Cooling Alpha Particles With Waves

Mark Christopher Herrmann

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

NOVEMBER 1998

©Copyright by Mark Christopher Herrmann, 1998. All rights reserved.

Abstract

DIVERTING ENERGY FROM ENERGETIC α-particles to waves (α-channelling) would be extremely beneficial for a magnetically confined deuterium tritium fusion reactor. If these waves were to damp on fuel ions, a hot ion mode would result, doubling the fusion power of the reactor at the same confined pressure. Alternatively, if these waves damp preferentially on electrons traveling in one direction, current would be driven. In both cases, the pressure profile could be modified and ash could be removed to advantage.

These potentially significant benefits motivate a detailed study of the implementation of α-channelling. This thesis identifies and explores issues in realizing α-channelling, making the following advances:

1. Calculating, through 0-dimensional reactor simulations, the substantial benefit in α-channelling which motivates the subsequent work.

2. Framing the α-channelling problem as a diffusion problem, with absorbing and reflecting boundaries. To develop insight, α-channelling is considered in a 2-dimensional phase space associated with a simple slab geometry. To solve realistic problems we pose α-channelling in the 3-dimensional constants of the motion space associated with particle orbits in tokamaks.

3. Developing a rapid Monte Carlo simulation in constants of motion space to keep track of wave-induced and collisional effects on the energetic particle distribution. This approach is equivalent to the full energetic particle dynamics in the limit of small changes during a single bounce time and diffusive wave-particle interactions.

4. Identifying the wave characteristics necessary to produce the channelling effect, which we discover are available in a combination of the mode converted ion-Bernstein wave

(MCIBW) and the toroidal Alfvén eigenmode (TAE).

5. Demonstrating how two waves can be combined in a reverse shear tokamak reactor to absorb 2/3 of the energy from the 93% of the α-particles ejected!

6. Showing how the basic building blocks of the α-channelling effect can be deduced from existing experimental data, including:

> *a. Reproducing, qualitatively, the results of TFTR experiments which show strong interaction of MCIBW with fast ions.*
>
> *b. Demonstrating the existence of the $k_{\|}$-flip of the MCIBW.*
>
> *c. Using the simulation to infer from experimental data a MCIBW diffusion coefficient, which significantly exceeds that which is predicted by geometrical-optics estimates.*

Taken together, the advances in this thesis show how experiments to date give us a measure of confidence in both the simulations themselves, the underlying physical assumptions, and ultimately the reasonableness of the application of these ideas to α-channelling in a tokamak reactor.

Acknowledgments

I have relied on many people in the writing of this thesis. Foremost, I would like to thank my advisor Nat Fisch who has been a constant source of wisdom, encouragement, and enthusiasm. His sage advice to "work as quickly as possible without sacrificing accuracy" has guided me through the completion of this thesis. On a more serious note, Nat has always been there at all hours of the day and night for discussions, and he has helped me stay focused on the problems that were important. I have learned a tremendous amount from him and will sorely miss both our personal and professional interaction.

This thesis has relied heavily on the resources available at the Princeton Plasma Physics Laboratory. In particular, I would like to thank Roscoe White and Yi Zhao for the spline and ORBIT codes which were used to incorporate general magnetic geometry and benchmark my simulation. Ernie Valeo provided a ray tracing code which was crucial to the comparisons with experiment. This thesis could not have been completed without the PPPL theory Unix cluster. Thanks to Ernie Valeo and Charles Karney for the many times I have interrupted their busy schedules to ask a question or get something fixed. The most wonderful resource of all was TFTR. To participate on experiments on this machine was truly an exhilarating experience. In particular, Doug Darrow and Dick Majeski designed and (with the help of the TFTR team) carried out the experiments which are discussed herein. Figures 6.3, 6.7, 6.8, 6.9, 7.1 were given to me by Doug Darrow. Most of all thanks to them for patiently explaining to me the differences between a theorist's naive perception and experimental reality.

This thesis has benefited from the comments of readers Dick Majeski, Charles Karney, as well as Mike Beer and Bob Heeter. Thanks for reading my thesis on such short notice.

I am sure I will not soon forget the joys of working in the theory wing cubicles and sleep-

ing under my desk. I would like to thank John Wright and Steve Smith for sharing many pointless arguments and conversations over the wall separating T191 from T193. Thanks to Phil Snyder who put up with my increasingly eccentric behavior during the final two months of my thesis. Thanks to Mike Beer, for sharing many laughs, reintroducing me to the game of chess, and seemingly never growing tired of being interrupted by my walking into his office to ask a question. Throughout my five plus years in graduate school I have shared many interesting conversations with Ben Chandran, Hans Herrmann, Jon Menard, Peter Stoltz and others.

To Barbara, thank you for helping me clear the many administrative hurdles. More importantly, you were always there to lend an ear, and to dote on Katie and Keith. You were like family to us. We will miss you but promise to keep in touch.

The patience of my wife, Kim, through these 5.75 years has been saintly. She has always, always been supportive even though these years have not been easy on her. In particular, for the last several months, when my presence at home was occasional at best, I say thank you. I hope that in some way I can pay her back over the course of our life together. Thanks to my daughter Katie who was always willing to drop whatever she was doing to be “Super Kitty” when I wanted to play “Superfriends” and fight off the Hall of Doom. Thanks to my son Keith whose constant cheer has helped me keep a smile on my face.

To my parents, words cannot describe the thanks I owe you. Simply put, this would not have been possible with out your love and support, from 1970 to today. To my brother Matt, thanks for being a great brother. The readers of this thesis can judge whether or not that tremendous rock that you hit my head with did any permanent damage.

Finally I would like to acknowledge the generous support of the Fannie and John Hertz Foundation.

Table of Contents

Abstract **iii**

1 Introduction **1**

1.1 Background 2

1.1.1 Fusion as an energy source 2

1.1.2 The need for concept innovation 3

1.2 Basic Elements of α-Channelling 4

1.2.1 Reactor implications of α-channelling 5

1.2.2 Requirements for α-channelling 7

1.3 Key Accomplishments 7

1.4 Outline of the Thesis 11

2 Reactor Implications of α-Channelling **13**

2.1 Current Drive 13

2.2 Hot Ion Mode 14

2.2.1 0-D power balance model 16

2.2.2 Operating point with no diversion 17

2.2.3 Operating point with diversion 22
2.2.4 Confinement times 25
2.2.5 Advanced fuels 29
2.3 Other Consequences of α-Channelling 29
2.4 Summary 32

3 Building Blocks of the Channelling Effect 35
3.1 Free Energy 35
3.2 Wave-Particle Interactions 36
3.3 Generic Requirements for α-Channelling 41
3.3.1 Tapping velocity space 42
3.3.2 Resonance conditions 42
3.3.3 Time scales 43
3.3.4 Deposition of channelled power 43
3.3.5 Power lost/wall loading 43

4 Constants of Motion Simulation 45
4.1 Orbits and Outline of the Simulation 47
4.2 Collisions 50
4.3 Stochastic Ripple Diffusion 57
4.4 Wave Particle Interactions 60
4.4.1 Action-angle formulation 61
4.4.2 Quasilinear diffusion in action-angle variables 62
4.4.3 Stochasticity 65

4.4.4 Quasilinear diffusion coefficient 68
4.4.5 Calculation of the diffusion coefficient for the MCIBW 70
4.4.6 MCIBW benchmark 71
4.4.7 Collisionless diffusion 72
4.4.8 Calculation of diffusion due to toroidal Alfvén eigenmodes . . . 79
4.5 Collisional Evolution of the Beam Ion Distribution 81
4.6 Timing 83
4.7 Summary 85

5 Reactor Scenario for α-Channelling with Two Waves 87

5.1 Basics 88
5.2 Cooling Scenario with Two Waves 92
5.3 Other Two Wave Scenarios 95
5.4 Cooled Distribution at the Wall 100
5.5 Two Key Questions 102
5.5.1 $k_{\|}$ flip of the MCIBW 103
5.5.2 Power for the toroidal Alfvén eigenmode 106
5.5.3 Power for the MCIBW 109
5.6 Summary 111

6 Modeling Fast Particle-MCIBW interactions on TFTR 113

6.1 Overview of Experiments 114
6.1.1 MCIBW physics 116
6.1.2 Fast ion species 119

6.1.3 Losses of fast particles and their detection 121

6.2 Fast Particle MCIBW Interactions on TFTR 125

6.3 Theoretical Considerations 130

6.3.1 IBW wave physics 131

6.3.2 Resonance condition 138

6.3.3 Trajectory in $v_\perp, v_\parallel$ space 141

6.3.4 Effect of n_ϕ on the particle trajectories 146

6.3.5 Theoretical summary 150

6.4 Simulation Results 151

6.4.1 Toroidal field scan 153

6.4.2 $n_{^3\mathrm{He}}/n_e$ scan 156

7 Key Implications of the TFTR experiments 163

7.1 The $k_\parallel$–flip 164

7.2 Diffusion Coefficient 167

7.2.1 Simple estimate 170

7.2.2 Simulation results 172

8 Directions for Future Work and Conclusions 177

8.1 Improvements to α-Channelling Scenarios 177

8.2 Further Experimental Validation 179

8.3 Applications to Alternate Concepts 180

8.4 Summary 182

A Energetic Particle Orbits in Tokamaks **185**

A.1 Constants of Motion Approach 186

A.2 General Geometry 189

A.3 Solving for Orbits 193

A.3.1 Orbit types 194

A.3.2 Bounds on orbits 198

A.3.3 Comparison with ORBIT code 202

A.4 ε, μ, and P_ϕ space 203

A.4.1 Basics 204

A.4.2 Stagnation point orbits 206

A.4.3 Classifying orbits 207

A.5 Averages Along the Particle Motion 209

Bibliography **213**

List of Figures

1.1 Flow of power under α-channelling 6

1.2 A child's game. 10

2.1 Fusion power density with no diversion 21

2.2 Fusion power density with diversion of 75% of the α-particle power . . 26

2.3 Fusion power density with diversion of 100% of the α-particle power and 100% ash removal 27

3.1 Diffusion path which extracts free energy from a inhomogeneous α-particle distribution 36

3.2 Optimal diffusion path for extracting α-particle energy. 38

3.3 Diffusion paths typical of waves in tokamaks. 39

3.4 Diffusion paths for a combination of two waves 40

3.5 Diffusion paths for a combination of two waves which could extract almost all of the α-particle energy. 41

4.1 Flow chart of COM simulation 49

4.2 Benchmark of slowing down due to collisions 55

4.3 Benchmark of energy diffusion due to collisions 56

4.4 Collisional evolution of a δ function in pitch. 56

4.5 Collisional evolution of a δ function in pitch, for longer times. 57

4.6 Comparison of predicted and actual stochastic loss domain 60

4.7 Benchmark of MCIBW constant D case 74

4.8 Benchmark of MCIBW nonconstant D case 75

4.9 Discontinuous D 76

4.10 Distribution function when D is constant 77

4.11 Distribution function when D is discontinuous 77

4.12 Losses versus time when D is discontinuous, explicit method. 79

4.13 Losses versus time when D is discontinuous, implicit method. 80

4.14 The density of beam particles vs. energy at 20 ms intervals 83

4.15 The density of particles vs. pitch 84

4.16 The density of particles vs. ψ 84

5.1 Orbits of an α-particle interacting with the MCIBW and being cooled . 90

5.2 Orbits of an α-particle interacting with the TAE and being cooled 91

5.3 Orbit of an α-particle interacting with both the MCIBW and the TAE . . 92

5.4 Energy extracted (MeV) by the MCIBW vs. initial location of particle in constants-of-motion space. 93

5.5 Energy extracted (MeV) by the TAE vs. initial location of particle in constants-of-motion space. 94

5.6 Location of waves for two wave scenario. 95

5.7 Histogram of the energy extracted from particles when the MCIBW is near the center. 96

5.8 Energy extracted (MeV) by the MCIBW vs. initial location of particle in constants-of-motion space. 98

5.9 Poloidal projection of a deeply trapped α-particle interacting with the TAE. 99

5.10 Energy extracted by the MCIBW when the n_ϕ/ω for the TAE is tailored. 100

5.11 Energy extracted by the TAE when n_ϕ/ω for the TAE is tailored. . . . 101

5.12 Velocity space position for particles leaving the tokamak after having their energy extracted. 102

5.13 Histogram of ejected α-particles vs. poloidal angle. 103

5.14 Effect of the $k_\parallel$ flip on particles transitioning from trapped to passing . . 105

6.1 Schematic of the D^{3}He mode conversion experiment 115

6.2 Location of the mode conversion layer vs. $n_{^3\mathrm{He}}$ and B_0 117

6.3 Detection of the first orbit loss of an α-particle and schematic of the lost alpha detector. 122

6.4 Examples of first orbit losses for 3.5 MeV α-particles 123

6.5 Distribution of first orbit losses 124

6.6 Examples of confined particles transitioning to lost orbits 125

6.7 Lost alpha signal vs. time for two discharges, one of which exhibits MCIBW-fast ion interactions. 127

6.8 The gyroradius distribution of the MCIBW related loss and first orbit loss at the 90° detector. 128

6.9 Losses of cogoing and countergoing deuterium beam ions in the presence of symmetrically phased RF. 130

6.10 Losses versus toroidal field at fixed $n_{^3\mathrm{He}}/n_e$. 131

6.11 Spectrum of coupled power for the TFTR bay M antenna, assuming symmetric phasing. 132

6.12 Spectrum of coupled power for the TFTR bay L and M antenna assuming symmetric phasing. 133

6.13 Projection of the poloidal field in the $\hat{R}$ direction 136

6.14 k_X vs. R for the MCIBW for $n_{^3\mathrm{He}}/n_e$ = 0.15, B = 5.0 T 137

6.15 k_X vs. $k_\parallel$ for the MCIBW for $n_{^3\mathrm{He}}/n_e$ = 0.15, B = 5.0 T 138

6.16 $k_\parallel$ vs. R for the MCIBW for $n_{^3\mathrm{He}}/n_e$ = 0.15, B = 5.0 T. 138

6.17 Electric field amplitude vs. $k_\parallel$ for the MCIBW for $n_{^3\mathrm{He}}/n_e$ = 0.15, B = 5.0 T 139

6.18 Passing/trapped boundary and the trajectories of a 100 keV D interacting with either the MCIBW or a fast wave vs. $v_\parallel$ and $v_\perp$. 144

6.19 Poloidal projection of the orbit of a D interacting with the MCIBW. . . 145

6.20 Passing/trapped boundary and trajectory of a 100 keV T interacting with MCIBW vs. $v_\parallel$ and $v_\perp$. 146

6.21 Poloidal projection of the orbit of a T interacting with the MCIBW. . . 147

6.22 A countergoing deuteron interacting with MCIBW for various values of $n_\phi < 0$ 148

6.23 A cogoing deuteron interacting with MCIBW for two values of $n_\phi > 0$. 149

6.24 Trajectory of deuteron interacting with MCIBW in $v_\parallel, v_\perp$ for different values of the toroidal field 154

