac{c^2}{\omega^2} \frac{\partial^2}{\partial x^2} E_z - i \frac{c^2 k}{\omega^2} \frac{\partial}{\partial x} E_x + P E_z = 0 \quad (\text{F.4})$$

Define the variable $\chi = \frac{\omega}{c} x$. Notice $\partial \chi = \frac{\omega}{c} \partial x$, i.e. $\frac{c}{\omega} \frac{\partial}{\partial x} = \frac{\partial}{\partial \chi}$ and $\frac{c^2}{\omega^2} \frac{\partial^2}{\partial x^2} = \frac{\partial^2}{\partial \chi^2}$. Finally, recall that $n \equiv \frac{ck}{\omega}$. So our equations become

$$x : -n^2 E_x - i n \frac{\partial}{\partial \chi} E_z + S E_x - i D E_y = 0, \quad (\text{F.5})$$

$$y : -n^2 E_y + \frac{\partial^2}{\partial \chi^2} E_y + S E_y + i D E_x = 0, \quad (\text{F.6})$$

$$z : \frac{\partial^2}{\partial \chi^2} E_z - i n \frac{\partial}{\partial \chi} E_x + P E_z = 0. \quad (\text{F.7})$$

The vacuum solution ($S = P = 1, D = 0$) is

$$E_z = C \text{Exp} \sqrt{n^2 - 1} \chi + D \text{Exp} -\sqrt{n^2 - 1} \chi \quad (\text{F.8})$$

$$E_x = \frac{-i n}{\sqrt{n^2 - 1}} \left(C \text{Exp} \sqrt{n^2 - 1} \chi - D \text{Exp} \sqrt{n^2 - 1} \chi \right) \quad (\text{F.9})$$

where we have imposed a wave polarization of $E_y = 0$.

The electric field of the wave within the plasma is given in general by

$$E_z = Ae^{m_{\perp}\chi} + Be^{-m_{\perp}\chi} \quad (\text{F.10})$$

However we can safely set $A = 0$ because the group velocity must be positive in order for the wave to continue into the plasma, and the phase velocity is opposite the group velocity, as will be shown in section D, so we must discard the wave with the positive group velocity.

The wave impedance at the plasma surface can be now be calculated (applying Faraday's law and F.5, and normalizing by a factor of $\mu_0 c$ for concision)

$$Z_p \equiv \frac{E_z}{H_y} = -\frac{n_{\parallel}^2 - S}{S} \frac{E_z}{E'_z} = -\frac{n_{\parallel}^2 - S}{n_{\perp} S} \quad (\text{F.11})$$

The impedance within the vacuum region is

$$Z = -\sqrt{n_{\parallel}^2 - 1} \frac{C + De^{-2\sqrt{n^2-1}\chi}}{C - De^{-2\sqrt{n^2-1}\chi}} \quad (\text{F.12})$$

Placing the antenna at $-\chi_a$, we use the plasma impedance and the antenna impedance as boundary conditions

$$Z_p = -\sqrt{n_{\parallel}^2 - 1} \frac{C + D}{C - D} \quad (\text{F.13})$$

$$Z_a = -\sqrt{n_{\parallel}^2 - 1} \frac{C + De^{2\sqrt{n^2-1}\chi_a}}{C - De^{2\sqrt{n^2-1}\chi_a}} \quad (\text{F.14})$$

Solving this first of these equations for D, and expanding for small ζ where

$$\zeta \equiv \frac{Z_p}{\sqrt{n_{\parallel}^2 - 1}} \quad (\text{F.15})$$

we finally arrive at the transcendental equation

$$\sqrt{n_{\parallel}^2 - 1} = Z_a \frac{1 + e^{-2\sqrt{n_{\parallel}^2-1}l_v\omega/c}}{1 - e^{-2\sqrt{n_{\parallel}^2-1}l_v\omega/c}} \left[1 - \frac{4Z_p l_s \omega e^{-2\sqrt{n_{\parallel}^2-1}l_v\omega/c}}{c\sqrt{n_{\parallel}^2-1}l_v\omega \left(1 - e^{-2\sqrt{n_{\parallel}^2-1}l_v\omega/c} \right)} \right] \quad (\text{F.16})$$

Finally we linearizing in terms of n and solve for the perturbed portion

$$n_{\parallel 1} = \frac{4(n_{\parallel 0}^2 - 1)^{1/2}}{n_{\parallel 0}} Z_p F(w_0) = -\frac{4(n_{\parallel 0}^2 - 1)^{1/2}}{n_{\parallel 0}} \frac{n_{\parallel 0}^2 - S}{n_{\perp} S} F(w_0) \quad (\text{F.17})$$

where

$$F(w_0) = \frac{e^{-2w_0}}{(1 - e^{-4w_0}) + 4w_0 e^{-2w_0}}, w_0 = \sqrt{n_{\parallel 0}^2 - 1} \omega l_v / c \quad (\text{F.18})$$

Our damping parameter is given by

$$|E| \sim \text{Exp}[-\theta/2\alpha_{coup}] = \text{Exp}\left[\frac{icn_{\parallel 1}r_{ap}\theta}{\omega}\right] \quad (\text{F.19})$$

i.e.

$$\alpha_{coup} = \frac{\omega}{cn_{\parallel 1}r_{ap}} = \frac{\omega n_{\parallel 0} n_{\perp} S}{4cr_{ap} \left(n_{\parallel 0}^2 - 1\right)^{1/2} \left(n_{\parallel 0}^2 - S\right) F(w_0)} \quad (\text{F.20})$$

which can be plugged into equation 2.15 to calculate a radiated power fraction.

