FAST ION STUDIES OF ION CYCLOTRON HEATING IN THE PLT TOKAMAK

Gregory Wayne Hammett

A DISSERTATION

PRESENTED TO THE

FACULTY OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE BY THE

DEPARTMENT OF

ASTROPHYSICAL SCIENCES

March 1986

Note: This online version of my thesis is missing the figures (which were put in by cut-n-paste in the good old days) and some appendix information. Some of the page numbering is also slightly different than the original 1986 version, which can be obtained by contacting the author or University Microfilms, www.umi.com.

Abstract

Fast Ion Studies of Ion Cyclotron Heating in the PLT Tokamak

Gregory Wayne Hammett

The most promising method for heating a tokamak reactor plasma to thermonuclear temperatures is the use of waves in the ion cyclotron range of frequencies. Measurements of the fast non-Maxwellian ions which are produced by this heating method provide a wealth of information about the physics of wave heating. Previous experiments have demonstrated that ion cyclotron heating tends to produce energetic ions whose banana tips are near the resonance layer. Cyclotron heating causes this "resonance localization" by imparting perpendicular energy to particles, and by imparting more energy to particles which spend more time in the resonance layer. A bounce-averaged quasilinear operator which properly includes these effects has been implemented in a Fokker-Planck code in order to make detailed comparisons with measurements. Good agreement is found with data from a horizontally-scanning, mass-resolving, neutral particle analyzer, although the assumed RF power deposition profile needed to match the data is broader than expected in some cases. Alternatively, radial diffusion of fast ions (which is ignored in the code) may make the RF power profile appear to be broader than it is. In addition to the usual charge exchange measurements of hydrogen and deuterium, double charge exchange measurements of ^{3}He have been made. Direct second harmonic heating of deuterium or tritium is a preferred technique for a reactor. The transition from hydrogen fundamental heating to deuterium second harmonic heating at low hydrogen concentrations has been clearly demonstrated. An isotropic model fit to the deuterium tail provides a direct measure of the central deuterium power density. These measurements are consistent with sawtooth broadening of the RF power profile and indicate that as much as 20% of the central RF power is directly absorbed by the deuterium in these experiments.

Dedicated to my parents,
Wilford and Mary Ann Hammett,
with thanks.

Acknowledgements

I could not have carried out this research without the help and advice of many people. Sid Medley, Lane Roquemore, George Gammel, and Bob Kaita were primarily responsible for the design, construction, and calibration of the charge exchange neutral analyzer. Bob Shoemaker, Jack Thomas, and Dick Yaeger helped install the analyzer on PLT. Bill Mycock and the PLT crew provided expert technical help on many occasions. Kenny Silber helped write the data analysis software. Tom Gibney provided the data acquisition software.

The PLT tokamak requires a small army of physicists to keep it running and to diagnose its behavior. Manfred Bitter, Ron Bell, Al Cavallo, Sam Cohen, Pat Colestock, Glen Greene, Joel Hosea, David Hwang, Bob Kaita, Steve Kilpatrick, Dennis Manos, Ernesto Mazzucato, Don McNeill, Dennis Mueller, Masayuki Ono, Cynthia Kieras Phillips, Charles Skinner, Jim Strachan, Szymon Suckewer, Schwick von Goeler, and Randy Wilson were all involved in the PLT ICRF experiments during my tenure. Al Cavallo's electron cyclotron emission measurements were particularly helpful for my experiments. Of course, the active interest and support of Joel Hosea were vital to the progress of my thesis.

My interaction with the ICRF group was both enjoyable and profitable. Pat Colestock, David Hwang, and Randy Wilson (who also served as one of my thesis advisors) not only helped run the experiments but also helped understand the results. Much of my day to day guidance came from them (and Bob Kaita). Their knowledge of the physics of ICRF heating is truly impressive, and they suggested many of the experiments done in this thesis. Cynthia Kieras Phillips also helped in modelling some of the ICRF heating experiments.

I had several useful discussions with Gary Kerbel, who explained some of

the finer points of the full bounce averaged quasilinear operator. Rob Goldston guided me through his original Fokker-Planck program which was the starting point for my computer simulations. Rob was also a great sounding board, helping me hear when some of my ideas might be out of tune with reality. I also benefited from useful discussions with Liu Chen, Sam Cohen, Charles Karney, Tom Stix, Jim Strachan, and Harry Towner at Princeton, and C.S. Chang, S.C. Chiu, and Bob Harvey in sunny southern California.

I would also like to take this opportunity to thank those who helped me get started in plasma physics: Ken Hill, Bill Mycock, Jim Strachan, Tom Stix, Kevin McGuire, Jim Glanz, and many others along the way, including Tony Favale and Bill Zarkowski at Grumman Aerospace. The people who gave me the most help over the years, and who deserve the most thanks, are my parents.

Finally, I wish to thank Bob Kaita for his dedicated service as a thesis advisor. His original measurements laid the foundation for this thesis. His time, advice, and guidance, in matters of physics and of life in general, are much appreciated. I could not have asked for a better advisor.

This work was supported in part by the United States Department of Energy Contract Number DE-A-AC02-76-CHO-3073.

Contents

Then the LORD answered Job out of the whirlwind and said,

"Who is this that darkens counsel
By words without knowledge?

"Now gird up your loins like a man,
And I will ask you, and you instruct Me!

"Where were you when I laid the foundation of the earth?
Tell *Me*, if you have understanding,

"Who set its measurements, since you know?
Or who stretched the line on it?

"On what were its bases sunk?
Or who laid its cornerstone,

When the morning stars sang together,
And all the sons of God shouted for joy?"

Job 38:1–7

Chapter 1

Introduction to ICRF Production of Fast Ions

1.1 Motivation

1.1.1 A Glimpse of the Future

A silver DeLorean roars out of nowhere and screeches to a halt. Out climbs Doc Brown, dressed in futuristic metallic garb, muttering, "I gotta have fuel." Taking a banana peel and a can of beer from a garbage can, he opens the trunk of his time-travelling sports car and deposits the garbage into a small device, about the size of a coffee machine, labelled "Mr. Fusion." Fusion energy now powers Doc's car—plutonium is no longer needed.

This fanciful scene from the movie "Back to the Future," written by Robert Zemeckis and Bob Gale, captures the dream of fusion scientists—that cheap and safe fusion will be the energy source of the future. Fusion energy would be much safer than the fission reactors of today, and would not produce any air pollution as fossil fuels do. The real fuel of fusion, deuterium, is practically as plentiful as garbage. The energy crisis of the 70's has faded from our memories, but, like a bad dream, it is sure to haunt us again in the future. If the fusion dream can be made a reality, then perhaps a future superpower war over scarce energy resources may be avoided. But we must keep fusion in perspective. It does not work yet. Even if it does work someday, it can not solve all of the problems of mankind.

1.1.2 A Look at the Past

The quest for controlled fusion energy began in the early 50's, long before the energy crises of the 70's. Although the progress of this quest has been slower than expected, much progress has been made. Many different approaches to fusion have been proposed and are being tested, but the tokamak has been the most successful to date. The tokamak concept was first proposed in 1950 by Igor E. Tamm and Andrei D. Sakharov[1] (the same Sakharov who is now a Soviet dissident). The minimum temperature of 75 million degrees C needed to make fusion work was first obtained in Princeton's PLT tokamak in 1978, while the requisite minimum confinement quality $n\tau > 10^{13}\text{cm}^{-3}\text{sec}$ was first reached in MIT's Alcator tokamak in 1975. The next step is to achieve both high temperature and high confinement quality simultaneously in the the same tokamak. The Princeton TFTR, the European JET, and the Japanese JT-60 tokamaks are designed to try to reach this goal, with the hope of demonstrating marginal fusion energy breakeven by the end of the 1980's. The tasks of igniting a self-burning plasma, and harnessing the resulting energy in an economically attractive way still remain further down the road.

Although it may be possible to heat a high field tokamak to ignition with Ohmic heating alone, a tokamak with more modest magnetic fields needs some form of auxiliary heating to reach ignition temperatures. Neutral beam injection (NBI) is one way to do this, and it has been studied extensively. Recent studies have focussed on using plasma waves, not only to heat the plasma but also to drive steady state currents. This thesis will concentrate on heating by waves in the ion cyclotron range of frequencies (ICRF), where the wave frequency ω resonates with the ion cyclotron frequency Ω (or with $2 \times \Omega$) of one of the ion species in the plasma. The wave energy is directly absorbed by the resonant ions, and then collisionally transferred from the energetic resonant ions to the rest of the plasma. While neutral beams have difficulty penetrating into a dense reactor plasma, ICRF heating should actually improve at higher densities. Access problems in a reactor should also be less severe for ICRF than for NBI. The potential usefulness of ICRF heating has long been recognized, with initial studies carried out in the early days of plasma physics research.[2,3] More recently, high power ICRF heating has been demonstrated on a number of tokamaks, including PLT at Princeton, Alcator C at MIT, TFR at Fontenay-aux-Roses in France, and JFT-2 at

Tokai in Japan. There are a number of good reviews of the present status of experimental and theoretical work on ICRF heating.[4,5]

The best heating results to date with ICRF have been obtained in the PLT tokamak (Fig. 1.1). A central ion temperature of 5 keV (60 million degrees C) was achieved with 4.3 MW of ICRF power. The heating efficiency was similiar to that of NBI.[6] In this case, the majority of the plasma ions were deuterium, while a small amount ($n_{^3\mathrm{He}}/n_e \sim 5 - 10$ %) of ^{3}He was puffed in before the ICRF was turned on. The magnetic field was set to 33 kG so that the 30 MHz ICRF waves would resonate with and heat the ^{3}He. These results are sure to be exceeded in the near future by the European JET tokamak, which will have $\sim$ 15 MW of ICRF power at its command by the summer of 1986.

1.1.3 The Focus of This Thesis

ICRF heating typically produces an energetic non-Maxwellian tail in the resonant ion velocity distribution function $f(\vec{v})$. The general goal of this thesis has been to measure this energetic tail to understand better the physics of ICRF heating. Figure 1.2 shows measurements at three different viewing angles of the energy distribution of fast hydrogen neutrals during hydrogen minority heating in the PLT tokamak. The viewing angles are shown in Fig. 1.3, and the spectra seen at each viewing angle are quite different. Maxwellian distributions ($f \propto \exp(-W/T)$) would be straight lines on the semilog plot of Fig. 1.2 and independent of viewing angle. Since cyclotron damping imparts primarily perpendicular energy to particles, one might naively expect to see the biggest tail in f while viewing perpendicular. Instead, the largest neutral flux is observed while viewing at an intermediate angle between perpendicular and parallel. Another interesting feature of the intermediate viewing angle data is the "negative temperature" region between 25 and 100 keV where $\partial f/\partial W > 0$. These peculiar spectra provide interesting subject material for this thesis.

Kaita et al.[7] showed that the large flux at the intermediate viewing angle is due to energetic trapped ions whose banana tips lie near the resonance layer and which charge-exchange in the high neutral density region at the edge of the plasma (Fig. 1.5). In Sec. 1.2 we will show why ICRF produces these energetic trapped particles with banana tips near the resonance layer, and will

Figure 1.1: Ion and electron temperature as a function of time in the PLT tokamak with 4.3 MW of ICRF heating. Triangles are charge exchange measurements (including a calculated 15% correction for profile effects) and squares are neutron measurements (assuming 30% depletion). $\bar{n}_e = 3.7 \times 10^{13}\text{cm}^{-3}$. Neutron production $\approx 10^{13}$ /sec.

Figure 1.2: Hydrogen energy spectra measured by the charge exchange neutral analyzer at three different viewing angles (shown in Fig. 1.3). Also shown are simulated curves to be discussed in Chapter 4.

Figure 1.3: Top view of the PLT tokamak and the horizontally scanning charge exchange neutral analyzer (CENA). The three CENA viewing angles corresponding to the data of Fig. 1.2 are marked PERP ($R_{\tan}$ = 13 cm), PEAK ($R_{\tan}$ = 70 cm), and PAR ($R_{\tan}$ = 102 cm).

Figure 1.4: Cross-sectional side view of PLT. All CENA sightlines lie in the horizontal midplane.

Figure 1.5: Sample banana orbits whose tips lie in the resonance layer and which charge exchange near the edge of the plasma, giving rise to the large charge exchange flux at the PEAK angle. Also shown are the positions viewed by other sightline angles, given by Eq. (1.2), assuming perfect resonance localization.

explore what may be learned from spectral measurements like Fig. 1.2. The physical ideas presented in Sec. 1.2 will be developed more quantitatively in the bounce averaged quasilinear theory of Chapter 2. This theory has been implemented in a Fokker-Planck computer program to allow detailed comparisons with experiments. These comparisons (shown in Chapters 3 and 4) are used to infer the RF power deposition profile from the fast neutral measurements. (We will find that this procedure is complicated if the tail is so energetic that unconfined orbit losses are important.) In many cases, the inferred RF power profile is significantly broader than theoretically expected. This may be because fast ions produced by the RF near the center of the plasma are quickly spread out over a larger region by some faster-than-neoclassical mechanism such as sawteeth, making the power profile appear to be broader than it really is.

Although second harmonic heating of deuterium or tritium may be preferred for a reactor, past attempts at deuterium second harmonic heating were thwarted by strong hydrogen fundamental absorption. This thesis presents clear evidence of direct second harmonic heating of deuterium. We have also developed an analytic model for the shape of the deuterium tail. This model, along with the full bounce averaged quasilinear program, provides a useful tool for analyzing the deuterium measurements.

We have also demonstrated the feasibility of double charge exchange measurements of ^{3}He minority heating (in a ^{4}He majority plasma). This technique could be used as an alpha particle diagnostic once the alpha particles have slowed down to ~ 400 keV. Our measurements are consistent with the higher efficiency of ^{3}He minority heating over hydrogen minority heating.

1.1.4 *A note about the length of this thesis*

One of my goals in writing this thesis was to avoid sacrificing clarity for brevity. I wanted to provide enough discussion of the assumptions and modelling behind my work that a future graduate student could find my mistakes. I also wanted to describe many of the interesting details I have come across.

As a result, this thesis is rather long, longer than even I expected. For the reader interested in a concise account of the highlights of this thesis, I suggest first reading the preceding section, Sec. 1.1.3, and then the final chapter, Chapter 5. The reader is also encouraged to use the detailed table of contents

to find the parts of the thesis which most interest him. Several aspects of this thesis research have already been reported,[6,8–11] but this dissertation is meant to be complete and self-contained.

1.1.5 The role of fast ions in ICRF heating

In minority ICRF heating, the resonant fast ions are the channel through which the RF power is transferred to the bulk plasma. In second harmonic ICRF, the bulk ions can directly absorb most of the RF power, but a substantial energetic tail may still be produced. These ICRF-produced fast ions play an important role in the propagation and damping of ICRF waves, in the heating of the bulk plasma, and in the transport and stability properties of the bulk plasma. Measurements of the ICRF-produced fast ions provide a useful tool for probing the physics of ICRF heating.

The physics of ICRF heating is interesting on a fundamental level as well. Standard quasilinear theory is only valid for a spectrum of incoherent waves, while ICRF heating is usually done with only a single monochromatic wave (or at most, a few discrete waves). Knowledge of the proper quasilinear theory for the ICRF case is necessary to predict $f(\vec{v})$ for the resonant ions and to calculate such basic quantities as the fusion reactivity enhancement[12,13] and the amount of electron heating vs. ion heating.

The direct effect of ICRF on transport is usually negligible. But there are a number of indirect ways in which ICRF may affect confinement. The average energy of the energetic resonant ions is usually determined by the balance between the input RF power and the collisional losses to the colder, bulk plasma. If the plasma current in the tokamak is not large enough, the energetic resonant ions will have such large banana widths that they will strike the outside limiter. These unconfined orbits not only represent a direct energy loss mechanism, but also enhance the sputtering of impurities from the limiter or wall[14] and may even enhance the inward neoclassical transport rate of impurities.[15] It is important to be able to identify when the tail is too energetic and to know how to control the tail. The TFR group put this loss mechanism to good use by pumping selected impurities out of the plasma.[16] Modelling of this ICRF-driven impurity pump-out has also been undertaken.[17]

The energetic trapped particles produced by ICRF will alter the RF power

deposition profile through finite gyroradius and Doppler-broadening effects, and to a lesser degree, through changes in the wave-particle correlation time. Although this thesis focusses on wave damping effects, the energetic trapped particles may also effect wave propagation itself.[18–20] Because the energetic ions produced by ICRF tend to be trapped to the low magnetic field side of the cyclotron layer, a poloidal electric field must be set up to maintain quasineutrality. There has been a suggestion that the poloidal electric field induced by ICRF may improve confinement while the field induced by electron cyclotron heating may degrade confinement.[21]

A fascinating possible application for ICRF is to try to stabilize the "fishbone" instability and to push a tokamak into the second stability regime of high β.[22–24] By positioning the cyclotron layer near the high field side of the $q = 1$ surface, it should be possibile to produce a large number of barely trapped, energetic particles within the $q = 1$ surface (Fig. 1.6). These barely trapped particles spend most of their time in the good curvature region of the tokamak and tend to stabilize the fishbone. (Recent calculations[25] indicate that it may be necessary to produce barely passing particles to improve stability. This can be accomplished by moving the cyclotron layer further to the high field side and relying on the the Doppler shift to heat ions inside the $q = 1$ surface.) If the cyclotron layer intersects the low field side of the $q = 1$ surface instead, the resulting energetic particles will spend most of their time in the bad curvature region of the tokamak and will tend to destabilize the fishbone. An experimental confirmation that the fishbone instability can be turned off and on by moving the cyclotron layer in and out would be most interesting.

1.2 Resonance Localization

Previous neutral particle measurements in the PLT tokamak indicated that ion cyclotron heating tends to produce energetic particles whose banana tips are near the resonance layer.[7] This effect has been termed "resonance localization",[21] referring to the localization of the banana tips near the resonance layer. In this section we will describe why resonance localization occurs, how it can be used to understand charge exchange spectra like Fig. 1.2, and the measurements on tokamaks which verify the resonance localization property of ICRF.

Figure 1.6: Sample fast ion orbits produced when the resonance layer intersects the (a) bad curvature side or the (b) good curvature side of the $q = 1$ surface.

It should be pointed out from the beginning that the resonance localization phenomenon itself is a simple process that has been known about in some form for a long time. Rothman[26] attributed the drop in coupling at high ICRF power levels in the model C stellarator to resonance localization. He pointed out that the ions gain perpendicular energy from the RF and are no longer able to make it up the magnetic beach, thus reducing the plasma density under the Stix coil and decreasing the loading. ICRF has been used in tandem mirrors, not only to heat the central cell, but also to trap escaping ions in the end plugs.[27,28] Electron cyclotron heating also exhibits the resonance localization effect.[29,30,69] Experimental evidence for ICRF resonance localization in tokamaks will be considered in Sec. 1.2.4 and 1.2.5, and the detailed theory behind it will be discussed in Chapter 2. First, however, we give a physical picture for why resonance localization should occur.

1.2.1 The Causes of Resonance Localization

The magnetic field in a tokamak varies approximately as $B \propto 1/R$, where R is the major radius. The cyclotron frequency is thus not constant in a tokamak, and a particle can resonate with the RF only at the major radius where its cyclotron frequency $\Omega = \frac{eB}{mc}$ equals the rf frequency ω. As a particle moves along a flux surface it will periodically pass through the resonance layer, R_{res}, where $\Omega = \omega$, and will gain or lose a small amount of energy depending on whether it is in or out of phase with the wave. A calculation of the energy kick $\delta W_\perp$ received by the particle is presented in Chapter 2, but all that is important now is that the interaction is confined to a narrow region around the resonance layer at R_{res}. (Chapter 2 will consider the case of general $k_\parallel \neq 0$ with the surprising result that the "resonance localization" phenomenon is independent of $k_\parallel$.)

Two mechanisms are important in producing energetic trapped particles whose banana tips are near the resonance layer. First, cyclotron heating is anisotropic. By this we mean that the wave imparts only perpendicular energy to the ions, while their parallel energy is unaffected (for $k_\parallel = 0$). Each time a particle passes through the resonance layer, its perpendicular velocity may increase. As the RF heats a particle up to high energies, it is also pushing it into the trapped region of velocity space. This process is illustrated by the sequence of orbits and corresponding points in $(v_{\parallel res}, v_{\perp res})$ phase space in

Figure 1.7: Sequence of orbits showing how the ICRF gives particles perpendicular energy, thus turning passing particles into trapped particles and pushing their banana tips closer to the resonance layer.

Fig. 1.7. (Note that $v_{\|}$ and $v_{\perp}$ are evaluated in the resonance layer rather than at the minimum of the magnetic well, which would be more traditional. This was chosen because $v_{\|res}$ is unchanged by the RF and because all of the orbits intersect the resonance layer at the same place.) Consider the initial 1 keV passing orbit in Fig. 1.7. After passing through the resonance layer many times, it may absorb enough perpendicular energy to become a barely trapped orbit. After many more passes through the resonance layer, a particle may be accelerated to even higher energies, pushed deeper into the trapped region of phase space, and its banana tips pushed closer to the resonance layer. Eventually a particle could become so energetic, and its banana width so wide, that it is scraped off by the outer limiter (the 120 keV orbit in Fig. 1.7). Although collisional pitch angle scattering becomes negligible at high energies, collisional drag with the electrons is still important

and usually serves to keep most of the ions from becoming so energetic that they are unconfined. In PLT, it is not difficult to produce a hydrogen tail so energetic that unconfined orbits are the dominant loss mechanism. In a higher current reactor designed to confine 3.5 MeV alpha particles, this should not be a problem.

Notice that the sequence of orbits in fig. 1.7 all intersect the resonance layer in the same place and all have their banana tips on the same flux surface. This is because the ICRF is unable to change $v_{\|res}$ (at least for $k_\| = 0$). From conservation of toroidal angular momentum this means that the ICRF is unable to cause any direct neoclassical-like radial transport. This subtle (and perhaps semantic) distinction between unconfined orbit losses and real radial transport has been the source of some confusion.

The anistropic nature of ICRF is one mechanism which produces energetic ions with banana tips near the resonance layer. The second mechanism is that cyclotron heating is preferential. Not all ions are heated uniformly, but ions which spend more time in the resonance layer will absorb more energy. A trapped particle spends a large fraction of its time near its banana tip, since $v_\| \to 0$ there. Thus a particle whose banana tip is already near the resonance layer will absorb more energy than a particle whose banana tip is far away (Fig. 1.8).

1.2.2 Unfolding Neutral Particle Measurements

By assuming perfect resonance localization, i.e., that all energetic ions have their banana tips exactly in the resonance layer, we are able to unfold the sightline integrated neutral particle measurements to obtain localized information. The pitch angle $\eta = v_\|/v$ accepted at different points along a sightline varies as $\eta = R_{tan}/R$, where R_{tan} is the tangency radius of the sightline, and R is the major radius where a fast neutral is born. This geometry is illustrated in the top view of a sightline through a tokamak in Fig. 1.9. The variation of a particle's pitch angle along a banana orbit can be worked out from conservation of energy W and magnetic moment μ:

$$\frac{1}{2}v_\|^2 = W - \mu B$$

$$\eta = \frac{v_\|}{v} = \pm\sqrt{1 - \frac{\mu B}{W}}$$

Figure 1.8: Particles whose banana tips are already near the resonance layer absorb more energy than a particle with the same initial energy but which quickly passes through the resonance layer.

Figure 1.9: Top view of tokamak and a detector sightline, showing the variation of the pitch angle $v_{\|}/v$ viewed along the sightline.

Taking the usual tokamak approximation $B \propto 1/R$, and knowing that $v_\parallel \to 0$ in the resonance layer where the banana tip is, we can write this as

$$\eta = \pm\sqrt{1 - \frac{R_{res}}{R}}. \tag{1.1}$$

But the neutral particle analyzer only sees particles from major radius R if it has the pitch angle $\eta = \pm R_{tan}/R$. Equating these two expressions for η yields an equation for $R(R_{tan}, R_{res})$ which determines where along a sightline the neutrals originated:

$$\sqrt{1 - \frac{R_{res}}{R}} = \frac{R_{tan}}{R}.$$

The physically meaningful root is

$$R(R_{tan}, R_{res}) = \frac{R_{res}}{2} + \sqrt{(\frac{R_{res}}{2})^2 + R_{tan}^2}. \tag{1.2}$$

With the resonance layer at 139 cm, we find that the $R_{tan} = 70$ cm sightline only accepts orbits which pass through the midplane at a major radius of $R = 168$ cm (2 of these orbits are shown in Fig. 1.5). This is near the edge of the plasma where the neutral density is high, thus producing the large neutral flux observed at $R_{tan} = 70$ cm in Fig. 1.2. On the other hand, the spectra measured at $R_{tan} = 13$ cm is from ions at $R = 140$ cm, near the center of the plasma. This radial information about the energetic ions can then be used to try to infer the radial power profile.

A striking feature of the $R_{tan} = 70$ cm spectra is that the 100 keV flux is 2 times bigger than the 30 keV flux. This "negative temperature" has two possible explanations. The first is simply that the ICRF power at point A in Fig. 1.5 is much higher than at point B, so that there are more 100 keV ions at point A than 30 keV ions at point B. But an alternative explanation is that the ICRF power at point C in Fig. 1.10 is large enough to produce a significant number of 500 keV, barely confined ions. As these ions pass through the cold edge plasma, they experience strong electron drag, but very little pitch angle scattering, and their banana tips move out as they slow down. Above 100 keV, the charge exchange cross section is negligible (Fig. 1.11). But as the ions slow down from 100 to 30 keV, charge exchange losses in the high neutral density edge plasma are large enough to cause the

Figure 1.10: Fast ions at point **A** may either be an indication of a large amount of RF power at point **A**, or of of a large amount of RF power at point **C**, if the fast ions at point **C** slow down in the cold edge plasma and move out in minor radius. The energies of these ions are: 500 keV at **C**, 100 keV at **A**, 30 keV at **B**.

Figure 1.11: Rate coefficient σv for protons charge exchanging with hydrogen atoms, as a function of energy. Taken from Freeman and Jones.[31]

observed 50% drop in flux. This edge drag process is probably the main cause of the "negative temperature" in the charge exchange spectra. However, we will find in Chapter 4 that the bounce averaged quasilinear program (which ignores radial transport processes and therefore this edge drag process) is also able to reproduce the "negative temperature" by putting much more ICRF power at point A than at point B.

1.2.3 Limitations on Resonance Localization

Of course not all ions have their banana tips exactly in the resonance layer, especially at low energies ($<$ 30keV) where pitch angle scattering is very rapid. In fact, the $R_{tan} = 102$ cm sightline views particles which must have pitch angle scattered out of the resonance layer, since Eq. (1.2) yields $R = 193$ cm for $R_{tan} = 102$ cm, which is well outside the vacuum vessel. Our primary goal here is to emphasize the fact that although the neutral particle spectra are sightline measurements, they do contain radial information because most of the high energy particles have their banana tips near the resonance layer. To be more quantitative about trying to unfold this radial information, one must be careful to include the effect of pitch angle scattering which tends to erase radial information. Pitch angle scattering and other important effects are included in the theory and computer program developed in Chapter 2 and used in Chapters 3 and 4. However, the pitch angle of an ion may be changed by sawteeth instabilities, a process ignored in our present computer modelling, as well as by collisions.

1.2.4 Neutral Particle Evidence for Resonance Localization

Experiments on TFR provided some of the earliest indications that resonance localization may play an important role in tokamaks. They observed a sharp rise in ripple losses of ions when the ICRF was turned on,[32] and in some cases, were able to pump selected impurities out of the plasma.[16] In this second paper, they specifically point out that ICRF should produce energetic trapped particles whose banana tips lie near the resonance.

Kaita's experiment[7] was more direct in showing that energetic ion banana tips actually lie in the resonance layer. While making routine fast neutral

Figure 1.12: Illustration of how the orbits which cause the peak neutral flux will change as the resonance layer is moved.

measurements as a function of viewing angle, Kaita found that the largest signal was observed at an intermediate viewing angle instead of at the perpendicular viewing angle as one might expect. Figs. 1.13–17 show measurements made in the early stages of my thesis research which reproduce their original results. They postulated that this large signal was due to fast ions whose banana tips were in the resonance layer and which charge exchange near the edge of the plasma where the neutral density is large. This model predicts that the largest neutral flux should be observed at an $\hat{R}_{tan}$ given by setting R in Eq. (1.2) to the outside edge of the tokamak:

$$\hat{R}_{tan} = \sqrt{R_{out}(R_{out} - R_{res})} \tag{1.3}$$

They tested this model by changing the position of the resonance layer (by changing the toroidal field) to see if the peak in the neutral flux moved as Eq. (1.3) would predict. The expected change in the orbits which cause the peak neutral flux is illustrated by Fig. 1.12. Figs. 1.13–14 show that the peak does move as this simple model predicts. (In these experiments R_{out} was taken to be 174 cm where the outer limiter was. In Kaita's original experiment there were only top and bottom limiters so R_{res} was taken to be 183, the position of the vacuum vessel wall.) Note that although the

Figure 1.13: Flux of 90 keV hydrogen neutrals as a function of analyzer sightline angle when $R_{\rm res} = 132$ cm.

magnitude of the energetic neutral flux has a strong density dependence, the position of the peak is independent of density. In addition to the density independence, Kaita's measurements showed that the position of the peak is independent of plasma current or the position of the ICRF antennas relative to the neutral particle analyzer. Figs. 1.15–17 show that although the peak at lower energies is less pronounced (because pitch angle scattering is larger) it still moves as the resonance layer is moved.