6.25 Loss energy vs. poloidal angle for five different values of toroidal field . 156

6.26 Pitch angle vs. poloidal angle for five different values of toroidal field . 157

6.27 Loss energy vs. poloidal angle, dependence on spot size 157

6.28 Loss energy vs. poloidal angle, dependence on spot size 158

6.29 Losses of beam deuterons for $n_{^3\mathrm{He}}/n_e$ = 0.1 (a), 0.15 (b) and B = 5.0 T . 159

6.30 Losses of beam deuterons for $n_{^3\mathrm{He}}/n_e$ = 0.2 (a), 0.25 (b) and B = 5.0 T . 160

6.31 Losses of beam deuterons for $n_{^3\mathrm{He}}/n_e$ = 0.15 (a), 0.2(b) and B = 4.86 T 161

7.1 Losses during phased RF. 166

7.2 Beam blip experiment 169

7.3 Beam blip experiment at low RF power 170

7.4 Collisionless and collisional loss of resonant ions vs. electric field amplitude 173

7.5 Fractional loss of deuterium ions vs. RF power. 174

7.6 Mean loss time of the exiting particles vs. power level 175

8.1 Simulation showing cooling of α-particles in a reverse shear TFTR discharge 180

A.1 Alternate representations of a poloidal cross section for a high aspect ratio, circular flux surface geometry 188

A.2 B_{orb} versus ψ along a particle orbit 189

A.3 Low energy orbits in both R,Z space and B, ψ space. 190

A.4 Alternate representations of a poloidal cross section for a reverse shear aspect ratio 3 reactor 192

A.5 Orbit type 1 195

A.6 Orbit type 2 196

A.7 Orbit type 3 196

A.8 Orbit type 4 197

A.9 Orbit type 5 197

A.10 Orbit type 6 198

A.11 Orbit type 7 198

A.12 Orbit type 8 199

A.13 Orbit type 8 blow up 201

A.14 Poloidal orbits generated by the COM approach and the ORBIT code . 203

A.15 A constant energy slice of ε, μ, and P_ϕ space. 205

A.16 Boundaries in ε, μ, and P_ϕ space 206

A.17 ε, μ, and P_ϕ blowup 208

A.18 $1/\dot{\psi}$ along a particle's orbit 210

A.19 ω_b vs. λ at high energy 212

A.20 ω_b vs. λ at low energy 212

List of Tables

2.1 Parameters for the ARIES-RS tokamak. 18

2.2 Aries-RS operating points, without diversion 19

2.3 Aries-RS operating points, with diversion 24

2.4 Aries-RS operating points, with diversion and high Z impurities 28

6.1 Some typical collisional times 120

6.2 Required $k_{||}$ for fast ions to resonate with MCIBW 140

A.1 Description of all the orbit types and their branches 202

Chapter 1

Introduction

IF THE POWER from energetic α-particles could be diverted to ions via waves (hereafter known as α-channelling), deuterium tritium magnetic confinement fusion reactors could be made far more attractive. For instance, the ions could be made significantly hotter than the electrons, doubling the output power of the reactor at the same confined pressure. At the same time, the wave power might be put to other beneficial uses, such as current drive. In some schemes to channel the α-particle power, ash would be coincidentally removed. Opportunities for pressure or current profile control also exist.

The main goal of this thesis is to advance α-channelling from an abstract idea in a simplified geometry towards a concrete implementation in a tokamak reactor. We lay a theoretical, numerical, and experimental foundation for the implementation. We quantify the advantages of α-channelling. We develop a theoretical framework for considering the problem in tokamak geometry, including the full energetic particle dynamics, simulate numerically α-particles interacting with the waves which might accomplish the α-channelling, find a scenario using two waves which accomplishes significant α-channelling in a reverse-shear tokamak reactor, and model TFTR experiments.

In the process of working towards the larger goal of advancing α-channelling, we developed a numerical simulation which is significantly faster than guiding center codes for a large class of problems. Using this numerical tool we managed to deduce

interesting and fundamental wave physics from TFTR experiments.

1.1 Background

World energy demand is growing rapidly, spurred on by the economic growth of developing countries. It is expected that, by the year 2030, the world consumption of energy will double from present levels [HOLDREN 1997]. At the same time, there is a limited supply of fossil fuels (although coal and natural gas reserves are probably sufficient into the twenty-second century). More serious, though, is a growing body of evidence which suggests that burning these fuels could change our environment by significantly increasing the levels of carbon dioxide in the atmosphere. Currently, there is uncertainty about the precise effects of increased carbon dioxide in the atmosphere and its consequences for humanity [MAHLMAN 1997]. Without regard to the nature of these consequences, there is no doubt that rising carbon dioxide levels represent a global risk. It is prudent to begin planning now on how society can mitigate this risk. In addressing both the limited supply of fossil fuels and global warming, fusion energy could play an important role.

1.1.1 Fusion as an energy source

Fusion offers the possibility of a high energy density power supply (as opposed to solar and wind energy, for instance) which can be used anywhere in the world to handle the baseline power load. Fusion is a carbon-free source of energy, and therefore it does not have an impact on the global climate (unlike fossil fuels). If a fusion reactor could be built out of low activation materials, the radioactive waste associated with thirty years of operation of a fusion power plant would qualify for shallow burial and be up to a million times less radioactive (after one year) than the corresponding waste from a fission plant [CONN et al. 1990]. Furthermore, while the dangers associated with fission power plants are frequently overstated, worst case scenarios for fusion involve doses at the site boundary (1 km) 100-500 times less than a fission plant [CONN et al. 1990] and would not require evacuations. Additionally the fuels used in fusion, unlike fission, do not pose a proliferation threat. Finally both the deuterium and lithium needed to make

the deuterium tritium cycle go are abundant, and found all over the world, providing a virtually limitless, apolitical supply of energy (unlike fossil fuels).

On the other hand, there are aspects of fusion which are not attractive; it is very complicated, it has not yet been demonstrated, and current projections suggest it will be an expensive energy source. Many of fusion's advantages over other energy sources are apparent but not readily quantifiable. For instance, when comparing to fission energy, how many cents per kilowatt hour is it worth to have a safer energy source which has less radioactive waste, and significantly reduced proliferation concerns?

1.1.2 The need for concept innovation

Much of the work of the last 47 years of fusion energy research focused on demonstrating fusion energy. Tremendous progress was made, not only in plasma performance, but also in the understanding of hot magnetically confined plasmas. Recently, experiments in the U. S. [McGuire et al. 1995] and Europe [Gibson and the JET Team 1997] achieved values of Q (fusion power out/power deposited in the plasma) approaching one (breakeven). A detailed engineering design was developed for the International Thermonuclear Experimental Reactor (ITER), a machine whose mission includes studying operation of a tokamak at high Q.

However, there has been a growing realization that "our present, conventional approach to deuterium tritium (DT) thermonuclear fusion may not lead to a truly attractive reactor product able to compete in the energy marketplace of the twenty-first century" [Perkins et al. 1995]. Galambos et al. (1995) show that, even under optimistic assumptions, power from a magnetic confinement fusion device will be significantly more expensive than the cost of power projected from better-experience fission power plants. While some cost differential is justified in light of fission's many external costs (safety, waste, proliferation), there is a widespread belief that, in order for fusion to compete, it will be necessary for the cost of electricity to be close to that of other energy sources.

Furthermore, the projected expense of a tokamak reactor is proving to be a barrier to development. While ITER is just an experiment, not a reactor or a reactor prototype,

the 8 billion dollar price tag (or even the 5 billion dollar price tag of ARIES-RS) has led many to reconsider the path towards fusion energy. Even if ITER were an experimental success (about which there is considerable uncertainty), it is not clear that such a success would point to an economically viable energy source.

Thus, the search for fusion energy finds itself today at a crossroads. Much of the effort leading up to the present focused on overcoming the obstacles in making fusion energy. However, while we may be able to create a fusion reactor, it is not yet known how to make a fusion reactor which can be economically competitive with other power sources. In light of these concerns, there is now a renewed focus on "alternate concepts", concepts that differ significantly from the conventional tokamak. At the same time, the tokamak is being reexamined in this light to see how much it might be improved itself.

This thesis is an investigation of α-channelling, a mechanism by which fusion reactors might be significantly improved. While the comprehensive control of the α-particle distribution necessary to accomplish α-channelling is a daunting task, the potentially large reward, along with the need for improvement in current concepts, makes this investigation worthwhile. The focus here is on α-channelling in tokamaks, although, in principle, α-channelling might also be useful in alternate concepts.

1.2 Basic Elements of α-Channelling

The main reactions by which fusion energy might be achieved in the foreseeable future are:

$$\mathrm{D} + \mathrm{T} \rightarrow {}^4\mathrm{He}\,(3.52\ \mathrm{MeV}) + \mathrm{n}\,(14.1\ \mathrm{MeV}) \tag{1.1}$$

$$\mathrm{D} + \mathrm{D} \underset{50\%}{\rightarrow} {}^3\mathrm{He}\,(0.82\ \mathrm{MeV}) + \mathrm{n}\,(2.45\ \mathrm{MeV}) \tag{1.2}$$

$$\underset{50\%}{\rightarrow} \mathrm{T}\,(1.01\ \mathrm{MeV}) + \mathrm{p}\,(3.02\ \mathrm{MeV}) \tag{1.3}$$

$$\mathrm{D} + \mathrm{He} \rightarrow {}^4\mathrm{He}\,(3.67\ \mathrm{MeV}) + \mathrm{p}\,(14.7\ \mathrm{MeV}) \tag{1.4}$$

In particular, the DT reaction is the most promising, as it has the highest reactivity per unit pressure. After the D and T fuse, typically 80% of the energy is carried off by a 14 MeV neutron which will deposit its energy somewhere in the wall of the reactor. The other 20% is carried off by the 3.5 MeV α-particle, which is confined by the strong magnetic field of the reactor, and slows down predominantly on the thermal electrons. It is this power from the α-particles which allows the fusion reaction to be self-sustaining, and it is this power that we seek to tap with α-channelling.

In contrast to the experiments of today, a DT magnetically confined fusion reactor will be dominated by energetic α-particles. For instance, in a reactor with a Q of 20, the α-particle power will be four times larger than the auxiliary heating power. Thus, the auxiliary heating flexibility of today's experiments will be in large part lost upon the transition to reactor sized machines. In view of this, any means of controlling the charged fusion products (CFP's) becomes quite valuable, as it may allow us to regain some control and substantially alter the operating characteristics of the reactor.

For example, while many experiments today achieve their best fusion power in the hot ion mode ($T_i > T_e$), in a reactor, the α-particles deposit their power on electrons, forcing $T_e \geq T_i$. Hot ion mode operation is advantageous since it allows more fusing ions to be present for the same total confined pressure. Clarke (1980) suggested that the hot ion mode could be attained in a reactor if, in some unspecified way, the power from the α-particles could be diverted to the ions. Fisch and Rax (1992a) suggested a mechanism using injected waves to accomplish Clarke's transfer of α-particle power to the ions, where waves diffuse α-particles both in space and in energy, thereby allowing the free energy of the α-particles to be tapped. In principle, all of the α-particle power could be diverted to waves in this way. The process is shown schematically in Fig 1.1.

1.2.1 Reactor implications of α-channelling

Through this diversion of α-particle power to ions, a reactor operating at $T_i = T_e = 20$ keV might instead be operated at $T_i = 20$ keV, $T_e = 12$ keV, if 75% of the α-particle power could be diverted to the ions via waves [FISCH and HERRMANN 1994; SNYDER, HERRMANN, and FISCH 1994]. This change would almost double the fusion power of a

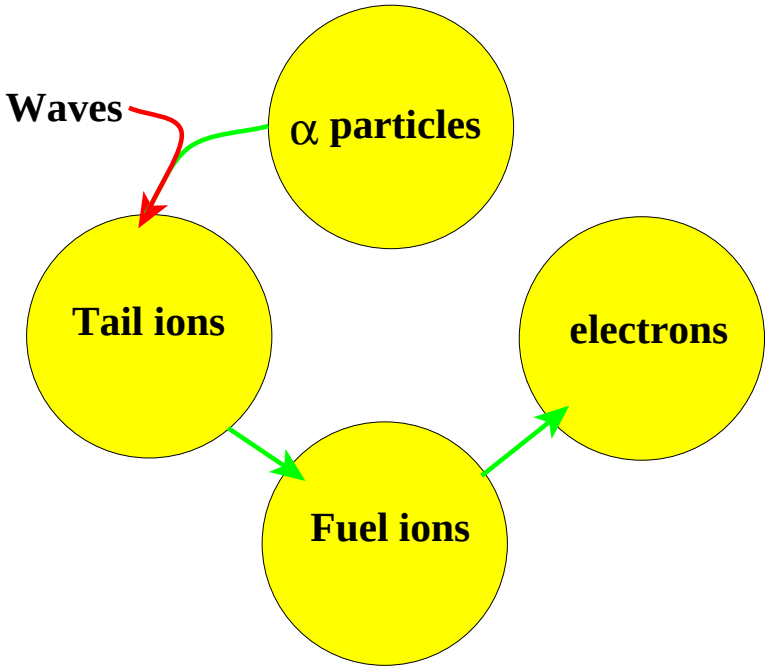

Figure 1.1: The goal of α-channelling is to redirect the α-particle power, which normally flows to electrons, to the ions thereby increasing the reactivity of the plasma.

reactor operating at the same magnetic field and pressure. Other benefits might include ash removal, a reduction in the fast particle pressure (reducing the drive for undesirable instabilities), and, in principle, the waves used to divert the power may damp in such a way as to drive a current [FISCH 1987]. Another effect is that α-channelling is typically optimized when the electron confinement time is very short (but above some minimum value), and the ion confinement is long. Such configurations [HERRMANN, FISCH, and SNYDER 1994] might be attained naturally, for example, in enhanced reverse shear plasmas, where the ion confinement is much better than the electron confinement [LEVINTON et al. 1995], or purposefully, for example, by injecting very high-Z impurities into the plasma to increase the radiated power, at the same time decreasing the heat load to the divertor. Experimental observations [SCOTT et al. 1994] and theoretical considerations [DORLAND et al. 1994] suggest yet another possible benefit, namely, that increasing the ratio of T_i/T_e significantly improves heat confinement. While most of the work in this thesis focuses on DT fuel, many of the benefits of α-channelling are also available for other fusion fuels.

1.2.2 Requirements for α-channelling

Thus, if α-channelling could be accomplished, it would have significant impact on the attractiveness of a magnetic fusion reactor. Since the benefits are clear, the challenge is to determine a way to implement α-channelling. The implementation will involve a search for the appropriate waves.

There are three effects that must occur so that significant diversion of power by waves is accomplished:
1. The waves must create a diffusion path in phase space, such that α-particles at high energy in the center diffuse to low energy near the edge of the tokamak.
2. The waves must interact strongly enough with the α-particles that the α-particle energy is diverted before the α-particle collisionally transfers its energy to the electrons. The wave is then convectively amplified.
3. The convectively amplified wave must then damp on ions.

1.3 Key Accomplishments

The key accomplishments of this thesis are:

1. We pose the α-channelling problem in a tokamak as a diffusion problem with absorbing and reflecting boundaries in the 3-dimensional constants of motion (COM) space. The α-particles in a tokamak exist, of course, in a 6 dimensional phase space, although making the guiding center approximation reduces this space to 5 dimensions. However, for α-channelling, we need not be concerned with the full 5 dimensional space, but rather just three dimensions. The constants-of-motion space [ROME and PENG 1979], using the particle's energy, ε, magnetic moment, μ, and canonical angular momentum, P_ϕ, [HSU and SIGMAR 1992] for the constants is a natural space in which to formulate the transport of the α-particles in the tokamak due to waves and collisions. Remarkably, the diffusion paths due to resonant wave-particle interactions in this space lay along a

straight line, giving us insight into the wave characteristics which are required to achieve α-channelling in a tokamak. This space is also ideally suited for a rapid Monte Carlo simulation of α-channelling as described below.