Appendix G: Some key error analysis

Key coupling parameters used in this work are

$$\alpha_{coup} = \frac{\alpha_{ap}}{\ln\left(\frac{F-R}{T}\right) - \alpha_{ant}/\alpha_{ohm}} \quad (G.1)$$

$$P_1 = (F - R)e^{-\theta_{ap1}/\alpha_{ohm}} \quad (G.2)$$

$$P_2 = (F - R)e^{-(\theta_{ap2}/\alpha_{ohm} - \alpha_{ap}/\alpha_{coup})} \quad (G.3)$$

$$\begin{aligned} P_{ohm} = & (F - R)(1 - e^{-\theta_{ap1}/\alpha_{ohm}}) \\ & + P_1 \frac{\alpha_{coup}}{\alpha_{coup} + \alpha_{ohm}} (1 - \text{Exp}[-\alpha_{ap}(1/\alpha_{ohm} + 1/\alpha_{coup})]) \\ & + P_2 (1 - e^{-(\alpha_{ant} - \theta_{ap2})/\alpha_{ohm}}) \end{aligned} \quad (G.4)$$

$$P_{rad} = P_1 \frac{\alpha_{ohm}}{\alpha_{ohm} + \alpha_{coup}} (1 - \text{Exp}[-\alpha_{ap}(1/\alpha_{ohm} + 1/\alpha_{coup})]) \quad (G.5)$$

The uncertainty in these values is not easily seen, and so the derivations are walked through here.

First, α_{coup} is a function of F, R, T, and α_{ohm} . Thus

$$\frac{\partial \alpha_{coup}}{\partial F} = -\frac{\alpha_{ap}}{[\ln\left(\frac{F-R}{T}\right) - \alpha_{ant}/\alpha_{ohm}]^2} \frac{1}{F - R} = -\frac{\alpha_{coup}^2}{(F - R)\alpha_{ap}} \quad (G.6)$$

$$\frac{\partial \alpha_{coup}}{\partial R} = \frac{\alpha_{coup}^2}{(F - R)\alpha_{ap}} \quad (G.7)$$

$$\frac{\partial \alpha_{coup}}{\partial T} = \frac{\alpha_{coup}^2}{T\alpha_{ap}} \quad (G.8)$$

$$\frac{\partial \alpha_{coup}}{\partial \alpha_{ohm}} = \frac{\alpha_{coup}^2}{\alpha_{ap}} \frac{\alpha_{ant}}{\alpha_{ohm}^2} \quad (G.9)$$

Therefore

$$\delta \alpha_{coup} = \alpha_{coup}^2 \sqrt{\left(\frac{\delta F}{(F - R)\alpha_{ap}}\right)^2 + \left(\frac{\delta R}{(F - R)\alpha_{ap}}\right)^2 + \left(\frac{\delta T}{T\alpha_{ap}}\right)^2 + \left(\frac{\delta \alpha_{ohm} \alpha_{ant}}{\alpha_{ap} \alpha_{ohm}^2}\right)^2} \quad (G.10)$$

For the intermediate calculated powers,

$$\delta P_1 = e^{-\theta_{ap1}/\alpha_{ohm}} \sqrt{\delta F^2 + \delta R^2 + (F - R)^2 \left(\frac{\theta_{ap1}}{\alpha_{ohm}^2} \right)^2} \delta \alpha_{ohm}^2 \quad (G.11)$$

$$\begin{aligned} \delta P_2 = & \text{Exp} [- (\theta_{ap2}/\alpha_{ohm} + \alpha_{ap}/\alpha_{coup})] \\ & \times \sqrt{\delta F^2 + \delta R^2 + (F - R)^2 \left(\frac{\theta_{ap2}}{\alpha_{ohm}^2} \right)^2 \delta \alpha_{ohm}^2 + (F - R)^2 \left(\frac{\alpha_{ap}}{\alpha_{coup}^2} \right)^2 \delta \alpha_{coup}^2} \end{aligned} \quad (G.12)$$