1.2.5 Other Evidence for Resonance Localization

Further evidence for resonance localization has been provided by a number of other diagnostics besides the neutral particle analyzer. Perhaps the clearest evidence is in edge probe measurements made by Manos.[33] They used a calorimeter probe with a rotable entrance aperture which only allows energetic ions with a selected pitch angle to enter the probe (Fig. 1.18). The probe was inserted into the edge of the plasma at the outside midplane, where unconfined banana orbits are lost (Fig. 1.7). The lower energy parti-

Figure 1.14: Flux of 90 keV hydrogen neutrals as a function of analyzer sightline angle when $R_{\rm res} = 152$ cm.

cles which are always present in the edge plasma have gyroradii smaller than the aperture, and they can enter the calorimeter whenever it is pointed along a field line. The peaks in Fig. 1.19 at 0 and 180 degrees are caused by the low energy edge plasma. The additional peak at 310 degrees in Fig. 1.19(a) occurs during H minority heating. This pitch angle corresponds to particles whose banana tips lie near the resonance layer, in agreement with Eq. (1.1). Reversing the direction of the plasma current reversed the peak angle in Fig. 1.19(b), as it should. No additional peak was observed during ^{3}He minority heating (Fig. 1.19(c)), consistent with expectations that unconfined orbit losses should be much weaker for ^{3}He minority heating than for H minority heating. Unconfined orbit losses were also observed during high power second harmonic hydrogen heating (Fig. 1.19(d)). In addition to the calorimeter probe at the outside midplane, a second probe was placed at the top of the plasma. It saw no fast ion peak during H or ^{3}He minority heating. This is also to be expected, since fast ions would be lost to the outside limiter long before their banana widths became large enough that the top probe would see them. The caps which cover the edge probes places at the outside

Figure 1.15: Flux of 30 keV hydrogen neutrals as a function of analyzer sightline angle when $R_{\rm res} = 132$ cm.

midplance suffer significant erosion after several hundred plasma discharges. The erosion is asymmetric, and is largest on the side which is bombarded by the ICRF produced energetic ions. This has also been interpreted as evidence for resonance localization.[14]

A number of innovative fusion product measurements also lend their weight to the argument for resonance localization. Spectral measurements by Heidbrink[34] of 15 MeV protons produced by D-^{3}He reactions show that the ^{3}He is highly anisotropic with $v_\perp \gg v_\parallel$ in the center of the plasma, indicating that they must be trapped particles. Measurements by Murphy and Strachan[35] of the major radius distribution of D-^{3}He reactions show that it is highly peaked near the resonance layer at R_{res}, and not necessarily near the magnetic axis.

Figure 1.16: Flux of 30 keV hydrogen neutrals as a function of analyzer sightline angle when $R_{\rm res} = 152$ cm.

1.3 Review of Previous Experimental Work

We have already discussed a large body of experimental work. Here we concentrate specifically on previous measurements of the energetic, non-Maxwellian ions produced by ICRF.

1.3.1 Deuterium or hydrogen absorption?

In the cold plasma limit, heating of a pure plasma at the ion cyclotron frequency is inefficient because the wave is circularly polarized and is rotating in the opposite directions as the ions. A wave at twice the cyclotron frequency does have a component which rotates in the same direction as the ions, and is able to heat the ions due to a finite $k_\perp \rho$ effect. In the early ST experiment, energetic ions were observed during what was thought to be second harmonic deuterium heating.[36] However, the charge exchange analyzer used for that measurement was incapable of distinguishing deuterium from hydrogen. The TFR group[37] pointed out that the dominant absorption mechanism may be

Figure 1.17: Flux of 60 keV hydrogen neutrals as a function of analyzer sightline angle when $R_{\rm res} = 152$ cm.

Figure 1.18: Schematic diagram of midplane probe rotated to high angle to accept co-going fast ions. (From Manos et al.[33]).

Figure 1.19: (a) Heat flux to probe as a function of rotation angle relative to plasma current. For D^+ plasma with H^+ minority heated by 1 MW of RF power at 42 MHz. (b) same as in (a) with all fields, including plasma current, reversed. (c) Same as for (a) with $^3He^{++}$ minority heated by 1.5 MW of RF power at 30 MHz. (d) same as in (a) for pure H^+ plasma heated by 2 MW of RF power at 2nd harmonic frequency. (From Manos et al.[33]).

the fundamental resonance with residual hydrogen, rather than the second harmonic resonance with deuterium. But it was Vdovin et al.[38] and Ivanov et al.[39,40] who used mass-resolving charge exchange analyzers to show that the energetic ions were in fact protons, not deuterons. The small amount of residual hydrogen always present in the plasma ($n_H/n_D \sim 10^{-2}$) was not enough to adversly affect the wave polarization, but was enough to absorb most of the RF power. This is because the hydrogen cyclotron frequency is twice the deuterium cyclotron frequency, and second harmonic damping is weaker than fundamental damping by a factor of $(k_\perp \rho)^2$. Further observations of an energetic tail (but still without mass resolution) were made in ATC.[41]

1.3.2 Comparisons with Stix's theory

The Russian experiments showed that the shape of the hydrogen energy distribution was in good agreement with the isotropic quasilinear theory of Stix[12] (which we will discuss in more detail in Chapter 2). The one free parameter in Stix's isotropic theory is the RF power density needed to sustain the non-Maxwellian tail against collisions. A number of experiments have tried to measure the RF power density by fitting Stix's theory to charge exchange spectra.

Mass-resolved charge exchange measurements on PLT[42–44] showed that a very energetic hydrogen tail (up to at least 80 keV) could be produced. The deuterium heated up (by collisions with the hydrogen), but remained Maxwellian, giving no indication of any direct second harmonic heating of the deuterium. In the early low power (90 kW) PLT experiments,[44] the radial power profile was inferred by fitting Stix's theory to charge exchange spectra from a vertically-scanning instrument, yielding $P(r/a = 0)/P(r/a = 15/40) \approx 1.4$. The TFR measurements[45,46] of the hydrogen tail indicate an RF power profile that is essentially flat out to $r/a \approx 12/18$. Odajima et al.[47] found that the central power density inferred from the energetic hydrogen tail seemed to drop by a factor of 3 as the hydrogen concentration was lowered from 10% to under 3%. More recent experiments on JFT-2[48,49] show evidence of a deuterium tail and direct second harmonic deuterium absorption which is stronger than expected. We discuss these recent JFT-2 results in more detail in Sec. 3.6. Ichimura et al.[50] measured the hydrogen tail on the small

JIPP T-II tokamak. They observed a saturation in the tail temperature, in agreement with a modified form of Stix's isotropic theory which includes unconfined orbit losses.

PLT demonstrated efficient second harmonic heating of hydrogen[51–53] (which, unlike deuterium, is not complicated by the simultaneous occurence of a fundamental resonance for another plasma species). They measured both the perpendicular and the *anti-parallel* (particles moving in the opposite direction as the plasma current) charge exchange spectra, finding that the perpendicular tail is much more energetic. (This is not in contradiction to Fig. 1.2, which shows the most energetic tail at an angle between perpendicular and *parallel*, viewing particles moving in the *same* direction as the plasma current.) They show that the energy dependence of the perpendicular tail can be fit with an isotropic Fokker-Planck calculation, but point out that a two-dimensional calculation must be carried out to describe the observed anisotropy. Vdovin[54,55] has numerically carried out two-dimensional calculations based on Stix's flux surface averaged quasilinear operator, but a comparison was made only with the perpendicular charge exchange measurements.

Chrien and Strachan[56] performed a comprehensive study of the scaling of the ICRF produced tail. They measured the magnitude of the D-^{3}He fusion reaction rate as a function of RF power, plasma density, plasma current, and electron temperature. These scalings were roughly consistent with a model based upon Stix's isotropic theory. Unconfined orbit losses were treated in their model by truncating the Stix distribution function at the loss energy, and the total reaction rate was found by integrating the $\langle \sigma v \rangle$ found from the Stix function over the plasma radius. They assumed a Gaussian shape for the power profile. The observed reaction rates did not increase as quickly as the model would predict. One interpretation of this is that the Gaussian width σ of the power profile broadened as the power increased, from σ of 9.5 cm at 0.4 MW to a σ of 12.5 cm at 1.0 MW for ^{3}He minority heating, and from a σ of 11.1 cm at 220 kW to a σ of 16.0 cm at 665 kW for D minority heating. They also observed correlations between the D-^{3}He reaction rate and sawteeth and $m = 2$ oscillations.

1.3.3 New observations

This thesis adds to previous experimental work in several ways. Our measurements provide clear evidence of direct second harmonic deuterium heating, and we have studied the scaling of the deuterium and hydrogen tails with hydrogen concentration and RF power. Until the work of Kaita et al.,[7] the angle dependence of the charge exchange spectra had not recieved much attention. We investigate this angle dependence further, study its scaling with RF power, and make detailed comparisons with the predictions of a bounce averaged Fokker-Planck program. Most previous measurements were compared with Stix's isotropic theory, which is valid at low energies where pitch angle scattering is dominant. Some two dimensional calculations have been done with Stix's flux surface averaged quasilinear operator, but this does not properly include the resonance localization effects necessary to explain spectra such as Fig. 1.2. All previous modelling of the charge exchange spectra during ICRF ignored radial profile effects. We compare our data with a model that not only includes resonance localization, but also includes finite $k_\perp \rho$, Doppler-broadening of the resonance layer, unconfined orbit losses, and radial profile effects, all of which can be important in real experiments.

1.4 Experimental Methods

1.4.1 PLT and the ICRF heating system

The experiments in this thesis were performed on the Princeton Large Torus (PLT), which has a major radius $R = 132$ cm, and a minor radius $a = 41$ cm. The ICRF generators operated at a frequency of 30 MHz, so for hydrogen minority heating the toroidal field was $B \sim 20$ kG, while for helium-3 minority heating the toroidal field was $B \sim 30$ kG. ICRF waves were launched from the low field side with a pair of out-of-phase antennas on adjacent ports which produce a spectrum of waves peaked around $k_\parallel = \pm 9/132\text{cm}^{-1} = \pm .068\text{cm}^{-1}$. (Higher harmonic waves at $\pm(2n+1)9/132\text{cm}^{-1}$ are evanescant in the low density edge plasma and carry little power into the main plasma.) Under the best conditions, up to 4.3 MW of ICRF power can be launched into the plasma, but most of our experiments were done with 1 to 2.5 MW. Typical plasma parameters for these experiments were: plasma current I_p

= 450–600 kA, line-averaged density $\bar{n}_e$ = 1–3×10^{13}cm^{-3}, central electron temperature T_e = 1–4 keV, and central bulk ion temperature T_i = 0.6–5 keV. The effective temperature (defined as two-thirds of the average energy) of the ion species which resonates with the ICRF can be as high as several hundred keV. Almost all of the discharges studied have sawtooth oscillations, with a typical inversion radius of 10–15 cm. Further information about PLT and the ICRF system can be found in a number of references.[53,57]

1.4.2 General Neutral Particle Analysis Techniques

There are a number of excellent reviews of diagnostic techniques based upon neutral particle analysis.[58] All of the measurements made for this thesis relied upon charge exchange of fast ions with background neutrals in the plasma. This is called passive charge exchange, as opposed to active charge exchange which uses a neutral beam injector to provide the source of neutrals. We use the conventions of Davis, Mueller, and Keane, where the quantity labelled as $f(W)$ in a charge exchange spectral plot is related to the true distribution function $f(W, \eta, \vec{x})$ (which depends on energy W, pitch angle $\eta = v_\parallel/v$, and position $\vec{x}$) by

$$f_{cx}(W) = \int ds f(W, \eta(s), \vec{x}(s)) n_0(\vec{x}(s)) \frac{\langle \sigma v \rangle_{cx}}{\sigma_{cx} v} e^{-\lambda}. \tag{1.4}$$

In general, passive charge exchange measurements are weighted averages over a sightline. The pitch angle varies along a sightline as described in Sec. 1.2.2. The thermal neutral density n_0 can drop 3 orders of magnitude from the edge to the center of the plasma and may also have strong poloidal and toroidal variation. $e^{-\lambda}$ represents the attenuation of the energetic neutrals as they travel through the plasma along the sightline. The factor $\langle \sigma v \rangle_{cx}/\sigma_{cx} v$ denotes the difference, which is usually small, between the true local charge exchange rate and the beam-target approximation made in the detector calibration. By looking at high energies (usually 3–10 $\times T_i(0)$) one can consider the $f_{cx}(W)$ integral to be weighted towards the hottest part of the plasma. A measure of the central ion temperature can then be obtained from the slope of $\log f_{cx}(W) \approx K - W/T_i(0)$. The charge exchange flux f_{cx} can also be interpreted as providing a localized measurement if the ions are highly anisotropic in pitch angle, as described in Sec. 1.2.2 for the case of ion cyclotron heating.

1.4.3 The $E\|B$ Neutral Particle Analyzer

The PLT neutral particle analyzer is of the same design as the TFTR $E\|B$ analyzers, and has been described in detail elsewhere.[59–62] This same instrument can be used to measure the bulk ion temperature and the energetic ion tail. The data in Figs. 1.1 and 1.2 were obtained with this instrument. As illustrated by Figs. 1.20-22, neutrals which escape the plasma are reionized in the stripping cell, separated by mass and energy by a region of parallel electric and magnetic fields, and detected by a set of double-layer microchannel plates. At the output of the microchannel plates is an array of 3 (mass) $\times$ 75 (energy) anodes which collect the resulting pulses of electrons and sends them on for electronic processing. This instrument is able to simultaneously measure the energy spectrum (covering a dynamic range of 30 in energy) of hydrogen, deuterium, and tritium neutrals with 1 msec time resolution, although the time resolution is usually count rate limited to 10–100 msecs. Ions with mass $\times$ energy of up to 600 keV $\times$ AMU can be analyzed. The instrument is connected to PLT with a flexible bellows and can be moved horizontally between discharges to look at tangency radii R_{tan} between 13 and 102 cm (Figs. 1.3–4).

The neutral particle analyzer originally had a problem with cross-talk between anodes which were adjacent in energy, but not between anodes adjacent in mass. During calibration of the instrument, it was found that with a monoenergetic beam focussed on a single anode, several adjacent anodes would trigger as well, but at a somewhat lower rate. When the instrument was first installed on PLT, the measured neutral particle spectra were not smooth but exhibited variations of a factor of 2–3 because the cross-talk problem was not uniform on the microchannel plate and was count-rate dependent. The original design used a stainless steel electrode (item C in Fig. 1.23) between the front and rear microchannel plates, which also served to separate anodes of different masses but not different energies. Roquemore[63] found that the cross-talk problem could be eliminated by replacing the intermediate electrode with one which blocked alternate anodes in a mass column. All of the measurements reported in this thesis are with the new intermediate electrode. An $\sim$ 10–25% variation in detection efficiency between anodes persists, a problem which can perhaps be solved by using improvements in microchannels plates recently developed by Armentrout.[64]

In the initial months of operating this neutral particle analyzer on PLT,

Figure 1.20: Cross sectional views in plan and elevation of the $E \| B$ charge exchange neutral analyzer.

Figure 1.21: Sketch illustrating how the $E \| B$ fields separate particles by energy and mass.

Figure 1.22: Sketch illustrating the operating principle of the multi-anode microchannel plate detector.

Figure 1.23: Isometric assembly drawing illustrating the components of the multianode chevron microchannel plate detector. A. Assembly hardware and front electrode; B. Front microchannel plate of the chevron assembly; C. Intermediate electrode and mask defining active plate area; D. Rear microchannel plate; E. Rear electrode; F. Pattern of gold deposition defining the 75 anode pads and surrounding guard electrode; G. Ceramic substrate on which the gold anode pad pattern is deposited. Electrical contact with the anode pads is by means of pins extending through the backside of the ceramic substrate; H. Mounting hardware.

several steps were taken to improve the signal-to-noise ratio. Two 1/4″ lead sheets were mounted around the microchannel plate vacuum feedthrus to shield against x-rays. A light dump was added to the straight-through detector port (Fig. 1.20), and the inside of the analyzer was painted black to cut down on stray light from the plasma. An RF pickup problem was solved by improving the shielding of cables between the microchannel plate vacuum feedthroughs and preamplifiers. The analyzer's magnet power supply also suffered from RF pickup. This was solved by operating it as a constant voltage source rather than a constant current source. Accurate background subtraction was accomplished by monitoring the background (x-rays, neutrons, and/or reflected light) with one of the unused anodes in the tritium mass column. The signal count rate was improved by raising the gauge pressure of the helium stripping cell from the original 0.5 mTorr to 2 mTorr, and by enlarging the stripping cell apertures (which are 25 cm apart) from the original 1.6 mm diameter to 4.64 mm × 1.5 mm. While the count rate of moderate energy (up to 2–5 $\times T_i$) deuterium was sometimes sufficient for 1 msec time resolution, good counting statistics at high energies very far out in the tail usually required 25-100 msec time averages. Data points are only plotted if the estimated standard deviation in $\ln f$ due to counting statistics and background subtraction is less than 0.25.

The analyzer's magnetic field was monitored by a Gaussmeter using a Hall effect probe. There is a $\pm \sim 0.05$ kG drift in the null of this Gaussmeter, which would lead to a $\pm \sim 10\%$ error in the energy of particles when the analyzer is set to a 1 kG reading, or an $\sim$3% error for a 3 kG field. The maximum energy × mass which can be detected by the analyzer scales as 20 keV × AMU / kG, and we typically set the analyzer field to 2–3 kG.

And they who dwell in the ends *of the earth* stand in awe of Thy signs;
Thou dost make the dawn and the sunset shout for joy.

Psalm 65:8

Chapter 2

Bounce-Averaged Quasilinear Theory

2.1 Review of Previous Theoretical Work

A general quasilinear operator which describes the evolution of the particle distribution function $f(v_\perp, v_\parallel, t)$ under the influence of an uncorrelated spectrum of electromagnetic waves in a uniform magnetic field was first derived by Yakimenko[65] and independently by Kennel and Engelmann.[66] The Kennel-Engelmann form of the quasilinear operator Q is given below in Eq. (2.15). It is necessary to assume a continuous spectrum of incoherent waves in these uniform field calculations to assure that particles undergo net stochastic diffusion rather than just oscillating in a phase space island created by a single wave.

2.1.1 Stix's flux surface averaged quasilinear theory

In a classic work, Stix[12] studied the quasilinear theory of ion cyclotron heating in a tokamak. In tokamak ICRF heating experiments, only one wave (or at most only a few discrete waves) at a single coherent frequency are launched into a plasma with an inhomogeneous magnetic field. At first glance, it may seem that quasilinear theory is inapplicable to this case, but Stix's insight into the problem showed how the earlier uniform field, continuous spectrum theories could be easily extended to the tokamak ICRF case. His argument

was essentially from the single particle point of view (which will be considered in more detail in Sec. 2.2). Because of the magnetic field inhomogeneity in a tokamak, as a particle moves along a field line, the local gyrofrequency Ω which it sees will vary. For a wave with a given frequency ω and parallel wavelength $k_\parallel$, the resonance condition $n\Omega = \omega - k_\parallel v_\parallel$ can only be satisified at selected points along a particle's orbit. As a particle passes through one of these resonances, it will absorb or lose a small amount of energy, depending on whether the particle and the wave are in or out of phase. The time between successive resonances is fairly long, and Stix showed that the relative phase between the particle and the wave is easily randomized during this time by even a small collision rate. A sequence of random, small changes in energy leads naturally to a quasilinear equation, although he did not write down the quasilinear equation which would result from the single particle viewpoint. Instead, he averaged the Kennel-Engelmann quasilinear operator Q over a flux surface to eliminate the δ function in Q (arguing that the energy absorbed in the resonance layer is quickly spread out over the flux surface by the parallel motion of the particles). He then showed that the absorbed RF power averaged over a flux surface was the same if calculated from the single particle viewpoint or from the flux surface average of Q. The flux surface averaged quasilinear operator $\bar{Q}$ was combined with the usual collision operator C in a Fokker-Planck equation

$$\frac{\partial f}{\partial t} = \bar{Q}(f) + C(f)$$

and used to calculate the non-Maxwellian shape of f caused by the RF. Stix showed that because of the energetic tail produced by the RF, it was possible to enhance the fusion reaction rates and therefore the fusion gain (which, unfortunately, is usually denoted by Q as well). Stix's steady state solution for f has been compared to measurements in a number of experiments and and his flux surface averaged $\bar{Q}$ has been used in a number of calculations.[67]

2.1.2 Bounce averaged quasilinear theory

Stix's paper lays the groundwork for applying quasilinear theory in tokamaks. But while the flux surface averaged $\bar{Q}$ is correct in certain limits and captures much of the essential physics, $\bar{Q}$ does not properly include the two effects (ICRF heating is anisotropic and preferential) needed to produce the

"resonance localizaton" observed in the experiments. The underlying reason for this is that Stix assumed that $f(v_\perp, v_\|)$ was constant along a field line in his flux surface averaging procedure. This is only true in the large aspect ratio limit, or in the large pitch angle scattering limit where $f = f(W)$ is a function of energy only. More generally, one should assume that $f(W, \mu)$ is a constant along a field line, where the energy W and the magnetic moment μ are constants of the motion. Furthermore, it is more accurate to weight the average of Q by the amount of time a particle spends at various points along a flux surface. This naturally leads to a bounce average,

$$\langle Q \rangle \propto \int Q d\tau \propto \int Q \frac{d\ell}{v_\|},$$

rather than a simple flux surface average $\bar{Q} \propto \int Q d\ell$. Particles which move slower through the resonance layer absorb more energy in the bounce-averaged theory, while all particles are treated equivalently in the flux-surface-averaged theory. Secs. 2.3–4 provide a formal derivation of the bounce-averaged quasilinear operator $\langle Q \rangle$ (for the completely general Kennel-Engelmann Q) and a discussion of the limits in which Stix's earlier results can be recovered.

Several researchers[68–70] have recently derived the proper quasilinear operator for a discrete number of waves in an inhomogenous plasma from first principles rather than by bounce averaging the Kennel-Engelmann Q. Perhaps not so surprisingly, the two approaches give equivalent results except for a small class of particles which spend a very long time near resonance. The underlying physical reason that our approach of bounce averaging the Kennel-Engelmann Q works is that in the particle's frame of reference, it sees waves at the Doppler-shifted frequency $\omega' = \omega - k_\| v_\| - \Omega(\vec{x}(t))$. Even for a single wave with fixed ω and $k\|$, the particle sees a whole spectrum of waves at various ω' as the particle moves to regions of various Ω. Unfortunately, our approach has a weak singularity for particles which spend a long time near a resonance (i.e., for particles whose banana tips are exactly at the resonance layer). One of the advantages of the first principles derivation is that it naturally leads to a finite result for all particles. However, the singularity in our $\langle Q \rangle$ is integrable, and the two approaches give the same answer for almost all particles. Because of this, I will suggest in Secs. 2.2 and 2.4 that it is sufficient to use our $\langle Q \rangle$ if the singularity is treated properly.

One of the advantages of our derivation is that it establishes the Kennel-Engelmann Q as the proper local operator, even for a single wave in a non-uniform field. This is useful for calculating power deposition from ray tracing or wave propagation codes. It will also show how Stix's widely used flux-surface-averaged $\bar{Q}$ can be properly extended to a bounce-averaged $\langle Q \rangle$ which includes resonance localization effects, and will show in what limits Stix's earlier results can be recovered. Another motivation for deriving $\langle Q \rangle$ by bounce averaging the Kennel-Engelmann Q is that the derivations from first principles tend to be fairly complicated. On the other hand, the first principles derivations can be used even when the relative phase between the particle and the wave is not sufficiently randomized between resonances. Bernstein and Baxter's[68] derivation is fully relativistic. Mauel's[69] derivation is non-relativistic, but his final answer is written in a much simpler form. Mauel's work is also interesting because it compares the predictions of the theory with measurements from electron cyclotron heating of a mirror. Kerbel and McCoy's[70] derivation is a non-relativistic version of Bernstein and Baxter's.

Kerbel and McCoy also developed a computer program which numerically solves the resulting equations on a single flux surface. This program was applied to a number of problems, one of which was the simulation of the charge exchange spectra discussed in this thesis. Although their single radius code reproduced some of the qualitative features of the charge exchange spectra, it could not reproduce the "negative temperature" region where $\partial f/\partial E > 0$ (Fig. 1.2). As discussed in chapter 1, this feature appears to be caused by particles interacting with the RF at various minor radii (or by radial transport of fast ions), while their code was initally for a single minor radius. (Their code has since been expanded to multiple radial zones.[71]) In order to provide more accurate comparisons with charge exchange measurements, it was decided to implement $\langle Q \rangle$ in a Fokker-Planck code with multiple radial zones originally written by Goldston[72,73] to model charge exchange spectra during neutral beam heating. This code will be described in Sec. 2.6, and comparisons with measured spectra will be shown in Chapters 3 and 4.

2.1.3 Other bounce averaged quasilinear theories.

The importance of many of the effects we have been discussing was recognized at an early date by Jaeger, Lichtenberg, and Lieberman[29,30] who were primarily concerned with determining when stochastic diffusion would, or would not occur, during electron cyclotron heating in mirror geometry. They were the first to derive the Airy function correction (which will be given in Sec. 2.2) for particles whose banana tips are very close to the resonance layer. They also derived a quasilinear operator for a few special limits. The Monte Carlo simulations of Whang and Morales[74] also show the resonance localization phenomenon.

Several previous authors have taken a similiar approach to the present one of bounce averaging the Kennel-Engelmann quasilinear operator. Kesner[75] performed the bounce average of the Kennel-Engelmann Q for simple mirror geometry but relied on Stix's flux surface averaged $\bar{Q}$ for tokamak geometry. Kesner points out one of the singularites of $\langle Q \rangle$ but does not say how it should be properly handled. (There is an error in Kesner's result. In the equations after his Eq. (A-3), he used $\int dz\, \delta(n\omega_{ci} - \omega + k_\| v_\|) = 1/(\partial n\omega_{ci}/\partial z)$ when the value of the integral should be $1/(\partial(n\omega_{ci} + k_\| v_\|)/\partial z)$, which will modify the effect of the singularity.) V.S. Chan et al.[76] and J.Y. Hsu et al.[21] have calculated the bounce average of the Kennel-Engelmann Q for the special case where Q acts on a Maxwellian. A framework for bounce averaging $Q(f)$ for general f (although not explicitly written down) and a discussion of the radial transport which may be caused by ICRF heating has been given by S.C. Chiu.[77] C.S. Chang[78] started with a bounce average approach, but ended by assuming that $\langle Q \rangle$ was independent of the speed with which a particle passed through resonance. S.C. Chiu and C.S. Chang reached rather different conclusions about the importance of ICRF driven transport. Recently, Anderson et al.[79] presented a derivation of $\langle Q \rangle$ which is very similiar to mine. However, they ignore the $\partial/\partial v_\|$ terms in the quasilinear operator, and they reach the incorrect conclusion that the bounce averaging will suppress the the acceleration of ions for energies $\frac{1}{2}mv_\perp^2 > \epsilon^{-1}T_e$.

2.1.4 Outline of the rest of Chapter 2

Although our derivation of the bounce averaged quasilinear operator $\langle Q \rangle$ in Sec. 2.3 is not revolutionary, we provide it in the interest of explicitly writing

down the correct bounce average of the complete Kennel-Engelmann quasilinear operator. In Sec. 2.4 we explore a number of interesting properties of $\langle Q \rangle$. In particular, an isotropic model for f during second harmonic heating is derived which will later be shown to be quite accurate in fitting the measurements. Almost all of the analytical and computational work to date on the RF driven non-Maxwellian tail has been based on balancing quasilinear diffusion with collisions, ignoring any effects of radial transport. Sec. 2.5 will explore a number of possible fast ion transport mechanisms which may be playing an important role in our experiments. Sec. 2.6 describes our implementation of the bounce averaged quasilinear operator in a Fokker-Planck code and explains the results of a sample simulation. But first, we will try to gain some physical insight into the problem at hand by looking at a single particle model.

2.2 Simple Single Particle Model

We will calculate the energy absorbed by a single particle as it passes through a resonance by integrating the equations of motion. Our approach is the same as Stix's,[12] but extended to consider trapped particles whose banana tips lie very close to the resonance layer. This leads to the Airy function correction of Jaeger et al.,[29] and helps resolve a mild singularity which arises in the bounce-averaged quasilinear operator.

Starting with the Lorentz force law:

$$\frac{d\vec{v}}{dt} = \frac{ze}{m}\vec{E} + \frac{ze}{mc}\vec{v} \times \vec{B}$$

consider particle motion in a static magnetic field $\vec{B}_0$ plus an oscillating wave field $\vec{E}_1$ and $\vec{B}_1$. Using Faraday's law

$$\nabla \times \vec{E} = -\frac{1}{c}\frac{\partial \vec{B}}{\partial t}$$

$$\vec{k} \times \vec{E}_1 = \frac{\omega}{c}\vec{B}_1$$

the equation of motion becomes

$$\frac{d\vec{v}}{dt} = \frac{ze}{m}\vec{E}_1 + \frac{ze}{mc}\vec{v} \times \vec{B}_0 + \frac{ze}{m}\vec{v} \times (\frac{\vec{k}}{\omega} \times \vec{E}_1)$$

We will be considering fast waves in the ion cyclotron range of frequencies which have a typical phase velocity of the order of the Alfvén speed, $\omega/k \sim v_A$. For thermal velocities, the last term on the right hand side will typically be smaller than the first term on the right hand side by a factor of $v/v_A \sim \sqrt{\beta} \ll 1$. For thermal particles, $kv/\omega = 0$ is a good approximation, but in PLT the ICRF can sometimes produce a very energetic tail for which this approximation breaks down. Although we will initially consider the $kv/\omega = 0$ limit in the single particle model and ignore the last term in this equation, the bounce-averaged quasilinear operator derived later is valid for general $kv/\omega \neq 0$.