2. We find that waves which satisfy the requirement that their diffusion path bring α-particles to edge while cooling them would be very difficult to excite in a reactor sized tokamak. However, we discovered that, by using waves with two very different frequencies, we are able to cool almost completely a single α-particle while bringing it to the edge [FISCH and HERRMANN 1995].

The two waves used to extract almost all the energy from a single α-particle are: one with $\omega \ll \Omega_\alpha$ (such as the toroidal Alfvén eigenmodes (TAE) [CHENG, CHEN, and CHANCE 1985; CHENG and CHANCE 1986]), and one with $\omega \sim \Omega_\alpha$ (such as the mode converted ion Bernstein wave (MCIBW) [PARK, LEE, PEEBLES, and LUHMANN 1985]). The high frequency wave extracts perpendicular energy, while the low frequency wave diffuses the α-particles to the tokamak periphery and extracts parallel energy. However, unlike the case of one wave only [FISCH and RAX 1992a], with two waves, there are no constraints on the particle motion, so that some ejected α-particles may be heated while others are cooled. This leaves open the question of how to find waves that successfully cool the full birth distribution of α-particles.

3. To expedite this search, we develop a rapid Monte Carlo simulation in constants-of-motion space to keep track of wave-induced and collisional effects on the energetic particle distribution. The simulation solves the orbit-averaged Fokker-Planck equation for the evolution of the energetic particle distribution function. The full energetic particle guiding center dynamics in general tokamak geometry are included. Additionally we include the effects of collisions, through the Landau collision operator, and effects of wave-particle interactions, assuming quasilinear diffusion is valid. This approach is equivalent to the full energetic particle dynamics in the limit of small changes of the constants of motion during a single bounce time and wave-particle interactions which satisfy the random phase criterion.

While in principle a guiding center code could be modified to simulate these interactions (indeed, the guiding center code ORBIT [WHITE and CHANCE 1984] already has a collision operator and TAEs, and incorporates the effect of ripple, but would need

to be modified to include resonant interactions of α-particles with waves which break the μ invariant), these codes are currently too slow to simulate the problems encountered in this thesis. In some of the examples used in this thesis COM simulations were able to simulate in 50 minutes of CPU time on DEC Alpha workstation what would take approximately 15 hours of CPU time on the CRAY C-90 computer at NERSC using ORBIT. Of course the development of the COM simulation requires significantly more overhead than guiding center codes such as ORBIT since each additional physics effect must be explicitly added. Additionally, the COM simulation is only applicable in the case where the changes in the constants of motion during a single poloidal orbit are small, and the random phase approximation is valid.

4. Using the Monte Carlo simulation, we find combinations of two waves which are able to extract a significant fraction of the energy from a birth distribution of α-particles. Accomplishing α-channelling in a reactor is like shaking particles out of a bottle through certain holes. The 3-dimensional volume here is the ε, μ, and P_ϕ space of the α-particles; the boundary of the bottle corresponds to values of these constants of the motion for orbits intersecting the physical boundary of the tokamak. Waves diffuse particles in this constants-of-motion space (ε, μ, and P_ϕ space). The trick is to devise plasma waves that shake most of the α-particles into "holes" in the bottle at low energy.

In devising such a bottle, the simultaneous satisfaction of criteria for different particles can be very frustrating, much as in the case of the analogous child's toy shown in Fig. 1.2. Here, while waves might be devised to extract energy from a single α-particle [FISCH and HERRMANN 1995], getting all of the particles to go into the lowest energy holes in response to the same set of waves is not simple. On the other hand, the parameter space of possible waves and magnetic bottle configurations is immense.

In the best case, based on an advanced tokamak reactor, we find that over 60% of the α-particle power can be diverted to a set of idealized waves [HERRMANN and FISCH 1997]. In this case, 93% of the α-particles were ejected from the plasma, after losing 2/3 of their initial energy (on average).

5. By comparing with experiments on TFTR, where a strong interaction of fast particles with the MCIBW was observed, we were able to validate, qualitatively, the simulation and uncover two results of critical importance to α-channelling. Experiments on

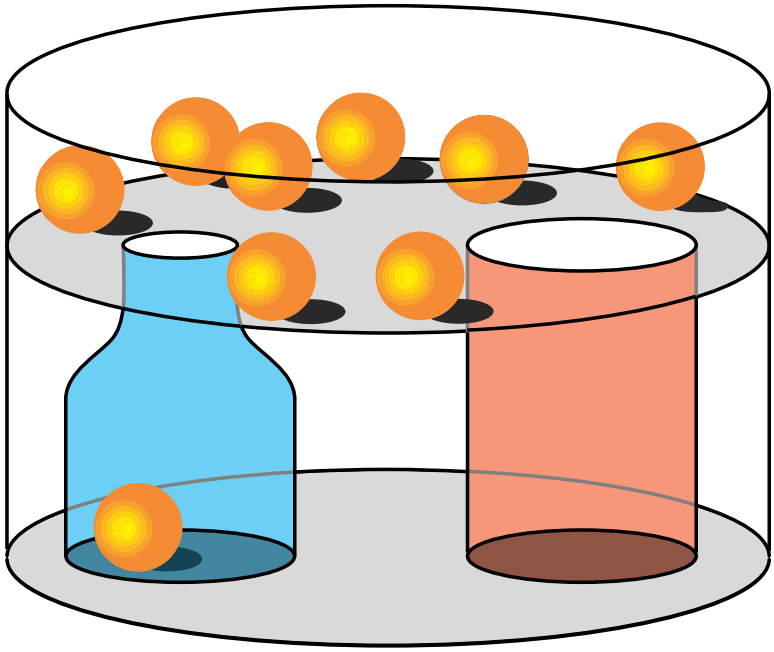

Figure 1.2: A child's game. By carefully shaking the game one tries to maximize the number of balls (α-particles) going in to the small hole (low energy) while minimizing the number of balls going in to the larger hole (high energy).

TFTR during 1994-1996 observed large losses of beam deuterons heated from 100 keV to 1-2 MeV by MCIBW in D^3He plasmas. At first blush, one may wonder what heated beam deuterons have to do with cooled α-particles and α-channelling. The straightforward answer is that these experiments give an incredibly rich database for validating at least the part of the COM simulation, which models the interaction of fast particles with MCIBW. We find that the COM simulation qualitatively agrees with the results of the experiments, namely, that countergoing passing beam particles are significantly heated and eventually cross the passing-trapped boundary and hit the wall. Note that the heating of beam particles is to be expected for the experimental setup on TFTR. The same simulation that predicts heating of beam deuterons, predicts cooling for the correctly phased combination of two waves in an advanced tokamak reactor.

We were able to go beyond the straightforward use of this data for validation of the code to infer two key pieces of MCIBW physics which are directly relevant for α-channelling:

The first is the reversal of the phase velocity of the MCIBW from the phase velocity of the launched fast wave at the antenna, i.e. the $k_{\|}$ flip. This effect was predicted by Valeo and Fisch (1994). It is very important to the implementation of α-channelling with the MCIBW, because it affects which α-particles can resonate with the MCIBW. It also affects the damping of the MCIBW on the background ions. This reversal can be inferred from the TFTR data by the observation of lost beam ions, which could only have resonated with a wave whose phase velocity had changed sign.

The second key inference using the simulation is the diffusion coefficient for the beam deuterons interacting with the MCIBW. This can be estimated by comparing the magnitude of the losses with the results from the simulations. We found that this diffusion coefficient significantly exceeds that which is predicted by geometrical-optics estimates. If we assume that we have the correct functional form for the diffusion coefficient, but that we are just underpredicting the electric field associated with the MCIBW, then the electric field in the experiment would need to be 5 to 9 times higher than what we are currently predicting. If this enhancement in fact occurs, it would also be beneficial to α-channelling in a reactor, since a larger electric field for the same power makes achieving the collisionless limit easier.

1.4 Outline of the Thesis

The outline of this thesis is as follows. In Chapter 2 we investigate the reactor implications of α-channelling, focusing on the attainment of the hot ion mode when significant fractions of the α-particle power can be channelled to the ions. In Chapter 3 we identify the basic elements of α-channelling in a simplified slab geometry. In Chapter 4 we formulate the theoretical problem of an energetic particle both colliding and interacting with various waves. We write the appropriate Langevin equations and explain the constants of motion simulation. Chapter 5 outlines two wave scenarios for cooling α-particles in an advanced tokamak reactor, including the 60% cooling scenario mentioned above. Chapter 6 gives an overview of the experiments in which fast ions were observed to interact with MCIBW on TFTR. We also give a theoretical picture of the losses as well as some simulation results. Chapter 7 shows how two key inferences can be made from the exper-

imental data. Chapter 8 summarizes the work presented here and suggests avenues for new research. Appendix A discusses the derivation of particle orbits in tokamaks from the constants of the motion, and describes the structure of the constants of motion space.

Chapter 2

Reactor Implications of α-Channelling

UNDERSTANDING THE IMPLICATIONS OF DIVERTING α-particle power to waves in a tokamak fusion reactor is important for appreciating the utility of α-channelling. Not only does such a study provide motivation for the investigation of α-channelling (if the potential benefits are not large then there is no need to go further), it also provides a context for evaluating implementation ideas. For example, important questions can be posed: should the implementation focus on current drive or the hot ion mode? What are the requirements for the fraction of wave power to be diverted in each case? In this chapter we attempt to answer these and other questions about implications of α-channelling on a future reactor.

2.1 Current Drive

A change in the current drive efficiency can strongly affect the economics of a tokamak power plant since, for fixed noninductive current, changing the efficiency changes the amount of radio-frequency power which must be launched into the tokamak. Thus, a change in efficiency changes both the recirculating power fraction ($\mathrm{P}_{aux}/\mathrm{P}_f \equiv 1/Q$) as well as the capital costs by reducing the number of antennae and generators required. It is an interesting historical footnote that the original motivation for the investigation

by Fisch and Rax (1992a) of the interaction of α-particles with lower hybrid waves was to address the concern [WONG and ONO 1984] that, in a reactor, α-particles might significantly damp the lower hybrid waves, thereby lowering the current drive efficiency.

As it turned out, Fisch and Rax (1992a) found it was possible, in fact, to use the spatial free energy in the α-particle distribution to amplify rather than damp the lower hybrid waves, and they suggested that the effective current drive efficiency could be enhanced by a factor 1.7 to 5, assuming 10 - 20% of the α-particle power could be diverted to waves and a Q of 20 in the absence of channelling.

Such a large increase in the current drive efficiency would be a notable achievement. However, as will be discussed later, in general it will be necessary to inject a significant amount of power in order to extract free energy from the α-particles before they slow down. This puts an upper limit on the increase in the effective current drive efficiency. Furthermore, external heating power will be required to start up the reactor. Thus, neither the capital costs of the RF system nor the recirculating power fraction can be taken to zero by the use of α-channelling. While there may indeed be a role for α-channelling to play in allowing a reactor to be designed with large current drive needs, using α-channelling to drive current is not likely, in itself, to make as large a change in an eventual fusion power plant as using α-channelling to obtain a hot ion mode.

2.2 Hot Ion Mode

It is important to note that present day experiments routinely operate and achieve their best discharges in hot ion modes, i.e. when $\mathrm{T}_i > \mathrm{T}_e$ (typical high power TFTR supershots had T_e = 10-12 keV, T_i = 30-40 keV). In fact, TFTR [MCGUIRE et al. 1995] and more recently JET [GIBSON and the JET Team 1997] have both achieved their highest fusion powers in the hot ion mode. There are three advantages of operating in the hot ion mode. First and foremost, since it is the ions that are fusing, it is desirable that they comprise as much of the plasma pressure as possible. Quantitatively, at a fixed confined pressure, $u_0 = n_i(\mathrm{T}_i + \mathrm{T}_e)$ (ignoring the hot α-particle pressure and impurities), the fusion power

can be written as

$$\mathrm{P}_f = 5\epsilon_\alpha f_D f_T {n_i}^2 \langle\sigma v\rangle(\mathrm{T}_i) = \frac{5\epsilon_\alpha {u_0}^2 \langle\sigma v\rangle(\mathrm{T}_i)}{4\mathrm{T}_i^2(1+\mathrm{T}_e/\mathrm{T}_i)^2}, \tag{2.1}$$

where ϵ_α is the α-particle birth energy, f_D and f_T are the fraction of ions which are deuterium and tritium respectively, and $\langle\sigma v\rangle$ is the reactivity of the DT reaction. If T_e could be made zero, the fusion power density would be four times higher than that in a T_e = T_i discharge at the same ion temperature and total pressure. Of course, it is impossible to maintain T_e at zero, since collisions will cause the ion and electron temperatures to equilibrate; however, current experiments, by preferentially heating ions with neutral beam heating, are able to maintain very disparate temperatures. Second, it has been observed both theoretically [DORLAND et al. 1994] and experimentally [SCOTT et al. 1994] that having $\mathrm{T}_i > \mathrm{T}_e$ can significantly reduce the turbulent heat transport, and better confinement leads to better discharges. Third, hot ion modes are achieved by neutral beam heating, which not only heats ions and increases reactivity, but also adds density to the center of the device, thereby peaking the fusion power (which goes like the density squared).

While hot ion modes are a very attractive operating point in present day experiments, deploying these modes in a future reactor is thought to be quite unlikely. In contrast to present day experiments which are dominated by injected power, a future tokamak power plant will have to get the majority of its heating from the 3.5 MeV α-particles, e.g. a reactor with Q of 25, will have α-particle power 5 times greater than the injected power. Below electron temperatures of 35 keV, when no impurities are present, the majority of the α-particle power goes directly to electrons; at 13.5 keV (i.e. the peak of $\langle\sigma v\rangle/\mathrm{T}_i^2$) only about a quarter of the α-particle power will flow into ions, leading to operation with $\mathrm{T}_e > \mathrm{T}_i$.

Clarke (1980) recognized that the hot ion mode might occur in a magnetically confined reactor, if the α-particle power could be diverted from the electrons into the ions and, at the same time, the ion energy confinement time exceeded the electron energy confinement time. The possibility that velocity space instabilities could provide the anomalous transfer of energy from the α-particles to the ions that Clarke (1980) was counting on has been examined [SUTTON et al. 1985; SIGMAR 1979; CHEN 1994],

but there is not enough free energy just in velocity space to achieve the hot ion mode. The advance of α-channelling is that, making use of diffusion in both velocity and space [FISCH and RAX 1992a], in principle 100% of the α-particle power could be diverted to waves. If those waves were to damp on ions, the hot ion mode might be retained in a reactor.

2.2.1 0-D power balance model

For the moment, let us assume that by means of α-channelling the α-particle power could be diverted to ions via waves and calculate the operating points which might be achievable in a future reactor both with and without α-channelling. For a more complete investigation the reader is referred to Fisch and Herrmann (1994) and Snyder, Herrmann, and Fisch (1994).