For P_{rad} ,

$$\frac{\partial P_{rad}}{\partial P_1} = \frac{P_{rad}}{P_1} \quad (G.13)$$

$$\begin{aligned} \frac{\partial P_{rad}}{\partial \alpha_{ohm}} = & -P_1 \frac{\alpha_{ap}}{\alpha_{ohm}(\alpha_{coup} + \alpha_{ohm})} \text{Exp}[-\alpha_{ap}(1/\alpha_{coup} + 1/\alpha_{ohm})] \\ & - P_1 \frac{\alpha_{ohm}}{(\alpha_{coup} + \alpha_{ohm})^2} (1 - \text{Exp}[-\alpha_{ap}(1/\alpha_{coup} + 1/\alpha_{ohm})]) \\ & \frac{P_1}{(\alpha_{coup} + \alpha_{ohm})} (1 - \text{Exp}[-\alpha_{ap}(1/\alpha_{coup} + 1/\alpha_{ohm})]) \end{aligned} \quad (G.14)$$

$$\begin{aligned} \frac{\partial P_{rad}}{\partial \alpha_{coup}} = & -P_1 \frac{\alpha_{ap}\alpha_{ohm}}{\alpha_{coup}^2(\alpha_{coup} + \alpha_{ohm})} \text{Exp}[-\alpha_{ap}(1/\alpha_{coup} + 1/\alpha_{ohm})] \\ & - P_1 \frac{\alpha_{ohm}}{(\alpha_{coup} + \alpha_{ohm})^2} (1 - \text{Exp}[-\alpha_{ap}(1/\alpha_{coup} + 1/\alpha_{ohm})]) \end{aligned} \quad (G.15)$$

It is not instructive to write this complete expression out, but of course

$$\delta P_{rad} = \sqrt{\left(\frac{\partial P_{rad}}{\partial P_1} \right)^2 \delta P_1^2 + \left(\frac{\partial P_{rad}}{\partial \alpha_{ohm}} \right)^2 \delta \alpha_{ohm}^2 + \left(\frac{\partial P_{rad}}{\partial \alpha_{coup}} \right)^2 \delta \alpha_{coup}^2} \quad (G.16)$$

Finally, P_{ohm} . Since this has a complicated dependence on several not-independent quantities, we break it down and analyze it in terms of smaller variables, specifically

$$\begin{aligned} \Omega_1 &= (F - R)(1 - e^{-\theta_{ap1}/\alpha_{ohm}}) \\ \Omega_2 &= \frac{\alpha_{coup}}{\alpha_{coup} + \alpha_{ohm}} (1 - \text{Exp}[-\alpha_{ap}(1/\alpha_{ohm} + 1/\alpha_{coup})]) \\ \Omega_3 &= (1 - e^{-(\alpha_{ant} - \theta_{ap2})/\alpha_{ohm}}) \\ \delta P_{ohm} &= \sqrt{\delta \Omega_1^2 + P_1^2 \delta \Omega_2^2 + \Omega_2^2 \delta P_1^2 + P_2^2 \delta \Omega_3^2 + \Omega_3^2 \delta P_2^2} \end{aligned} \quad (G.17)$$

Proceeding,

$$\begin{aligned}\frac{\partial \Omega_1}{\partial F} &= (1 - e^{-\theta_{ap1}/\alpha_{ohm}}) \\ \frac{\partial \Omega_1}{\partial R} &= -(1 - e^{-\theta_{ap1}/\alpha_{ohm}}) \\ \frac{\partial \Omega_1}{\partial \alpha_{ohm}} &= \frac{\theta_{ap1}}{\alpha_{ohm}^2} (F - R) e^{-\theta_{ap1}/\alpha_{ohm}}\end{aligned}$$

$$\begin{aligned}\frac{\partial \Omega_2}{\partial \alpha_{coup}} &= \frac{1}{(\alpha_{coup} + \alpha_{ohm})^2} \\ &\quad \left(\alpha_{ohm} - \left[\alpha_{ohm} + \frac{\alpha_{ap}}{\alpha_{coup} (\alpha_{coup} + \alpha_{ohm})} \right] \text{Exp}[-\alpha_{ap} (1/\alpha_{coup} + 1/\alpha_{ohm})] \right) \\ \frac{\partial \Omega_2}{\partial \alpha_{ohm}} &= \frac{-\alpha_{ohm}^2 + (\alpha_{ohm}^2 - \alpha_{ap}\alpha_{coup} - \alpha_{ap}\alpha_{ohm}) \text{Exp}[-\alpha_{ap} (1/\alpha_{coup} + 1/\alpha_{ohm})]}{\alpha_{ohm}^2 (\alpha_{coup} + \alpha_{ohm})^2} \alpha_{coup}\end{aligned}$$

and

$$\frac{\partial \Omega_3}{\partial \alpha_{ohm}} = \frac{\alpha_{ant} - \theta_{ap2}}{\alpha_{ohm}^2} e^{[-(\alpha_{ant} - \theta_{ap2})/\alpha_{ohm}]}$$

The convenience of having done error analysis in this way is that the backplane measurements or the 0-D calculations can be used to determine α_{coup} , and the error analysis holds either way.