Consider a local frame of reference near the resonance layer with $\hat{z} \| \vec{B}$. The equations for motion perpendicular to the magnetic field $\vec{B}$ are:

$$\begin{aligned} \frac{dv_x}{dt} - \Omega(t) v_y &= \frac{ze}{m} E_x cos(\omega t) \\ \frac{dv_y}{dt} + \Omega(t) v_x &= -\frac{ze}{m} E_y sin(\omega t) \end{aligned} \tag{2.1}$$

where $\Omega(t)$ is the local cyclotron frequency seen by the particle as it moves along a field line. Defining $E_\pm = \frac{1}{2}(E_x \pm E_y)$, $u = v_x + i v_y$, these two equations can be combined into:

$$\frac{du}{dt} + i\Omega(t) u = \frac{ze}{m}(E_+ e^{-i\omega t} + E_- e^{+i\omega t}) \tag{2.2}$$

This is just the equation for a driven harmonic oscillator. Ignoring the non-resonant E_- term, the solution is

$$u(t_1) = e^{-i\int_{t_0}^{t_1} \Omega dt} [u(t_0) + \frac{ze}{m} E_+ \int_{t_0}^{t_1} e^{-i\int_{t_0}^{t} (\omega - \Omega) dt'} dt] \tag{2.3}$$

Particles will gain or lose energy depending on the relative phasing of $u(t_0)$ and E_+. If $u(t_0)$ and E_+ are randomly phased, then the average change in perpendicular energy is:

$$\begin{aligned} \langle \delta W_\perp \rangle &= \frac{m}{2} \langle u(t_1) u^*(t_1) - u(t_0) u^*(t_0) \rangle \\ &= \frac{z^2 e^2}{2m} |E_+ \int_{t_0}^{t_1} e^{-i\int_{t_0}^{t} (\omega - \Omega) dt'} dt|^2 \end{aligned} \tag{2.4}$$

The dominant contribution to this integral is in the vicinity of a resonance where $\omega = \Omega$. Stix evaluates this integral by expanding Ω near the resonance as $\Omega \approx \omega + \Omega' t$, where $\Omega' = v_\parallel \frac{\partial \Omega}{\partial \ell}$, and $\frac{\partial}{\partial \ell}$ is the derivative along the field line. In this limit, the average energy absorbed is:

$$\begin{aligned} \langle \delta W_\perp \rangle &= \frac{\pi z^2 e^2}{m} \frac{|E_+|^2}{|\Omega'|} \\ &= \left(\frac{\pi r B}{|v_\parallel| B_\theta} \right)_{res} \frac{\mathrm{P_{Stix}}}{\mathrm{n_{min}}} \frac{R_x}{R_{res}} \end{aligned} \tag{2.5}$$

where

$$P_{\mathrm{Stix}} = \left(n_{min} \frac{zec}{B} |E_+|^2 \frac{R}{r|sin\theta|} \right)_{res} \frac{R_{res}}{R_x} \tag{2.6}$$

is the absorbed RF power density averaged over a flux surface. The minority species which is resonant with the RF has a density of n_{min} and a charge of ze. Stix calculated P_{Stix} by summing Eq. (2.5) over all particles. (Eqs. (2.5–2.6) include a small R_{res}/R_x correction to Stix's original results.) As we will see later, the bounce-averaged quasilinear operator also gives P_{Stix} in the small kv/Ω limit, even for a non-Maxwellian, anisotropic plasma, as long as all parameters in Eq. (2.6) are evaluated in the resonance layer—including n_{min}. In deriving Eq. (2.5) it was assumed that the flux surfaces were concentric circles and that $\Omega \propto B \propto 1/R \propto 1/(R_x + r cos\theta)$. The geometry is shown in Fig. 2.1. Assuming axisymmetry, the parallel derivative can be expressed as

$$\frac{\partial}{\partial \ell} = \frac{B_\theta}{rB} \frac{\partial}{\partial \theta}$$

so that

$$\frac{\partial \Omega}{\partial \ell} = \frac{\Omega B_\theta \sin\theta}{RB} \tag{2.7}$$

Equation (2.5) reveals some interesting physics. In the large aspect ratio limit, ignoring toroidal effects, then $\pi r B/(B_\theta |v_\parallel|) \approx \pi R q/|v_\parallel|$ is the time it takes a particle to travel from one resonance to the next. P_{Stix}/n_{min} is the average power per particle. Multiplying these two gives the average energy absorbed during a single pass through the resonance layer. The RF power is not uniformly deposited among all particles but is weighted toward particles which spend the most time in the resonance layer, i.e. particles which move slowest through the resonance layer and have small $v_\parallel$ at resonance. Since

Figure 2.1: Cross sectional view of tokamak with definitions of minor radius r, poloidal angle θ, major radius $R_{\rm res}$ of the resonance layer, and major radius $R_{\rm x}$ of the magnetic axis.

$v_\parallel = 0$ at the turning point of a trapped particle, this implies that particles whose banana tips lie closer the resonance layer are preferentially heated.

Although Eq. (2.5) is valid for almost all particles, it breaks down for very small $|v_\parallel|$, basically because a particle does not spend an infinite amount of time at its banana tip even though $v_\parallel = 0$ there. We will consider the Airy function corrections to $\langle \delta W_\perp \rangle$ near $v_\parallel = 0$ below, but the corrections are only important for such a small region of phase space that other effects are usually more important in removing the singularity. The singularity in Eq. (2.5) is integrable, yielding a finite total power density P_{Stix} when $\langle \delta W_\perp \rangle$ is integrated over all particles. In practice, the singularity is quickly averaged over a finite region of phase space by collisions. Although the 90 deg scattering rate may be fairly slow, the 1 deg scattering rate is 90^2 times faster. A more quantitative analysis of the role of collisions is presented in Sec. 2.4.13, but the basic idea is that pitch angle scattering prevents any particles from remaining exactly at $v_\parallel = 0$ for long. After a time δt, a group of particles initially at $v_\parallel = 0$ will have spread out over a region of width $\langle \delta v_\parallel^2 \rangle^{1/2} \approx v(\nu_\perp \delta t)^{1/2}$. The energy absorbed by the ensemble of particles is now seen to be finite:

$$\Delta W = \int_0^{t'} dt \frac{dW_\perp}{dt} \sim \int_0^{t'} dt/v_\parallel \sim \int_0^{t'} dt/t^{1/2}$$

Just as the simple expression for $\langle \delta W_\perp \rangle$ contains a weak, integrable singularity, P_{Stix} contains an integrable singularity for $r|sin\theta| \to 0$. This $r|sin\theta|$ factor is a geometric effect related to the fraction of a flux surface which intersects the resonance layer (Fig. 2.1). In the $k_\parallel = 0$ limit, all the power is absorbed in a resonance layer of zero width, and a singularity occurs when a flux surface is tangent to the resonance layer (i.e., where $r|sin\theta| \to 0$). Despite this singularity, integrating P_{Stix} over the whole plasma volume yields a finite total power. The $r|\sin\theta|$ singularity is eliminated in the real world by Doppler-broadening which gives the resonance layer a finite width. The bounce averaged quasilinear operator derived in the Sec. 2.3 includes Doppler-broadening along with other $k_\parallel \neq 0$ effects. In some cases, Doppler-broadening plays a major role in determining the central power density, and a calculation of this is presented in Sec. 2.4.8. Nevertheless, the simplicity of the $k = 0$ limit is an aid in understanding some of the important physics involved, and helps establish the relationship between the bounce-averaged quasilinear approach and the single particle approach.

The fact that these singularities are integrable is exploited when numerically solving the bounce-averged quasilinear equation. $\langle \delta W_\perp \rangle$ is always averaged over the finite velocity grid size, while P_{Stix} is always averaged over the finite radial grid size. This prevents the rise of any numerical singularities, and insures that the total integrated power is correct.

To resolve the singularity in $\langle \delta W_\perp \rangle$ for $v_\parallel \to 0$, we must go to higher order in the expansion of Ω in Eq. (2.4) and use $\Omega = \omega + \Omega' t + \frac{1}{2}\Omega'' t^2$. The second order term accounts for the fact that a trapped particle spends only a finite amount of time at its banana tip, even though Ω' vanishes there. This second order expansion of Ω also means that we will be calculating the energy absorbed during two passes through the resonance layer. A particle passes through resonance, is reflected by the inhomogeneous magnetic field, and passes through resonance a second time, absorbing a total energy of :

$$\begin{aligned} \langle \delta W_\perp \rangle_2 &= \frac{z^2 e^2}{2m} |E_+|^2 | \int_{-\infty}^{\infty} dt\, e^{i(\Omega' \frac{1}{2} t^2 + \Omega'' \frac{1}{6} t^3)} |^2 \\ &= \frac{z^2 e^2}{2m} |E_+|^2 \, 2 | \frac{6}{\Omega''} |^{1/3} \int_0^{\infty} d\tau \, cos(\tau^3 - \frac{\Omega'^2}{12} | \frac{6}{\Omega''} |^{4/3} \tau) \end{aligned} \tag{2.8}$$

which can be written in terms of the Airy function as

$$\langle \delta W_\perp \rangle_2 = \frac{z^2 e^2}{2m} |E_+|^2 [2\pi | \frac{2}{\Omega''} |^{1/3} Ai(x)]^2 \tag{2.9}$$

where

$$x = -\frac{\Omega'^2}{4} | \frac{2}{\Omega''} |^{4/3} \tag{2.10}$$

A similiar Airy function behaviour was first obtained in the study of RF heating in mirror geometry.[29] The behaviour of the Airy function is shown in Fig. 2.2. To compare with Stix's result in Eq. (2.5), consider the $\Omega'' \to 0$ limit of Eq. (2.9). Using the asymptotic expansion of $Ai(-x)$ for $x < -1$ yields[80]

$$\langle \delta W_\perp \rangle_2 = \frac{\pi z^2 e^2 |E_+|^2}{m|\Omega'|} 4 sin^2 (\frac{2}{3}|x|^{3/2} + \frac{\pi}{4}) \tag{2.11}$$

The highly oscillatory sin^2 term represents coherence effects between the two passes through the resonance layer made by a trapped particle. Stix showed that collisions are usually sufficient to decorrelate successive passes. So we can recover Eq. (2.5) by taking an average value for sin^2 of $\frac{1}{2}$ and by

Figure 2.2: $Ai^2(x)$, where Ai is the Airy function. Also shown is an approximation discussed in the text.

multiplying by another factor of $\frac{1}{2}$ since Eq. (2.11) is for two passes through resonance.

A particle which comes very close to the resonance layer but is reflected before actually reaching it can still absorb some power. Expanding Ω in Eq. (2.4) near the turning point as $\Omega = \Omega_t + \frac{1}{2}\Omega_t'' t^2$ yields an equation identical to Eq. (2.9) but with the argument of the Airy function given by

$$x = (\omega - \Omega_t)|\frac{2}{\Omega_t''}|^{\frac{1}{3}} \tag{2.12}$$

x is seen to be a measure of the distance between the resonance layer and the particle's turning point. Note that Ω and its derivatives are evaluated at the particle turning point in Eq. (2.12), while they are evaluated at the resonance layer in Eq. (2.10). The definition of x in Eqs. (2.10) and (2.12) are equivalent for particles which pass through resonance, while Eq. (2.12) extends the definition of x to particles which are reflected before reaching the resonance layer. For these particles which don't reach resonance ($x > 0$), the absorbed energy $\langle \delta W_\perp \rangle$ vanishes exponentially fast, as $exp(-\frac{4}{3}x^{3/2})$ for $x > 1$. If we define an effective resonance layer width δ_A due to the Airy function correction by how close a particle's banana tip must come to the

resonance layer in order for $|x| = 1$, then we find

$$\delta_A = (\rho_\perp^2 r \frac{r sin^2\theta}{4Rq^2})^{1/3} \tag{2.13}$$

where $\rho_\perp = v_\perp/\Omega$ is the gyroradius, $q = rB/(RB_\theta)$ is the inverse rotational transform, B_θ and B are the poloidal and total magnetic fields, R is the major radius of the resonance layer, and θ is the poloidal angle where the resonance layer and the flux surface intersect (Fig. 2.1). In deriving Eq. (2.13), Ω_t'' was evaluated for a tokamak with circular flux surfaces and a magnetic field $B \propto 1/R$, so that

$$\begin{aligned}\Omega_t'' &= \frac{dv_\parallel}{dt}\frac{\partial\Omega}{\partial\ell} \\ &= -\frac{1}{2}v_\perp^2(\frac{B_\theta sin\theta}{BR})^2\Omega\end{aligned}$$

The effective resonance layer width due to the Airy function correction is quite small. For typical PLT parameters, $r \sim 40$ cm, $R \sim 132$ cm, $q \sim 2$, and $B \sim 20$ kG, a 100 keV H ion with a gyroradius $\rho \sim 1.6cm$ sees an effective resonance layer width $\delta_A \sim 1.2$ cm. Only a small amount of pitch angle scattering is necessary to change the location of the turning point of a particle by an amount δ_A. As discussed earlier, the Doppler-broadening of the resonance is usually much larger than δ_A. (One must take care in considering the Doppler effect, since it seems to disappear at the banana tip where $k_\parallel v_\parallel \to 0$. As we will see in the Sec. 2.3.3, for $k_\parallel \neq 0$ the singularity in $\langle \delta W_\perp \rangle$ no longer appears at $v_\parallel = 0$ but at a slightly Doppler-shifted velocity). The highly oscillatory coherence nature of the Airy function is easily averaged over, not only by collisions and Doppler-broadening, but also by effects which we have ignored so far, such as collionless RF-induced stochasticity[29,74] finite banana widths, finite gyroradii, toroidal precession, toroidal asymmetries in the magnetic field, or even nonlinear interactions. In practice, we may approximate the quantity $Ai^2(x)$ which appears in Eq. (2.9) by

$$\begin{aligned}Ai^2(x) &\to \frac{1}{2\pi\sqrt{|x|}} \text{ for } x < -.3079 \\ &\to .28683 \quad \text{for} \quad -.3079 < x < 0 \\ &\to 0 \qquad \text{for} \quad x > 0\end{aligned}$$

A comparison of $Ai^2(x)$ and its approximation is show in Fig. 2.2. This approximation is equivalent to Stix's result in Eq. (2.5) for most particles, removes the singularity by setting an upper limit for $\langle \delta W_\perp \rangle$, and ignores the small amount of power transferred to particles which don't reach the resonance layer. We see that Stix's formula only breaks down for trapped particles which turn within a distance of the order of a gyroradius from the resonance layer. This represents such a small region of phase space, usually smaller than the velocity grid spacing used in numerical solutions, that it is frequently acceptable to use Stix's formula everywhere. In this case, one must rely on the rapidity of pitch angle scattering over small regions of $v_\parallel$ and the integrabilty of the $1/v_\parallel$ singularity in $\langle \delta W_\perp \rangle$ to prevent numerical singularities and to produce realistic results. The justifiability of this approach has been checked by running the code with extremely small grid spacing in order to resolve the Airy function upper limit for $\langle \delta W_\perp \rangle$. The results were virtually identical to runs with coarser grids.

2.3 Derivation of a Bounce-Averaged Quasilinear Operator

2.3.1 The Kennel-Engelmann quasilinear operator

Quasilinear theory describes the velocity space diffusion of particles caused by interaction with plasma waves. The quasilinear operator for an infinite, spatially uniform, nonrelativistic plasma was worked out by Kennel and Engelmann.[66] Representing the wave electric field by its Fourier transform,

$$\vec{E}(\vec{x},t) = Re[\int d^3k e^{i\vec{k}\cdot\vec{x}-i\omega t}\vec{E}_{\vec{k}}], \tag{2.14}$$

they found

$$\frac{\partial f(v_\perp, v_\parallel, t)}{\partial t} = \lim_{V\to\infty}\int \frac{d^3k}{V}\sum_n \frac{\pi e^2}{2m^2} L[v_\perp^2 \delta(\omega - k_\parallel v_\parallel - n\Omega)|\Theta_{n,\vec{k}}|^2 Lf] \tag{2.15}$$

$$\equiv Q(f).$$

The limit involving the plasma volume V provides the proper normalization for the the wave energy density. $\Theta_{n,\vec{k}}$ has units of electric field per wavelength k and is defined by

$$\Theta_{n\vec{k}} = E_{\vec{k}+}e^{-i\psi}J_{n-1} + E_{\vec{k}-}e^{+i\psi}J_{n+1} + \frac{v_\parallel}{v_\perp}E_{\vec{k}\parallel}J_n$$

$$E_{\vec{k}\pm} = \frac{1}{2}(E_x \pm iE_y)_{\vec{k}}$$

$$k_x = k_\perp \cos\psi$$

$$k_y = k_\perp \sin\psi$$

$$k_z = k_\parallel$$

and the argument of the Bessel functions is $(k_\perp v_\perp/\Omega)$. The phase factors $e^{\pm i\Psi}$ insure that $|\Theta_{n\vec{k}}|^2$ is independent of the choice of x and y coordinates. The operator L is defined by

$$L = (1 - \frac{k_\parallel v_\parallel}{\omega})\frac{1}{v_\perp}\frac{\partial}{\partial v_\perp} + \frac{k_\parallel}{\omega}\frac{\partial}{\partial v_\parallel}$$

Note that L takes the derivative along circles of constant energy in the wave's frame of reference, i.e. $LC = 0$, where

$$C = v_\perp^2 + (v_\parallel - \omega/k_\parallel)^2 = \text{constant} \tag{2.16}$$

L appears twice in Eq. (2.15) making it a diffusion type equation. As we can see from the form of the derivatives in L, Q conserves particles as it should:

$$\frac{\partial n}{\partial t} = \frac{\partial}{\partial t} 2\pi \int_0^\infty dv_\perp v_\perp \int_{-\infty}^\infty dv_\parallel f = 2\pi \int_0^\infty dv_\perp v_\perp \int_{-\infty}^\infty dv_\parallel Q = 0$$

2.3.2 The bounce averaging procedure

The presence of the delta function in the quasilinear operator means that only particles which resonate with the wave (so that $\omega - k_\parallel v_\parallel = n\Omega$) will be affected by the wave. Of course, delta functions only make sense when used in integrals. Usual quasilinear theory is derived for a continuous spectrum of incoherent waves. The $\int d^3k$ in Eq. (2.15) integrates over this wave spectrum and the delta function to yield a finite diffusion coefficient. Stix's contribution was to note that because of the inhomogenous magnetic field, the resonance condition $\omega - k_\parallel v_\parallel = n\Omega(\vec{x})$ is only satisfied at certain positions in the plasma and one can integrate over the delta function in Q by averaging over a flux surface. By bounce averaging, we are merely weighting Stix's flux surface average by the amount of time a particle spends at each point on the flux surface. Although the physical idea behind this is simple, the formal derivation of bounce averaging is a little more involved.

The bounce-averaging procedure I will outline here is a specific example of the annihilation technique discussed by Kruskal in his discourse on the art of asymptotics.[81] This annihilation technique finds frequent application in plasma physics. Our starting point is the gyro-averaged drift kinetic equation[82] to which we have added a quasilinear operator:

$$\frac{\partial f}{\partial t} + v_\parallel \frac{\partial f}{\partial \ell} = C(f) + Q(f) \tag{2.17}$$

where the gyroaveraged distribution function $f(W, \mu, \sigma, \vec{x}, t)$ is a function of the constants of motion, the energy $W = \frac{1}{2}m(v_\perp^2 + v_\parallel^2)$ and the magnetic moment $\mu = mv_\perp^2/(2B)$, and of $\sigma = \text{sign}(v_\parallel)$, position $\vec{x}$, and time t. The

parallel velocity is given by $v_\parallel = \sigma\sqrt{2(E-\mu B)/m}$, while $\frac{\partial}{\partial \ell} = \frac{\vec{B}\cdot\nabla}{|\vec{B}|}$ denotes the spatial derivative along a field line. The key difference between our work and Stix's is that he ignored the $v_\parallel \partial f/\partial \ell$ term.

We have made the small banana width approximation and have ignored cross-field drifts. C is the usual collision operator. We will assume that both the confining magnetic field and f are axisymmetric, $\frac{\partial}{\partial \phi} = 0$. For the usual circular flux surface model of the tokamak (Fig. 2.1), this means that $\frac{\partial}{\partial \ell} = \frac{B_\theta}{rB}\frac{\partial}{\partial \theta}$, where θ is the poloidal angle, B_θ is the poloidal field, and r is the minor radius. f is periodic in ℓ since it must be periodic in θ.

In present and future tokamak experiments, the effect of collisions and quasilinear heating during one bounce of a particle is small. We will assume that $v_\parallel \frac{\partial f}{\partial \ell}$ is the largest term in Eq. (2.17), and that all other terms are order ϵ smaller. Expanding $f = f_0 + \epsilon f_1 + \ldots$, we find to lowest order that $\frac{\partial f_0}{\partial \ell} = 0$, i.e. that $f_0 = f_0(W, \mu, \sigma, r, t)$ is constant on a flux surface. We must go to next order in ϵ to find how f_0 evolves in time:

$$\frac{\partial f_0}{\partial t} + v_\parallel \frac{\partial f_1}{\partial \ell} = C(f_0) + Q(f_0) \tag{2.18}$$

The f_1 term can be annihilated by integrating Eq. (2.18) over an orbit with $\int \frac{d\ell}{v_\parallel}$, since f_1 is periodic in ℓ. This results in a bounce-averaged kinetic equation involving f_0 only:

$$\frac{\partial f_0}{\partial t} = \langle C(f_0)\rangle + \langle Q(f_0)\rangle \tag{2.19}$$

where the bounce average of some quantity X is defined by

$$\langle X\rangle = \frac{1}{\tau_B}\int \frac{d\ell}{|v_\parallel|} X$$

$$\tau_B = \int \frac{d\ell}{|v_\parallel|}$$

The bounce time τ_B is the time it takes a trapped particle to travel from one banana tip to the other, while for passing particles it is the time a particle takes in making one poloidal circuit of a flux surface. Cordey[83] has calculated the the bounce-averaged collision operator $\langle C\rangle$ (which is given in Eq. (2.45) for completeness) and has given the appropriate boundary conditions which

f_0 must satisfy. For notational convenience we will assume a discrete wave spectrum:

$$\vec{E}(\vec{x},t) = Re[\sum_{\vec{k}} e^{i\vec{k}\cdot\vec{x}-i\omega t}\vec{E}_{\vec{k}}]$$

so that the limit as the plasma volume goes to infinity is no longer necessary in Eq. (2.15). Then $\langle Q\rangle$ is simply:

$$\langle Q\rangle = \sum_{\vec{k},n}\frac{\pi e^2}{2m^2}\frac{1}{\tau_B}\int\frac{d\ell}{|v_\parallel|}L[v_\perp^2\delta(\omega-k_\parallel v_\parallel-n\Omega)|\Theta_{n,\vec{k}}|^2 Lf] \tag{2.20}$$

2.3.3 Evaluating the bounce average

It seems reasonable that Eq. (2.20) reduces to Eq. (2.25) by application of the basic delta function property, Eq. (2.22). But because the L operator acts on the delta function, we must take care to first change the order of operations in this integral. We use the chain rule to transform the L operator from $(v_\perp, v_\parallel)$ space to (W,μ) space. Using

$$\frac{\partial}{\partial v_\perp} = \frac{\partial W}{\partial v_\perp}\frac{\partial}{\partial W} + \frac{\partial\mu}{\partial v_\perp}\frac{\partial}{\partial\mu}$$

$$= mv_\perp\frac{\partial}{\partial W} + \frac{mv_\perp}{B}\frac{\partial}{\partial\mu}$$

and

$$\frac{\partial}{\partial v_\parallel} = \frac{\partial W}{\partial v_\parallel}\frac{\partial}{\partial W} + \frac{\partial\mu}{\partial v_\parallel}\frac{\partial}{\partial\mu}$$

$$= mv_\parallel\frac{\partial}{\partial W}$$

we find

$$L = m\frac{\partial}{\partial W} + (1-\frac{k_\parallel v_\parallel}{\omega})\frac{m}{B}\frac{\partial}{\partial\mu}$$

Note that the operator L depends on position through B and $v_\parallel$. Denoting the expression in brackets in Eq. (2.20) by G, we can use the identity

$$\frac{1}{|v_\parallel|}LG = m\frac{\partial}{\partial W}(\frac{G}{|v_\parallel|}) + m\frac{\partial}{\partial\mu}((1-\frac{k_\parallel v_\parallel}{\omega})\frac{G}{B|v_\parallel|})$$

and can pull the $\frac{\partial}{\partial W}$ and $\frac{\partial}{\partial \mu}$ operations out of the ℓ integral to get

$$\langle Q \rangle = \frac{m}{\tau_B}\frac{\partial}{W}\sum_{n,\vec{k}}\frac{\pi e^2}{2m^2}\int\frac{d\ell}{|v_\parallel|}[v_\perp^2\delta(\omega - k_\parallel v_\parallel - n\Omega)|\Theta_{n\vec{k}}|^2 Lf]$$

$$+\frac{m}{\tau_B}\frac{\partial}{\mu}\sum_{n,\vec{k}}\frac{\pi e^2}{2m^2}\int\frac{d\ell}{|v_\parallel|}(1-\frac{k_\parallel v_\parallel}{\omega})\frac{1}{B}[v_\perp^2\delta(\omega - k_\parallel v_\parallel - n\Omega)|\Theta_{n\vec{k}}|^2 Lf] \quad (2.21)$$

The δ function now makes the bounce average integrals easy. The integrals are of the form:

$$\int\frac{d\ell}{|v_\parallel|}\delta(\omega - k_\parallel v_\parallel - n\Omega)G(\ell) = \sum_{\ell_{res}}\frac{G(\ell_{res})}{|v_\parallel\frac{\partial}{\partial\ell}(k_\parallel v_\parallel + n\Omega)|_{res}}$$

$$= \sum_{\ell_{res}}\frac{G(\ell_{res})}{|(nv_\parallel - k_\parallel\mu\frac{c}{e})\frac{\partial\Omega}{\partial\ell}|} \quad (2.22)$$

where the sum is over all positions where a resonance occurs. The factor $(1-\frac{k_\parallel v_\parallel}{\omega})/B$ which appears in Eq. (2.21) must be evaluated at resonance, yielding $n\Omega_{res}/(\omega B_{res}) = ne/(\omega mc)$ which is a constant indepedent of $k_\parallel$. L must also be evaluated at resonance, giving:

$$L_{res} = m\frac{\partial}{\partial W} + \frac{ne}{\omega mc}m\frac{\partial}{\partial\mu} \quad (2.23)$$

which is independent of $k_\parallel$ and the position of the resonance. One of the consequences of this surprising result is that the banana tips of all particles approach the $k_\parallel = 0$ resonance layer as the particles are accelerated to high energy. To see this, note that the L_{res} operator causes diffusion to occur only along characteristics defined by

$$K = W - \frac{\omega mc}{ne}\mu = \text{constant} \quad (2.24)$$

i.e., $L_{res}K = 0$. The magnetic field at a particle's turning point is given by

$$B_{tp} = \frac{W}{\mu} = \frac{W}{W-K}\frac{\omega mc}{ne}$$

As a particle is accelerated to high energies so that $W >> K$, the particle's banana tip approaches the $k_\parallel = 0$ resonance layer where $\omega = n\Omega$.

The general, bounce averaged, Kennel-Engelmann operator can thus be written as:

$$\langle Q\rangle = \frac{1}{\tau_B}\sum_{n\vec{k}}\frac{\pi e^2}{2m^2}L_{res}\sum_{res}[\frac{v_\perp^2|\Theta_{n\vec{k}}|^2}{|(nv_\parallel - k_\parallel\mu\frac{c}{e})\frac{\partial\Omega}{\partial\ell}|}]L_{res}f \tag{2.25}$$

The positions of the resonances are found by solving the simultaneous equations:

$$\omega - k_\parallel v_\parallel = n\Omega = n\Omega_0\frac{B}{B_0}$$

$$v_\parallel = \pm\sqrt{\frac{2}{m}(W-\mu B)}.$$

B_0 is the minimum magnetic field on the the flux surface. From these equations, the parallel velocity at resonance is found to be

$$v_{\parallel res} = \pm\sqrt{\frac{2}{m}(W-\mu\frac{B_0}{n\Omega_0}\omega) + (\frac{\mu B_0 k_\parallel}{mn\Omega_0})^2} + \frac{\mu B_0 k_\parallel}{mn\Omega_0} \tag{2.26}$$

There is no resonance if the argument of the square root is negative. In order for this solution to be meaningful, the resonance must occur somewhere on the flux surface which the particle is constrained to move on. A particle must actually attain this value of $v_{\parallel res}$ somewhere on the flux surface, i.e., $W - \mu B_0 > \frac{1}{2}mv_{\parallel res}^2 > W - \mu B_{max}$. Using the expression for $v_{\parallel res}$, the denominator in Eq. (2.25) can be written as:

$$|(nv_\parallel - k_\parallel\mu\frac{c}{e})\frac{\partial\Omega}{\partial\ell}|_{res} = n|\frac{\partial\Omega}{\partial\ell}|_{res}\sqrt{\frac{2}{m}(W-\mu\frac{B_0}{n\Omega_0}\omega) + (\frac{\mu B_0 k_\parallel}{mn\Omega_0})^2} \tag{2.27}$$

In the $k_\parallel = 0$ limit, particles in the simple circular cross-section tokamak have either 2 resonances or no resonances. (Trapped particles may pass through 4 resonances during a complete bounce, but our bounce average is over only half a bounce, from one banana tip to the other. Since f must be symmetric in $v_\parallel$ for trapped particles, one must take care to symmetrize $\langle Q\rangle$ by averaging over the two directions $\sigma = \pm 1$ for trapped particles.)

2.4 Properties of the Bounce-Averaged Quasilinear Operator

2.4.1 Recovering Stix's results

The general expression for $\langle Q \rangle$ is sufficiently complicated to suggest the use of computers to evaluate it. In the $k_\parallel = 0$ limit, however, the position of the resonance is identical for all particles, and further analytic progress can be made. If we further consider a simple limit used by Stix, keeping just $n = 1$, setting $k_\parallel = k_\perp = E_- = E_\parallel = 0$, and using the definition of $\langle \delta W_\perp \rangle$ in Eq. (2.5), we can write $\langle Q \rangle$ in the form

$$\langle Q \rangle = \frac{1}{\tau_B} L_{res} [v_{\perp res}^2 \frac{\langle \delta W_\perp \rangle}{m} L_{res} f] \tag{2.28}$$

This relates the quasilinear operator to the average energy absorbed by a single particle found in the previous section. (It is postulated that $\langle Q \rangle$ can always be written in a form similiar to Eq. (2.28) for general $k_\parallel \neq k_\perp \neq E_- \neq E_\parallel \neq 0$ if the single particle energy gain $\langle \delta W_\perp \rangle$ was calculated for the more general case.) In this $k_\parallel = 0$ limit, we can identify L_{res} as $(\frac{1}{v_\perp}\frac{\partial}{\partial v_\perp})_{res}$. Writing $\langle \delta W_\perp \rangle$ in terms of P_{Stix}, $\langle Q \rangle$ becomes:

$$\langle Q \rangle = \Big[\{ \frac{1}{\tau_B} (\frac{2\pi r B}{B_\theta |v_\parallel|}) H \} \frac{1}{v_\perp} \frac{\partial}{\partial v_\perp} (v_\perp^2 \frac{P_{Stix}}{2mn_{min}} \frac{R_x}{R_{res}} \frac{1}{v_\perp} \frac{\partial f}{\partial v_\perp}) \Big]_{res} \tag{2.29}$$

The Heaviside step function $H = H(W - \mu B_{res})$ is zero for trapped particles which never reach the resonance layer. The res subscript emphasizes that everything is evaluated at resonance. In particular, $v_\perp$ and $v_\parallel$ depend on position and are to be evaluated at resonance. If we take the limit of large aspect ratio or $v_\parallel >> v_\perp$, then $\tau_B \to 2\pi R q / |v_\parallel|$ and the factor in braces in Eq. (2.29) approaches unity so that $\langle Q \rangle$ reduces to the flux surface averaged $\bar{Q}$ found by Stix (with a small R_0/R_{res} correction). The factor in braces weights Stix's $\bar{Q}$ by the fraction of time a particle spends in the resonance layer.