To calculate the operating point of a fusion plasma, a 0-D model of the power balance is used,

$$\frac{du_e}{dt} = \nu(\zeta u_i - u_e) + (1-\eta)\mathrm{P}_\alpha + (1-\eta_{Hi})\mathrm{P}_H - u_e/\tau_{Ee} - \mathrm{P}_{br} - \mathrm{P}_{sync} \tag{2.2}$$

$$\frac{du_i}{dt} = \nu(u_e - \zeta u_i) + \eta\mathrm{P}_\alpha + \eta_{Hi}\mathrm{P}_H - u_i/\tau_{Ei}, \tag{2.3}$$

where $u_e = 3n_e\mathrm{T}_e/2$ is the electron energy density, u_i the ion energy density, $\nu \propto n_e/\mathrm{T}_e{}^{3/2}$ is the energy equilibration rate, τ_{Ee} and τ_{Ei} are the electron and ion transport energy confinement times, P_α is the charged fusion product (CFP) power, η is the fraction of α-particle power deposited on ions which depends on T_e, P_H is the injected heating power, η_{Hi} is the fraction of P_H deposited on the ions, P_{br} is the power radiated away by bremsstrahlung, and P_{sync} is the power radiated away by synchrotron radiation [MCNALLY 1982], P_α is equal $\mathrm{P}_f/5$ and P_f is given in Eq. (2.1), finally $\zeta \equiv \sum_j n_j Z_j / \sum_j n_j$, the ratio of electron to ion densities, with Z_j defined as the ion charge state for the jth species. Note that by adding the two equations in steady state, we find the relation:

$$\mathrm{P}_\alpha + \mathrm{P}_H - \mathrm{P}_{br} - \mathrm{P}_{sync} = u_e/\tau_{Ee} + u_i/\tau_{Ei} \equiv (u_e + u_i)/\tau. \tag{2.4}$$

The total confined pressure is taken to be a constant,

$$u_0 = u_i + u_e + u_\alpha, \tag{2.5}$$

where u_α includes the pressure taken up by the hot charged fusion products. The amount of helium ash, n_α, is self-consistently calculated by taking the α-particle production rate (P_α/E_α) times τ_{He}, where τ_{He} is assumed to be 10 τ, and τ is the transport confinement time defined in Eq. (2.4).

We now have a complete set of equations which we can solve to determine a self-consistent operating point, including the effects of ash and impurities, hot α-particle pressure, and different confinement times for electron and ion heat. To solve these equations, we specify a T_i and T_e, and then solve self-consistently for the fusion power in steady state, the density, and the confinement times. A solution is obtained using an initial guess for the helium ash, the amount of helium is then altered so as to be consistent with the α-particle power and τ_{He}, and an operating point is again solved for. By iterating in this way, a solution with self-consistent levels of helium ash is found.

Note that no scaling law is assumed for the confinement times, nor is any assumption made about the ratio of τ_{Ei} to τ_{Ee}; rather, specifying T_i, T_e, and the total pressure allows us to solve Eq. (2.3) in steady state for the values of τ_{Ei} and τ_{Ee}. Also, note that steady state at fixed heating power will not be attainable at arbitrary values of T_i and T_e, even with no constraints on the confinement times. Finally, neglecting radiation losses for the moment, if for total pressure u_0 a solution of Eq. (2.3) with parameters T_i, T_e, n_i, n_e, P_f, τ_{Ei}, τ_{Ee} exists then for pressure λu_0 there exists a solution with parameters T_i, T_e, λn_i, λn_e, $\lambda^2 P_f$, τ_{Ei}/λ, τ_{Ee}/λ, where λ is an arbitrary constant.

2.2.2 Operating point with no diversion

In this section we introduce a variety of plots and examples loosely based on ARIES-RS [NAJMABADI et al. 1997], a reverse shear tokamak reactor designed as a 1000 MW (electric) power plant. In Table 2.1 we list the nominal operating parameters of ARIES-RS [The ARIES Team 1996].

R_0(m)	5.52
a (m)	1.38
A	4.00
κ (95%)	1.70
B (T)	7.98
I (MA)	11.3
β (%)	4.87
n_i (10^{14}/cm^3)	1.72
n_e (10^{14}/cm^3)	2.11
$n_{D/T}/n_e$	0.33
n_α/n_e	0.15
n_{imp}/n_e	0.0045
T_i (keV)	18.0
T_e (keV)	18.7
f_{rad}	0.18
τ (sec)	1.39
P_f (W/cm^3)	6.22
P_{aux}/P_α	0.19
f_{bs}	0.88

Table 2.1: Parameters for the ARIES-RS tokamak.

Using the stored pressure, impurity concentrations, and operating temperature from Table 2.1 as input parameters, the output of the 0-D power balance model is calculated, the result is shown in Table 2.2. Our 0-D power balance model reproduces the ARIES-RS operating point within about 10% of all major parameters.

An alternative operating point is shown in Table 2.2(b). This operating point is at significantly lower temperature than the nominal ARIES-RS operating point and at the same Q, but has significantly higher fusion power (almost 50% higher). It is illustrative to examine the differences between these two scenarios, since when we consider the affect of α-channelling many similar issues will arise. Of course, the ARIES-RS operating point, Table 2.1, was chosen by considering many constraints and effects, such as minimization of recirculating power (by maximizing current drive efficiency), neutron wall loading, and many others which have not been included in our simplified analysis, so the following is just a didactic exercise, rather than an argument for designing ARIES-RS to operate at a lower temperature.

u_0 (10^{14}keV/cm^3)	124.6
u_i/u_0	0.40
u_e/u_0	0.52
$u_{\alpha H}/u_0$	0.08
T_i(keV)	18.0
T_e(keV)	18.7
η	0.30
η_{Hi}	0.10
ν(sec^{-1})	2.14
τ_i(sec)	1.57
τ_e(sec)	1.47
τ(sec)	1.52
n_i(10^{14}/cm^3)	1.86
n_e(10^{14}/cm^3)	2.28
$n_{D/T}/n_e$	0.33
n_α/n_e	0.14
n_{imp}/n_e	0.0045
P_f(W/cm^3)	6.05
P_{aux}/P_α	0.2
f_{rad}	0.17

(a) Operating Point Based on ARIES-RS

u_0(10^{14}keV/cm^3)	124.6
u_i/u_0	0.44
u_e/u_0	0.52
$u_{\alpha H}/u_0$	0.04
T_i(keV)	11.4
T_e(keV)	11.6
η	0.21
η_{Hi}	0.10
ν(sec^{-1})	7.02
τ_i(sec)	1.55
τ_e(sec)	0.94
τ(sec)	1.15
n_i(10^{14}/cm^3)	3.20
n_e(10^{14}/cm^3)	3.70
$n_{D/T}/n_e$	0.38
n_α/n_e	0.096
n_{imp}/n_e	0.0045
P_f(W/cm^3)	8.84
P_{aux}/P_α	0.2
f_{rad}	0.20

(b) ARIES-RS at Lower Temperature

Table 2.2: Operating point based on the ARIES-RS design. Operating point based on the ARIES-RS design, except for T_i,$T_e \approx 11.5$ keV.

Note that, Table 2.2(b) has fusion power 46% higher than Table 2.2(a) at the same confined pressure. This can be attributed to three main reasons. First, the reactivity per unit pressure (Eq. (2.1)) is about 5% higher at 11.4 keV than at 18.0 keV (assuming Maxwellian distributions with $T_D = T_T = T_i$ then $\langle\sigma v\rangle(T_i)/T_i^2$ peaks at about 13.5 keV and remains above half of its peak value between 5 and 35 keV). Second, operation at the lower T_e means that the α-particles will slow down more quickly (the slowing down time is proportional to $T_e^{3/2}$), and, therefore, take up less pressure. Hot α-particles take up one half as much pressure in (b) as in (a). This results in a 10% increase in the fusion power. Third, as mentioned above, the α-particle ash is proportional to 10 $P_\alpha\ \tau$. We also know that $P_\alpha\ \tau \approx u_0$, from Eq. (2.4), implying the amount of α-particle ash is about the same without regard to operating point. However, what matters for the dilution

of the fuel is the ratio of n_α to n_i. Since we are at lower temperatures, n_i and n_e are higher and therefore the dilutive effect of the α-particle ash is significantly reduced at lower temperatures. By this argument the ratio of α-particle ash divided by ion density between the two scenario's should be about the inverse of their temperatures and this is verified in Table 2.2. This effect gives the final 25% increase in the fusion power. Of course, this final effect depends on an assumption, that τ_{He} is a fixed multiple of τ independent of other variables like density and temperature, which is unlikely to be valid in a reactor. These three effects are multiplicative, yielding a 46% gain in P_f overall.

While lowering the electron temperature leads to an increase in the fusion power density, the lower temperature makes it more difficult to obtain a hot ion mode. For instance, it is now much harder to maintain a difference between the ion and electron temperatures, since ν has more than tripled from case (a) to case (b). Since the fusion products slow down more quickly on electrons, less power is going to the ions and η has gone from 0.3 to 0.21. Also, a very serious concern for steady state reactors is the maintenance of the current. Operating point (b) would have the same bootstrap current as the design point, within the context of a 0-D simulation, since the electron pressure is the same in both cases. However, at lower electron temperature and higher density the amount of power necessary for the RF driven currents goes up substantially. Assuming the waves were operated at the same phase velocity relative to the thermal velocity, the RF power requirements go up by a factor of 2.6. This would reduce the nominal Q of the reactor to about 14. If, on the other hand, the wave frequency were consider fixed, the ratio of phase velocity to thermal velocity would increase, increasing the efficiency (but possibly lowering the per pass absorption of the RF power) and moving Q back towards its original value of 25.

In Table 2.2 we compared two different solutions of Eq. (2.3) for an ARIES-RS-like reactor. Actually there exists a whole space of operating points for the reactor at a given total pressure. Figure 2.1 shows one representation of this space, with contour plots of three quantities, τ_{Ee}, τ_{Ei}, and P_f versus T_i and T_e in the range of 8 keV to 25 keV. The shaded regions indicate where steady state operation is impossible, i.e. it is impossible to sustain the temperature differences, or radiation losses are greater than the amount of energy going into the electron channel from the α-particles and the auxiliary power. The operating points in Table 2.2 are marked with an "X". The dark line represents the values

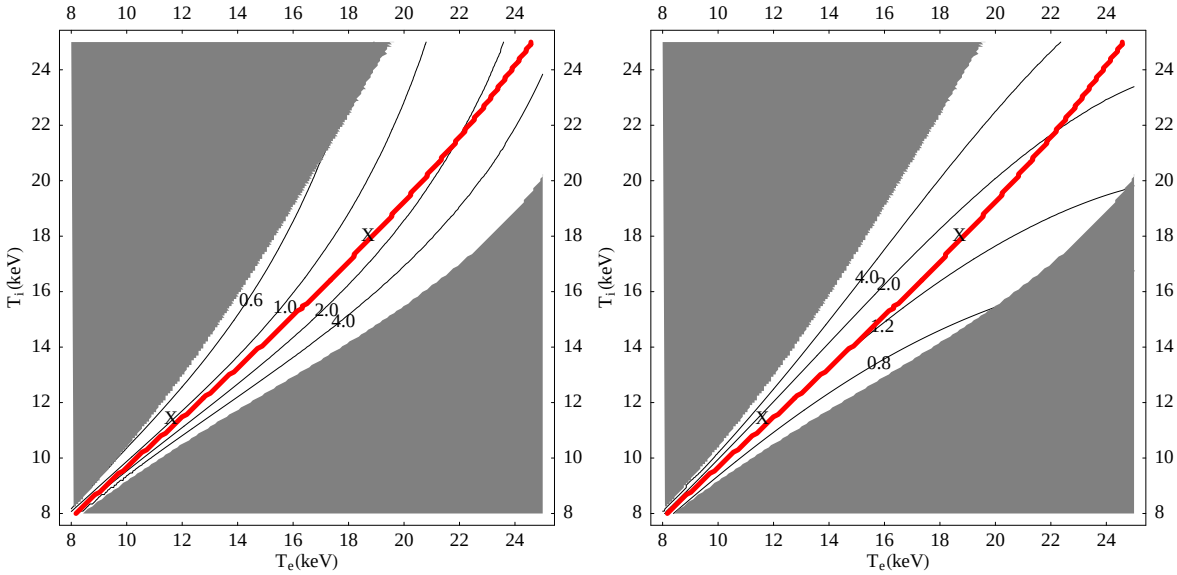

(a) τ_{Ee} (s)

(b) τ_{Ei} (s)

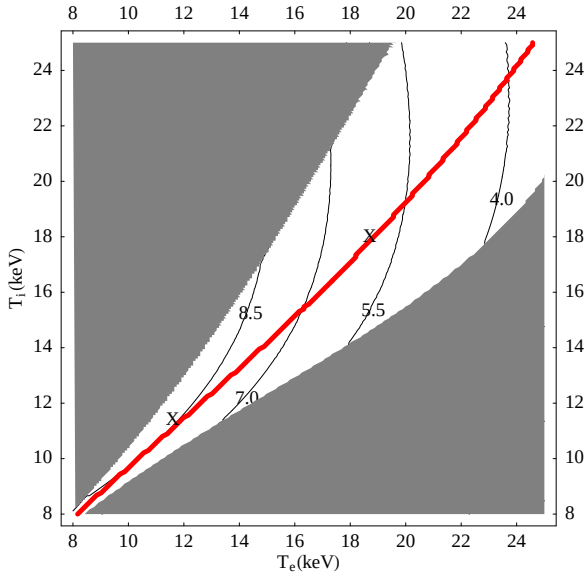

(c) P_f (W/cm^3)

Figure 2.1: Contours of electron transport confinement time (a), ion confinement time(b), and fusion power density(c) for Aries-RS like reactor with no diversion of α-particle power versus T_e and T_i. The grey area represents values of the T_e and T_i where a steady state solution to Eq. (2.3) does not exist for the given value of P_{aux}/P_α. The solid curve is where $\tau_{Ei} = \tau_{Ee}$. The operating points in Table 2.2 are each marked by an "X".

of T_i and T_e where $\tau_{Ei} = \tau_{Ee}$.

Note that everything in the upper diagonal has $T_i > T_e$, and conversely for the lower diagonal. As expected, as T_e increases, the range of T_i increases, since the equilibration rate depends strongly on T_e. Also note that the highest fusion power density is achieved in the hot ion mode, although large differences between T_i and T_e cannot be maintained if the ratio of τ_{Ei} to τ_{Ee} is close to one. Although the "conventional wisdom" that the more confinement the better is valid for current experiments, when operating at a fixed pressure in steady state we see that while the direction of increasing fusion power coincides with the direction of increased ion confinement, it also coincides with decreasing electron confinement. One can see that, in some cases, it might be desirable to purposefully spoil the electron confinement to achieve higher fusion power, if it could be done without spoiling the ion heat confinement. One way this might be accomplished is discussed below.

As a final note in this section on operating points without diversion of α-particle power, it is worth mentioning that occasionally operation of a reactor in a hot ion mode is proposed. This would be achieved by means of large electron temperatures, and an electron heat confinement time (including both transport and radiation losses) significantly shorter than the ion heat confinement time. From Fig. 2.1 we see that such regimes do exist. Unfortunately, these operating points do not share the high fusion power density per unit pressure which is a hallmark of the hot ion mode and one of its most attractive features. In order to achieve these high T_e, hot ion modes one must typically be far from the fusion power maximum in T_e, T_i space due to the ion temperature being far from the peak reactivity per unit pressure and the large pressure taken up by the slowing down distribution of the fast particles. As we will see in the next section, when α-particle power is diverted to the ions, one can operate in the hot ion mode near the maximum of the fusion power density.

2.2.3 Operating point with diversion

In the presence of α-channelling a few modifications to the standard picture are introduced. The wave diverts a fraction, η_w, of the α-particle power directly to the thermal

ions which, for the purposes of this discussion, is taken as a free parameter. Although there is some gain if this energy goes to suprathermal ions [FISCH and RAX 1992b; FISCH and HERRMANN 1994; SNYDER, HERRMANN, and FISCH 1994], it is not a large effect and so, for the purposes of this discussion, we neglect it. Since the wave must damp the α-particles' energy in a time short compared with their slowing down in order to achieve α-channelling, the portion of plasma pressure taken up by the fast α-particles is decreased by the fraction $1-\eta_w$. Once the α-particles get to $1-\eta_{wave}$ of their birth energy, we can make two different assumptions, the first, modeling *in situ* cooling, is that the α-particles are then assumed to slow down naturally. On the other hand, the most promising methods for α-channelling suggest that the way to tap the free energy is by extracting the α-particles from the tokamak while cooling them. If the α-particles leave while still significantly above the thermal energy, we can entertain the possibility that these α-particles might be made to implant themselves somewhere and remain outside the plasma. Under this scenario, no α-particles would be left for ash. (Of course we would have to deal with the new problem of energetic α-particles impinging on the first wall.) The case with no α-channelling corresponds to η_{wave}=0.

In Table 2.3 we give two nominal operating points of ARIES-RS with 75% of the α-particle power diverted to waves which then heat ions. In both cases the cooled α-particles are assumed to slow down naturally. In these cases almost 90% of α-particle power eventually goes to ions, since, as the α-particles slow down, the fraction of their energy going to ions increases. By the time the α-particles have an energy of $\approx$ 40 T_e, the amount of α-particle power going to ions is equal to that going to electrons. Note that the fusion power density of Table 2.3(a) is more than twice that of Table 2.2(a), and the thermal ion pressure has gone up by 30%. While the ion confinement times are about the same between the case with no diversion and the case with diversion, the electron confinement time is significantly shorter ($\tau_{Ee} < 1/3\ \tau_{Ei}$) when there is diversion. With such a large equilibration rate, a short electron confinement time is necessary in order to maintain T_i more than 30% larger than T_e.

Table 2.3(b) shows an operating point with 75% diversion but with T_i = 1.5 T_e. This operating point is at a higher electron temperature with about the same ion and electron confinement time as Table 2.3(a). Note that the fusion power density is about 10% lower than (a), but still almost 90% higher than the case with no diversion. While case

u_0 (10^{14}keV/cm^3)	124.6
u_i/u_0	0.52
u_e/u_0	0.47
$u_{\alpha H}/u_0$	0.01
T_i(keV)	14.4
T_e(keV)	11.0
η	0.88
η_{Hi}	0.80
η_{wave}	0.75
ν(sec^{-1})	7.22
τ_i(sec)	1.58
τ_e(sec)	0.46
τ(sec)	0.74
n_i(10^{14}/cm^3)	3.05
n_e(10^{14}/cm^3)	3.53
$n_{D/T}/n_e$	0.38
n_α/n_e	0.10
n_{imp}/n_e	0.0045
P_f(W/cm^3)	12.7
P_{aux}/P_α	0.20
f_{rad}	0.12

(a) Operating Point of ARIES-RS with 75% Diversion

u_0(10^{14}keV/cm^3)	124.6
u_i/u_0	0.55
u_e/u_0	0.44
$u_{\alpha H}/u_0$	0.01
T_i(keV)	20.3
T_e(keV)	13.5
η	0.89
η_{Hi}	0.80
η_{wave}	0.75
ν(sec^{-1})	4.09
τ_i(sec)	1.70
τ_e(sec)	0.48
τ(sec)	0.80
n_i(10^{14}/cm^3)	2.26
n_e(10^{14}/cm^3)	2.69
$n_{D/T}/n_e$	0.36
n_α/n_e	0.12
n_{imp}/n_e	0.0045
P_f(W/cm^3)	11.4
P_{aux}/P_α	0.20
f_{rad}	0.10

(b) ARIES-RS at with 75% diversion at higher temperature

Table 2.3: Operating point based on the ARIES-RS design but diverting 75 % of the α-particle power to the fuel ions (a). Operating point like(a), but with a larger difference between T_i and T_e (b).

(b) has more pressure in the ions than case (a), the ion temperature is now significantly off of the fusion power maximum. The transport confinement times in both cases are significantly lower than the case with no diversion. This is a necessary consequence of higher fusion power and Eq. (2.4).

In Fig. 2.2 the operating space with 75% diversion is shown. With this large a fraction of the α-particle power going to the ions, steady state can only be reached in the hot ion mode. The fusion power maximum has now moved up to higher temperatures, in part due to the ease with which significant ion and electron temperature differences can

be maintained there.

Finally, while it may be impossible to achieve, it is at least interesting to consider the maximal extension of α-channelling. That is, what would the reactor consequences be if we could divert all of the α-particle power to the ions, and in doing so remove all of the α-particles from the plasma in such a way that they would not return. In Fig. 2.3 and Table 2.4 we show the operating regime for this "maximal" extension. The fusion power density at $T_i \approx 20$ keV, $T_e \approx 10$ keV is almost 4 times the ARIES-RS base case and more than 2.5 times the maximum fusion power for ARIES-RS without diversion.

Note that the cases investigated here have focused on diverting large fractions (75 - 100%) of the α-particle power to waves. Snyder, Herrmann, and Fisch (1994) found that the increase in fusion power density at the optimal electron confinement time (with fixed ion confinement time) is approximately linear in the amount of power diverted.

2.2.4 Confinement times

From the curves in Figures 2.2 and 2.3 the fusion power density is typically maximized when the electron confinement time is as small as possible, consistent with steady state. Surprisingly, when τ_{Ee} is too large, operation could be significantly enhanced by lowering τ_{Ee} while keeping τ_{Ei} fixed. This may seem problematic as, historically, not much has been known about altering the relative ratio of ion and electron confinement times. However recent experiments [LEVINTON et al. 1995] have discovered regimes, called enhanced reverse shear (ERS), where the ion thermal diffusivity is significantly lower than the electron thermal diffusivity within the confinement barrier. The ion heat confinement time is much longer than the electron heat confinement time inside this barrier. Actually, the ERS regime, if attained in a reactor, would be quite compatible with α-channelling. Since it satisfies $\tau_{Ei} \gg \tau_{Ee}$, operation near the fusion power maximum should be possible. Also, since in ERS operation the thermal ash will be very well but inconveniently confined, α-channelling could provide a natural way for these α-particles to be pulled out past the confinement barrier and discharged, avoiding the poisoning which might eventually choke off the plasma.

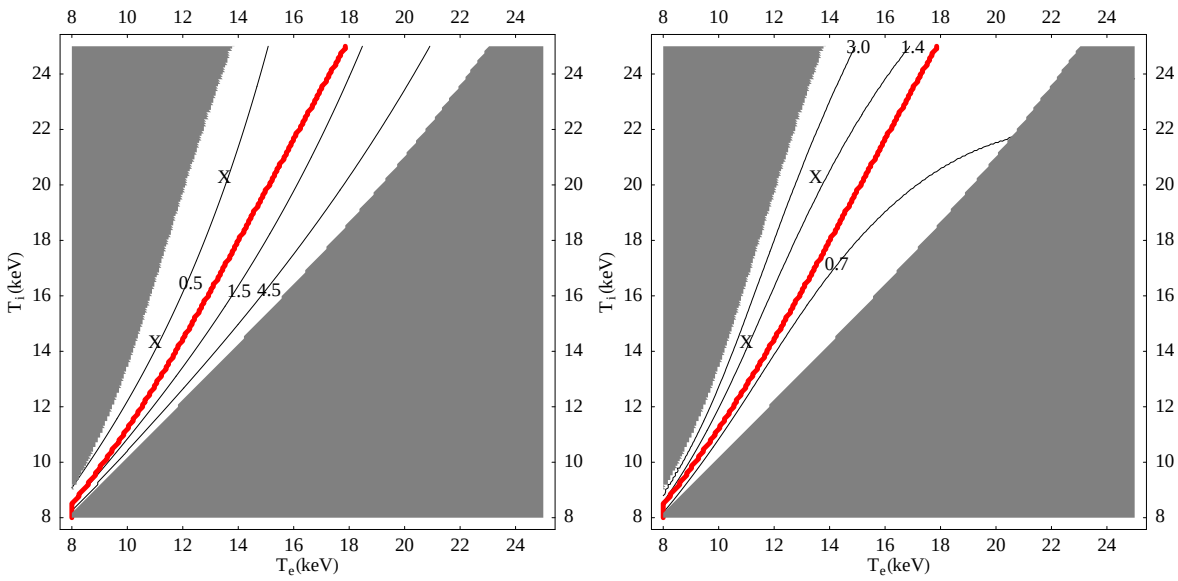

(a) τ_{Ee} (s)

(b) τ_{Ei} (s)

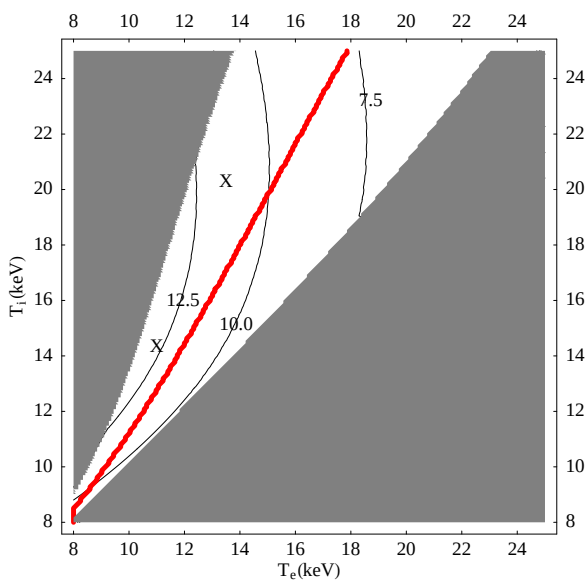

(c) P_f (W/cm^3)

Figure 2.2: Contours of fusion power density(a), electron transport confinement time(b), and ion confinement time(c) for Aries-RS like reactor with diversion of 75 % of the α-particle power to the fuel ions plotted versus T_e and T_i. The operating points of Table 2.3 are marked by "X".

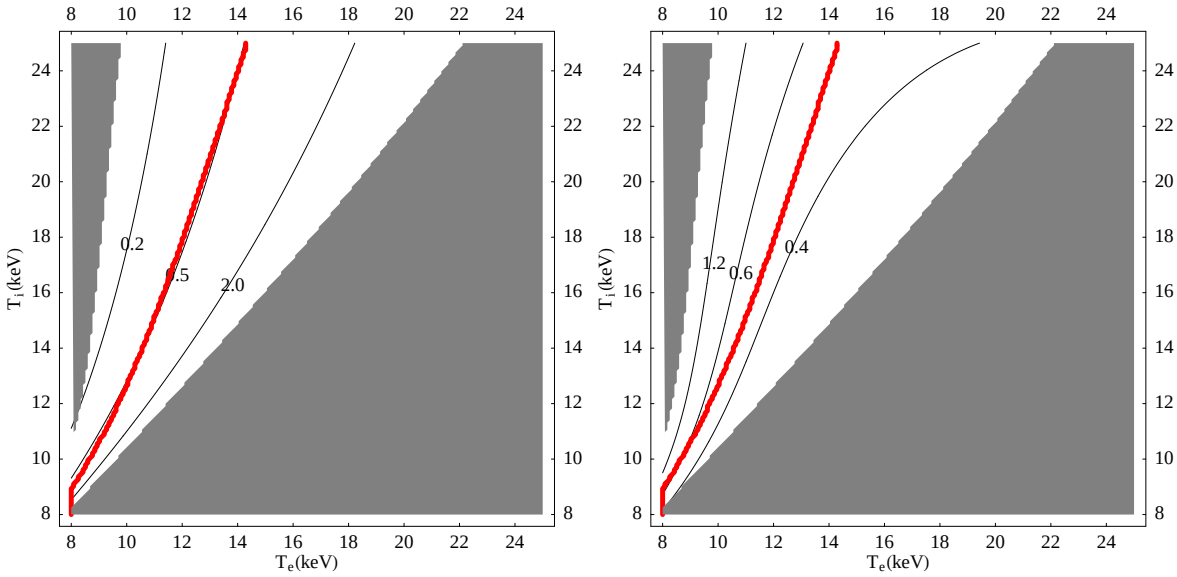

(a) τ_{Ee} (s)

(b) τ_{Ei} (s)

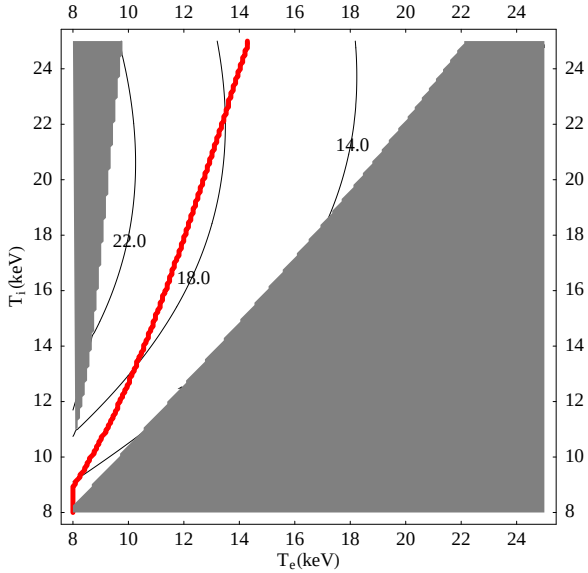

(c) P_f (W/cm^3)

Figure 2.3: Contours of fusion power density(a), electron transport confinement time(b), and ion confinement time(c) for Aries-RS like reactor with diversion of 100 % of the α-particle power to the fuel ions and zero helium ash plotted versus T_e and T_i.

u_0 (10^{14}keV/cm^3)	124.6
u_i/u_0	0.66
u_e/u_0	0.34
$u_{\alpha H}/u_0$	0.00
T_i(keV)	20.0
T_e(keV)	10.0
η	1.00
η_{Hi}	1.00
η_{wave}	1.00
ν(sec^{-1})	6.48
τ_i(sec)	1.44
τ_e(sec)	0.16
τ(sec)	0.39
n_i(10^{14}/cm^3)	2.73
n_e(10^{14}/cm^3)	2.85
$n_{D/T}/n_e$	0.48
n_α/n_e	0.0
n_{imp}/n_e	0.0045
P_f(W/cm^3)	22.3
P_{aux}/P_α	0.20
f_{rad}	0.04
$Z_{\rm eff}$	1.41

(a) Operating Point of ARIES-RS with 100% Diversion

u_0(10^{14}keV/cm^3)	124.6
u_i/u_0	0.65
u_e/u_0	0.35
$u_{\alpha H}/u_0$	0.0
T_i(keV)	20.0
T_e(keV)	10.0
η	1.00
η_{Hi}	1.00
η_{wave}	1.00
ν(sec^{-1})	6.51
τ_i(sec)	1.60
τ_e(sec)	1.55
τ(sec)	1.58
n_i(10^{14}/cm^3)	2.72
n_e(10^{14}/cm^3)	2.87
$n_{D/T}/n_e$	0.47
n_α/n_e	0.0
n_{imp}/n_e	0.0045
n_{Bi}/n_e	2.6×10^{-4}
P_f(W/cm^3)	22.14
P_{aux}/P_α	0.20
f_{rad}	0.76
$Z_{\rm eff}$	2.19

(b) ARIES-RS at with 100% diversion and line radiation

Table 2.4: Operating point based on the ARIES-RS design but diverting 100% of the α-particle power to the fuel ions and all α-particle ash removed. (a). Operating point like(a), but τ_{Ei} = τ_{Ee} because of the injection of bismuth (Z=83) impurities.