2.4.2 Compared with Mauel and with Kerbel and McCoy

Eq. (2.29) is singular for particles which have $v_{\|res}$ exactly equal to zero, and one is sorely tempted just to use the Airy function correction to $\langle \delta W_\perp \rangle$ found in the previous single particle model section. In fact, our general quasilinear operator in Eq. (2.25) is essentially the same as the one derived by Mauel and by Kerbel and McCoy, except that they leave their $\langle Q \rangle$ in terms of a correlation time which automatically includes the Airy function correction. In a simple limit, their correlation time τ_c is defined by

$$\tau_c^2 = \left| \int_0^{\tau_B} d\tau e^{-i \int_0^\tau d\tau' (\omega - k_\| v_\| - n\Omega)} \right|^2 \tag{2.30}$$

This is similiar to the integral in Eq. (2.4), with the addition of the Doppler shift. In almost all of phase space it is valid to expand the argument of the exponential near resonance keeping terms only to second order in τ, yielding

$$\tau_c^2 \simeq \frac{4\pi}{|v_\| \frac{\partial}{\partial \ell}(k_\| v_\| + n\Omega)|_{res}} = \frac{4\pi}{|(n v_\| - k_\| \mu \frac{c}{e}) \frac{\partial \Omega}{\partial \ell}|_{res}} \tag{2.31}$$

so that Mauel's and Kerbel and McCoy's operator reduces to the $\langle Q \rangle$ given here by Eq. (2.25). In a small region of phase space, the approximations leading to Eq. (2.31) break down and Eq. (2.25) has a singularity. We will show in Sec. 2.4.6 that this is an integrable singularity. We will further suggest in Secs. 2.4.13–2.4.15 that, because of collisions and Doppler-shifts, it is usually sufficient to ignore the corrections of Mauel and Kerbel and McCoy and just use the $\langle Q \rangle$ found by bounce averaging the Kennel-Engelmann Q.

It must be stressed that this general bounce-averaged quasilinear operator given by Eq. (2.25) is applicable to a wide class of problems. In this thesis it is used for studying ICRF heating. It can also be used to study heating or current drive by lower hybrid waves or waves in the electron cyclotron range of frequencies. Preferential interaction of particles which stay in resonance longer, and creating of trapped particles, may play important effects in other applications of RF besides ICRF heating. It is conceivable that it may be used to determine the velocity space stability properties of a particular f (i.e., if there exists a wave which absorbs energy from f rather than giving energy to f, then an instability exists). It is often desirable to

include relativistic effects when considering electron behaviour. If this same bounce-averaging procedure were applied to the relativistic, uniform plasma quasilinear operator of Karney and Fisch,[84] the calculation of $\langle Q \rangle$ from first principles by Bernstein and Baxter should be recovered.

2.4.3 Bounce averaged particle conservation

The local particle density is:

$$n_{min} = \int d^3v f = \sum_\sigma \int_0^\infty dW \int_0^{W/B} d\mu \frac{2\pi}{m^2} \frac{B}{|v_\parallel|} f \tag{2.32}$$

Note that although f is independent of position (using only the lowest order $f = f_0$), n_{min} depends on position through B and $|v_\parallel| = \sqrt{\frac{2}{m}(W - \mu B)}$. The local density n_{min} is conserved by the local collision and quasilinear operators. It is not relevant to ask if the bounce-averaged collision and quasilinear operators conserve local density. For example, a trapped particle unable to reach $\theta = \pi/2$ in Fig. 2.1 may pitch angle scatter so that it can reach that point, adding to the local density there. The bounce-averaged quasilinear operator must conserve the total number (or equivalently, the average density) of particles on a flux surface. Consider the average density in a tube of flux $\delta\Phi = B\delta A$ (Fig. 2.3):

$$\langle n_{min} \rangle = \frac{\int d\ell \delta A n_{min}}{\int d\ell \delta A} = \frac{\delta\Phi \int \frac{d\ell}{B} n_{min}}{\delta\Phi \int \frac{d\ell}{B}}$$

Inserting the velocity space integral for n_{min}, we have:

$$\langle n_{min} \rangle = \frac{1}{\int \frac{d\ell}{B}} \int \frac{d\ell}{B} \sum_\sigma \int_0^\infty dW \int_0^\infty d\mu \frac{2\pi}{m^2} \frac{B}{|v_\parallel|} H(W - \mu B) f$$

$$= \sum_\sigma \int_0^\infty dW \int_0^\infty d\mu \frac{2\pi}{m^2} f \frac{1}{\int \frac{d\ell}{B}} \int \frac{d\ell}{|v_\parallel|} H(W - \mu B)$$

$$= \frac{2\pi}{m^2} \sum_\sigma \int_0^\infty dW \int_0^\infty d\mu f \frac{\tau_B}{\int \frac{d\ell}{B}}$$

The bounce time τ_B is defined as before, with the explicit incorporation of the Heaviside step function H in τ_B meaning that $\tau_B = 0$ for nonphysical

Figure 2.3: Sketch of a tube of constant flux $\delta\Phi$ along which particles are free to move. ℓ is the distance along the field line.

particles which have $\mu B > W$ everywhere. Since $\langle Q\rangle$ must conserve the total number of particles on a flux surface,

$$\frac{\partial\langle n\rangle}{\partial t} = \frac{2\pi}{m^2}\sum_\sigma \int_0^\infty dW \int_0^\infty d\mu \frac{\tau_B}{\int \frac{d\ell}{B}}\langle Q\rangle$$

$$= 0$$

we must be able to write $\langle Q\rangle$ as the divergence of a flux in (W,μ) space:

$$\langle Q\rangle = \frac{1}{\tau_B}\left(\frac{\partial\Gamma_W}{\partial W} + \frac{\partial\Gamma_\mu}{\partial\mu}\right) \tag{2.33}$$

$-\Gamma_W$ and $-\Gamma_\mu$ are the fluxes in the directions of increasing W and μ, respectively. From Eqs. (2.25) and (2.23) we see that $\langle Q\rangle$ can indeed be written in conservative form. Note that $\Gamma_\mu = \frac{ne}{\omega mc}\Gamma_W$ so that the direction of the quasilinear flow in phase space is independent of $k_\parallel$.

2.4.4 Transforming $\langle Q\rangle$ to energy and pitch angle space

Rather than working in (W,μ) space, it is convenient to transform $\langle Q\rangle$ to (W,ξ) space, where $\xi = \sigma\sqrt{1-\mu B_0/W}$. B_0 is the minimum magnetic field

along the field line. In the usual circular flux surface model of the tokamak, the minimum magnetic field is at the outside midplane, ($\theta = 0$ in Fig. 2.1). $\xi = (v_\parallel/v)_{B_0}$ is the cosine of the particle's pitch angle at the outside midplane. Denoting the old coordinates as (W', μ), we use the chain rule to transform the derivatives:

$$\frac{\partial}{\partial W'} = \frac{\partial}{\partial W} + \frac{\partial \xi}{\partial W'}\frac{\partial}{\partial \xi}$$
$$= \frac{\partial}{\partial W} + \frac{1-\xi^2}{2\xi W}\frac{\partial}{\partial \xi}$$

and

$$\frac{\partial}{\partial \mu} = \frac{\partial \xi}{\partial \mu}\frac{\partial}{\partial \xi}$$
$$= -\frac{B_0}{2\xi W}\frac{\partial}{\partial \xi}$$

So that L_{res} transforms to

$$L_{res} = m\frac{\partial}{\partial W} - (\xi^2 - \xi_*^2)\frac{m}{2\xi W}\frac{\partial}{\partial \xi}$$

ξ_* is defined by $\xi_*^2 = 1 - \frac{neB_0}{\omega mc}$ and is the pitch angle of particles whose banana tips lie in the un-Doppler-shifted resonance layer where $\omega = n\Omega$. The flux-surface averaged particle density is given in (W, ξ) coordinates by:

$$\langle n \rangle = \frac{4\pi}{m^2 B_0 \int d\ell/B}\int_0^\infty dW \int_{-1}^1 d\xi |\xi| W \tau_B f \tag{2.34}$$

Again using particle conservation, we must be able to write $\langle Q \rangle$ as the divergence of a flux in (W, ξ) space:

$$\langle Q \rangle = \frac{1}{|\xi| W \tau_B}\left(\frac{\partial \Gamma_W}{\partial W} + \frac{\partial \Gamma_\xi}{\partial \xi}\right) \tag{2.35}$$

The fluxes in the new (W, ξ) coordinate system are related to the fluxes in the old (W', μ) coordinate system by

$$\Gamma_W = W|\xi|\Gamma_{W'}$$

$$\Gamma_\xi = \sigma\frac{1-\xi^2}{2}\Gamma_{W'} - \sigma\frac{B_0}{2}\Gamma_\mu$$

More explicitly, the fluxes in Eq. (2.35) can be written as

$$\Gamma_W = W|\xi| \sum_{\vec{k},n} \frac{\pi e^2}{2m} \sum_{res} [\frac{v_\perp^2 |\Theta_{\vec{k},n}|^2}{|v_\parallel \frac{\partial}{\partial \ell}(k_\parallel v_\parallel + n\Omega)|}]_{res} L_{res} f \tag{2.36}$$

$$\Gamma_\xi = -\frac{(\xi^2 - \xi_*^2)}{2W\xi} \Gamma_W$$

2.4.5 $k_\parallel \neq 0$ resonance localization

We have already pointed out that $\langle Q \rangle$ only causes diffusion along certain characteristics, Eq. (2.24), so that as a particle is accelerated to high energy, its banana tip approaches the $k_\parallel = 0$ resonance layer. We also see this from the form of Γ_ξ in Eq. (2.36). The fascinating thing about this is that although the position of the true resonance ($\omega - k_\parallel v_\parallel = n\Omega$) is **different** for various particles and waves with different $k_\parallel v_\parallel$, all particles have their banana tips approach the **same** position as they are accelerated to high energy.

To understand the origins of this effect, consider the constraint in Eq. (2.16). As a particle passes through a resonance, the RF induced change in $v_\parallel$ is related to the change in $v_\perp^2$ by

$$\delta v_\perp^2 + 2(v_\parallel - \omega/k_\parallel)\delta v_\parallel = 0.$$

Defining the cyclotron frequency at resonance as Ω_r, and using the resonance condition $\omega - k_\parallel v_\parallel = n\Omega_r$, we have

$$\delta v_\parallel = \frac{k_\parallel}{2n\Omega_r} \delta v_\perp^2 \tag{2.37}$$

This change in $v_\parallel$ and $v_\perp$ will move the position of the banana tip. Since $B \propto 1/R$, the banana tip position can be parameterized by the magnetic field at the banana tip, defined by $B_{tip} = E/\mu = B_r(1 + v_\parallel^2/v_\perp^2)$. The change in B_{tip} is simply

$$\delta B_{tip} = B_r(\frac{2v_\parallel \delta v_\parallel}{v_\perp^2} - \frac{v_\parallel^2 \delta v_\perp^2}{v_\perp^4}) = B_r \frac{\delta v_\perp^2}{v_\perp^2}(\frac{k_\parallel v_\parallel}{n\Omega_r} - \frac{v_\parallel^2}{v_\perp^2})$$

But using the resonance condition and the definition of B_{tip}, this can be written in a form independent of $k_\parallel$ and B_r:

$$\delta B_{tip} = \frac{\delta v_\perp^2}{v_\perp^2}(\frac{\omega mc}{ne} - B_{tip})$$

where $\frac{\omega mc}{ne}$ is the magnetic field in the un-Doppler-shifted resonance layer. As particles are accelerated to higher energy ($\delta v_\perp^2 > 0$) their banana tips approach the $k_\parallel = 0$ resonance layer, despite the wave having $k_\parallel \neq 0$. Even a particle so deeply trapped that it does not reach the $k_\parallel = 0$ resonance layer (i.e., its banana tips lie to the low field side of the $k_\parallel = 0$ resonance layer), but which resonates with the wave because of the Doppler shift, will have its banana tips pushed towards the $k_\parallel = 0$ resonance layer. In Eq. (2.37), it is interesting to note that $\delta v_\parallel$ has the same sign as $k_\parallel$ so that the position of the real resonance $n\Omega = \omega - k_\parallel v_\parallel$ shifts outward to larger R each time the particle is accelerated to higher energy, for either sign of $k_\parallel$.

2.4.6 Bounce averaged power deposition

We now turn to calculating the power absorbed by the plasma from the wave. First we will calculate the general formula for the flux-surface averaged power, and then consider a simple limit where Stix's formula $P_{\rm Stix}$ can be recovered. Next, the local power deposition is calculated from the original Kennel-Engelmann Q, and evaluated in a simple limit to show the importance of Doppler-broadening in determining the central power deposition. By analogy with Eq. (2.34), we see that the flux-surface-averaged power is given by:

$$\langle P \rangle = \frac{2\pi}{m^2}\int_0^\infty dW \int_{-1}^1 d\xi \frac{2|\xi|W^2}{B_0}\frac{\tau_B}{\int \frac{d\ell}{B}}\langle Q \rangle$$

Using the general bounce-averaged quasilinear operator $\langle Q \rangle$ found in Eqs. (2.35) and (2.36), and integrating by parts twice, we can write $\langle P \rangle$ as

$$\langle P \rangle = \frac{2\pi}{m^2 B_0 \int \frac{d\ell}{B}}$$

$$\int_0^\infty dW \int_{-1}^1 d\xi 2|\xi| W f L_{res} \sum_{\vec{k},n}\sum_{res}[\frac{\pi e^2}{2m}\frac{v_\perp^2 |\Theta_{\vec{k},n}|^2}{|v_\parallel \frac{\partial}{\partial \ell}(k_\parallel v_\parallel + n\Omega)|}]_{res}$$

We will evaluate this in Stix's limit, $k_\parallel = k_\perp = E_\parallel = E_\perp = 0$ and $n = 1$, but retain f as an arbitrary function of W and ξ. In the $k_\parallel = 0$ limit, $v_{\parallel res}$ and $v_{\perp res}$ are given in terms of W and ξ by:

$$v_{\parallel res} = \pm\sqrt{\frac{2}{m}W(\frac{\xi^2 - \xi_*^2}{1 - \xi_*^2})}$$

$$v_{\perp res}^2 = \frac{2}{m}W\frac{1 - \xi^2}{1 - \xi_*^2} \tag{2.38}$$

For a tokamak with concentric circular flux surfaces where $B \propto B_\theta \propto 1/R$, we can calculate the integral $\int \frac{d\ell}{B} = r\int d\theta/B_\theta = 2\pi r/B_{\theta x}$, where $B_{\theta x}$ is the magnitude of the poloidal field at $\theta = \pi/2$ (Fig. 2.1). Carrying out the L_{res} operation and doing a little algebra leads to

$$\langle P\rangle = (\frac{P_{Stix}}{n_{min}})_{res}[\frac{2\pi}{m}\int_0^\infty dW\int_{-1}^1 d\xi|\xi|\sqrt{\frac{2}{m}W}\frac{H(\xi^2 - \xi_*^2)}{\sqrt{1 - \xi_*^2}\sqrt{\xi^2 - \xi_*^2}}f]$$

As mentioned before, the integrand has a singularity at $\xi^2 \to \xi_*^2$ (i.e., where $v_{\parallel res} \to 0$ and $\langle\delta W_\perp\rangle \to \infty$). Integrating over all pitch angles ξ, we see that this is an integrable singularity as long as f is finite at ξ_*. In fact, the expression in brackets is the definition of $(n_{min})_{res}$ (as can be seen by transforming the integral in Eq. (2.32) to (W, ξ) space). We are left with $\langle P\rangle = P_{Stix}$.

The simplicity of the expression for P_{Stix} in Eq. (2.6) is deceiving. P_{Stix} depends on the arbitrary $f(W, \xi)$ through $(n_{min})_{res}$. As the ICRF accelerates particles to high energies, it makes them trapped particles whose banana tips approach the resonance layer. This increases $(n_{min})_{res}$ and increases the RF damping. Pitch angle scattering prevents all the particle banana tips from being exactly in the resonance layer, and so prevents $(n_{min})_{res}$ from rising indefinitely. For typical experimental parameters, $(n_{min})_{res}$ usually reaches an equilibrium value only 10 to 30 % greater than its initial value. Preferential heating of some particles may make only a small difference in the total absorbed power, but as we can see from Fig. 1.2 in chapter 1, it makes a large difference in the shape of f.

2.4.7 Local power deposition

Quasilinear theory determines how f evolves for a given wave field $\vec{E}$. To be self consistent, this f should then be used to calculate the dispersion relation which determines how the waves propagate. If one includes the magnetic field inhomogeniety and the rotational transform (so that $\frac{\partial\Omega}{\partial\ell} \neq 0$) in the wave equations, one is typically left with a nonlocal dispersion relation. The issue of whether or not the dispersion relation is strongly affected by particles whose turning points are very near the resonance layer is under investigation.[18–20]

It is well known that the damping calculated from the imaginary part of $\vec{k}$ from the dispersion relation is identical to the damping found from quasilinear theory — as long as the same f and $\vec{E}$ are used in each calculation. We just calculated the flux surface averaged damping $\langle P\rangle$ in terms of the bounce-averaged quasilinear operator $\langle Q\rangle$Now we will calculate the local power P to allow more direct comparison with wave propagation studies. Knowledge of the local power also helps elucidate the importance of finite $k_\parallel$ in determining the central power density. The local power P is difficult to extract from the first principles theories[68–70] because they average over bounce-motion from the start. But in our approach, the local damping rate is immediately attainable from the local Kennel-Engelmann Q as long as we use the proper $f(W,\mu) = f(\frac{m}{2}(v_\perp^2 + v_\parallel^2), \frac{mv_\perp^2}{2B(\ell)})$. The power absorbed by the particles is

$$\begin{aligned} P &= \frac{\partial}{\partial t}\int d^3v W f \\ &= \int_0^\infty dv_\perp 2\pi v_\perp \int_{-\infty}^\infty dv_\parallel \frac{m}{2}(v_\parallel^2 + v_\perp^2) Q \end{aligned}$$

Integrating by parts leads to

$$\begin{aligned} P &= -\sum_{n\vec{k}} \frac{\pi e^2}{2m^2}\int_0^\infty dv_\perp 2\pi v_\perp \int_{-\infty}^\infty dv_\parallel m v_\perp^2 \delta(\omega - k_\parallel v_\parallel - n\Omega)|\Theta_{n\vec{k}}|^2 L f \\ &= -\sum_{n\vec{k}} \frac{\pi e^2}{2m^2}\int_0^\infty dv_\perp 2\pi v_\perp m v_\perp^2 |\Theta_{n\vec{k}}|^2 \\ &\quad \times \frac{1}{|k_\parallel|}\left[(1 - \frac{k_\parallel v_\parallel}{\omega})\frac{1}{v_\perp}\frac{\partial f}{\partial v_\perp} + \frac{k_\parallel}{\omega}\frac{\partial f}{\partial v_\parallel}\right]_{v_\parallel = \frac{\omega - n\Omega}{k_\parallel}} \end{aligned}$$

This is the general expression for the local damping rate.

2.4.8 Doppler-broadening of the power profile

To emphasize the effects of $k_\parallel$ on the power deposition, we will use a simple limit similiar to Stix's, taking $k_\perp = E_\parallel = E_\perp = 0$ and $n = 1$ while allowing $k_\parallel \neq 0$. To make analytic progress, we will assume that f is Maxwellian,

$$f = n_{min} \frac{e^{-\frac{v_\perp^2 + v_\parallel^2}{2v_t^2}}}{(2\pi v_t^2)^{3/2}}$$

where $v_t^2 = T/m$. Using $Lf = -f/v_t^2$ and performing the $v_\perp$ integral leads to

$$P = n_{min} \frac{\pi e^2}{m} |E_+|^2 \frac{e^{-\frac{(\omega-\Omega)^2}{2k_\parallel^2 v_t^2}}}{\sqrt{2\pi k_\parallel^2 v_t^2}}$$

Assuming that $\Omega = \omega R_{res}/R$ gives

$$P = n_{min} \frac{\pi e^2}{m} |E_+|^2 [\frac{e^{-\frac{\omega^2(R-R_{res})^2}{2k_\parallel^2 R^2 v_t^2}}}{\sqrt{2\pi k_\parallel^2 v_t^2}}] \tag{2.39}$$

The damping occurs in a narrow region around the $\omega = \Omega$ resonance layer of width $\Delta \sim 2k_\parallel R v_t/\omega$. In the limit of $k_\parallel \to 0$, the expression in brackets can be written as a δ function, so that

$$\lim_{k_\parallel \to 0} P = n_{min} \frac{\pi e^2}{m} |E_+|^2 \delta(\omega \frac{R - R_{res}}{R_{res}}) \tag{2.40}$$

The power averaged over a flux surface is defined by

$$\langle P \rangle = \frac{\int_0^{2\pi} d\theta \, r 2\pi (R_x + r\cos\theta) P}{\int_0^{2\pi} d\theta \, r 2\pi (R_x + r\cos\theta)}$$

If the limits of integration extend to either side of the resonance by more than the resonance layer width , then we may use the δ function approximation for P in Eq. (2.40) to find

$$\langle P \rangle = [n_{min} \frac{ec}{B} |E_+|^2 \frac{R}{r|\sin\theta|}]_{res} \frac{R_{res}}{R_x}$$

and Stix's result in Eq. (2.6) is again recovered. From Eq. (2.39) the maximum value P attains is for $R = R_{res}$, and therefore the maximum possible $\langle P \rangle$ is

$$\langle P \rangle = n_{min} \frac{\pi e^2}{m} |E_+|^2 \frac{1}{\sqrt{2\pi k_\parallel^2 v_\parallel^2}} \tag{2.41}$$

$$= n_{min} \frac{ec}{B} |E_+|^2 \frac{R\sqrt{2\pi}}{\Delta}$$

This the actual maximum value of $\langle P \rangle$ only when the resonance layer intersects the magnetic axis so that $|\sin\theta| = 1$. If $R_{res} \neq R_x$, the actual maximum value of $\langle P \rangle$ will be even less. For typical PLT parameters, $k_\parallel R = 9$ and $B \sim 20$ kG, the resonance layer width for thermal particles with $T \sim 1$ keV is only $\Delta \sim 3$ cm and can usually be ignored (except within ~ 3 cm of the place where $r|\sin\theta| \to 0$). In some heating modes, however, the minority species may become very energetic, sometimes reaching an effective temperature (define as two-thirds of the average energy) of ~ 100 keV in PLT. Most of these particles have their banana tips near the resonance layer, and $v_\parallel$ goes to zero at the banana tip, but pitch angle scattering may lead to an average parallel energy at resonance $\sim 10keV$. The resulting resonance layer width $\Delta \sim$ 10–15 cm plays a dominant role in determining the central RF power density.

2.4.9 The bounce averaged collision operator

It is only possible to obtain analytic solutions of the general bounce-averaged Fokker-Planck equation

$$\frac{\partial f}{\partial t} = \langle Q \rangle + \langle C \rangle + \langle S \rangle$$

in special cases. Numerical solutions for the general case will be discussed in the section III.F. Here we will derive analytic solutions for fundamental and 2^{cd} harmonic heating in the $k_\parallel = 0$, small $k_\perp \rho$ limit. $\langle S \rangle$ represents sources and sinks, such as neutral beam injection, charge exchange losses, and unconfined orbit losses, which are included in the numerical solutions but will be ignored here. The linearized collision operator appropriate for

energetic ions moving faster than thermal ions but slower than the electrons (i.e. $T_i/m_i \ll W/m \ll T_e/m_e$) is

$$C = \frac{1}{\sqrt{W}}\frac{\partial}{\partial W}[\frac{2}{\tau_s}(W^{3/2}+W_c^{3/2})f + \frac{2}{\tau_s}(W^{3/2}T_e + W_c^{3/2}T_i)\frac{\partial f}{\partial W}] \qquad (2.42)$$

$$+\frac{\nu_{ii}}{2}\frac{\partial}{\partial \eta}[(1-\eta^2)\frac{\partial f}{\partial \eta}]$$

This standard collision operator has been extensively used to study the slowing down of energetic ions from neutral beam injection.[85,83] In terms of the notation used in the *NRL Plasma Formulary*,[87] $1/\tau_s = \nu_s^{i/e}$, $\nu_{ii} = 0.5 \times \nu_\perp^{i/i'}$, and $(W_c^{3/2}T_i)/(W^{3/2}W\tau_s) = \nu_\parallel^{i/i'}$.

In Eq. (2.42), the first term proportional to f on the right hand side represents drag due to collisions with electrons and ions. The second terms represents energy diffusion due to collisions with electrons and ions. The last term represents pitch angle scattering off of the ions. Here, $\eta = v_\parallel/v$ is the local pitch angle. In this equation, the slowing down time due to electrons is given by

$$\tau_s = \frac{6.32\times 10^8 A T_e^{3/2}}{n_e Z^2 \log\Lambda_{ie}}\text{sec} \qquad (2.43)$$

The terms proportional to W_c and ν_{ii} are due to collisions with the background ions. Ion drag exceeds electron drag for energies $W < W_c$. The critical energy W_c and the ion-ion collision frequency ν_{ii} are given by

$$W_c = A14.8T_e\left\langle\frac{Z_i^2}{A_i}\right\rangle^{2/3}$$

$$\nu_{ii} = \frac{9.009\times 10^{-8} Z^2 Z_{eff} n_e \log\Lambda_e}{A^{1/2}E^{3/2}} sec^{-1}$$

$$Z_{eff} = \langle Z_i^2\rangle = \frac{\sum_i n_i Z_i^2 \log\Lambda_i}{n_e\log\Lambda_e}$$

$$\left\langle\frac{Z_i^2}{A_i}\right\rangle = \frac{\sum_i n_i (Z_i^2/A_i)\log\Lambda_i}{n_e\log\Lambda_e} \qquad (2.44)$$

In these formulas, all energies and temperatures are in eV, while masses A_i are in AMU. Note that for $T_e = T_i$, the steady state solution $C(f) = 0$ is

$\partial f/\partial\eta = 0$ and $\partial f/\partial W = -f/T_i$. C forces f to relax to a Maxwellian. The use of a linearized collision operator means that collisions between the energetic ions are ignored. This is valid for minority heating when $n_{min} \ll n_i \sim n_e$. It is also justified for majority second harmonic heating if the distribution function can be separated into a Maxwellian plus an energetic tail, $f = f_{Max} + f_{tail}$, and if $n_{tail} \ll n_{Max}$. The bounce average of C calculated by Cordey is

$$\langle C\rangle = \frac{1}{\sqrt{W}}\frac{\partial}{\partial W}[\frac{2}{\tau_s}(W^{3/2}+W_c^{3/2})f + \frac{2}{\tau_s}(W^{3/2}T_e + W_c^{3/2}T_i)\frac{\partial f}{\partial W}]$$
$$+\frac{\nu_{ii}}{2|\xi|\tau_b}\frac{\partial}{\partial\xi}\left[(1-\xi^2)\frac{\tau_b}{|\xi|}\left\{\left\langle\frac{B_0}{B}\right\rangle - (1-\xi^2)\right\}\frac{\partial f}{\partial\xi}\right] \qquad (2.45)$$

2.4.10 $\langle Q\rangle$ in the isotropic limit

If pitch angle scattering is rapid enough to force f to be isotropic, $\frac{\partial f}{\partial\xi} = 0$, but f is still allowed to be an arbitrary function of energy, then we can simplify $\langle Q\rangle$ by averaging it over pitch angle. When f is independent of pitch angle, then the flux surface averaged density is the same as the local density

$$n_{min} = 4\pi\int_0^\infty dv\, v^2 f = \frac{\sqrt{2}}{m^{3/2}}4\pi\int_0^\infty dW\, W^{1/2} f$$

Comparing this with Eq. (2.34) indicates that the proper weighting for the pitch angle average is :

$$\langle\ldots\rangle_\xi = \frac{\int_{-1}^1 d\xi|\xi|\tau_b\ldots}{\sqrt{\frac{2m}{W}}B_0\int\frac{d\ell}{B}}$$

Note that $\langle f\rangle_\xi = f$ for isotropic f. We use the conservative form for $\langle Q\rangle$ in Eqs. (2.35) and (2.36). The Γ_ξ term integrates to zero, leaving

$$\langle\langle Q\rangle\rangle_\xi = \frac{1}{\sqrt{2m}B_0\int\frac{d\ell}{B}}\frac{1}{\sqrt{W}}\frac{\partial}{\partial W}$$
$$\times\left\{W\sum_{n\vec{k}}\frac{\pi e^2}{2}\sum_{res}\int_{-1}^1 d\xi|\xi|\left[\frac{v_\perp^2|\Theta_{n\vec{k}}|^2}{|(nv_\parallel - k_\parallel\mu\frac{c}{e})\frac{\partial\Omega}{\partial\ell}|}\right]_{res}\frac{\partial f}{\partial W}\right\} \qquad (2.46)$$

In order to do the ξ integral analytically, we restrict our attention to the $k_\parallel = 0$ limit and consider fundamental ($n = 1$) heating with $|\Theta|^2 = |E_+|^2$. Using Eq. (2.38) to express $v_{\parallel res}$ and $v_{\perp res}$ in terms of W and ξ, using Eq. (2.5) to express $|E_+|^2$ in terms of P_{Stix}, and doing a little algebra leads to

$$\langle\langle Q\rangle\rangle_\xi = \frac{1}{\sqrt{W}}\frac{\partial}{W}[W^{3/2}\frac{2P_{Stix}}{3n_{min}}\frac{\partial f}{\partial W}] \tag{2.47}$$

This is identical to the isotropic, flux surface averaged quasilinear operator found by Stix. The equivalence of the flux surface averaging and the bounce averaging techniques in the isotropic limit can be traced back to Eq. (2.18) where $\partial f_1/\partial \ell \to 0$ in this infinite pitch angle scattering limit.