In the event that an ERS-like regime is not attained in a reactor, the requisite lowering of τ_{Ee} might be achieved without lowering τ_{Ei}, e.g., by the injection of high-Z impurities whose presence will not disturb the plasma too much, but whose line radiation will significantly cool the electrons. For example, consider Table 2.4 which shows two operating points for the case of 100% α-channelling with ash removal. Case 1 has $\tau_{Ei} \gg \tau_{Ee}$, while Case 2 has $\tau_{Ei} \approx \tau_{Ee}$, but has added line radiation losses (which are calculated using the tables of Post et al. (1977)). Here, bismuth (any high-Z material could have

been substituted for the purposes of this analysis) is injected into the plasma at the level of $n_{Bi}/n_e = 2.6 \times 10^{-4}$. Note the ratio of ion to electron density has hardly changed between the two plasmas, but f_{rad} has gone from 10% to 76%.

2.2.5 Advanced fuels

Operation in the hot ion mode is also advantageous for fuels other than DT [SCHULTZ, BROMBERG, and COHN 1980]. The α-channelling mechanism might also be applied to attain hot ion modes using fuels other than DT. The advantage of these fuels over DT is their significantly smaller neutron output, however, they also have lower power density and higher operating temperatures.

In considering an ARIES-III-like tokamak, a reactor design [BATHKE et al. 1992] that burns deuterium helium-3 (D^3He), Fisch and Herrmann (1994) and Snyder, Herrmann, and Fisch (1994) showed that the gain in P_f is again about a factor of two over the case with no diversion. This might be somewhat unexpected since there is more power to divert here (almost 100% comes out in charged fusion products), and a high electron temperature favors substantial temperature differences. However the high densities required for sufficient fusion power density and the necessity of two electrons for every ^3He limit the maximum temperature difference and the amount of pressure that can be taken up by ions.

2.3 Other Consequences of α-Channelling

The previous section discussed how α-channelling facilitates operation in a hot ion mode thereby increasing the fusion power density of a reactor. Other benefits do accrue from diverting the energy of the α-particles to waves.

If the free energy of the α-particles is used to amplify waves which damp preferentially on electrons traveling in one direction or the other, then we could drive larger currents with the same power thereby increasing the current drive efficiency. We already

addressed the benefits which arise from having the ions heated by the α-particle power via waves and thereby attaining the hot ion mode. Note that it may also be possible to drive currents in the case where ions are heated, if they are are also heated preferentially with respect to their direction, using a minority ion species current drive technique [FISCH 1987]. These schemes do suffer from lower efficiency then electron current drive techniques, so the diversion of considerable power is important.

The mechanism, elucidated in Chapter 1, for tapping the free energy of the α-particles depends on ejecting the α-particles from the plasma after significantly cooling them. Note that schemes which just eject fast α-particles from the plasma do not significantly reduce the ash levels, because most cold He ions recycle back into the plasma as thermal particles about 10 times [SYNAKOWSKI et al. 1995] before being removed from the system. Thus, at most, these schemes would reduce the residence time of the He ions by 10%. If, however, the α-particles are ejected while they are still energetic, and then implant into a structure or are otherwise prevented from returning, they will not contribute to the ash at all, thereby reducing the dilution of the fuel and allowing more pressure for the reacting ions.

If the short electron times, which are desirable to achieve the maximum difference in electron and ion temperatures, discussed above were achieved through introduction of impurities, the fraction of power radiated might be significantly increased. For example, in Table 2.4 the radiated power fraction goes from 10% to 76%. This effectively reduces the heat load to the divertor by about a factor of four. Note that recent experiments suggest that large radiated power fractions can be compatible with good plasma performance [HILL et al. 1998]. Of course, the table above is a case with maximal α-channelling, which is unlikely to be achieved. Also, profile effects, which have been entirely neglected so far, would be particularly important when it comes to the introduction of impurities.

If a significant fraction of the α-particle power is diverted to waves, then there will be a large amount of RF power present in the device. For instance, if 75% of the power is diverted, and Q is 25, then 19% of P_f is present as wave power in the plasma. This wave power might be put to many uses, for example the formation of transport barriers [ONO et al. 1994] might be contemplated, affecting both the density and temperature profiles.

We have not addressed at all the issue of where the wave heat is deposited, and how that will differ from where the α-particles are born. We also have not addressed the shape of the density or temperature profiles. While these will in large part be determined by transport processes, it is possible to think of ways in which the profiles are affected by α-channelling. For instance, α-channelling extracts energy from the α-particles by bringing them out and cooling them, at the same time, the waves will be damping on ions, heating them and bringing them to the center of the device. It may be possible to use such fluxes advantageously in a reactor.

Now we enumerate some of the costs. The 0-D operating points which maximize the impact of α-channelling have low electron temperatures. Thus these plasmas have low current drive efficiency. They also have large ion pressures, and therefore correspondingly smaller electron pressure. This leads to a reduction in the bootstrap current. Both of these effects tend to reduce the driven current. Of course, there is a large amount of RF power present which might mitigate this current shortfall, but it is not clear that a self-consistent equilibrium can be maintained without going to lower toroidal current, which would give lower β and thus lower fusion power. Thus, it is important to demonstrate that self-consistent equilibria, i.e. those where the channelled α-particle power is sufficient to drive the required current, can be maintained at the lower electron temperatures which are favorable for maximizing the power increase due to the channelling effect.

Other costs will be associated with α-channelling. For instance it will be necessary to inject power into the plasma, possibly lowering the Q. It may be necessary to use several different wave systems, or have a large amount of RF power for the startup of the channelling effect, both of which will affect the capital cost. Further, α-channelling may induce significant, localized losses of energetic α-particles to the first wall of the reactor. In all likelihood the damage caused by such losses would far outweigh any benefits of α-channelling, unless a clever way of dealing with this problem can be found.

2.4 Summary

If α-channelling were to work (and the problems mentioned above could be avoided or ameliorated), it might be possible to double the fusion power output of a reactor at the same pressure. However, just because the fusion power output could be doubled does not mean that doing so would be the most economical way to take advantage of this innovation. In particular, both the first wall heat load and neutron wall load constraints may prevent one from taking advantage of this doubling of the fusion power density. Instead, one may contemplate just lowering the current or the field in the device. Emmert et al. (1994) have done a systems analysis of a reactor using a highly idealized model of α-channelling. They found that by lowering the toroidal field, a reduction in the cost of electricity on the order of 15% could be achieved. Alternatively, one could consider a smaller device. Note that both of these operating points would have significantly higher current drive requirements than the ARIES-RS base case, however, these plasmas have several hundred megawatts of wave power coursing through them. With so much RF power it is also possible to contemplate other beneficial uses, for instance, the formation of transport barriers, or profile control.

To achieve a hot ion mode with α-channelling requires comprehensive control of the α-particles in the plasma, i.e. the goal is to take all the energy out of all of the α-particles to achieve a step change in the eventual reactor. This may not be possible. However, it may be possible to control a suitable portion of the α-particle phase space in order to accomplish some more limited objective like pressure profile control, seed current generation, ash removal, and the like, all of which would be valuable additions to an advanced tokamak. While these have not been specifically addressed in this thesis, many of the tools developed here would be equally applicable to these problems.

Finally, while much of this has been written with the tokamak specifically in mind, most of the issues raised here are generic. Thus, the benefits of α-channelling are likely to transfer to other magnetically confined fusion plasmas, such as stellarators and reverse field pinches, provided, of course, a method can be found for diverting the free energy of the α-particles to waves. In that respect, the insights developed here for how combinations of waves can be applied in concert to produce the channelling effect should

be valuable in developing similar schemes on related toroidal confinement devices.

Chapter 3

Building Blocks of the Channelling Effect

EXTRACTING THE FREE ENERGY of the α-particles with waves in a tokamak is in principle possible but in practice challenging. Significant insight may be had by examining α-channelling in a simpler phase space. This chapter lays out the issues involved in extracting the α-particles energy (see Fisch, Fruchtman, Karney, Herrmann, and Valeo (1995), Fisch and Herrmann (1995), and Fisch (1997)).

3.1 Free Energy

In Fig. 3.1, the distribution function of the α-particles in a device which is hot in the center and cold at the edge is plotted versus energy and space. Although the α-particles are born at 3.5 MeV, eventually they slow down on the background plasma and develop a distribution function which is monotonically decreasing in energy. If the waves which were used only diffused particles in velocity space, no free energy could be extracted from this distribution function. However, waves which diffuse particles in energy and space can tap the population inversions which exist along their diffusion paths, diverting

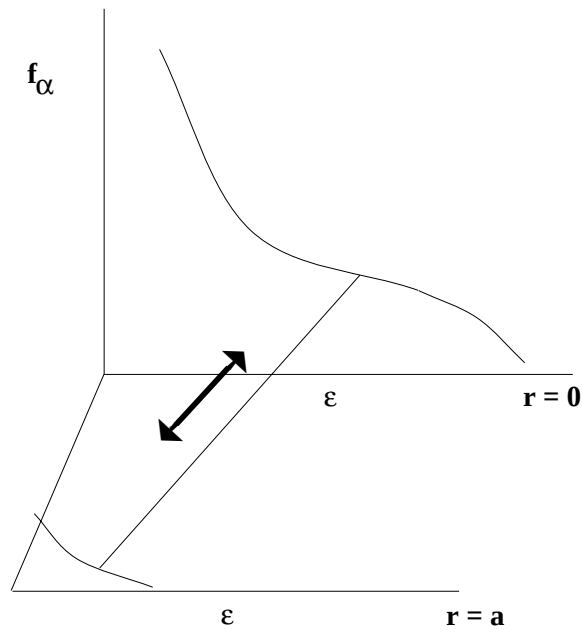

Figure 3.1: Schematic of a diffusion path in velocity and real space which would, on average extract energy from the α-particle distribution even though it is monotonically decreasing in energy at any fixed radius.

energy from the particles to the waves. Normally particles diffuse until the distribution function is flattened along this path (neglecting background damping which would tend to leave the distribution with a small but finite slope). If, however the diffusion path is in contact with a sink of particles, e.g., the wall, then the steady state solution will be that the distribution function is zero along the path, i.e., all of the particles will hit the wall, and in doing so each irrevocably gives up its energy to the wave. In principle, it is possible to extract all of the energy of the α-particle distribution. Our goal is to construct the diffusion path, or paths, which maximize the amount of energy extracted from the energetic α-particles.

3.2 Wave-Particle Interactions

To begin, it is necessary to understand how the diffusion path in energy and space is constructed. Typically we think of a particle resonating with a wave and either gaining

or losing energy, but particles also undergo a spatial displacement upon interacting with the wave. Consider the equations of motion for a particle in a uniform magnetic field interacting with an electrostatic wave which has potential $\Phi = \Phi(\mathbf{k} \cdot \mathbf{x} - \omega t)$,

$$m\frac{d\mathbf{v}}{dt} = -q\boldsymbol{\nabla}\Phi + q\frac{\mathbf{v}}{c} \times \mathbf{B} \tag{3.1}$$

We are interested in the change of the guiding center position, $\mathbf{r}_{g.c.\perp}$, where

$$\mathbf{r}_{g.c.\perp} = \frac{\mathbf{B} \times \mathbf{r} \times \mathbf{B}}{B^2} - \frac{\mathbf{B} \times \mathbf{v}}{\Omega B} \tag{3.2}$$

In the absence of the wave the time derivative of $\mathbf{r}_{g.c.\perp}$ is zero. The equations of motion can be rewritten as

$$\frac{d}{dt}(m\mathbf{v} - \frac{q\mathbf{r} \times \mathbf{B}}{c}) = -q\mathbf{k}\Phi' \tag{3.3}$$

where we have taken advantage of the form of Φ. Taking $c\mathbf{B}/(qB^2)$ crossed with this we find

$$\frac{d\mathbf{r}_{g.c.\perp}}{dt} = -\frac{c\mathbf{k} \times \mathbf{B}}{qB^2}q\Phi' \tag{3.4}$$

Using the relation $d\varepsilon/dt = q\mathbf{v} \cdot \mathbf{E}$ write

$$\frac{d(\varepsilon + q\Phi)}{dt} = -q\mathbf{v} \cdot \boldsymbol{\nabla}\Phi + \frac{dq\Phi}{dt} = q\frac{\partial\Phi}{\partial t} = -q\omega\Phi' \tag{3.5}$$

Substituting for Φ' from above we find

$$\frac{d\mathbf{r}_{g.c.\perp}}{dt} = \frac{c\mathbf{k} \times \mathbf{B}}{q\omega B^2}\frac{d(\varepsilon + q\Phi)}{dt} \tag{3.6}$$

Integrating this equation through the interaction, and evaluating where Φ is zero we see that

$$\frac{\Delta\mathbf{r}_{g.c.\perp}}{\Delta\varepsilon} = \frac{\mathbf{k} \times \mathbf{B}}{m\omega\Omega B} \tag{3.7}$$

Thus, when the particle receives a kick in energy, it will receive a correlated kick in position, creating a diffusion path in phase space. Note that the kick is proportional to

the magnitude of k, and inversely proportionally to the frequency of the wave.