Stix's isotropic quasilinear operator can be extended to include what is commonly called second harmonic ($n = 2$) heating by making a small $k_\perp \rho$ expansion in Eq. (2.46) for $|\Theta|^2 = |E_+|^2 J_1^2 \approx \frac{|E_+|^2}{4}(\frac{k_\perp v_\perp}{\Omega})^2$ to find

$$\langle\langle Q\rangle\rangle_\xi = \frac{1}{\sqrt{W}}\frac{\partial}{\partial W}[W^{3/2}\frac{k_\perp^2 W}{5\Omega^2 m}\frac{2P_{Stix}}{3n_{min}}\frac{\partial f}{\partial W}] \tag{2.48}$$

This equation is valid for general 2^{cd} harmonic heating, even when there is no minority fundamental heating. P_{Stix}/n_{min} is independent of n_{min} and is just used as a mnemonic for the constants in Eq. (2.6). In chapter 3 we will be considering deuterium second harmonic heating. There is always a small amount of hydrogen in the plasma, and since $\Omega_H = 2\Omega_D$, second harmonic deuterium heating must compete with fundamental hydrogen heating. From Eq. (2.47) we find that the power absorbed by the hydrogen $P_H = P_{Stix}$, while from Eq. (2.48) we find that the power absorbed by the deuterium is

$$P_D = \int d^3v W \langle\langle Q\rangle\rangle_\xi = P_H \frac{n_D}{n_H}\frac{k_\perp^2 \rho_D^2}{2} \tag{2.49}$$

where $\rho_D^2 = T_D/(m_D\Omega_D^2)$. Using a simple Alfvén wave dispersion relation, $k_\perp \approx \omega/v_A$, leads to the convenient formula[86]

$$\frac{P_D}{P_H} = \frac{\beta_D}{\eta_H} \tag{2.50}$$

The hydrogen concentration $\eta_H = n_H/n_D$ must be less than the deuterium beta $\beta_D = 8\pi n_D T_D/B^2$ in order for deuterium second harmonic heating to

exceed hydrogen fundamental heating. In most previous experiments, the hydrogen absorption was so strong that no evidence of direct RF heating of the deuterium could be obtained. It is useful to rewrite Eq. (2.48) in terms of P_D by using Eq. (2.49), so that the isotropic quasilinear operator for second harmonic deuterium heating can be written as:

$$\langle\langle Q\rangle\rangle_\xi = \frac{1}{\sqrt{W}}\frac{\partial}{\partial W}[W^{3/2}\frac{2W}{5T_D}\frac{2P_D}{3n_D}\frac{\partial f_D}{\partial W}] \tag{2.51}$$

2.4.11 Stix's isotropic solution for minority heating

Stix combined Eqs. (2.47) and (2.42) to find the steady state distribution function reached when fundamental minority heating is balanced by collisions. He then calculated the fusion reaction rate enhancement due to the energetic RF produced tail. We define f_1 as the isotropic, steady state solution for fundamental heating, and f_2 (to be considered in the next section) as the isotropic, steady state solution for second harmonic heating. Stix found that f_1 must satisfy

$$(-\frac{d\log f_1}{dW})^{-1} = T_1(W) = \frac{W^{3/2}T_{tail} + W_c^{3/2}T_i}{W^{3/2} + W_c^{3/2}}$$

$$T_{tail} = T_e + \frac{P_{Stix}}{3n_{min}}\tau_s \tag{2.52}$$

It is possible to analytically integrate $d\log f_1/dW$ to find

$$\log f_1(W) = -\int^W \frac{dW}{T_1(W)}$$

$$= \text{const} + \frac{W}{T_{tail}}\left\{1 + \left(\frac{T_{tail}}{T_i} - 1\right)H\left(\frac{W}{W_c}\left(\frac{T_{tail}}{T_i}\right)^{2/3}\right)\right\} \tag{2.53}$$

where, as Stix found,

$$H(x) = \frac{1}{x}\int_0^x \frac{du}{1+u^{3/2}}$$

$$= \frac{1}{x}\left\{\frac{1}{3}\log[\frac{x-\sqrt{x}+1}{x+2\sqrt{x}+1}] + \frac{2}{\sqrt{3}}\arctan(\frac{2\sqrt{x}-1}{\sqrt{3}}) + \frac{\pi}{3\sqrt{3}}\right\}$$

In the high energy limit, $W >> W_c$ and $T_{\rm Tail} >> T_e$, where the RF input power is balanced by drag on the electrons, f_1 has a particularly simple form:

$$f_1(W) \propto e^{-W/T_{\rm Stix}} \tag{2.54}$$

$$T_{\rm Stix} = \frac{P_{\rm Stix}}{3n_{\rm min}}\tau_s$$

(Stix's paper defines a parameter ξ, which is equivalent to our $T_{\rm Stix}/T_e$.) The physical interpretation of this is clear. The tail temperature is proportional to the energy absorbed in a slowing down time. Energy goes into the minority species at a rate P_{Stix}. Energy is lost from the minority species at the rate $\frac{3}{2}n_{min}T_{\rm Stix}\frac{2}{\tau_s}$. Balancing the two just gives Eq. (2.54). It must be emphasized that the analytic solutions given by Eqs. (2.53) and (2.54) are valid only if unconfined orbit losses can be ignored. It is easy to make the hydrogen minority in PLT so energetic that unconfined orbit losses carry away most of the RF power. The effects of unconfined orbit losses will be considered in more detail in Sec. 2.5.

2.4.12 Second harmonic isotropic solution

With Eq. (2.51) we can extend Stix's steady state solution to second harmonic heating, finding that

$$(-\frac{d\log f_2}{dW})^{-1} = T_2(W) = \frac{W^{3/2}T_{tail} + W_c^{3/2}T_i}{W^{3/2} + W_c^{3/2}}$$

$$T_{tail} = T_e + (\frac{2W}{5T_D})\frac{P_D}{3n_D}\tau_s \tag{2.55}$$

where T_{tail} is now energy dependent. Chapter 3 will show that this simple formula describes the shape of the observed fast neutral spectra quite well. We will obtain a measure of the central deuterium power density by adjusting P_D in this formula to obtain a best fit to the data. It would be useful to have an analytic solution for f_2, but to date we have had to rely on numerically integrating Eq. (2.55). The symbolic manipulation program MACSYMA is able to find an analytic solution only if the T_e contribution to T_{tail} is ignored, an approximation which is probably valid for most cases of interest. The symbolic manipulation program SMP (version 1.5.0) claims to be able to find

the analytic solution including the T_e term, but unfortunately SMP's solution is wrong. Simple solutions for f_2 can be obtained in several interesting limits. At high energies, $W >> W_c$, the solution is

$$f_2 \propto \frac{1}{W^\alpha}$$

$$\alpha = \frac{15 T_D n_D}{2 P_D \tau_s}$$

and f_2 is normalizable in the sense of containing a finite amount of energy only for $\alpha > 5/2$. (Using Eqs. (Eq. (2.49)) and (Eq. (2.52)), and assuming $k_\perp = \omega/v_A$, we find that $\alpha = \frac{5}{2} v_A^2 m_H / T_{\mathrm{Stix,H}}$. This curious result says that a normalizable deuterium distribution function exists if the average hydrogen speed is less than the Alfvén speed.) Actually, Eq. (2.55) breaks down at very high energies because the small $k_\perp \rho$ expansion of J_1 which led to Eq. (2.48) is no longer valid. Inclusion of the full J_1^2 term would prevent "runaway" at high energies, and would always lead to normalizable f. At lower energies, Eq. (2.55) accurately describes the shape of f_2. For $W \ll W_c$ and $T_{tail} \gg T_e$, we have

$$T_2(W) \approx T_i + \frac{2}{5} \frac{W^{5/2}}{W_c^{3/2} T_D} \frac{P_D}{3 n_D} \tau_s \tag{2.56}$$

which is independent of electron temperature. Notice that W_c for deuterium is twice W_c for hydrogen. This means that for $T_e \sim 1$ keV and $\langle Z_i^2/A_i \rangle \sim .5$, Eq. (2.54) is valid for hydrogen above ~ 10 keV, while Eq. (2.56) is valid for deuterium below ~ 20 keV. These two analytic solutions are very useful for qualitative analysis of the data.

2.4.13 Singularities and pitch angle scattering

At high energies, the measured distribution functions are highly anisotropic and we must resort to numerical solutions (discussed in the next section) to provide any quantitative comparison with theory. Several workers have attempted to find analytic solutions which model the anistropic nature of f at high energies,[77,78] but they seem to gloss over the handling of the singularity in $\langle Q \rangle$ at $v_\parallel \to 0$ (in the $k_\parallel = 0$ limit). Although $\langle Q \rangle \propto \frac{1}{v_{\parallel res}} \propto \frac{1}{\sqrt{\xi^2 - \xi_*^2}}$ is singular at $\xi \to \xi_*$, it is apparent that when $\langle Q \rangle$ is averaged over any finite

region in pitch angle space, a well behaved answer will result:

$$\int_{\xi_*}^{\xi_*+\Delta} d\xi \xi \langle Q \rangle \sim \sqrt{(\xi_* + \Delta)^2 - \xi_*^2}$$

At low energies, pitch angle scattering is rapid enough to force f to be completely independent of ξ, justifying the averaging of $\langle Q \rangle$ over all pitch angles in Eq. (2.46). At higher energies, pitch angle scattering is no longer sufficient to keep f completely isotropic, but it is rapid enough to smooth f over small regions of ξ, thus justifying some sort of averaging of $\langle Q \rangle$ over small regions of ξ.

2.4.14 Model anisotropic solution

To understand better the importance of collisions in leading to a well behaved solution, consider the model equation

$$\frac{\partial f}{\partial t} = \frac{2TP_{Stix}}{3n_{min}} \delta(\xi) \frac{\partial^2 f}{\partial W^2} + \frac{2T}{\tau_s} \frac{\partial f}{\partial W} + \frac{\nu_{ii}}{2} \frac{\partial^2 f}{\partial \xi^2}.$$

The first term on the right hand side models the integrable singularity in $\langle Q \rangle$, replacing the $1/\sqrt{\xi^2 - \xi_*^2}$ singularity with a δ function singularity. The second term models collisional drag, while the third term models collisional pitch angle scattering. Despite the singular RF diffusion coefficient, it is easy to show that a steady state solution exists which does not run away to infinite energy:

$$f = e^{-|\xi|/\sigma_\xi} e^{-W/T_{tail}} \tag{2.57}$$

$$\sigma_\xi = \sqrt{\frac{\nu_{ii}\tau_s}{4}}$$

$$T_{tail} = \frac{P_{Stix}}{3n_{min}} \tau_s$$

Although this is only a model equation, it captures the general idea that including collisions will lead to reasonably behaved solutions. Any attempt to derive an analytic solution to the full bounce averaged Fokker-Planck equation should probably follow a similiar approach. A typical steady state numerical solution to the full equation $0 = \langle Q \rangle + \langle C \rangle$ is shown in Figs. 2.8–2.10. Its qualitative similiarity to Eq. (2.57) is evident. $f(\xi)$ exhibits peaks

near the singular pitch angle ξ_*. The width of f in ξ is fairly broad at low energies where pitch angle scattering is quite rapid and $\nu_{ii}\tau_s$ is large. At higher energies, ν_{ii} gets smaller, and the width σ_ξ also gets smaller. We suggest that a good model of the general solution is

$$f(W,\xi) = f_{iso}(W)K[e^{-|\xi-\xi_*|/\sigma_\xi} + e^{-|\xi+\xi_*|/\sigma_\xi}]$$

where $f_{iso}(W)$ is the solution of Eq. (2.52) for fundamental heating, or Eq. (2.55) for 2^{cd} harmonic heating. The pitch angle dependence is given by the expression in brackets, with the width $\sigma_\xi = \sigma_\xi(W)$ generalized to include ion drag:

$$\sigma_\xi = \sqrt{\frac{\nu_{ii}}{4}\tau_s(1+W_c^{3/2}/W^{3/2})^{-1}} = \sqrt{\frac{Z_{eff}}{4A\langle Z_i^2/A_i\rangle}(1+W^{3/2}/W_c^{3/2})^{-1}} \quad (2.58)$$

and the constant K chosen to give the proper normalization

$$f_{iso}(W) = \int_{-1}^{1} d\xi f(W,\xi)$$

It must be emphasized that this is only a model solution and has not been strictly derived from the equations. It is similiar to an analytic formula given by C.S. Chang, although his uses a Gaussian $\exp(-(\xi-\xi_*^2)^2/2\sigma_\xi^2)$ for the pitch angle dependence, and uses only the high energy fundamental heating limit for $f_{iso}(W) = exp(-W/T_{tail})$.

2.4.15 Singularities and Doppler-shifts

Doppler shifts also help lead to well behaved solutions despite the singularities in $\langle Q\rangle$. To understand how this occurs, consider the particle orbits and resonance curves plotted in Fig. 2.5 (this useful figure is due to Kerbel and McCoy[70]). As a particle moves along a field line to regions of different $\Omega \propto 1/R$, its parallel velocity varies because of energy and magnetic moment conservation according to:

$$\frac{v_\parallel}{v} = \sqrt{1-\frac{\mu B}{W}} = \sqrt{1-\frac{R_{tip}}{R}}$$

The solid line in Fig. 2.5 show some of these orbits for various values of R_{tip}. A resonance occurs wherever $\omega - k_\parallel v_\parallel = n\Omega$. Putting the resonance layer

Figure 2.4: Solid curves are particles orbits in (R,$v_\parallel/v$) phase space, where R is the major radius and $v_\parallel/v$ is the local pitch angle. Dotted curves show the location of resonances for different values of $k_\parallel v/\omega$.

through the center of the plasma so that $n\Omega = \omega R_x/R$, we can write the resonance condition as

$$\frac{v_\parallel}{v} = \frac{\omega}{k_\parallel v}(1 - \frac{R_x}{R})$$

The dashed lines in Fig. 2.5 show this resonance curve for various values of $\omega/(k_\parallel v)$. Points where the orbit curves and the resonance curves are tangent (called "tangent resonances" by Kerbel and McCoy) are points where $\langle Q \rangle$ has a singularity (unless the Airy function correction is included). For $k_\parallel = 0$, the orbit labelled by **B** in Fig. 2.5 (corresponding to a particle whose banana tip lies exactly in the resonance layer) experiences infinite RF acceleration and would runaway to infinite energies if it were not for pitch angle scattering which migh cause it to scatter into orbit **C**, for example. When $k_\parallel \neq 0$, the singularity in $\langle Q \rangle$ no longer appears at $v_\parallel = 0$, but instead appears at $v_\parallel = k_\parallel \mu c/(en)$ (orbit **A** in Fig. 2.5). For $k_\parallel \neq 0$, it is no longer necessary to call upon collisional pitch angle scattering to save particles from the fate of being accelerated to infinite energy. As orbit **A** is acclerated to higher energy, its resonance curve shifts out to higher $k_\parallel v_\parallel/\omega$, while its orbit shifts in towards orbit **B**. The particle originally at orbit **A** no longer has a tangent resonance.

We have discussed the importance of pitch angle scattering and $k_\parallel \neq 0$ in leading to well behaved solutions for f despite the singularity in $\langle Q \rangle$. The third effect which leads to well behaved solutions is the Airy function correction (Eqs. (2.30) and (2.9)) which prevents $\langle Q \rangle$ from actually becoming inifinite. But as argued after Eq. (2.13), the correction is only important for pitch angles extremely close to ξ_*, usually much closer than the width σ_ξ found in Eq. (2.58). It is therefore usually acceptable to ignore the Airy function corrections.

2.5 Obtaining Numerical Solutions

We have presented analytic solutions for the distribution function in several simple limits. These solutions help to clarify the physics, and can be directly compared with measurements in some cases. But we must use numerical solutions to treat the general bounce averaged quasilinear operator $\langle Q \rangle$. This also allows us to simultaneously include such important effects as less-than-perfect resonance localization, unconfined orbit losses, and radial profile effects on the charge exchange spectra. We have upgraded an existing bounce averaged Fokker-Planck computer program to include $\langle Q \rangle$. The original program was written by Rob Goldston[72,73] and has been used to study a number of aspects of neutral Beam injection.[102–104] We have also upgraded the program to model adiabatic compression experiments.[105,106]

2.5.1 The Basic Equations

The distribution function $f(W, \xi, r, t)$, a function of energy W, pitch angle ξ, radial position r, and time t, is found by solving the equation

$$\frac{\partial f}{\partial t} = \langle Q \rangle + \langle C \rangle + \langle S \rangle + \langle V \rangle.$$

The bounce averaged quasilinear operator $\langle Q \rangle$ is given by Eqs. (2.35–2.36). $E_+(r)$ and $|k_\parallel|$ ($\pm k_\parallel$ are used in the calculation) are specified as inputs to the calculation. E_- and $k_\perp$ are calculated from the cold plasma wave equation (Eqs. (3) and (5) from Stix[12]). $E_\parallel$ is ignored. The full Bessel function effects are included, so for fundamental heating $\langle Q \rangle \propto |E_+ J_0 + E_- J_2|^2$. Several authors have improperly treated the $E_- J_2$ term. Because E_- is usually much bigger than E_+, it can actually cause the tail absorption to increase in some cases. E_+ is held fixed in the calculation, while in the experiments it is the RF power which is held fixed. In some cases the calculation must undergo two or three iterations to produce the proper RF power.

To simulate the theoretically predicted power profile, we typically used $E_+^2(r) \propto (1 - r^2/a^2)$ in the quasilinear operator. As can be seen from the form of $P_{\rm Stix} \propto n_{\rm min}(r) E_+^2(r)/r$ in Eq. (2.6), this modestly peaked, parabolic E_+^2 profile leads to a very peaked power profile (even after including the Doppler-broadening effects of Sec. 2.4.8). Full wave[107] and ray-tracing[108] calculations

have been carried out for our experimental parameters, but they usually lead to power profiles which are comparable to, or more peaked than, the power profile obtained from a parabolic E_+^2 profile. We will find in the next two chapters that it is frequently necessary to assume a hollow $E_+^2(r)$ profile in order to produce a good fit to the charge exchange data.

The bounce averaged collision operator $\langle C \rangle$ is given by Eq. (2.45). The measured electron temperature, electron density, and ion temperature profiles, are used to calculate the various collision frequencies. (Detailed profile information was frequently unavailable, in which case we would usually assume $n_e \propto (1 - r^2/a^2)$ and $T_e \propto T_i \propto (1 - r^2/a^2)^2$.) A single impurity model is used to relate the hydrogenic depletion to $Z_{\rm eff}$, and to calculate the impurity contributions to the collision operator. $\langle V \rangle$ is the bounce average of $(e/m)E_\parallel(1 - Z/Z_{eff})\partial f/\partial v_\parallel$, and includes the combined effects of the toroidal electric field and the drift of the electrons.[109] This term is usually very small. Sources and sinks of particles in this problem are included in the $\langle S \rangle$ operator. Charge exchange losses are modelled by $-f/\tau_{\rm cx}$ and recapture of escaping neutrals is ignored. A source of low energy ions is arbitrarily included to replace charge exchange and unconfined orbit losses, maintaining the resonant ion density at a specified level. The original neutral beam injection source can also be turned on if desired.

A complete discussion of the boundary conditions can be found in the papers by Cordey[83] or Kerbel and McCoy.[70] f is symmetric in ξ in the trapped region of velocity space. Conservative boundary conditions are used for the flow across the trapped-passing boundary, and at $\xi = \pm 1$. Conservative boundary conditions are usually used at $W = 0$, following McCoy's[110] treatment in accounting for the finite density and energy of the zone at $W = 0$. (During neutral beam injection simulations, the boundary condition at $W = 0$ is usually $\partial f/\partial W = 0$ to remove paritcles from the simulation once they have slowed down to thermal energies.) f is independent of ξ at $W = 0$.

Unconfined orbit losses are incorporated by forcing $f = 0$ along a loss boundary at high energies. The loss region is found from the excursion of a particle orbit from its average flux surface, i.e., co-passing orbits are approximately circular but shifted outward in major radius from their average flux surface, and counter-passing orbits are shifted to smaller major radius, while trapped orbits execute the usual banana orbits. Barely trapped orbits have

the largest excursion from their average flux surface and become unconfined at the lowest energy. The present loss model is an improvement over the model in the original program which ignored losses for co-passing orbits (unless they were unconfined when first injected into the plasma by the neutral beam or they pitch angle scattered onto an unconfined trapped orbit), and underestimated the orbit shift of counter-passing orbits (unless $\xi = -1$).

2.5.2 Finite Banana Width Effects

Although finite banana width effects are included in calculating unconfined orbit losses and the charge exchange *spectra* (see below), the actual Fokker-Planck calculation is carried out in the zero-banana width limit on the average flux surface of a particle orbit. Particles may undergo pitch angle scattering, but they remain on the same average flux surface. There is no radial transport of particles. The calculation of f is done independently on a number of radial zones, and the radial dependence of f comes from the radial dependence of the RF power profile and the collision frequencies. Charge exchange *losses* are calculated using the neutral density on the average flux surface. The zero banana width assumption can be traced back to the starting point of our derivation of the bounce averaged equations, Eq. (2.17), where we used only the parallel particle motion $v_\parallel \partial f/\partial \ell$ and ignored cross-field drifts which would lead to $(v_\parallel \hat{b} + \vec{v}_D) \cdot \nabla f$. Given the importance of the radial transport mechanisms outlined in Sec. 2.5, an interesting—and challenging—area for future research is the inclusion of radial transport in a Fokker-Planck program.

The finite banana width effect on charge exchange detection is perhaps best illustrated with a simple figure, Fig. 2.7, although we will give the explicit formula below. The particle charge exchanges at point **B**, not **A**. Not only is the neutral density different at the two points, but the pitch angle of the particle at the two points is different as well. These finite banana width effects are essential for reproducing the "negative temperature" feature of Fig. 1.2. The charge exchange flux $f_{cx}(W, R_{tan}, t)$ is related to the $f(W, \xi, r, t)$

Figure 2.5: Sketch of the zero banana width orbit (**A**) used in the bounce averaged Fokker-Planck calculation, and the equivalent banana orbit (**B**) used to calculate the charge exchange spectrum and to include unconfined orbit losses.

calculated by the Fokker-Planck code by

$$f_{cx}(W, R_{tan}, t) = \int ds\, f(W, \xi(R(s), R_{tan}, \Delta), r(R(s), \Delta), t) \times n_0(r(R(s), \Delta)) \frac{\langle \sigma v \rangle_{cx}}{\sigma_{cx} v} e^{-\lambda}.$$

This is the same as Eq. (1.4), except we have explicitly noted that the position $R(s)$ and pitch angle $R_{tan}/R(s)$ of the particle when it charge exchanges is different from the average radius r and the ξ (defined as the pitch angle at $R_{maj} + r$) of the particle. The shift of the particle orbit is Δ, so that $r(R(s), \delta) = R(s) - \Delta - R_{maj}$. The particle's pitch angle at $R_{maj} + r$ can easily be found from the pitch angle where it is detected by conservation of energy and magnetic moment.

The FRANTIC[111] subroutine package is used to calculate the neutral density profile $n_0(r)$. FRANTIC assumes toroidal and poloidal symmetry. In addition to the T_e, T_i, and n_e profiles, one must specify an edge neutral density and temperature. The edge neutral density was typically adjusted to give a particle confinement time of ~ 30 msecs, but this is not critical as it only affects the magnitude and not the shape of $n_0(r)$. The edge neutral

temperature does effect the shape of the neutral density near the plasma edge. But the edge neutrals charge exchange with ions near the edge of the plasma, and the overall shape of the neutral density profile is not affected very much by changes in the edge neutral temperature.

2.5.3 Numerical Methods

Adding the bounce averaged quasilinear operator to an existing Fokker-Planck program was not as simple as originally thought. It was quite a challenge to find a method of obtaining physically meaningful solutions within a reasonable amount of computer time. Many of the techniques we used are discussed in the literature or in textbooks, but some were found only by trial and error.

We differenced the quasilinear and collision operators in conservative form,[110] which leads to a convenient discrete analog of particle and energy conservation.[112] In order to represent both the initial low energy Maxwellian and the tail which extends to very high energies for a small region of pitch angles, we used a variably spaced energy grid. With a grid spacing of $\delta W = 0.1$ keV at low energies, and $\delta W/W = .1$ at higher energies, it is possible to span 0 keV to 1 MeV with just 83 grid points. The convective (drag in the collision operator) term can easily lead to unphysical results if not treated in a special way. (See Patankar[113] and Karney[112] for a discussion of the relative merits of various treatments of the convective term). We used upstream differencing of the convective term $(\partial f_i/\partial t = \partial Af/\partial W = (A_{i+1}f_{i+1} - A_i f_i)/(W_{i+1} - W_i))$.

The original program used the ADI (alternating direction implicit) method[110,114] for solving the discrete equations. Briefly explained, the ADI methods treats all W derivatives implicitly and all ξ derivatives explicitly on the first time step. This is reversed on the second time step, with implicit treatment of ξ derivatives and explicit treatment of W derivaties. The ADI method provides a fast (because the standard tridiagonal inversion method can be used) and accurate solution to the original problem, which includes only the collision operator. However, the cross-derivative $(\partial^2 f/\partial W \partial \xi)$ terms in $\langle Q \rangle$ can not be treated implicitly, unless the whole 9-banded matrix is inverted. So, we first tried what was most straightforward: use the ADI method for all of the terms in $\langle Q \rangle$ except for the cross-derivative terms, which were treated explicitly on each step. We found that this approach was unstable

unless a very small time step was used, thus leading to large cpu requirements. We next tried a number of modified explicit schemes, including the method proposed by Saul'yev[115] None of these methods were any better, and some were even unconditionally unstable.

We can gain insight into this numerical instability problem by considering a simple example. $\langle Q \rangle$ causes diffusion only along certain characteristics (see Eq. (2.24)), meaning that there exists a coordinate system (x, y) such that

$$\left(\frac{\partial f}{\partial t}\right)_{\rm RF} = \langle Q \rangle = \frac{\partial}{\partial x} D \frac{\partial f}{\partial x},$$

which can be easily solved by the standard implicit method. Taking D to be a constant, and transforming this equation to (w, z) space by $w = x + y$ and $z = x - y$, yields a diffusion equation with cross-derivatives:

$$\left(\frac{\partial f}{\partial t}\right)_{\rm RF} = D\frac{\partial^2 f}{\partial w^2} + 2D\frac{\partial^2 f}{\partial w \partial z} + D\frac{\partial^2 f}{\partial z^2}.$$

Using the von Neumann method for stability analysis,[116] it is easy to show that the ADI method, with explicit treatment of the cross derivative term, is unstable unless $\delta t < D(\delta x)^2$. However, the implicit operator splitting method[115] is found to be unconditionally stable. In this method, the time advancement is done in three steps. On the first step, the cross-derivative term is advanced explicitly, while the other two terms are ignored. On the second step, the $\partial^2 f/\partial w^2$ is carried out implicitly and the other two terms are ignored. On the final step, the $\partial^2 f/\partial z^2$ terms are implicitly advanced with the other two terms ignored. We found that the implicit operator splitting method works not only for this model equation with constant D and straight line characteristics, but also for the full bounce averaged quasilinear operator.

Another numerical instability arose when finite $k_\parallel$ effects were added to the program. We were able to solve this problem by replacing the original differencing scheme for the cross derivatives (which was the same as Eq. (58) of McCoy et al.,[110] with a slightly more accurate scheme:

$$\frac{\partial}{\partial W} D \frac{\partial f}{\partial \xi} \approx \left\{ D_{i+1/2,j} \left[\frac{(f_{i+1,j+1} + f_{i,j+1})/2 - (f_{i+1,j-1} + f_{i,j-1})/2}{2\delta\xi_j} \right] \right.$$
$$\left. - D_{i-1/2,j} \left[\frac{(f_{i,j+1} + f_{i-1,j+1})/2 - (f_{i,j-1} + f_{i-1,j-1})/2}{2\delta\xi_j} \right] \right\} / \delta W_i.$$

This is the differencing method used in the more recent work of Kerbel and McCoy (see their Eq. (B5)).

The last topic to be considered in this numerical methods section is our treatment of the singularities in $\langle Q \rangle$ which occur whenever the expression in Eq. (2.27) is zero. As noted before, all of these singularities are integrable, leading to finite power absorption. We have also argued that the effects of these singularities are reduced by pitch angle scattering and Doppler broadening. In the program, we deal with these singularities by *analytically* averaging the singular expression Eq. (2.27) over the finite grid size. The singularity for $\partial\Omega/\partial\ell = 0$ (see Eq. (2.7)) is eliminated by integrating over the radial grid width. The singularity which occurs when the argument of the square root in Eq. (2.27) is zero is eliminated by averaging over the pitch angle grid width. This procedure avoids the numerical problems which occur when a singularity happens to fall exactly on a grid point, provides reasonably accurate answers even for large grid spacing, and converges to the exact answer in the limit that the grid spacing is made very small.

2.5.4 Sample Simulation

$f(W, \xi, r, t)$ is usually calculated on an 83 (energy) $\times$ 50 (pitch angle) grid for 10 radial shells. A typical simulation of 60 msecs of physical time is done with a 0.25 msec time step and takes $\sim$ 8 cpu minutes on a VAX 8600 computer. The time advancement algorithm requires $\sim$ 42 μs/meshpoint on the VAX 8600, which is $\sim$ 22 times slower than McCoy's[110] 1.9 μs/meshpoint on a CRAY-1. McCoy[117] has recently developed a fully implicit algorithm which inverts the full 9-banded (plus a few extra terms) matrix. This can speed up the search for equilibrium solutions by a factor of $\sim$ 30.

In Chapters 3 and 4 we will compare our bounce averaged quasilinear simulations with a large number of measurements. Here we will briefly look at one of those simulations, the one contained in Figs. 4.5, 4.9, 4.11, 4.16, and 4.17. For completeness, Appendix B provides the actual input files used for this simulation. (Much of the information there will be meaningless to most readers, or may not even be used in the simulation, but is given to provide a flavor of how it is performed.) Fig. 4.5 gives the main plasma parameters. Fig. 4.9 shows the $E_+(r)$ and $k_\perp(r)$ used in the simulations. The plots labelled Prf and Ttail in Fig. 4.9 are based on the simple Stix formula, and are not

Figure 2.6: Calculated f as a function of pitch angle $\xi = v_\parallel/v$ at $\theta = 0$, for various energies.

the actual RF power or minority temperature profiles. The actual RF power profile calculated by the full bounce averaged quasilinear code is shown in Fig. 4.17, along with a plot of where that power eventually goes (most of it unconfined orbit losses). Fig. 4.16 shows the final "temperature" (defined as two thirds of the average energy) of the hydrogen minority.

In Figs. 2.8-2.10 we show the equilibrium distribution function $f(W, \xi)$ calculated at the radius $r/a = .55$. The pitch angle dependence of f is qualitatively consistent with our model solution of sec. 2.4.14. At low energies it is isotropic in pitch angle, while at high energies it is very peaked around ξ_*, the pitch angle of particles whose banana tips lie exactly in the resonance layer. Integrating over a sightline through all radii produces the simulated charge exchange spectra of Fig. 4.11, which agrees fairly well with the data. The assumptions necessary to produce this good fit are described in Chapter 4.

Figure 2.7: Calculated f as a function of energy W, for various pitch angles.

Figure 2.8: Contour plot of f as a function of $v_\parallel/v_0$ and $v_\perp/v_0$ at $\theta = 0$, where the normalizing velocity v_0 is equivalent to 956 keV.