Consider a slab geometry, with $\mathbf{B}$ in the z direction, the wave vector $\mathbf{k}$ in the y direction and a wall at x = a. The phase space of this configuration is shown in Fig 3.2. The α-particles are born near the center (x = 0) with 3.5 MeV of energy, as depicted by the contour curves. In the absence of waves, eventually the α-particles will develop a slowing down distribution due to collisions, and this distribution will be monotonically decreasing in energy as shown in the figure. Also shown in the figure is the diffusion path of particles interacting with a wave. From Eq. (3.7) we know that the slope of this line depends on the parameters of the wave, which in this case have been chosen so that $dx/d\varepsilon \approx a/\varepsilon_0$, with ε_0 the birth energy of the α-particles. If an α-particle exchanges energy with this wave it must move along the diffusion path shown. In the absence of

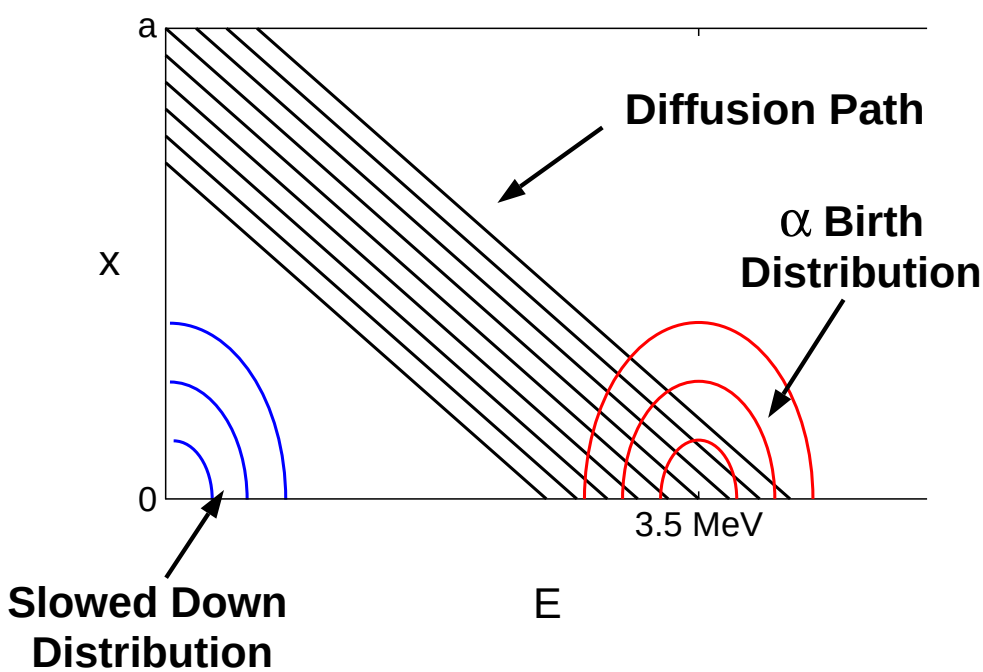

Figure 3.2: The optimal diffusion path for extracting α-particle energy. The horizontal axis is energy and the vertical axis is radius, x.

collisions, assuming a reflecting wall at x = 0,(corresponding to the center of the device), and an absorbing wall at x = a, (corresponding to the first wall of the device), eventually the particle must diffuse out near x = a and be absorbed. In doing so the particle will have lost almost 3.5 MeV of energy which in turn amplifies the wave. For this to work

in the presence of collisions, the α-particles must exit in a time short compared to the slowing down time.

Unfortunately, it appears to be difficult to excite waves in a reactor size tokamak which have this optimal slope. More typical are waves with slopes illustrated in Fig. 3.3. The wave with the large slope ($\Delta x/\Delta\varepsilon \gg 1$) is analogous to a low frequency wave such as the toroidal Alfvén eigenmode (TAE) in a tokamak. (Note that throughout this thesis we use TAE as a generic term for any low frequency eigenmode). While these waves do extract some energy from the α-particle distribution, they eject the particles with most of their energy remaining, causing damage to the first wall and a loss of heating to the plasma. The wave with the small slope ($\Delta x/\Delta\varepsilon \ll 1$) is analogous to the mode converted ion Bernstein wave (MCIBW) in the tokamak. It can extract energy from a birth distribution of α-particles by diffusing them to lower energies. However, as a slowing down distribution develops, the gradient of the distribution function along the diffusion path changes from $\partial f/\partial l > 0$ to $\partial f/\partial l < 0$ and the wave will tend to be damped rather than amplified.

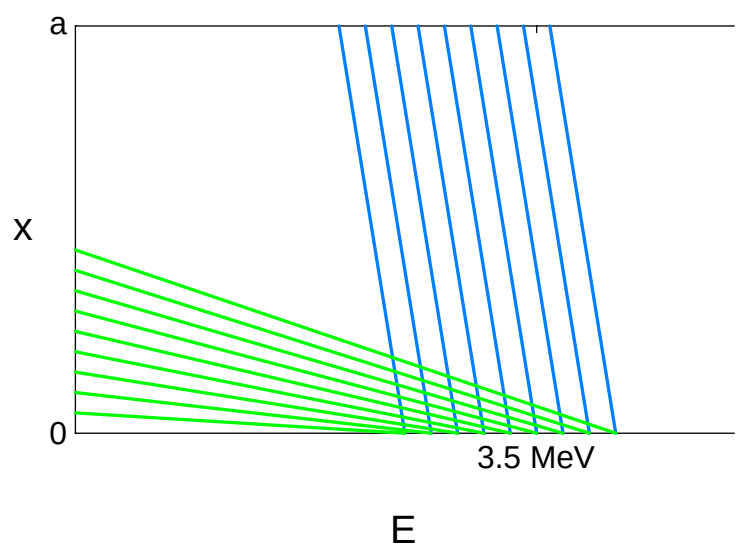

Figure 3.3: Diffusion paths which are typical of waves in tokamaks.

A qualitatively different picture emerges when the α-particles interact with two waves. To begin with, the diffusion is no longer constrained to one dimension, but now is fully two dimensional. Consider Fig. 3.4 which shows the phase space described above

with two waves present. One wave diffuses particles only in energy, and might arise from a Landau resonance, i.e., from Eq. (3.7) above, when $\mathbf{k} \parallel \mathbf{B}$. The other "wave" diffuses particles spatially without any energy exchange. This could occur due to the presence of stochastic magnetic fields or via stochastic ripple diffusion which arises from the ripple of the toroidal field. Immediately we see that the constraints that occur with only one

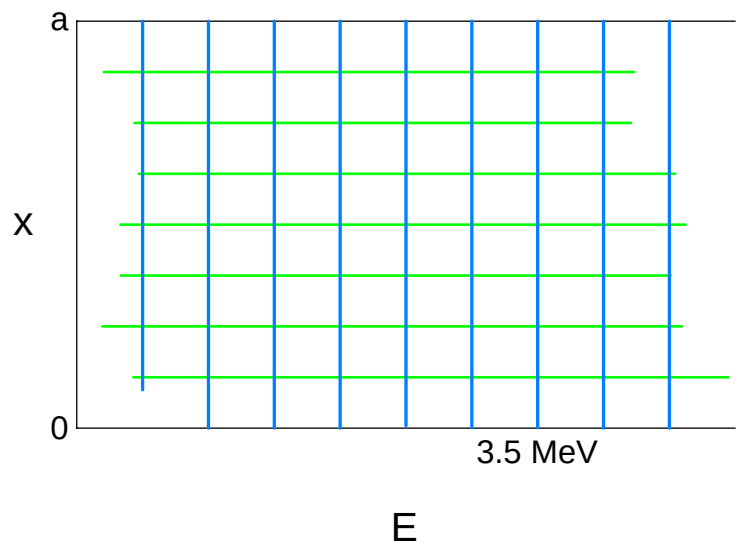

Figure 3.4: The diffusion paths of two waves, one which diffuses only in energy, one which diffuses only in space.

wave are removed. For instance, the concept of a fixed $\Delta x/\Delta\varepsilon$ which is true for one wave is not true for two. A particle could move from the center to the edge and be cooled, without any requirement on the slopes of the waves. However, this relaxation of constraints also allows at least some α-particles to be heated and ejected from the plasma.

Unlike the case with one wave, where cooling was assured, we now may have both heating and cooling. Fortunately, more control can be exercised over the system. By selectively choosing where the diffusion paths exist in phase space (by means of resonance conditions and spatial localization of the waves), configurations of two waves can be created that predominantly involve cooling. One such example is shown in Fig. 3.5. Since the only way for particles to leave this system is by going to low energy and being radially transported out, we are assured that this system of waves will lead to significant cooling. Of course the actual system we are dealing with, α-particles in a tokamak, is

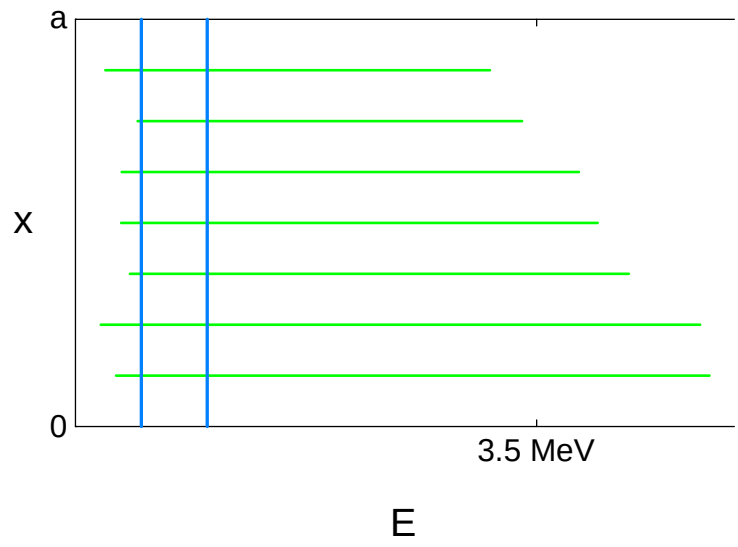

Figure 3.5: The diffusion path for a combination of two waves which could extract almost all of the α-particle energy. By choosing carefully where these waves exist in phase space (by means of resonance conditions and location of the waves) one can guarantee significant energy extraction.

more complicated than this simple two dimensional system. While waves with diffusion paths are still relevant, as we shall see, this system is three dimensional (it can be characterized in terms of the constants of the particles' motion, ε, μ, and P_ϕ) with complicated boundaries. Nevertheless, the insight gained in the simple two dimensional system is applicable to α-channelling in tokamaks.

3.3 Generic Requirements for α-Channelling

While creating a diffusion path which connects 3.5 MeV α-particles at the center with cooled α-particles at the edge is one of the fundamental requirements for α-channelling, there exist several other requirements which must be satisfied in order for α-channelling to be realized.

3.3.1 Tapping velocity space

Magnetically confined fusion plasmas typically have a strong magnetic field which creates significant anisotropies in the velocity space of the particles. A distribution of α-particles in a tokamak will have about $1/3$ of its energy parallel to the magnetic field and $2/3$ perpendicular to the field. In order to channel almost all of the energy, it is necessary to tap into both types of energy. However, because of the strong magnetic field, different waves are required for tapping perpendicular energy and parallel energy. For example, waves with $\omega \ll \Omega$, where Ω is the cyclotron frequency of the particles, will be unable to break the μ invariant, and thus will be more efficient at tapping parallel energy. (Note, though, that by moving particles from regions of high field to regions of lower field, e.g. from the center to the edge, a particle will lose perpendicular energy, even though μ is conserved because of changes in the magnetic field). On the other hand, waves with $\omega \approx \Omega$ will be able to break the μ invariant and extract perpendicular energy from the particles. These waves may also extract or give parallel energy to the particles as they interact, as discussed extensively in Chapter 6. Note also that the form of Eq. (3.7) suggests that the low frequency waves will push particles much further than waves with $\omega \approx \Omega$, given similar values for k.

3.3.2 Resonance conditions

In order for the α-particles to interact with the waves, they must satisfy a resonance condition. In particular, since we are interested in moving particles out and cooling them, it is necessary that the α-particle be resonant over a wide range of positions and velocities as they are slowed down by the waves and moved out. Furthermore, as discussed in the two wave scenario above, it would be useful to use those waves over which we might exert some control, both of the resonance conditions and the spatial localization.

3.3.3 Time scales

The α-particle must move out in a time short compared to a slowing down time, τ_{sd}. For a diffusive process, the time it takes for the particles to get out can be estimated as $\tau \approx (\Delta\varepsilon)^2/D_\varepsilon$. Since the diffusion coefficient is typically proportional to the power in the wave, the requirement $\tau < \tau_{sd}$, can be converted to a requirement on the electric field amplitude, which in turn is a constraint on the minimum intensity of the waves. To maximize the diffusion coefficient, it would be desirable to minimize the area over which the power is spread, while maximizing the power in the waves. However, the wave which is launched may also be amplified by the α-channelling effect. The wave may be convectively damped, so long as the power from the α-particles goes into the wave before the wave is significantly damped. This effect may reduce the required power, but the issue of startup will still remain, i.e. before the α-channelling power is flowing into the waves, enough power must be going to the waves in order to channel the α-particles. Thus, while the recirculating power requirements may not be large, it may be necessary to have a significant power capability to start up the α-channelling. This requirement might be alleviated if there were a way to bootstrap the α-channelling effect with small amounts of power.

3.3.4 Deposition of channelled power

Once the power gets into the waves from the α-particle, it must, in some way, be made to damp on the ions, preferably at the center. Thus it is necessary to choose waves which are likely to do this.

3.3.5 Power lost/wall loading

While much of the power may be diverted to waves, inevitably some energy will be left in the channelled α-particles when they leave the plasma. If this flux is localized, it might cause damage to the first wall. Furthermore, power lost in this way is not available for heating the plasma, thereby limiting the amount power available for sustaining a hot ion

mode. On the other hand, α-particles lost in this way (i.e. while still energetic), may be embedded in the walls and then would not be available as helium ash to choke the plasma. Note that it may be possible to take advantage of a localized loss by some clever means, that, for example, extracted further energy.

Chapter 4

Constants of Motion Simulation

ALPHA PARTICLES IN TOKAMAKS EXIST IN A phase space that is considerably more complicated than that discussed in the previous chapter. The guiding center approximation, which is usually quite well satisfied, allows us to neglect the gyromotion and reduces the phase space of the α-particles to a 5-dimensional one. Typically, one would simulate the evolution of the α-particle distribution using a guiding center code such as ORBIT [WHITE and CHANCE 1984]. However, there are two drawbacks with this approach for the problem we are considering here. The first is that the ORBIT code can simulate 1000 particles for 100 toroidal transits, where a toroidal transit is $2\pi qR/v_\alpha \sim \mathcal{O}(1\ \mu\text{second})$ in TFTR, in about 30 CPU seconds on the A machine (CRAY C-90) at the National Energy Research Scientific Computing Center (NERSC). For the problems we are interested in, with 1000 particles for something on the order of the slowing down time of about a quarter of a second this would translate to approximately 20 hours of CPU time.

Furthermore, much of the motion that takes place in the guiding center code is uninteresting to us. For instance, for the study of α-channelling, one is interested in whether or not a particle can diffuse to the edge of the plasma, not in the details of the ion motion within each orbit.

The approach we have taken is to study the transport of the energetic particles in constants-of-motion (COM) space, in particular, the ε, μ, and P_ϕ space introduced by

Hsu and Sigmar (1992). This representation retains all the physics of the unperturbed guiding center motion (except for the inconsequential phase of the particle's motion on its orbit and an equally irrelevant initial toroidal angle). In this 3-d space, it is possible to draw boundaries which indicate when a particle is on an orbit which intersects the wall and what the orbit topology is. It is even possible to draw boundaries for resonant interactions with the waves in this space (i.e. if a particle has a certain ε, μ, and P_ϕ, it can be determined whether or not it will resonate with a particular wave). There is also a precise direction in which particles move in COM space upon resonant interaction with a wave of given parameters. Thus, particles interacting with many waves in a tokamak (under the assumption of a random phase interaction) can be simulated without computing their real space orbits.

This simulation is equivalent to a Monte Carlo solution of the orbit averaged Fokker-Planck (FP) equation in the three dimensional COM space. This approach to simulating energetic particles in tokamaks is not, in itself, new; for instance, Rome and Peng (1979) suggested that the COM approach they took to finding energetic α-particles orbits in tokamaks would have applications to an orbit averaged FP code. The theoretical framework was provided even earlier; an elegant description of quasilinear diffusion in an axisymmetric torus formulated in terms of action angle variables, was given by Kaufman (1972). More recently, Eriksson and Helander (1994) calculate Monte Carlo operators for solving the orbit-averaged Fokker-Planck equation for arbitrary constants of motion (i.e. not just the action angle variables, which are inconvenient since the second adiabatic invariant, J_2, depends on an integral over the orbit). Eriksson and Helander (1994) also briefly review previous work along these lines. This approach was used to model ion cyclotron resonant heating (ICRH) in tokamaks, keeping the effects of full orbit width and spatially varying collision frequencies [CARLSSON, HELLSTEN, and ERIKSSON 1996; CARLSSON, ERIKSSON, and HELLSTEN 1994; CARLSSON, ERIKSSON, and HELLSTEN 1997]. Many of the calculations which follow are similar to the work of Eriksson and Helander (1994) and Carlsson, Hellsten, and Eriksson (1996).