Chapter 3

The Transition to Deuterium Second Harmonic Heating

3.1 Introduction

As discussed in the exerimental review, Sec. 1.3.1, the original experiments on what was thought to be second harmonic deuterium heating turned out to be dominated by fundamental absorption by residual hydrogen. There was no evidence that any RF power was directly absorbed by the deuterium. It is theoretically expected (Eq. (2.50)) that the ratio of deuterium absorption to hydrogen absorption is given by $P_{\mathrm{D}}/P_{\mathrm{H}} = \beta_{\mathrm{D}} n_{\mathrm{D}}/n_{\mathrm{H}}$. We have not been able to reduce the hydrogen concentration to such a low level that deuterium damping dominates hydrogen damping. But we have been able to reduce it enough that $\sim 22\%$ of the central RF power is absorbed by the deuterium. We have seen a substantial deuterium tail, providing clear evidence of direct second harmonic deuterium heating. We have also studied the scaling of this tail with hydrogen concentration and total power.

The hydrogen concentration plays a critical role in hydrogen minority heating. It determines how energetic the hydrogen ions become, and therefore determines the relative magnitudes of electron heating, ion heating, and unconfined orbit losses. It also governs how much RF power flows directly to the deuterium, and how much power is mode converted into the ion Bernstein wave. Sections 3.2–3.4 describe a set of discharges from the same day in which the hydrogen concentration was varied from 1.6% to 55% while

monitoring the deuterium and hydrogen charge exchange spectra. All of the measurements from this day were made with the analyzer viewing perpendicular ($R_{tan} = 13$ cm in Fig. 1.3). No simultaneous scans of concentration and viewing angle were performed. At sufficiently high energies, the perpendicular view should give us a measure of the central hydrogen and deuterium energy distributions.

This concentration scan covers important unexplored territory because most of the experiments for this thesis were performed at very low hydrogen concentrations ($n_H/n_D \sim 1 - 3\%$) in order to avoid the complications of mode conversion, and maximize the amount of second harmonic deuterium heating. In addition, it was originally thought that the measurement of the hydrogen tail would be improved by operating at the lowest possible hydrogen concentration and making the tail as energetic as possible. The results of this chapter, however, show that the tail was already more energetic than necessary because of the large power available, and that by increasing the hydrogen concentration up to 10% the hydrogen signal could be dramatically improved.

This chapter begins with a discussion of our methodology for measuring the hydrogen concentration and the difficulties involved with this measurement. This is followed by a comparison of the bounce averaged Fokker-Planck code with the measured spectra from three discharges with low (2.8%), moderate (14%), and high (33%) hydrogen concentrations. The deuterium tail is found to increase as the hydrogen concentration is lowered, but not as quickly as the theory would predict. While it is necessary to invoke effects outside the scope of the present modelling at the lowest and highest hydrogen concentrations, the RF-produced tails observed at moderate hydrogen concentrations are consistent with the bounce averaged quasilinear simulations which ignore fast ion transport and assume a peaked RF power profile.

In order to be more quantitative about the scaling of the deuterium tail with hydrogen concentration, we have developed a new method of fitting the measured deuterium tail with an analytic model. This provides a quick, direct measure of the central deuterium power density, P_D. This measured P_D scales as $(n_D/n_H)^{\sim.59}$ in disagreement with the theoretical scaling which is proportional to n_D/n_H. This is subject to a number of possible explanations. It may be an indication that the RF power profile is broadening as the concentration is lowered. Or it may be that fast ion transport is increasing,

making the RF power profile appear to be broad. The broadest RF power profiles observed are consistent with sawtooth mixing of fast ions out to an $\sim$ 15 cm radius. If sawtooth mixing always occurs, then the apparently peaked profiles would actually be evidence of anomalous enhancement of second harmonic deuterium damping. Another alternative is that is is insufficient to model the time dependent sawtooth transport process as a time independent broadening of the RF power profile.

The scaling of the deuterium tail with total RF power is studied in Sec. 3.5. Three different power scans from three different days are presented, two of which show roughly linear scaling, while the third shows very nonlinear scaling. Plasma conditions did vary from day to day, and vary with power as well, but it is not known why the scaling was linear in some cases and nonlinear in others.

3.2 Measuring the Hydrogen Concentration

PLT would frequently run for weeks on end without any hydrogen gas being deliberately puffed into the machine. The residual hydrogen in the plasma probably came from recycling with the walls, leaks in the vacuum vessel, or water inside the vacuum vessel. Measuring the hydrogen concentration in these cases is difficult because it is so small, with $n_H/(n_D + n_H)$ in the range of 1 to 3%. The hydrogen concentration varies during a discharge, and may vary in minor radius, further complicating attempts to measure it.

There are several traditional methods for measuring the hydrogen concentration. Spectroscopic determination of the ratio of H_α to D_α line radiation provides a measure of the edge source of hydrogen and deuterium. The residual gas immediately after a discharge can be mass analyzed. Another traditional method is to measure the ratio of hydrogen to deuterium charge-exchange neutrals emitted by the plasma. We relied on this last technique, and used the measured density rise during hydrogen gas puffing as a cross-check.

3.2.1 Density rise measurements

In these experiments, the discharge was initiated in a prefill of deuterium, and the line average electron density was brought up to $1.2 \times 10^{13}\text{cm}^{-3}$ by

Figure 3.1: Relative electron density rise versus the "open" duration of the gas valve during hydrogen gas puffing.

deuterium gas puffing. Hydrogen would be puffed into the tokamak ~ 350 msec into the discharge. The amount of hydrogen puffed in was controlled by the amount of time the gas valve was open. The resulting density rise should be proportional to the amount of hydrogen introduced into the plasma. Figure 3.1 shows the size of the density rise versus the "gas valve duration." This relation is roughly linear, with the 18 msec offset due to the finite time the gas valve takes to open once its control voltage has been set. Unfortunately, the density rise is not a direct measure of the amount of hydrogen now in the plasma. Two effects are important. First, in the energy range of interest, the charge exchange rate exceeds the electron impact ionization rate by roughly a factor of 3. A deuterium ion can be replaced by a hydrogen ion via charge exchange without increasing the electron density. This effects tends to make Δn_e an underestimate of the hydrogen density. However, this tends to be offset by the second effect, in which the hydrogen can end up in the wall instead of in the plasma. The freshly puffed hydrogen ions in the edge of the plasma will not survive there for long, but will charge exchange in the high edge neutral density. Depending on an ion's direction at the time of its charge exchange, it can either penetrate further into the plasma, or leave the plasma and bury itself in the wall. Since the wall was already saturated with deuterium, this extra hydrogen in the wall forces the desorption

of deuterium from the surface of the wall. The result of this second effect is that puffing hydrogen into the tokamak could actually increase the amount of deuterium in the plasma. The relative importance of these two competing effects is unknown. Nevertheless, the density rise should be linearly proportional to the amount of hydrogen which finally ends up in the plasma and can be used as a check of the the charge exchange measurement of the hydrogen concentration.

3.2.2 Charge exchange measurements

In these experiments, two separate neutral particle analyzers were used. The horizontally scanning analyzer (CENA) described in chapter 2 was usually set to look at high energy ions (5 to 200 keV) produced by the RF, while a vertically scanning instrument (MACE) was set to look at lower energy ions (.75 to 5 keV) during the ohmic phase of the plasma to measure the hydrogen concentration. MACE (described in more detail in Ref. [58]) has 10 channels to look at a range of energies in a single shot, but only looks at 1 ion species at a time. On consecutive shots the instrument was set to measure the hydrogen distribution, the deuterium distribution, and the background level. A typical measurement with the verically scanning instrument in a case where n_H is of the same order as n_D is shown in Fig. 3.2. As described in Chapter 2, the quantity labelled as f on a charge exchange plot is related to the true distribution function f by

$$f_{cx} = \int d\ell f n_0 \frac{\langle \sigma v \rangle_{cx}}{\sigma_{cx} v} e^{-\lambda}.$$

We ignore the small $\langle \sigma v \rangle / (\sigma v)$ correction and the difference in attenuation between deuterium and hydrogen (a $\sim$ 30 % effect), and assume that the temperature profile of the hydrogen and deuterium are similiar. If the shape of the hydrogen and deuterium density profiles are similiar, then by taking the ratios of the hydrogen and deuterium charge exchange measurements, $f_{cx,H}/f_{cx,D}$, the neutral density n_0 cancels and we are left with $f_H/f_D = n_H/n_D$.

Figure 3.2: Typical hydrogen and deuterium spectra measured by MACE at high hydrogen concentration during the ohmic phase of the plasma.

Figure 3.3: Typical hydrogen and deuterium spectra measured by MACE at low hydrogen concentration during the Ohmic phase of the plasma.

3.2.3 Mass rejection problems at low hydrogen concentration

Although this measurement is straightforward when $n_H \sim n_D$, it requires some additional interpretation when $n_H \ll n_D$. Figure 3.3 shows a measurement of the hydrogen and deuterium spectra in this low hydrogen concentration limit. Taking the ratio of the charge exchange fluxes at 1.0 keV would give $n_H/n_D = e^{-5} = .7\%$, while the ratio at 3.5 keV would give $n_H/n_D = e^{-1.5} = 22\%$. The hydrogen appears to be much hotter than the deuterium. The slope of a straight line fit to the points above 1 keV gives a temperature 322 eV for the deuterium and 676 eV for the hydrogen. The intercept of this straight line fit is supposed[58] to be proportional to $n_i n_0/T_i^{3/2}$. Factoring out the temperature dependence and taking the ratio gives $n_H/n_D \sim .2\%$. It is difficult to conceive of how the hydrogen could be twice as hot as the deuterium in the ohmic phase of the plasma and we believe that the higher energy hydrogen signal is not real but is caused by a small amount of deuterium which was not rejected by the mass analyzing

magnetic and electric fields of the instrument. When the MACE instrument was originally calibrated, it was able to discriminate between hydrogen and deuterium with a rejection ratio of ∼1000:1. This means that when a beam of pure deuterium was fired at the analyzer, only .1% would be improperly detected as hydrogen. The small amount of deuterium of energy W which is not properly mass rejected appears to the detector to be hydrogen of energy $2 \times W$ and is significant relative to the low signal from the hydrogen tail ions at these energies. This is the main reason why the hydrogen temperature appears to be roughly 2 times the deuterium temperature.

It was hypothesized that although the higher energy hydrogen signal may be dominated by deuterium, the lower energy signal may still be giving an accurate hydrogen measurement. Ignoring the difference in stripping cross sections, the real hydrogen signal is proportional to $n_H \exp(-W/T)$, while the small fraction of deuterium which contaminates the hydrogen produces a signal proportional to $(n_D/R)exp(-W/(2T))$, where R is the rejection ratio and the factor of 2 in the exponential dependence accounts for deuterium of energy E masquerading as hydrogen of energy $2 \times E$. The ratio of the deuterium contamination to the real hydrogen signal is $(n_D/(n_H R))exp(+E/(2T))$, and we see that the contamination problem gets more severe at higher energies. For $n_D/n_H \sim 100$, $R \sim 1000$, $T \sim 300$ eV, we find that deuterium contamination swamps the real hydrogen signal above 1.4 keV. At energies under 1.4 keV, one might hope to be seeing mostly "real" hydrogen, especially since it is easier to strip low energy hydrogen than deuterium at even lower energy. The stripping efficiency is a function of E/m and begins to drop dramatically below a few keV/AMU. Even if the low energy hydrogen channels are partially contaminated, they at least provide an upper bound on the hydrogen concentration. This is useful because an upper bound on the hydrogen concentration provides a lower bound on the expected deuterium power density.

A further complication is that there were systematic differences in the n_H/n_D ratio measured by the .75 keV detector and the 1.0 keV detector of the MACE instrument. Because the .75 keV signal seemed to be less reproducible, and in order to provide a consistent definition of the charge exchange measurement of n_H/n_D, we always used the ratio as measured by the second detector at 1 keV.

Figure 3.4: Hydrogen concentration measured by MACE immediately after a hydrogen gas puff versus the relative electron density rise of the gas puff.

3.2.4 Comparing charge exchange and density rise measurements

One way to check our hypothesis that the 1 keV hydrogen signal was free from deuterium contamination is to compare the hydrogen concentration measured by this method with the density rise during a hydrogen gas puff. The results are shown in Fig. 3.4 (for the same discharges shown in Fig. 3.1). The linear relationship between the charge exchange hydrogen concentration measurement and the size of the density rise provides evidence for the usefulness of this measurement. The lowest points on this graph, with no hydrogen gas puff at all, had a charge exchange measured hydrogen concentration of $n_H/(n_H + n_D) = .6$–$.7\%$. During periods of strong deuterium gas puffing, hydrogen concentrations as small as .5–.6% have been measured. Thus it appears possible to measure hydrogen concentrations at least as small as .5–.6%, and that any determination of higher values is probably not affected by deuterium contamination errors. Getting a sufficiently high number of counts to measure such small hydrogen concentrations frequently requires the signal to be averaged over a fairly long time period ($\sim$50 msecs) or the measurements to be made at low plasma density (which usually occurs near the end of the discharge).

Figure 3.5: Hydrogen concentration measured by MACE after the RF versus that measured before the RF.

The hydrogen gas puff occurs at 350 msecs. The concentration measurements in Fig. 3.4 were made after the gas puff and before the RF began at 400 msecs. The RF power reached its peak power level at 450 msecs and stayed at that level until 650 msecs when the RF power began to decline. The RF power was completely off by 700 msecs. The concentration measurements were repeated after the RF, at about 750 msecs, and are compared with the pre-RF measurement in Fig. 3.5. For large hydrogen gas puffs, it is seen that 60% of the hydrogen has left the plasma by the time of the post-RF measurement. Without any hydrogen gas puff, the post-RF concentration is actually higher than the pre-RF concentration. This is probably because the pre-RF plasma is fueled by deuterium gas puffing while the post-RF plasma is fueled by recycling from the walls. Figure 3.5 shows that the hydrogen concentration can vary significantly in time during the discharge. We actually want to know the hydrogen concentration during the RF, but it is not possible to use straightforward charge exchange measurements because the RF causes the hydrogen to become highly non-Maxwellian. The measurements of the deuterium and hydrogen tails which will be presented in the following sections are made near the end of the RF (averaged over 575 to 625 msecs), so the post-RF concentration values are used instead of the pre-RF concentration. No attempt is made to interpolate between the pre

Figure 3.6: Hydrogen and deuterium spectra as measured by CENA at the same time as Fig. 3.2.

and post-RF values because this would tend to make the disagreement with theory (to be discussed in the following sections) worse than it already is.

3.2.5 Comparing two different charge exchange analyzers

The hydrogen concentration measurement is complicated further by the disturbing fact that the ratio of n_H/n_D measured by the CENA instrument was consistently 2 to 3 times bigger than the ratio measured by the MACE instrument. Most of the time, CENA was set to look at high energy neutrals and was unable to measure the hydrogen concentration in the ohmic phase of the plasma. Occasionally, CENA was set to look at lower energies and was able to make the concentration measurement. Figure 3.6 shows one of these measurements, made in the same discharge as Fig. 3.1. The CENA spectra give $n_H/n_D = \exp(1.2)$, while the MACE spectra give $n_H/n_D = \exp(.5)$, a factor of $\exp(.7) \sim 2$ lower. The horizontally scanning CENA instrument was looking through the center of the plasma (r=0) at a perpendicular viewing

Figure 3.7: Hydrogen to deuterium ratio measured by CENA versus that measured by MACE.

angle (Rtan=13 cm in Fig. 1.3). The vertically scanning MACE instrument was not looking through the center of the plasma, but was aimed at a tangency radius of r=20 cm, half way out in minor radius. This is assumed to be the reason the temperature in Fig. 3.6 is 500 eV, while Fig. 3.1 only gives 300 eV. MACE and CENA were mounted on adjacent ports, spaced 63 cm apart toroidally. MACE was 63 cm toroidally from the top limiter, while CENA was twice as far away. But unless the radial profile shape of $n_H(r)$ is different from the shape of $n_D(r)$, there should not be any difference in the ratios n_H/n_D measured by the two instruments.

Figure 3.7 shows that the difference between the two instruments was not an irreproducible anomaly on one shot, but the factor of 2 to 3 discrepancy persisted over a wide range of hydrogen concentrations. We had known about the discrepancy at low hydrogen concentration for some time, but had dismissed it as due to the poorer mass rejection of CENA. That the discrepancy persisted even at high concentrations was not discovered until well after the experiment was originally done, and MACE had been removed from the tokamak to make room for electron cyclotron heating ports. If the discrepancy had been known at the time of the experiments, one check would be to compare the hydrogen concentration measured by the vertically view-

ing MACE instrument at two different viewing angles, one with a tangency radius of r=20 cm, and one looking through the plasma center at r=0 cm. This would verify that the problem was indeed instrumental, and not due to a real radial variation of n_H/n_D. Another potential cause of the discrepancy is an error in one of the analysis programs for the two instruments. This was eliminated from consideration by looking at the raw count rate of 1 keV hydrogen and deuterium by the two instruments. This showed the same difference of a factor of 2. Other possible sources of the discrepancy could be incorrect stripping cell pressure measurements, or errors in the measured particle energy due to drifts in the analyzer magnetic field monitoring circuits.

Because CENA was calibrated more recently than MACE, and because an upper bound on n_H/n_D provides a useful lower bound on the expected deuterium power density, all MACE measurements of n_H/n_D quoted in the rest of this chapter have been scaled up by a factor of 2 to give agreement with CENA. The readings from the two instruments are linearly related to each other (Fig. 3.7), and the charge exchange measurement is linearly proportional to the density rise during hydrogen gas puffing (Fig. 3.4). For these reasons, the measurements can used to monitor relative changes in the hydrogen concentration if not the absolute magnitude as well.

3.3 Bounce Averaged Quasilinear Simulations of Three Discharges

In this section we make detailed comparisons between bounce averaged quasilinear theory and actual hydrogen and deuterium spectra from three discharges on February 26, 1985. These three discharges have hydrogen concentrations of $n_H/(n_H + n_D) = 2.8\%$, 14%, and 33%, and exhibit very different characteristics. Although there is good aggreement between the theory and the measurements for the intermediate concentration (14%), there are significant differences for very low or very high concentrations and effects not included in the present modelling must be invoked. In this section we will also compare the full bounce averaged quasilinear predictions with a simple isotropic model for the second harmonic deuterium tail. The agreement is very good and establishes the isotropic model as an accurate and quick way

to analyze measured deuterium spectra.

3.3.1 Assumed plasma parameters

A number of assumptions went into the bounce averaged quasilinear simulations. The main plasma parameters for these three discharges are listed in the table below.

Shot #	$\frac{n_H}{n_D+n_H}$	$\bar{n}_e$ $\frac{10^{13}}{\mathrm{cm}^3}$	P_{RF} MW	T_{e0} keV	T_{i0} keV	T_{neut} keV	P_{D} $\frac{mW}{cm^3}$	$\frac{Z_{\mathrm{eff,RF}}}{Z_{\mathrm{eff,OH}}}$
27063	0.028	1.25	1.10	1.20	2.10	3.5	78.	5.5/1.5
27085	0.142	1.15	1.20	1.60	1.67	2.0	29.	4.0/1.25
27107	0.332	1.10	1.12	1.32	1.23	1.5	13.	3.5/1.2

The quoted hydrogen concentrations were measured immediately after the RF using the MACE instrument (including the factor of 2 correction discussed in the preceding section) for the lowest 2 concentrations and the CENA instrument for the highest concentration measurement. Thomson scattering measurements were unavailable on this day. The electron density profile was assumed to be parabolic, $n_e(r) = n_{e0}(1 - r^2/a^2)$, normalized to the line averaged density measured by the microwave interferometer. PLT's electron cyclotron emission system[118] provides $T_e(t)$ at 10 different radii, but for simplicity we assumed that the electron temperature profile was parabolic squared, $T_e(r) = T_{e0}(1-r^2/a^2)^2$, and used only the peak electron temperature from the electron cyclotron emission.

The ion temperature profile was also assumed to be parabolic squared, with the peak ion temperature determined by the fitting the deuterium charge exchange spectra with the isotropic model discussed below. It is hoped that this model fit allows one to measure the ion temperature even in the presence of a non-Maxwellian tail. It is because of this tail that the temperature inferred from the neutron flux is significantly higher than the charge exchange measurement at the lowest concentrations. The values quoted for 'neutron T_i' assume not only a Maxwellian plasma but also assume a pure deuterium plasma with $n_D = n_E$. In most of our experiments, no other measure of T_i was available. In one experiment where FeXXV Doppler broadening mea-

surements were available, the impurity ion temperature was measured to be 1900 eV, the charge exchange ion temperature was 1500 eV, and the neutron ion temperature was 2700 eV. (This low charge exchange measurement is not unique to deuterium second harmonic heating, but is observed during other heating methods as well.) Due to lack of time, we have not systematically studied the neutron flux enhancement during deuterium second harmonic heating, which could be an interesting topic of future research.

Relative changes in Z_{eff} were inferred from an uncalibrated measurement of visible bremsstrahlung radiation.[119] Z_{eff} was observed to rise by a factor of $\sim$2.3–3.6 during RF heating, with the larger increases corresponding to lower hydrogen concentrations and lower electron densities. Because the results are not very sensitive to Z_{eff}, and there is some debate over the reliability of the visible Bremsstrahlung technique, a constant value of $Z_{eff} = 4$ was used for all of the simulations. Detailed spectroscopic measurements were not available on this day, so the plasma composition was unknown. Previous experiments[120] have shown that a significant fraction of the rise in Z_{eff} during ICRF heating is due to metallic impurities. An average impurity model was used in these simulations with $Z_{im} = 16$, $A_{im} = 32$. The hydrogenic ion density in the average impurity model is

$$\frac{n_H + n_D}{n_e} = \frac{Z_{im} - Z_{eff}}{Z_{im} - 1} = .8$$

which was assumed to be the same for all of the simulations. If the hydrogenic fraction actually does decrease as the hydrogen concentration is lowered and Z_{eff} rises, then the discrepancy between theory and experiment to be discussed later would be worsened. A .25 eV edge neutral temperature was used to predict the neutral profile, and the resulting profile was normalized to give a particle confinement time of $\sim$ 30 msec.

Unless noted otherwise, the simulations in this section assumed a very modest amount of focussing of the wave, $E_+ \propto (1 - r^2/a^2)^{1/2}$, roughly consistent with ray-tracing[108] and full wave calculations.[107] The coupling efficiency was assumed to be 85%, with 15% of the total input power lost to resistive heating of the antenna, the Faraday shield, and the walls. Electron damping was ignored, and E_+ for the simulations was iterated (if needed) so that the total RF power to the hydrogen and the deuterium added to 85% of the total input power. The outer flux surface had a minor radius of 41 cm and a major radius of 132 cm. The bounce averaged code uses a concentric circular flux

surface model. In an attempt to model the Shafranov shift, R =135 cm and a =38 cm was used in the simulations. The toroidal field at 135 cm was 20.12 kG, which placed the resonance layer at 138.4 cm. The plasma current for these discharges ranged from 420 to 470 kA.

3.3.2 14% hydrogen concentration

We will first discuss the 14% hydrogen concentration case, where the agreement between theory and experiment is best. Figures 3.8 and 3.9 show the deuterium and hydrogen spectra (measured simultaneously in the same discharge with a perpendicular view of the plasma), along with predictions of the bounce averaged quasilinear code for 3 different assumed values of the hydrogen concentration. The best fit to the data is obtained for a hydrogen concentration of 14%, in agreement with value measured by charge exchange techniques shortly after the RF was turned off. Lowering the assumed hydrogen concentration increases the second harmonic heating of the deuterium (in accord with Eq. (2.50)), and increases the theoretically expected deuterium tail. All three theoretical curves approach each other at low energies, where ion-ion collisions dominate and the slope of the theoretical curves is determined by the ion temperature.

Figure 3.8 also shows a fit to the data with the isotropic second harmonic model given in Eq. (2.55), which was numerically integrated to find $\log f(W)$. Assuming that T_e and W_c are known (using $< Z_i^2/A_i >= .5$ in the expression for W_c), there are 3 free parameters in the model: the usual intercept $\log f(0)$ and ion temperature T_i, plus a new RF parameter,

$$RFK = \frac{2W_c}{5T_D}\frac{P_D}{3n_D}\tau_s$$

A nonlinear least squares fitting routine using a grid search minimization algorithm[121] was used to find the values of these three parameters which provide a best fit to the data. From the expression for RFK, one obtains a measure of the second harmonic deuterium power density P_D without having to make any assumptions about $k_\perp$ or about n_H (except to the extent that it depletes deuterium). This is because one is directly measuring the amount of RF power needed to sustain the tail against the collisional drag forces. Only data points above 6 keV were used in the fit, in the hope that the

Figure 3.8: Deuterium spectrum during ICRF with 14% hydrogen. Solid lines are bounce averaged quasilinear simulations for different assumed hydrogen concentrations. The peaked $E_+(r)$ of Fig. 3.15 was assumed. Dashed line is isotropic model fit.

Figure 3.9: Hydrogen spectrum during ICRF with 14% hydrogen. Solid lines are bounce averaged quasilinear simulations for different assumed hydrogen concentrations. The peaked $E_+(r)$ of Fig. 3.15 was assumed.

neutral spectra above 6 keV is representative of the distribution function in the central hot region of the plasma. This fitting procedure typically took 30 seconds of CPU time on a DECsystem-10 computer. The computing time could probably be reduced by using a gradient search algorithm. The fit could also be speeded up if an analytic solution to Eq. (2.55) could be found. The parameters which best fit the data in Fig. 3.8 are $\log f(0) = 40.3$, $T_i = 1.67$ keV, and $RFK = 10.7$ keV, which corresponds to a central value for the deuterium second harmonic heating power density of $P_D = 29 mW/cm^3$. The full bounce averaged quasilinear simulation for 14% hydrogen gave a central $P_D = 34 mW/cm^3$, in very good agreement with the isotropic model. The isotropic model fit and the full bounce averaged quasilinear simulation for 14% hydrogen are virtually indistinguishable at higher energies. At lower energies the isotropic model begins to diverge from the experimental data and from the bounce averaged simulation because it does not include the contribution to the sightline integrated signal from the lower T_i and higher n_0 regions of the plasma near the edge.

Figure 3.9 compares the measured hydrogen spectra with the predictions of bounce averaged quasilinear theory. The parameters assumed for the hydrogen simulations in Fig. 3.9 are the same as assumed for the deuterium simulations of Fig. 3.8. For a fixed shape of the neutral density profile, the theoretical curves are linearly proportional to the magnitude of the neutral density. The neutral density not only has strong radial variation but also toroidal and poloidal asymmetries. Because the magnitude of the neutral density in front of the charge exchange analyzer is not known, one is usually free to adjust the theoretical curves up or down to obtain the best fit to the data, thus determining the magnitude of the neutral density. In this case, the magnitude of the neutral density has already been determined by adjusting the deuterium theoretical curves to match the deuterium spectra at low energies, so the amplitude of the hydrogen theoretical curves is now fixed. As in the case of the deuterium spectrum, the best fit is for a hydrogen concentration of $\sim$14%. The fit is not as good as the deuterium fit, but is probably within the error bars of the data. Errors in the neutral profile predicted by FRANTIC, or those due to the small banana width approximations made in the simulation may also be playing a role.

Note that the main effect of varying the hydrogen concentration is to shift the hydrogen simulations uniformly up and down without changing

their slope significantly. This is because the tail is so energetic that unconfined orbit losses have become significant and the simple scaling one might expect from $T_{\rm Stix} = \frac{P_{\rm H}}{3n_{\rm H}}\tau_s$ is no longer relevant. As discussed in Sec. 2.5.1, the shape of the distribution function becomes independent of RF power in the large orbit loss limit. The maximum energy hydrogen that could be confined in PLT for these experiments was $\sim$ 710 keV. As the hydrogen concentration drops from 24% to 7.6%, $T_{\rm Stix}$ rises from $\sim$120 to 360 keV and $T_{\rm Stix}/W_{\rm loss}$ climbs from $\sim$ .17 to .51. According to the model in Fig. 2.6, the fraction of power lost on unconfined orbits will rise from 10% to 60%, in good agreement with the full bounce averaged simulations which indicate that this fraction should rise from 28% to 72%. Because such a large fraction of the RF power is lost on unconfined orbits, the bounce averaged simulations shown in Fig. 3.9 indicate that the 'temperature' of the hydrogen (defined as 2/3 of the average energy of the hydrogen) only rises from 65 to 93 keV as the hydrogen concentration drops from 24% to 7.6%.

3.3.3 2.8% hydrogen concentration

We have been discussing why, in this parameter regime, the slope of the charge exchange spectra is not expected to change much as the hydrogen concentration is varied. Experimental evidence of this is provided in Fig. 3.10, which shows the hydrogen spectra for concentrations of 14%, 7.2%, and 3.5%. The tail temperature does increase somewhat as the concentration is dropped from 14% to 7.2%, but there is no change as the concentration is dropped further to 3.5%. One must be wary of comparing the absolute magnitudes of the raw data from different shots because the neutral density may vary from shot to shot. It is more reliable to compare the deuterium and hydrogen spectra from the same shot, which is done in Fig. 3.11 and 12 for an even lower concentration case. The concentration measured by MACE soon after the RF was turned off was 2.8%. But the hydrogen spectra during the RF has dropped well below the simulation which assumed 1% hydrogen. Note that the high energy ($\sim$ 50 keV) hydrogen flux in Fig. 3.12 has dropped more than the low energy ($\sim$ 10 kev) hydrogen flux when compared with Figs. 3.9 and 3.10. This may be an indication that the central hydrogen density has dropped more than the edge hydrogen density, so that the hydrogen density profile no longer has the same shape as the electron density profile. It may

Figure 3.10: Hydrogen spectra for 14%, 7.2%, and 3.5% hydrogen concentrations.

Figure 3.11: Deuterium spectrum during ICRF with 2.8% hydrogen. Solid lines are bounce averaged quasilinear simulations for different assumed hydrogen concentrations. The peaked $E_+(r)$ of Fig. 3.15 was assumed. Dashed line is isotropic model fit.

Figure 3.12: Hydrogen spectrum during ICRF with 2.8% hydrogen. Solid lines are bounce averaged quasilinear simulations for different assumed hydrogen concentrations. The peaked $E_+(r)$ of Fig. 3.15 was assumed.

be that unconfined orbit losses are so large that whatever particle sources exist in the plasma are unable to maintain the central hydrogen density at 2.8%, so that it has actually dropped to less than 1%. In this extremely low concentration limit, there are not enough counts to see any signal above 75 keV.

The deuterium spectra of Fig. 3.11 seem to tell a different story. Lowering the hydrogen concentration from Fig. 3.8–3.11 did cause the deuterium tail to rise, but not as much as the theory would predict. The deuterium tail seems more consistent with an 8% hydrogen concentration (equivalent to a central $P_D = 92\text{mW/cm}^3$) than with the post RF measurement of $\sim 3\%$ (or $P_D = 330\text{mW/cm}^3$) or the even lower estimate of 1% (or $P_D = 1300\text{mW/cm}^3$) from the hydrogen tail during RF.

Figures 3.13 and 3.14 show that it is possible to fit the hydrogen and deuterium spectra of Figs. 3.10 and 3.11 simultaneously if one assumes the RF $E_+(r)$ profile is hollow. The choice of the shape of a hollow $E_+(r)$ profile is somewhat arbitrary. The one used here is the same as the one which will be used in Sec. 4.1 to match the spectra measured at different viewing angles. This hollow E_+ profile is shown in Fig. 3.15 for comparison with the moderately peaked E_+ profile used before. The RF power profile to the hydrogen in the peaked and hollow cases are shown in Fig. 3.16, while the RF power profile to the deuterium is shown in Fig. 3.17. Different hydrogen concentrations are used for the two profile shapes (8% for the peaked profile, .9% for the hollow profile) in order to get similiar maximum values for P_D. Note that they predict similiar deuterium spectra as well. The isotropic model fit for this discharge gave $P_D = 83mW/cm^3$, showing that the isotropic model provides a good measure of the peak deuterium power, even if the radial variation of P_D is unknown.

3.3.4 33% hydrogen concentration

We finally turn our attention to the high concentration case in Figs. 3.18 and 3.19. Although the deuterium tail did fall as the concentration was raised from 14% in Fig. 3.8 to 33% in Fig. 3.18, it did not fall quite as much as expected. The central deuterium power is 7 mW/cm^3 for the 33% theoretical curve and 18 mW/cm^3 for the 16% curve. The truth lies between these two curves, and the isotropic fit gives a central P_D of 13 mW/cm^3. Both

Figure 3.13: Same deuterium spectrum as Fig. 3.11, compared with a bounce averaged quasilinear simulation assuming 0.9% hydrogen and the hollow $E_+(r)$ of Fig. 3.15.

Figure 3.14: Same hydrogen spectrum as Fig. 3.12, compared with a bounce averaged quasilinear simulation assuming 0.9% hydrogen and the hollow $E_+(r)$ of Fig. 3.15.

Figure 3.15: Comparison of the moderately peaked and hollow $E_+(r)$ profiles used in the bounce averaged quasilinear simulations.

Figure 3.16: Hydrogen RF power profiles corresponding to the peaked and hollow $E_+(r)$ profiles of Fig. 3.15. 8% hydrogen was assumed for the peaked $E_+(r)$ power profile, and 0.9% hydrogen was assumed for the hollow $E_+(r)$ power profile.

Figure 3.17: Deuterium RF power profiles corresponding to the peaked and hollow $E_+(r)$ profiles of Fig. 3.15. 8% hydrogen was assumed for the peaked $E_+(r)$ power profile, and 0.9% hydrogen was assumed for the hollow $E_+(r)$ power profile.

Figure 3.18: Deuterium spectrum during ICRF with 33% hydrogen. Solid lines are bounce averaged quasilinear simulations for different assumed hydrogen concentrations. The peaked $E_+(r)$ of Fig. 3.15 was assumed. Dashed line is isotropic model fit.

Figure 3.19: Hydrogen spectrum during ICRF with 33% hydrogen. Solid lines are bounce averaged quasilinear simulations with different assumptions: Peaked $E_+(r)$ and 33% hydrogen, peaked $E_+(r)$ and 16% hydrogen, hollow $E_+(r)$ and 33% hydrogen, and, finally, 33% and peaked shape for $E_+(r)$ but reduced in magnitude to give a total hydrogen power of 190 kW—down a factor of 5 from the 950 kW used for the other simulations.

theoretical curves in Fig. 3.18 assumed a moderately peaked E_+ profile. The equivalent curves for the hydrogen spectra are shown in Fig. 3.19, along with additional curves which assume a 33% concentration but have the central power density reduced by either using the hollow E_+ profile or by dropping the total RF power by a factor of 5. At such high concentrations it is unlikely that unconfined orbit losses are significant. The large drop in the hydrogen tail between Figs. 3.9 and 3.19 is probably evidence for a decline in the RF power absorbed by the hydrogen, with most of the RF power being diverted to some other channel such as mode conversion.

This section provided a very detailed analysis of the spectra from 3 discharges. Despite all of the questions that are raised by this detailed analysis, three conclusions can be drawn. First, at moderate hydrogen concentrations ($\sim$ 15%) the deuterium and hydrogen spectra are in good agreement with the predictions of the bounce averaged quasilinear code ignoring fast ion transport and using a peaked power profile. Second, the observed scaling of the deuterium tail with hydrogen concentration is not as strong as expected. There is more deuterium tail than expected when the concentration was raised to 33%, and less than expected when the concentration was lowered to 2.8%. Third, the central deuterium power density obtained by fitting the spectra with the isotropic model is very similiar to that obtained by fitting the spectra with the full bounce averaged quasilinear code. We will make use of this last conclusion to study the first two conclusions in more detail in the next section.

3.4 Scaling of the Deuterium Power With Concentration

The isotropic model provides a useful tool for quickly measuring the central deuterium power density from the deuterium spectra. Figure 3.20 shows this measured P_D for 45 discharges which have hydrogen concentrations ranging from 1.6% to 33% plotted versus $\langle P_{\rm D}\rangle_{15}$, the theoretically expected deuterium power density averaged over a 15 cm minor radius. More specifically, the quantity identified as the theoretical $\langle P_{\rm D}\rangle_{15}$ is defined as:

$$\langle P_{\rm D}\rangle_{15} = .64 \times .85 \times \frac{P_{Total}}{2\pi R \pi 15^2}\left(\frac{.5k_\perp^2\rho_D^2 n_D/n_H}{1+.5k_\perp^2\rho_D^2 n_D/n_H}\right) \qquad (3.1)$$

Figure 3.20: Measured versus theoretical deuterium RF power density, as the hydrogen concentration is varied from 1.6% to 33%. In this and later plots, the deuterium RF power density is measured by fitting the deuterium tail with the isotropic second harmonic model, Eq. (2.55), while quantity called "theoretical $\langle P_{\mathrm{D}} \rangle$" is given by Eq. (3.1) and is an average over a 15 cm minor radius.

Figure 3.21: Ratio of the measured to theoretical deuterium RF power density versus hydrogen concentration, for the same discharges as Fig. 3.20.

Figure 3.22: Ratio of measured to theoretical deuterium RF power density versus T_{Stix} for the hydrogen.

The factor in parenthesis is $P_D/(P_H + P_D)$ from Eq. (2.49). The cold plasma dispersion relation is used to find $k_\perp$, and central plasma parameters are used to calculate v_A and ρ_D. P_{Total} is the total RF power into the vacuum vessel, with an estimated coupling efficiency of .85. A rough estimate of the average RF power absorbed inside the 15 cm minor radius is obtained by integrating the expression for $P_{Stix} \propto n_H|E_+|^2/r$ from Eq. (2.6) over a 15 cm minor radius volume. Assuming that $n_H(r)$ and $|E_+(r)|^2$ are proportional to $(1 - r^2/a^2)$ gives 64% of the RF power absorbed inside of $r/a = 15/40$. Note that averaging P_{Stix} over such a large volume is an underestimate of the peak deuterium power density that the wave physics predicts. According to Eq. (2.41), the Doppler broadening of the resonance layer determines the maximum power density. For 2 keV deuterium, the Doppler width is only $\sim$ 3 cm and gives a power density which is equivalent to averaging P_{Stix} over a $\sim$ 1 cm minor radius volume. This would result in a power density 15 times bigger than the one used here. Averaging $P_{\rm Stix}$ over a 15 cm minor radius allows for a true RF power density which may be much higher on axis but assumes that some rapid transport process mixes the fast ions over a 15 cm radius. Although the choice of the mixing radius r_{mix} is somewhat arbitrary, it only leads to a multiplicative constant in the expression for $\langle P_{\rm D}\rangle \sim 1/r_{mix}$ and does not effect the scaling of $\langle P_{\rm D}\rangle$ with hydrogen concentration.

3.4.1 Interpretations of the nonlinear scaling

The data in Fig. 3.20 supports the conclusion of the previous section that the deuterium power density does not increase as fast as expected as the hydrogen concentration is dropped. The same data in Fig. 3.20 is plotted in Figs. 3.21 and 3.22 to show explicitly how the ratio of the measured to expected deuterium power density scales with hydrogen concentration and with T_{Stix}. A complete list of plasma parameters for the discharges shown in Figs. 3.20–22 is given in the appendix. P_D should be proportional to n_D/n_H, but is instead found to only scale as $(n_D/n_H)^{.59\pm.04}$. This is subject to a number of possible interpretations. It may be that that P_D scales as expected with hydrogen concentration, but that fast ion transport is increasing as the concentration is lowered. At the highest concentrations, this transport only causes mixing over a 3 cm minor radius, but for some of the lowest concentration discharges the mixing radius is 18 cm. A number of possible

fast ion transport mechanisms are discussed in Sec. 2.5. In the center of the plasma, the dominant transport mechanism is sawteeth. The edge q values of the discharges analyzed in this section varied by less than 10%. The inversion radius as measured by the ECE system was $\sim 14cm$ for the low concentration discharges. At the highest concentrations, the inversion radius may have moved inward a small amount, but the shift is less than the $\sim 5cm$ spacing of the ECE channels. Although ion sawteeth transport would thus be consistent with the broadest power profiles observed, it can not explain the variation of the power profile with concentration.

If sawteeth transport of fast ions is always occuring, then the question is no longer why the apparent power profile is so broad at the lowest concentrations, but why the deuterium power density is larger than expected at the highest concentrations. It has been suggested[53] than in the case of pure hydrogen second harmonic heating, $k_\perp$ is increased over its cold plasma value by Bernstein wave effects. Calculations[122,108] for deuterium second harmonic/hydrogen fundamental heating show no such increase in second harmonic damping.

Sawteeth transport is a time-dependent process, while we are treating it as equivalent to a time-independent broadening of the RF power profile. The sawtooth period for these discharges is typically ~ 5 msec, while the charge exchange signal was averaged over 50 msec to obtain an adequate count rate in the tail. If the time scale for the RF to produce a small tail is short compared to the sawtooth period, while the time scale to extend the tail out to high energies is long compared to the sawtooth period, then the size of the tail may have a nonlinear relationship to the RF power density. Perhaps this explains why the deuterium tail is larger than expected at the highest hydrogen concentrations. Further investigation of this possibility is reserved for future work. A time dependent sawtooth model needs to be implemented in the quasilinear simulations. And higher time resolution charge exchange measurements need to be made.

Another interpretation would be that there is no significant transport of the tail ions, but that the RF power profile really is broadening (and the central power density dropping) as the hydrogen concentration falls. For example, if the hydrogen that is pumped out of the center of the plasma ends up in the edge of the plasma, then the damping rates will increase near the edge and less power will be available at the center of the plasma to drive

the deuterium tail. If the central ions become sufficiently collisionless, then the particle-wave interaction may become superadiabatic and no stochastic heating will occur. Particles which mirror in the resonant layer may interact so strongly with the wave as to change its polarization or cause reflection of the wave.

If the single pass absorption becomes sufficiently weak, a substantial fraction of the RF power may never be absorbed in the plasma at all. Measurements of the Q of the PLT vacuum vessel[53] indicate that the indicate that the "reflection" losses can be as large as a few percent per bounce. These losses may be caused by resistive damping in the walls, absorption of the wave by unused antennas, or escape of the wave through vacuum vessel ports. But even for the lowest hydrogen concentrations of 1.6%, the estimated[122] single pass absorption is $\sim 20\%$, and reflection losses should not be severe. At the lowest hydrogen concentrations where E_+ is largest, there may be some nonlinear process at the edge of the plasma which is diverting RF power from the plasma center.[123]

All 45 discharges plotted in Figs. 3.20–3.22 had sawteeth. There were 2 sawtoothless dicharges which are not shown. The values of P_D measured for these two discharges were only slightly above (by 0–30%) comparable sawtooth discharges. This may seem to be a refutation of our claim that sawteeth are the dominant transport process. However, tokamak plasma usually fail to sawteeth only because something else is wrong with the plasma. Frequently there are m=2 oscillations present if sawteeth are absent. These two particular discharges had much colder electron temperatures (~ 950 eV) than the usual sawtooth discharges (1200–1500 eV).

Most of the discharges in Fig. 3.20–3.22 had fairly low line averaged electron densities ($< 1.3 \times 10^{13}$). The few discharges at higher densities ($1.5 - -2.2 \times 10^{13}$), marked with an 'X', tended to outperform the lower density discharges with equivalent theoretical $\langle P_{\rm D} \rangle_{15}$. This may be because of improved focussing of the wave at higher densities. Although the most dramatic tails observed with the charge exchange analyzer (i.e., the deuterium spectra that extended to the highest energies and had the largest values of RFK) were at fairly low densities, where the collisionality was low and it was easy to make a non-Maxwellian tail, the highest values of P_D were actually obtained at the highest densities where $k_\perp$ was largest.

In the low density ($< 1.3 \times 10^{13}$), low concentration ($< 15\%$), high power

regime where most of the data was taken, the electron density was very difficult to control. The line averaged density would frequently show large ($\sim 10--20\%$) erratic fluctuations and would sometimes begin to rise extremely rapidly towards the end of the RF power flattop. The visible bremsstrahlung Z_{eff} would drop about as fast as the density rose. The biggest source of scatter in the data of Figs. 3.20–3.22 may be that many of the discharges were not really in steady state at the time of the measurement.

3.5 Scaling of the Measured Deuterium Power With Total Power

The deuterium tail should increase not only as the hydrogen concentration is lowered but also as the total RF power is increased. Figures 3.23–3.34 show the results of three different power scans. A complete list of plasma parameters for these discharges is given in the appendix. The first power scan, in Figs. 3.23–3.27, shows that the increase in the deuterium tail is in good agreement with the expected scaling. The measured deuterium and hydrogen spectra for a range of power levels are shown in Figs. 3.23 and 3.24. Also shown are theoretical curves from the bounce averaged quasilinear code for the lowest and highest power levels. A hydrogen concentration of $n_H/n_e = 1\%$ and a hollow $E_+(r)$ similiar to Fig. 3.15 were used to provide a good fit to the lowest power data. The same hydrogen concentration and E_+ profile shape were then used to predict the high power data, with only measured parameters (total RF power, T_e, n_e, T_i, Z_{eff}) changed in the simulation. The observed increase in the deuterium tail is in very good agreement with the bounce averaged quasilinear code. The hydrogen tail is already so energetic that the simulation predicts that the hydrogen charge exchange flux should drop as the power is increased because of the simultaneous rise in n_e and fall in n_0. The observed hydrogen tail actually drops slightly faster than the code predicts. As in the previous section, we fit the deuterium spectra with the isotropic model to obtain a direct measure of P_D. Figs. 3.25–3.27 show that this measured P_D increases linearly with total RF power for this particular set of discharges.

A second power scan (performed a week after the first power scan and a few days before the data of Sec. 3.4) is shown in Figs. 3.28–3.31. P_D is

Figure 3.23: Deuterium spectra for a range of ICRF power levels. Solid curves are bounce averaged quasilinear simulations of the lowest and highest power data.

Figure 3.24: Hydrogen spectra for a range of ICRF power levels. Solid curves are bounce averaged quasilinear simulations of lowest and highest power data.

Figure 3.25: Measured versus theoretical deuterium RF power density, as total RF power is varied from 680 kW to 1150 kW.

Figure 3.26: Ratio of measured to theoretical deuterium RF power density versus total RF power, for the same discharges as Fig. 3.25.

Figure 3.27: Ratio of measured to theoretical deuterium RF power density versus T_{Stix} for the hydrogen, for the same discharges as Fig. 3.25.

Figure 3.28: Measured versus theoretical deuterium RF power density, as total RF power is varied from 120 to 1250 kW.

Figure 3.29: Ratio of measured to theoretical deuterium RF power density versus total RF power, for the same discharges as Fig. 3.28.

Figure 3.30: Ratio of measured to theoretical deuterium RF power density versus T_{Stix} for the hydrogen, for the same discharges as Fig. 3.28.

Figure 3.31: Line averaged electron density as a function of total RF power, for the same discharges as Fig. 3.28.

observed to increase linearly with $P_{\rm total}$ from 500 to 1300 kW, but the scaling is nonlinear below 500 kW. The low power discharges marked with closed circles had deuterium gas puffing during the RF which maintained $\bar{n}_e$ at $1.3\times10^{13}{\rm cm}^{-3}$, while the other low power discharges had $\bar{n}_e \sim 1.0\times10^{13}{\rm cm}^{-3}$ during the RF. This additional deuterium gas puffing may have reduced the source of hydrogen from the walls, or it may be some other effect associated with higher density which caused P_D to be higher for these two discharges.

A third power scan (performed six months later) shows a strong nonlinearity up to the highest power levels. The data in Figs. 3.32–3.34 were taken at three power levels: 660 kW, 1210 kW, and 2300 kW. A least squares fit the the data yields $P_D \propto \langle P_{\rm D}\rangle_{15}^{.38\pm.04}$, an even weaker scaling than found from the hydrogen concentration scan. The discharges in Figs. 3.32–3.34 tended to be at higher $\bar{n}_e$ and lower Z_{eff} than those of Sec. 3.4. No erratic oscillations in $\bar{n}_e$ were observed, although the density did begin to climb rapidly at the highest power level. This data also presents a test of the isotropic model. At each power level, measurements were made at three different viewing angles (Fig. 1.3). The variation of T_i, RFK, and P_D measured at the three viewing angles is less than 30%, providing further evidence for the utility of the isotropic model. Sec. 4.2 will compare the power scan data of Figs. 3.32–3.34 with the bounce averaged quasilinear code.

No hydrogen was puffed into the tokamak during any of the three

Figure 3.32: Measured versus theoretical deuterium RF power density, as total RF power is varied from 660 kW to 2300 kW.

Figure 3.33: Ratio of measured to theoretical deuterium RF power density versus total RF power, for the same discharges as Fig. 3.32.

Figure 3.34: Ratio of measured to theoretical deuterium RF power density versus T_{Stix} for the hydrogen, for the same discharges as Fig. 3.32.

power scans discussed in this section. The residual hydrogen concentration $n_H/(n_D + n_H)$ used to calculate the theoretical $\langle P_{\rm D}\rangle$ was assumed to be a constant 3% for all of the discharges, although no accurate measurement of the concentration was attempted for these discharges. The hydrogenic concentration $(n_H + n_D)/n_e$ was assumed to be a constant 80% for the first two power scans, while the variation of Z_{eff} was included in estimating the deuterium depletion for the third power scan. Assuming that $Z_i = 8$ was the main impurity, the hydrogenic fraction only varied from .79 to .93 for this last power scan.

Science is not a democracy. Although the first two power scans showed that P_D increases linearly above 500 kW of total RF power, the third power scan showed that this is not always true. Part of the complication may be that the density varies with RF power, even though we try to control for the density variation through the definition of $\langle P_{\rm D}\rangle$. Fig. 3.31 shows the natural variation (except for two of the low power discharges where the density feedback system was enabled) of $\bar{n}_e$ with total power for the second power scan. The density did not vary as much in the third power scan. At the highest power it was $2.2 \times 10^{13}{\rm cm}^{-3}$, dropping to $1.7 \times 10^{13}{\rm cm}^{-3}$ at medium power, and staying there via deuterium gas puffing at the lowest power. Of course, T_i and T_e vary with RF power as well, and separating out the importance of various parameters is inherently difficult. Future experiments should try to resolve the present ambiguity by performing independent scans of power and density. Simultaneous measurements of the hydrogen concentration would also be helpful.

3.6 Conclusions

We have presented clear evidence of direct second harmonic heating. A deuterium tail has been observed out to energies in excess of 50 keV (equivalent to $25 \times T_i$ or $2 \times W_c$). The slope of the distribution function at the highest energies was 1/(16 keV), well above the bulk ion temperature of 2.1 keV. The highest measured value of P_D for which hydrogen concentration measurements exist was 220 mW/cm^3 (or $\sim 22\%$ of the total central RF power) and is consistent with theoretical expectations (Eq. (3.1)) assuming a 15 cm mixing radius. The highest P_D achieved in PLT to date is 300 mW/cm^3 at $\bar{\rm n}_{\rm e} = 2.9 \times 10^{13}/{\rm cm}^3$. P_D does not scale linearly with hydrogen concentration.

The higher-than-expected values of P_D found at higher concentrations may be evidence of enhanced deuterium damping or reduced fast ion transport.

The first experimental observation of direct second harmonic deuterium heating was made on the Japanese JFT-2 tokamak.[48,49] They observed a deuterium tail up to 6 keV (equivalent to $10 \times T_i$ or $.75 \times W_C$). Ion-ion collisions are very rapid at these lower energies and keep the distribution function nearly Maxwellian. Because of this, the slope of the distribution function at the highest energies they observed was only 1/(1.2 keV), only 2 times the bulk ion temperature of 550 eV. Based on charge exchange and power balance analysis, they claimed that the deuterium absorption was stronger than predicted by theory. For low hydrogen concentrations, $n_H/n_D \sim$ 2–4%, they obtained $P_D/P_H \sim 3$, while theory would predict $P_D/P_H = \beta_D n_D/n_H \sim .7$. At higher concentrations, $n_H/n_D = 10\%$, they observed $P_D/P_H = 1$ while theory would predict .078. Although the absolute values of P_D/P_H from their analysis is much higher than ours, note that they observe that $(P_D/P_H)_{Exp}/(P_D/P_H)_{Theory}$ gets larger at higher concentrations, which is similiar to our results. Perhaps the differences between our results and theirs is that JFT-2 uses a high field side antenna, while the PLT experiments were done with a low field side antenna.

There are a number of ways in which future experiments could expand upon these results. Perhaps the most obvious is to improve the measurement of the hydrogen concentration. In addition to the charge exchange method described here, spectroscopic H_α/D_α and after-the-shot residual gas analysis measurements would provide corroborating evidence for the true concentration. The time variation of the hydrogen concentration during the discharge could be reduced by always fuelling the plasma with a premixed supply of hydrogen and deuterium gas, rather than initiating the discharge in deuterium and puffing in a small amount of hydrogen later. Operation at a fixed concentration for tens of discharges in a row may be necessary to allow the hydrogen concentration to come to equilibrum with the walls. Because unconfined orbit losses were probably playing a large role in these experiments, it would be useful to repeat them at lower power densities, higher electron densities, and higher plasma currents to reduce these losses. Systematic documentation of the enhancement of the neutron production rate over that expected from spectroscopic measurements of T_i is another interesting area for future research. Our present experiments have concentrated on the low density

regime where the non-Maxwellian deuterium tail was largest. To reach the highest values of P_D, future experiments should try to maximize the density, perhaps with the aid of a pellet injector. The new technique of fitting charge exchange spectra with an isotropic model can probably be fruitful in the the study of pure second harmonic hydrogen heating where there is no complication of a competing fundamental heating mechanism which produces an extremely energetic tail. Finally, on the theoretical front, there is a clear need to incorporate radial transport into models of the RF-produced tail.

Chapter 4

Pitch Angle Dependence of RF Generated Fast Ions.

While the last chapter focussed only on perpendicular charge exchange measurements, this chapter will explore the angle dependence of the charge exchange spectra. The introductory chapter showed that the resonance localization model provides a good, qualitative explanation of the dramatic viewing angle dependence of the charge exchange spectra, and argued that this may provide radially local information. This chapter will make detailed comparisons of the measured spectra with predictions of the bounce-averaged quasilinear code, and will investigate the uncertainties in uncovering radial information.

In Sec. 4.1, we compare the measured spectra from a single high power case with the predictions of the code based on a variety of assumptions. In Sec. 4.2, we compare the predictions of the code with spectra taken at RF power levels ranging from 200 kW to 2300 kW. There is good qualitative agreement between the predicted and measured spectra, although the best fits seem to require an RF power profile which is much broader than theoretically expected. This is consistent with the more quantitative results of the last chapter, and is subject to the same interpretation: the true RF power profile is probably quite peaked, but radial transport (which is ignored by the code) leads to a significantly broader distribution of fast ions. It should be emphasized that whenever we speak in this chapter, as in the last, of the RF power profile being broader than expected, we mean only that it *appears*

to be broad, perhaps because of radial fast ion transport.

In Sec. 4.3, we present the first charge exchange measurements of ^{3}He minority heating. These measurements, made possible by operating with a ^{4}He majority plasma, have similiar features to the hydrogen spectra, but do not require as broad an RF power profile to obtain a good fit.

4.1 Angle Dependence at High RF Power

In this section we will compare the deuterium and hydrogen spectra shown in Figs. 4.1 and 4.2 with bounce-averaged quasilinear code predictions for a variety of assumptions. This hydrogen data is the same as discussed in the introduction (Fig. 1.2) and the viewing angles are shown in Fig. 1.3. The parallel deuterium spectra were not measured in this experiment. The theoretical deuterium spectra predicted for a variety of assumed values of the hydrogen concentration are shown in Fig. 4.1. The corresponding predictions of the hydrogen spectra are shown in Figs. 4.2–4. The main plasma parameters for these spectra are shown in Fig. 4.5–6. A modestly peaked $E_+(r)$ profile proportional to $(1 - r^2/a^2)^{1/2}$ was assumed for the simulations in Figs. 4.1–4. The total RF power to the deuterium and hydrogen, after 15% coupling losses, was 1310 kW. Other parameters used for these simulations include: $Z_{eff} = 3.0$, main plasma impurity $Z_I = 8$, minor radius a = 41 cm, major radius $R_{maj} = 132$ cm, magnetic field $B_{tor} = 20.61$ kG, resonance layer $R_{res} = 138.7$ cm, plasma current $I_p = 450$ kA, loop voltage $V_\ell = 1.0$ V, and $k_\| = \pm.07\text{cm}^{-1}$. The neutral density profile was calculated using an edge neutral temperature of .25 eV, and the toroidally local magnitude of the neutral density was adjusted, as in Chapter 3, to obtain the best fit to the lower energy deuterium spectra.

Accurate measurements of the true hydrogen concentration were not made for the discharges shown in Figs. 4.1 and 4.2, although the rough measurements indicate that it was less than or of the order of 2–3% during the ohmic phase of the plasma. For the assumed peaked E_+ profile, the observed deuterium tail is best fit by a higher concentration of $\sim 6\%$. On the other hand, the best fit to the hydrogen spectra requires a lower concentration of $\sim 0.7\%$. These observations are similiar to those made in the previous chapter for low hydrogen concentrations. Note that the simulated hydrogen spectra do not change their shape much as the assumed concentration is

Figure 4.1: Deuterium energy spectra measured by the charge exchange neutral analyzer at two different viewing angles (shown in Fig. 1.3). Also shown are bounce averaged quasilinear simulations for different assumed hydrogen concentrations and a moderately peaked $E_+(r)$ profile (Fig. 4.6).

Figure 4.2: Hydrogen energy spectra measured by the charge exchange neutral analyzer at three different view angles (shown in Fig. 1.3). Also shown is a bounce averaged quasilinear simulation assuming 6% hydrogen and a moderately peaked $E_+(r)$ (Fig. 4.6).

Figure 4.3: Same as Fig. 4.2, but assuming 1.5% hydrogen for the simulation.

Figure 4.4: Same as Fig. 4.2, but assuming 0.7% hydrogen for the simulation.

Figure 4.5: Main plasma parameters for the data of Figs. 4.1–4.4.

Figure 4.6: $E_+(r)$ and $k_\perp(r)$ profiles assumed for the simulations of Figs. 4.1 to 4.4. The plots of $P_{\rm RF}(r)$ and $T_{\rm tail}(r)$ in this figure, and in Figs. 5.7, 5.9, 4.20, and 4.25 are based upon Stix's simple formulas and are not the actual power or temperature profiles calculated by the program.

dropped from 1.5% to 0.7%, although the relative amplitudes of the spectra at different viewing angles do change. This is consistent with the result of Sec. 2.5.1 that the shape of $f(W)$ becomes independent of RF power when orbit losses are dominant.

The bounce-averaged quasilinear code includes the effects of pitch angle scattering and can be used to determine when the assumption of resonance localization is valid. According to Eq. (1.2), which assumes perfect resonance localization, the perpendicular sightline ($R_{tan} = 13$ cm) should see energetic neutrals which originate near $R = 140$ cm. The code indicates that only above 50 keV is this a valid approximation, and unfortunately there is not much perpendicular signal above 50 keV. Below 50 keV, pitch angle scattering is rapid enough, and there are enough fast ions at larger minor radii, so that over 50% of the flux comes from outside $R = 148$ cm. The peak angle ($R_{tan} = 70$ cm) should view ions which charge exchange near $R = 168$ cm (~ 5 cm from the edge), according to Eq. (1.2). The assumption of resonance localization is valid down to much lower energies at this angle, because pitch angle scattering is less and the neutral density is higher towards the edge.

One of the interesting features of the data is the "negative temperature", or $\partial f/\partial W > 0$, observed between 30 keV and 100 keV at the peak angle. As discussed in the introduction, one possible explanation is that there is much more RF power at point A in Fig. 1.10 than at point B. Indeed, keeping all other parameters the same as for the Fig. 4.4 simulations, but reducing E_+ significantly at $r/a = .75$ (Fig. 4.7), we are able to partially reproduce this negative temperature (Fig. 4.8). The good fit in Fig. 4.8 does not rule out the possibility of alternative explanations, such as the "edge drag" mechanism discussed in Secs. 1.2.2 and 2.5.2. The energetic hydrogen signal at other viewing angles is unaffected, as is the predicted deuterium spectra (not shown here) which still disagrees (assuming 0.7% hydrogen) with the observed spectra. It is necessary to keep the same E_+ for $r/a > .75$ in order to produce enough charge exchange flux below 30 keV. The $E_+(r)$ profile of Fig. 4.7 may be actually be close to reality, if there is a gap inbetween the large evanescent fields at the edge and propagating waves which are focussed towards the center.

To fit the hydrogen and deuterium spectra simultaneously, it is necessary to reduce the central deuterium power density by making $E_+(r)$ hollow (fig. 4.9). Except for raising the hydrogen concentration slightly to 0.9%, all other

Figure 4.7: $E_+(r)$ and $k_\perp(r)$ profiles assumed for the simulation of Fig. 4.8. Note the drop in E_+ at $r/a = .75$.

Figure 4.8: Same data as Fig. 4.2, but using the $E_+(r)$ profile of Fig. 4.7 and assuming 0.7% hydrogen.

parameters are kept fixed. Unlike the simulations of Fig. 4.1, the hollow E_+ deuterium simulation of Fig. 4.10 produces more deuterium tail at the peak angle than at the perpendicular angle, in agreement with the measurements. The hollow E_+ hydrogen simulation of Fig. 4.11 is not much different than the previous one. In both cases, good agreement is found with the perpendicular and peak measurements, but not with the parallel measurements.

In an earlier analysis[10] of the hydrogen spectra of Fig. 4.11, I concluded that the RF power profile is broader than theoretically expected. While I reach the same conclusion in this thesis on the basis of the deuterium and hydrogen spectra in Figs. 4.10 and 4.15, it now appears difficult to draw any unique conclusions based only on the hydrogen spectra of Fig. 4.11. Part of the problem is due to uncertainty about the hydrogen concentration. If the concentration really is as high as 3%, then a peaked power profile cannot produce enough energetic ions at point A in Fig. 1.10 to fit the data, as concluded in my earlier analysis. If the concentration is lower, then less power at point A is needed because $T_{\rm Stix} \propto P/n_H$. But the central problem is that the tail is so energetic in this case that unconfined orbit losses make the shape of f independent of T_{Stix}. In the next section, we will analyze a low power case where orbit losses are small and the hydrogen spectra alone provide some indication of a broad RF power profile. (The earlier version of the code assumed $k_\| = 0$ and $E_- = 0$, but this difference does not have a major effect on our conclusions.)

In addition to the drop in E_+ at $r/a = .75$, another ingredient necessary to produce the "negative temperature" in the simulations is a very hollow neutral density profile to maximize the edge contribution to the peak angle signal. If a flatter neutral profile is used (by raising the edge neutral temperature significantly above the .25 eV assumed here), then it becomes more difficult to obtain good fits, although broadening the RF power profile further does introduce a slight improvement. Many different values of $E_+(r)$ and n_H/n_e, as well as other plasma parameters, have been tried in the computer simulations in an attempt to improve the fit. Slightly better fits than shown here can be obtained, but no fit is substantially better.

The hydrogen "temperature" (defined as two thirds of the average energy) calculated by the code for the case of 0.7% hydrogen and $E_+ \propto (1-r^2/a^2)^{1/2}$ is shown in Fig. 4.12. The corresponding power balance in Fig. 4.13 indicates that over half of the hydrogen power is lost to unconfined orbits.

Figure 4.9: $E_+(r)$ and $k_\perp(r)$ profiles assumed for the simulations of Figs. 4.10 and 4.11.

Figure 4.10: Same data as Fig. 4.1, with a bounce averaged quasilinear simulation using the hollow $E_+(r)$ of Fig. 4.9 and assuming 0.9% hydrogen.

Figure 4.11: Same data as Fig. 4.2, with a bounce averaged quasilinear simulation using the hollow $E_+(r)$ of fig. 4.9 and assuming 0.9% hydrogen.

Figure 4.12: Initial and final hydrogen temperature calculated by the bounce averaged quasilinear program for the parameters of Figs. 4.5 and 4.6.

Fig. 4.13 shows that the steady state hydrogen power profile is significantly broader than the initial power profile, even though $E_+(r)$ is held fixed. As the hydrogen heats up, Doppler broadening reduces the peak power density (Eq. (2.41)), while the $E_- J_1$ term increases the damping elsewhere (Eq. (2.15)). For completeness, Figs. 4.14–15 show the deuterium "temperature" and the deuterium RF power profile for the same peaked E_+, 0.7% case, even though this is a poor fit to the deuterium spectra. The "temperature" and power profiles for the hollow E_+, 0.9% case which fits the data fairly well, is shown in Figs. 4.16–19.

4.2 Scaling with Total RF Power

The hydrogen and deuterium spectra at the perpendicular, peak, and parallel angles were measured for RF power levels of 660, 1210, and 2300 kW. The analysis of these deuterium spectra using the isotropic second harmonic model has already been discussed in Sec. 3.5 (Figs. 3.32–34). In this section, we provide comparisons of the hydrogen and deuterium spectra with the full bounce-averaged quasilinear code. A number of interesting discrepancies be-

Figure 4.13: Initial (0) and final (R) hydrogen RF power profile calculated by the bounce averaged quasilinear program for the parameters of Figs. 4.5 and 4.6. Also shown is the power flowing out of the hydrogen due to ion heating (i), electron heating (e), unconfined orbit losses (L), and charge exchange losses (C).

Figure 4.14: Initial and final deuterium temperature calculated by the bounce averaged quasilinear program for the parameters of Figs. 4.5 and 4.6.

Figure 4.15: Initial (0) and final (R) deuterium RF power profile calculated by the bounce averaged quasilinear program for the parameters of Figs. 4.5 and 4.6.

Figure 4.16: Initial and final hydrogen temperature calculated by the bounce averaged quasilinear program for the parameters of Figs. 4.5 and 4.9.

tween the data and the modelling will be pointed out. Hydrogen spectra were also measured on this same day at even lower RF power levels (205 and 320 kW) where unconfined orbit losses are negligible. This low power data provides further evidence that the RF power profile is broader than expected.

Rather than trying to obtain the best fit at each power level by tediously adjusting the E_+ profile and the hydrogen concentration, we adjusted these parameters only once, to provide a decent match to the medium power data, and then kept them fixed to simulate the high and low power data. Only measured quantities (total RF power, T_e, n_e, Z_{eff}, and T_i) were varied for the simulations of different power levels. A hollow E_+ profile (shown in Fig. 4.20 along with other plasma parameters used for the medium power simulation), somewhat different than the previous one (Fig. 4.9), was used for these simulations. The hydrogen concentration n_H/n_e was assumed to be 3%. The fit to the medium power spectra, and a comparison with the higher and lower power spectra, are shown in Figs. 4.21 and 4.22.

Many interesting features of this data will be discussed below, but let us begin with the observation that the deuterium tail does not increase with RF power as quickly as it should. This result is stated more quantitatively in Figs. 3.32–34, where it is shown that the measured deuterium power density

Figure 4.17: Initial (0) and final (R) hydrogen RF power profile calculated by the bounce averaged quasilinear program for the parameters of Figs. 4.5 and 4.9. Also shown is the power flowing out of the hydrogen due to ion heating (i), electron heating (e), unconfined orbit losses (L), and charge exchange losses (C).

Figure 4.18: Initial and final deuterium temperature calculated by the bounce averaged quasilinear program for the parameters of Figs. 4.5 and 4.9.

Figure 4.19: Initial (0) and final (R) deuterium RF power profile calculated by the bounce averaged quasilinear program for the parameters of Figs. 4.5 and 4.9.

Figure 4.20: Plasma and RF parameters for the medium power bounce averaged quasilinear simulations in Figs. 4.21 and 4.22.

Figure 4.21: Measured and calculated deuterium spectra at different viewing angles for total RF powers of 660 kW, 1210 kW, and 2300 kW. The hollow $E_+(r)$ of Fig. 4.20 was assumed.

Figure 4.22: Measured and calculated hydrogen spectra at different viewing angles for total RF powers of 660 kW, 1210 kW, and 2300 kW. The hollow $E_+(r)$ of Fig. 4.20 was assumed.

increases much slower than linearly with total RF power. One explanation is that the RF power profile is broadening as the total power increases. However, the results of the two other power scans in Sec. 3.5 (Figs. 3.25–27 and Figs. 3.28–30) show a scaling which is much closer to linear. The reason for the different behavior found here is unknown.

There are qualitative similiarities between the data and the simulations, but the agreement is not particularly good at any power level nor at any viewing angle. Many different E_+ profiles and hydrogen concentrations were tested to try to improve the fits, but no satisfactory fit to all of the details was ever achieved. A fundamental reason for this may be the lack of radial transport in our simulations. Modelling it as a simple broadening of the RF power profile may be inadequate. Even if this was an acceptable model, there is no reason that the effective E_+ needed to model transport for deuterium is the same as that needed for hydrogen.

Another reason for the poor agreement between theory and data is toroidal asymmetry in the neutral density profile. The simulations assume a toroidally symmetric neutral density and predict that the parallel deuterium flux should always be smaller than the flux at other angles. In reality, the opposite occurs. The parallel deuterium flux is larger even at low energies where f should be isotropic. (Note that the vertical distance between the parallel data and the perpendicular data is fairly large at low energies.) The theoretical curves already account for the longer path length through the plasma, and the large attenuation of signal, for the parallel sightlines. Even the parallel hydrogen flux is larger than the flux at other angles up to 40 keV, and it is larger than the predicted flux up to even higher energies. Merely increasing Z_{eff} in the simulations fails to improve the fits. This data seems to give clear evidence of important toroidal variations in the neutral density. The data of Sec. 4.1, however, is much more isotropic at low energies and is in better agreement with the simulations, indicating that the neutral density is more symmetric in that case (at least near the analyzer). In principle, the variation of the neutral density is irrelevant to our purposes. At sufficiently high energies, information about the radial RF power profile is contained in the slope of the charge exchange spectra, $\partial f/\partial W$, as a function of angle, and is independent of the relative amplitude of the charge exchange flux at different angles. Thus, one should be able to adjust the magnitudes of the simulated curves by different amounts at different angles to obtain the best

fit to the data.

As in Secs. 3.3 and 4.1, the ion temperature used here for the full bounce-averaged quasilinear simulations was found by first fitting the perpendicular deuterium spectra with the isotropic second harmonic model. This T_i was slightly lower than spectroscopic measurements (as usual with passive charge exchange methods), but more importantly, it produced a final bounce-averaged simulation which fit the data well. This same procedure did not work as well, however, for the present data (Fig. 4.21). Lowering the T_i used in the simulations by $\sim 30\%$ (or from 1270 eV to 960 eV for the medium power case) does improve the fits somewhat, but it does not change the conclusion that the deuterium tail increases slower than it should with total RF power.

Another interesting difference from earlier data is that the hydrogen signal level increases in Fig. 4.22 as the RF power increases, while the opposite occured in the power scan of Fig. 3.24. The earlier result was roughly consistent with a drop in the central neutral density as the electron density rose with RF power. The present data were fit with a higher hydrogen concentration (3%) than the earlier data (1%), and perhaps unconfined orbit losses have not yet caused the tail temperature to saturate. Note that the negative temperature region ($\partial f/\partial W > 0$ at the peak angle) appears only at the highest powers where orbit losses are largest and edge drag transport is most important. Unlike the fits in Figs. 4.8 and 4.11, the simulation in Fig. 4.22 does not reproduce the negative temperature. The E_+ profile must have very strong gradients at precisely the right place in order to do this.

4.2.1 Low power hydrogen spectra

Measured and predicted hydrogen spectra for even lower power levels (320 and 205 kW) from this same series of discharges are shown in Fig. 4.23. The perpendicular and peak fits are fairly good at high energies, but not at low energies. One way to increase the simulated low energy signal may be to raise the RF power near the edge, $r/a > .8$. The simulated low energy signal is already large enough at the highest power levels (1210 and 2300 kW in Fig. 4.22), but increasing the edge RF power would probably not change the high power fit because of orbit losses.

At the lowest power level (205 kW), unconfined orbit losses are rather

Figure 4.23: Measured and calculated hydrogen spectra at different angles for total RF powers of 205 kW and 320 kW. The hollow $E_+(r)$ of Fig. 4.20 was assumed.

small ($\sim 3.5\%$), and one should be more successful at inferring the RF power profile from the hydrogen spectra alone. The 205 kW data is simulated in Fig. 4.24 with $E+ \propto (1 - r^2/a^2)^{1/2}$ and hydrogen concentrations n_H/n_e of 3% and 6%. All other parameters are the same as the hollow E_+, 3% simulation of Fig. 4.23. The peaked power simulations of Fig. 4.24 do not do as well in predicting the relative amplitude of the charge exchange signal at different angles as the broad power simulation of Fig. 4.23. We attribute this to unaccounted-for variations in the neutral density, and instead focus our attention on the slopes of f at different angles. The 3% simulation in Fig. 4.24 produces the approximately correct slope at the peak and parallel angles, but it underestimates the slope of the perpendicular spectra. Increasing the hydrogen concentration to 6% makes the tail less energetic and improves the perpendicular fit, but now the peak and parallel slopes are not as good. It seems that a peaked power profile can not fit the data, regardless of the hydrogen concentration, providing further evidence for a broad power profile.

To make accurate measurements of the power profile, unconfined orbit losses must be negligible. For hydrogen minority heating in PLT, this means one must operate in a low power regime. The low power data presented here provides evidence that the power profile is broader than expected, but additional low power experiments need to be carried out to confirm this. A number of ways in which future experiments could improve upon the present ones are discussed in Sec. 4.4.

4.3 ^{3}He Minority Heating in a ^{4}He Majority Plasma

4.3.1 The double charge exchange process

The most successful ICRF heating results in PLT have been obtained in the ^{3}He minority, D majority regime. It is difficult to observe the RF produced ^{3}He tail using standard charge exchange techniques because most of the neutrals in the plasma are D^0 which have only one electron and cannot neutralize the $^3He^{++}$. Post, Grisham, and Medley[124] proposed to neutralize the energetic $^3He^{++}$ by *double charge exchange* with $^4He^0$ from standard neutral beam injectors. A variant of this method using multi-MeV, $Z \geq 2$

Figure 4.24: Same as the 200 kW data of Fig. 4.23, but simulations are with 3% and 6% hydrogen concentrations assuming a modestly peaked $E_+(r)$ profile.

neutral beams has been proposed for detection of 3.5 MeV alpha particles.[125] The double charge exchange process

$$^3\mathrm{He}^{++} + {}^4\mathrm{He}^0 \rightarrow {}^3\mathrm{He}^0 + {}^4\mathrm{He}^{++}$$

has a fairly high cross section (only a factor of ~ 10 less than the usual $\mathrm{H}^+ + \mathrm{H}^0 \rightarrow \mathrm{H}^0 + \mathrm{H}^+$ cross section) because it is a resonant process. Our charge exchange analyzer was not oriented to look across the sightline of the existing neutral beam injectors on PLT, so we decided to use a passive charge exchange version of the Post-Grisham-Medley scheme. By making a ^{4}He majority plasma instead of a D majority plasma (both species have the same e/m), we were able to get a large enough $^4\mathrm{He}^0$ neutral background to produce a measurable flux of energetic $^3\mathrm{He}^0$ neutrals to our detector. The cross sections for the FRANTIC neutral profile subroutine were modified to allow it to calculate the ^{4}He neutral density profile. The ionization cross section included not only electron impact ionization but also single charge exchange $^4\mathrm{He}^{++} + {}^4\mathrm{He}^0 \rightarrow {}^4\mathrm{He}^+ + {}^4\mathrm{He}^+$. The ^{4}He neutral density profile shown in Fig. 4.25 assumes a 5 eV edge neutral temperature and is similiar to the profiles calculated earlier for a deuterium plasma.

4.3.2 Analyzing ^{3}He neutrals

The energetic $^3\mathrm{He}^0$ must be reionized before it can be analyzed. As suggested by Post, Grisham, and Medley, H_2 was used in the stripping cell instead of the usual He because it has a higher cross section for ionizing He. Operation at high stripping cell pressures (gauge reading of 3.5 mTorr with a gauge factor of 2.0 for hydrogen) also improved the stripping efficiency. Measurements were also made with a 50 Åcarbon foil[126] but the count rate was lower than for a gas stripping cell, probably because of large scattering losses.

It was necessary to run the charge exchange analyzer at very high magnetic fields (~ 4.7 kG) to measure ^{3}He up to 150 keV because the gyroradius of singly charged $^3\mathrm{He}^+$ (which is what comes out of the stripping cell) is larger than the same energy proton. Although the magnet is water cooled, it can not run steady state at such a high field setting, and it was necessary to turn the magnet on and off just before and after each discharge. On a few occasions when the magnet was not turned off manually, the temperature safety interlock would shut down the magnet to prevent overheating

damage. Setting the analyzer to look at the highest possible energies had the added advantage of increasing the energy range ΔW, and therefore the count rate, of each individual anode in the microchannel plate detectors. Obtaining an adequate signal to noise ratio for the ^{3}He measurements was difficult not only because the charge exchange cross section is a factor of 10 less and the stripping efficiencies slightly worse than for hydrogen, but also because the x-ray background noise was fairly high, perhaps because of the larger breakdown voltages needed to start up a helium plasma. One of the disadvantages of operating at such a high magnetic field is that the lower energy particles fail to strike the proper mass anodes because the snubber (an iron cylinder between the stripping cell and the analyzing magnetic field which is supposed to shield out and compensate for fringe fields) saturates at high magnetic fields. This effect was well documented during the calibration of the instrument, where it was found that although all of the channels were good at 3.5 kG, the lowest sixth of the channels are unreliable at 4.0 kG, while the lowest third are unreliable at 4.5 kG. For the ^{3}He spectra taken at 4.7 kG, slightly more than a third of the data points have been deleted as inaccurate.

4.3.3 Measured spectra compared with theory

The ^{3}He spectra measured at three different viewing angles are compared in Fig. 4.26 to the bounce-averaged quasilinear calculations for a peaked power profile ($E_+ \propto (1 - r^2/a^2)^{1/2}$) and in Fig. 4.27 to calculations for a hollow power profile ($E_+(r)$ given by Fig. 4.9). The corresponding power balances are shown in Figs. 4.28–29. The main plasma parameters used for the simulations are shown in Fig. 4.25. The total RF power to the ^{3}He was 2100 kW after 15% coupling losses. The ^{3}He concentration $n_{^3\mathrm{He}}/n_e$ as measured by the density rise during the ^{3}He gas puff is 6%. Unlike the case of hydrogen minority heating, one might expect this to be a good measurement because the walls are less likely to absorb helium than hydrogen, a conjecture verified by the measurements of Chrien et al.[127] Other parameters used for these simulations include: $Z_{eff} = 3.4$, $Z_I = 6.7$ (which may seem low because it includes contributions from the ^{3}He), $a = 37$ cm, $R = 136$ cm, $B_{tor} = 31.04$ kG, $R_{res} = 143.4$ cm, $I_p = 450$ kA, $V_\ell = 2.3$ V, and $k_\| = \pm.07\mathrm{cm}^{-1}$.

Comparing the ^{3}He spectra of Fig. 4.26 to the H spectra of Fig. 4.2

Figure 4.25: Plasma and RF parameters for the simulation of ^{3}He minority heating in a ^{4}He majority plasma in Fig. 4.26.

Figure 4.26: Measured ^{3}He spectra compared with predictions assuming a hollow $E_+(r)$.

Figure 4.27: Measured ^{3}He spectra compared with predictions assuming a modestly peaked $E_+(r)$ profile.

Figure 4.28: Calculated power balance for the simulation shown in Fig. 4.26.

Figure 4.29: Calculated power balance for the simulation shown in Fig. 4.27.

leads to a number of interesting observations. As in the hydrogen case, the largest charge exchange flux is observed not at the perpendicular angle, but at an intermediate angle between perpendicular and parallel. The anisotropy between different viewing angles is not as strong as the hydrogen case, consistent with the higher pitch angle scattering rate of ^{3}He. The ^{3}He tail is not as energetic as the H tail, and no "negative temperature" is observed at the the peak angle. (The scatter in the data near 60 keV is thought to be instrumental. The ^{3}He data was taken at an early stage in the present researc, when a number of steps were being taken to improve the uniformity of the detectors.) The fraction of power lost on unconfined orbits is much less for ^{3}He than for H. The plasma current for this ^{3}He case was the same as for the H case. If it were increased by a factor of 1.5, as allowed by the higher toroidal magnetic field used for ^{3}He, unconfined orbit losses would decline further.

The peaked power simulation (Fig. 4.26) predicts a perpendicular spectra with a slightly hotter tail than observed, although the calculated peak angle slope is approximately correct. The broad power simulation (Fig. 4.27) produces the proper perpendicular slope but the peak slope is much too energetic. These results suggest that the ^{3}He RF power profile is not as broad as the hydrogen RF power profile. Not only are unconfined orbit losses smaller for the ^{3}He case, but the radius of the $q = 1$ surface is also smaller (measured to be $r = 9$ cm from the ECE inversion radius, versus $r = 15$ cm for the H case). This is because the toroidal field is higher while the plasma current is the same. These observations do not provide solid evidence that the broad hydrogen RF power profile is due to sawtooth and edge drag transport, but they are consistent with this interpretation. On the other hand, the low power hydrogen spectra of Fig. 4.3 seemed to indicate a broad RF power profile in a case of negligible unconfined orbit losses. A more definite conclusion which can be drawn from these data is that ^{3}He minority heating is more efficient than H minority heating because ^{3}He is more collisional, transfers more of its energy to ions instead of electrons, and suffers less from unconfined orbit losses.

4.4 Future Directions

In retrospect, it is clear that accurate measurement of the RF power profile from the hydrogen spectra alone is best done in a regime where unconfined orbit losses are negligible ($T_{\rm Stix} < .1 \times W_{\rm loss}$). In a higher current device this is easily satisfied, but in PLT this restricts the method to relatively low power levels (for hydrogen minority heating) where $T_{\rm Stix} < 70$ keV. It may even be best to operate with $T_{\rm Stix} \sim$ 5–20 keV on axis, because it is difficult to distinguish a 50 keV tail from a flat, infinite energy tail if measurements exist only up to 100 keV where the charge exchange reaction rate begins to drop. The hydrogen concentration should be increased to the 5–15% range, both to provide a good charge exchange count rate during the RF and to allow accurate measurements of the concentration. Another advantage of the high concentration is that the power profile may actually be as peaked as theoretically expected (Figs. 3.8 and 3.9). Increasing the concentration further may introduce the complications of wave reflection and mode conversion.

Measuring the RF power profile with good radial resolution requires that charge exchange data from many different viewing angles be analyzed, not the just the three angles studied in this work. Of course, measurements of the neutral density profile are always helpful for passive charge exchange analysis. D_α measurements of the poloidal and toroidal variation of the edge neutral density could be used to set the proper boundary conditions for a neutral profile calculation.

Future experiments to measure the RF power profile should investigate other approaches as well. A slight variation of the technique presented here is a perpendicular viewing, vertically scanning charge exchange analyzer. It always sees particles at their banana tips, which are in the resonance layer according to the resonance localization model. A vertical scan would thus give a direct radial profile measurement. A diagnostic neutral beam could be used to improve the spatial localization of the charge exchange measurements. During ^{3}He minority heating, other diagnostics which might also be useful include radial fusion product measurements[56,35] and charge exchange recombination spectroscopy.[128] Although multi-MeV, $Z \geq 2$ neutral beams are necessary to diagnose 3.5 MeV alpha particles, the double charge exchange technique we have demonstrated here could use standard neutral beam in-

jectors to measure the alpha particles once they have slowed down to the 50–400 keV range.

But beyond this, my son, be warned: the writing of many books is endless, and excessive devotion *to books* is wearying to the body. The conclusion, when all has been heard, *is:* fear God and keep His commandments, because this *applies to* every person. For God will bring every act to judgment, everything which is hidden, whether it is good or evil.

Ecclesiastes 12:12–14

Chapter 5

Summary

We have measured the energy and angle dependence of fast ions produced by ICRF, and compared these measurements with a bounce averaged Fokker-Planck program. This chapter will briefly summarize our results, and will suggest interesting areas for future research.

5.1 Summary of Results

Chapter 1 described the basic idea of resonance localization: ICRF heating produces energetic trapped particles whose banana tips are near the resonance layer. Resonance localization occurs not only because cyclotron heating gives particles perpendicular energy, but also because it gives more energy to particles which mirror near the resonance layer and so spend more time in resonance. Using this idea, we presented a qualitative framework for understanding peculiar charge exchange spectra such as Fig. 1.2.

The physics of resonance localization is quantitatively described by bounce averaged quasilinear theory, the topic of Chapter 2. We showed how the results of Bernstein and Baxter,[68] Mauel,[69] and Kerbel and McCoy[70] could be recovered with an extension of Stix's flux surface averaging technique. Our bounce averaged Fokker-Planck program solves for $f(W, \xi, r, t)$ as a function of energy W, pitch angle ξ, minor radius r, and time t. It integrates over a sightline to simulate charge exchange spectra. This program includes the complete effects of the Bessel functions and Doppler-shifts in the bounce averaged quasilinear operator. Finite banana width effects are

incorporated in unconfined orbit losses and in simulating charge exchange spectra. The usual collisional processes of energy drag, energy diffusion, and pitch angle scattering are also included. The program does not include radial transport of fast ions, which is indistinguishable in our model from a broadening of the RF power profile. There is a need for future work to properly include fast ion transport.

Chapter 3 presents clear evidence of direct second harmonic deuterium heating. Previous attempts at second harmonic deuterium heating have been thwarted by strong absorption at the fundamental resonance of the residual hydrogen in the plasma. Operating at low hydrogen concentration to reduce hydrogen absorption, and at low density to slow the rate of collisional relaxation to a Maxwellian, we were able to observe a very energetic deuterium tail (Fig. 3.11). The central RF power density, P_D, needed to sustain this tail could be found equally well by fitting the data with a simple isotropic quasilinear model, or with the complete bounce averaged Fokker-Planck program. We have studied the scaling of this measured P_D with total RF power (Figs. 3.25, 3.28, and 3.32) and hydrogen concentration (Fig.3.20). The first two power scans showed P_D increasing linearly with P_{total} between 500 kW and 1250 kW, while the third power scan indicated P_D increased less than linearly for P_{total} between 600 kW and 2300 kW. The dependence of P_D on hydrogen concentration was also weaker than expected. As discussed in Sec. 3.4.1, this nonlinear scaling is subject to a number of possible interpretations, one being that fast ion transport is increasing. The highest deuterium power densities were consistent with averaging the theoretically expected power profile (which is very peaked) over an $\sim$15 cm minor radius. Up to 22% of the total central RF power was directly absorbed by the deuterium.

In principle, the central RF power to the hydrogen can be inferred from the shape of the perpendicular hydrogen spectrum. However, in most of our experiments the hydrogen tail was so energetic that unconfined orbit losses made the shape of $f(W)$ independent of RF power (Sec. 2.5.1). Furthermore, the bounce averaged Fokker-Planck program suggested that, in many cases, the hydrogen charge exchange signal was dominated by energetic ions from the outer half of the plasma. A few low power experiments were done where unconfined orbit losses should be negligible. These data are more consistent with a hollow power profile (Fig. 4.23) than a peaked power profile (Fig. 4.25).

Chapter 4 focussed on the angular dependence of the charge exchange

spectra (Fig. 4.11), showing that the bounce averaged Fokker-Planck program could (with the proper assumptions) reproduce many interesting features of the data. The program produces the largest hydrogen signal at the proper "peak" angle. It is possible to reproduce the "negative temperature" observed at the peak angle by assuming a large dip in the RF power profile at $r/a = 0.75$, so that there are more 100 keV ions at point **A** in Fig. 1.5 than 30 keV ions at point **B**. In order to simultaneously fit the hydrogen and deuterium spectra, it is necessary to assume RF power profiles (Figs. 4.17 and 4.19) which are much broader than theoretically expected. The point made clear by the power scan of Chapter 4 is that there are many subtleties in the data, and that it is difficult to simultaneously fit all of the data with the same set of assumptions.

Finally, we have demonstrated the feasibility of a new double charge exchange diagnostic of ^{3}He. The ^{3}He spectra (Fig. 4.26) show resonance localization features similiar to the H spectra, and are consistent with ^{3}He heating being more efficient because unconfined orbit losses are less and ion heating is better.

5.2 Suggestions for Future Research

We need to understand why the RF power profile appears to be much broader than theoretically expected in many cases. We also need to understand why the central deuterium power density does not always scale as expected. There is a clear need to incorporate fast ion transport (such as sawtooth oscillations, and neoclassical, ripple, and RF-driven mechanisms) into quasilinear models. The role (or lack thereof) of RF-driven transport needs to be clarified. We have assumed that the minority density is proportional to the electron density in all of our modelling. This may not be true when unconfined orbit losses are substantial (and may not even be true in general), and improved modelling and measurements would be useful.

Our measurements of the deuterium and hydrogen tails, and their dependence on hydrogen concentration and RF power, contain a number of puzzling results. It would be useful to repeat these experiments on other machines, or even on PLT. Specific suggestions on how these experiments could be improved are given at the ends of Chapters 3 and 4.

The theoretically predicted RF power profile is very peaked. We observed RF power profiles which are much broader than this, but which are still comparable to typical neutral beam heating profiles. As a result, ICRF heating still works fairly well (Fig. 1.1). Major ICRF heating experiments are beginning around the world, and the next few years should be an exciting time for learning more about the physics of ICRF. Unconfined orbit losses should be less severe in higher current tokamaks. Sawtooth transport will remain, unless a way to stabilize the sawtooth instability can be found. The RF power may be better focussed in a larger, higher density tokamak. It will be interesting to test second harmonic deuterium heating in a high beta regime where it should be most efficient. We have observed a substantial hydrogen tail even at hydrogen concentrations as high as 15%, and future experiments in the high minority concentration regime would be of interest. Experiments with neutral beams and ICRF may provide another useful approach to learning more about the details of quasilinear theory.

Appendix A

Databases from the isotropic second harmonic model fits of Secs. 3.4 and 3.5.

Definitions of parameters:
DATE of discharge.
SHOT number identifying the discharge.
NEBAR line averaged density ($1/cm^3$).
PTOT total RF power (Watts).
TE0 central electron temperature (eV).
TI0 central ion temperature from second harmonic model fit (eV).
RFK defined in Chapter 3 (eV).
HTOD $\log(n_H/n_D)$.
FRHD $(n_H + n_D)/ne$.
FRH n_H/n_e.
TAUS slowing down time on electrons (s).
WCRIT critical energy above which electron drag dominates (eV).
DEND central n_D ($1/cm^3$).
PDCX Measured central deuterium power density (W/cm^3).
PAV estimated RF power density inside 15 cm (W/cm^3).
KPER2 $k_\perp^2$ from Stix's cold plasma equations (cm^{-2}).
PDTHR Theoretical deuterium power density averaged over 15 cm (W/cm^3).
THSTX T_{Stix} (eV).

Parameters for discharges in Figs. 3.20–3.22.
Parameters for discharges in Figs. 3.23-3.27.
Parameters for discharges in Figs. 3.28-3.31.
Parameters for discharges in Figs. 3.32–3.34.

Appendix B

Input files for the sample simulation of Sec. 2.6.4.