When is this approach most useful? The COM approach is valid when the changes in the constants due to waves, collisions, or ripple in one bounce are small, i.e., when the orbit changes so little in one bounce that using the old orbit to determine the kicks is a good approximation. For instance, the COM approach would not be good at

following particles trapped in a ripple well, although the particles most likely to become ripple trapped can be identified by their COM. Also, the COM approach is most useful when any phase information can be thrown out, i.e. we are not interested in the phase of the particle on its orbit, or its phase relative to a wave. In this sense, the COM approach is not useful for determining the correct nonlinear saturation of a single TAE mode, or the details of the approach to stochasticity for stochastic ripple diffusion. Rather, it is ideally suited to studying energetic particles in the presence of many perturbations, such as waves, collisions, and toroidal field ripple, when it can be assumed that each interaction is diffusive or is diffusive with a drift term, i.e., when the system can be modeled with a Fokker-Planck equation.

Apart from a more rapid simulation, an additional benefit is the understanding that goes along with reducing the problem to three dimensions. It is very easy to visualize 2-d slices of this space (with the relevant boundaries). A particle's location in this phase space tells all there is to know about this particle. (Does it hit the wall? Will it interact with the wave?) In an α-channelling scenario, one can determine which groups of particles have done well, which have done poorly, and one can see how that changes as wave or plasma characteristics are varied. An example of this is given in Chapter 5.

This chapter is organized as follows: Sec. 4.1 describes the general outline of the simulation. The Monte Carlo increments for a particle interacting with collisions (Sec. 4.2), stochastic ripple diffusion (Sec. 4.3), MCIBW (Sec. 4.4), and toroidal Alfvén eigenmode (Sec. 4.4.8) are then discussed. The code is benchmarked, where possible, against analytical calculations and the existing guiding center code, ORBIT. In Sec. 4.5 the code is applied to the case of a neutral beam distribution colliding with the background plasma, to illustrate some of the simulation's capabilities. Finally, the ways in which the code performance could be improved are discussed.

4.1 Orbits and Outline of the Simulation

As shown in Appendix A, given the ε, μ, and P_ϕ (and the sign of $v_\parallel$) of a particle the details of its guiding center orbit can be calculated in a rapid and numerically efficient

way. The normalizations used throughout this thesis are discussed in Appendix A. A three dimensional space, ε, μ, and P_ϕ space (which is two-sheeted in some places due to the sign of $v_\parallel$) can be constructed. Once the basics of this space are understood we can start to formulate a Monte Carlo code in this space.

An overview of the simulation is given in Fig. 4.1. Given the details of the orbit, we compute the mean and the standard deviation of the changes in the COM which occur over one poloidal orbit of the particle. For instance, we compute the changes in ε, μ, and P_ϕ, which would occur due to collisions by averaging the effect of the Landau collision operator on ε, μ, and P_ϕ over the orbit (as detailed in the next section), or we determine if the particle passes through the IBW region, if it resonates with the wave in that region, and then calculate the size of the kick it receives. Next, we apply these changes for the number of poloidal orbits necessary to achieve the desired time step. Note that the changes due to drift terms are linear in the timestep and are deterministic, whereas the diffusive changes go like the square root of the time step and are multiplied by an appropriately chosen random number. We then apply the kick, and determine the characteristics of the new orbit.

We also determine whether the kick was so large that it moved the particle out of the physical portion of ε, μ, and P_ϕ space, or it caused the particle to change its topology in an unphysical way. Of course, this only occurs because of the orbit-averaged way in which the changes are calculated; in reality the particle's COM cannot be kicked into unphysical regions of ε, μ, and P_ϕ space. If an unphysical kick has been given, the time step is halved, the kick is applied again, and the orbit characteristics are determined. This process is repeated until a physical orbit is realized. A check is also made to see if a particle has hit the wall or if the energy has dropped below a threshold value which is on the order of the ion thermal temperature for the plasma. We are not interested in simulating the dynamics of thermal particles, and so particles below this energy are removed from the simulation. If the desired timestep was not achieved due to a kick resulting in an unphysical position, the particle is iterated upon until the desired timestep is achieved. In this way, each particle in a multi-particle simulation is recorded at the same value of time so that the distribution function at a given time step can be determined for diagnostic purposes.

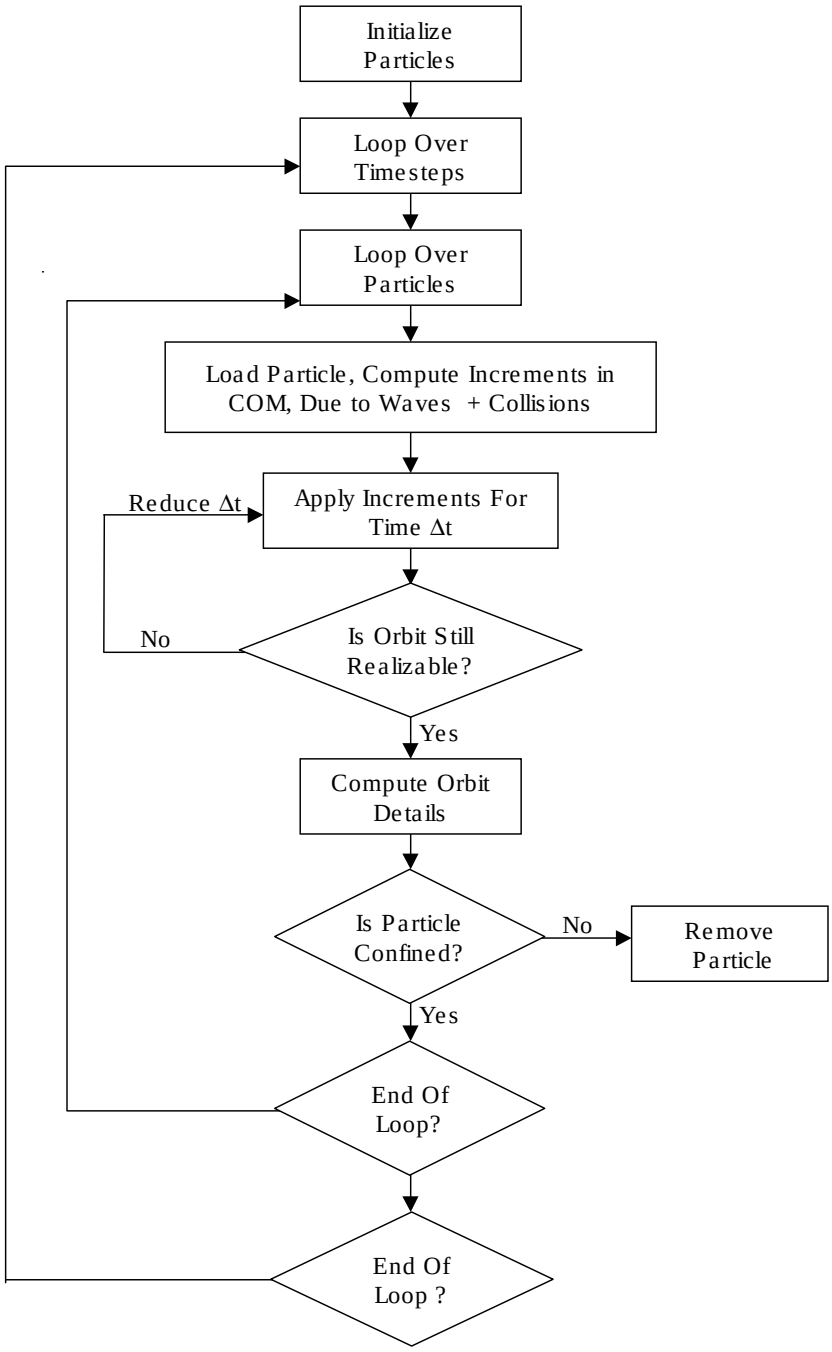

Figure 4.1: Flow chart of constants-of-motion Monte Carlo simulation.

While this seems straightforward there are subtleties which arise. For instance what if a particle crosses from trapped to passing? Does it become cogoing or countergoing? If it crosses the passing trapped boundary due to an interaction with a wave in the ion cyclotron range of frequencies this question is relatively easy to answer. Since the particle had to be resonant with the wave in order for it to get the push across the passing trapped boundary, it must be that the particle continues going in the direction it was going when it resonated with the wave. For a wave like the toroidal Alfvén eigenmode or collisions this question is more difficult to answer. The answer can be thought of in terms of the continuity of phase space density across the boundary. If this did not hold the fluxes would be infinite. Furthermore there also has to be reciprocity, i.e. the particles must be just as likely to cross from A to B as to cross from B to A. Together these two constraints allow the determination of the flux.

Further subtleties arise in ensuring that the wave particle interactions are microscopically reversible, i.e. that a particle which receives a kick from the wave remains resonant with the wave after the kick. This is discussed in greater detail in Sec. 4.4.6.

4.2 Collisions

As a first approximation, the 3.5 MeV α-particles in a reactor can be treated as collisionless (as is done in the investigation of α-channelling in a reactor in Chapter 5). However, the addition of collisions to simulations of α-channelling is important, as the channelling effect must take place on a time scale short compared with the slowing down time. Thus adding collisions will allow us to determine the wave amplitudes and powers necessary to accomplish α-channelling, although the power levels have been estimated in previous work [HERRMANN and FISCH 1997]. Collisions are even more important for the modeling of experiments conducted on TFTR [DARROW et al. 1996] which saw a strong interaction of beam ions with MCIBW (see Chapters 6 and 7). In some experiments, it was possible to determine that this interaction continued for several beam particle slowing down times, eliminating the possibility of a collisionless interaction. Thus collisional effects are necessary to model these experiments.

Consider the Landau collision operator for a weak beam of species α interacting with a Maxwellian background plasma species β. Then

$$\frac{df^\alpha}{dt} = \boldsymbol{\nabla}_v \cdot \left(\frac{m_\alpha}{m_\alpha + m_\beta} \, {\nu_s}^{\alpha/\beta} \, \mathbf{v} \, f^\alpha + \frac{1}{2} \, {\nu_\parallel}^{\alpha/\beta} \, \mathbf{v}\mathbf{v} \cdot \boldsymbol{\nabla}_v f^\alpha + \frac{1}{4} \, {\nu_\perp}^{\alpha/\beta} \, \left(v^2 \mathbb{I} - \mathbf{v}\,\mathbf{v} \right) \cdot \boldsymbol{\nabla}_v f^\alpha \right) \tag{4.1}$$

with the following definitions,

$$ {\nu_0}^{\alpha/\beta} = \frac{4\pi {e_\alpha}^2 {e_\beta}^2 \lambda_{\alpha\beta} n_\beta}{{m_\alpha}^2 {v_\alpha}^3} \tag{4.2}$$

$$ {\nu_s}^{\alpha/\beta} = \left(1 + \frac{m_\alpha}{m_\beta} \right) \psi \left(x^{\alpha/\beta} \right) \nu_0^{\alpha/\beta} \tag{4.3}$$

$$ {\nu_\parallel}^{\alpha/\beta} = \frac{\psi \left(x^{\alpha/\beta} \right)}{x^{\alpha/\beta}} \nu_0^{\alpha/\beta} \tag{4.4}$$

$$ {\nu_\perp}^{\alpha/\beta} = 2 \left(\left(1 - \frac{1}{2x^{\alpha/\beta}} \right) \psi \left(x^{\alpha/\beta} \right) + \psi' \left(x^{\alpha/\beta} \right) \right) \nu_0^{\alpha/\beta} \tag{4.5}$$

$$ x^{\alpha/\beta} = \frac{m_\beta {v_\alpha}^2}{2T_\beta} = \frac{m_\beta}{m_\alpha} \frac{\varepsilon_\alpha}{T_\beta} \tag{4.6}$$

$$ \psi(x) = \frac{2}{\sqrt{\pi}} \int_0^x dt \sqrt{t} e^{-t} \tag{4.7}$$

If we transform this expression to spherical coordinates, assume that the distribution function is gyrophase independent, and sum over all the thermal species, we can rewrite this as:

$$\frac{\partial f^\alpha}{\partial t} = \frac{1}{v^2} \partial_v \left(\alpha(v) v^2 f^\alpha \right) + \frac{1}{2v^2} \partial_v^2 \left(\beta(v) v^2 f^\alpha \right) + \frac{\gamma(v)}{4v^2} \partial_\lambda (1 - \lambda^2) \partial_\lambda f^\alpha \tag{4.8}$$

$$ \alpha(v) = \sum_\beta \frac{1}{v^2} \left(\left(\frac{1}{2x^{\alpha/\beta}} - \frac{m_\alpha}{m_\beta} \right) \psi \left(x^{\alpha/\beta} \right) + \psi' \left(x^{\alpha/\beta} \right) \right) \nu_0^{\alpha/\beta} \tag{4.9}$$

$$ \beta(v) = \sum_\beta \frac{\psi \left(x^{\alpha/\beta} \right)}{v x^{\alpha/\beta}} \nu_0^{\alpha/\beta} \tag{4.10}$$

$$ \gamma(v) = \sum_\beta \frac{2}{v} \left(\left(1 - \frac{1}{2x^{\alpha/\beta}} \right) \psi \left(x^{\alpha/\beta} \right) + \psi' \left(x^{\alpha/\beta} \right) \right) \nu_0^{\alpha/\beta}. \tag{4.11}$$

Now we would like to determine an equation for the evolution of ε, μ, and P_ϕ

which corresponds to this Fokker-Planck equation. It will have a general form [RISKEN 1989a]:

$$X_i(t+\Delta t) - X_i(t) = \Delta X_i = h_i \Delta t + g_{ij}\zeta_j\sqrt{\Delta t}, \tag{4.12}$$

where i and j run from 1 to 3 and the ζ_j's are random variables satisfying $\langle\zeta_j\rangle = 0$ and $\langle\zeta_i\zeta_j\rangle = \delta_{ij}$, and summation over repeated indices is assumed. Taking the average of this equation, we find $h_i = \langle\Delta X_i/\Delta t\rangle \equiv \langle dX_i/dt\rangle$. Multiplying this by the equation for ΔX_k and taking the average, we find

$$\langle\Delta X_i \Delta X_k\rangle = h_i h_k \Delta t^2 + g_{ij} g_{kl}\langle\zeta_j\zeta_l\rangle\Delta t \tag{4.13}$$

$$\frac{\langle\Delta X_i \Delta X_k\rangle}{\Delta t} = g_{ij} g_{kj} + h_i h_k \Delta t, \tag{4.14}$$

where we have used the definition of ζ_i. A little algebra expanding the ΔX's shows that

$$\lim_{\Delta t\to 0}\frac{\langle\Delta X_i \Delta X_k\rangle}{\Delta t} = \frac{d\langle X_i X_k\rangle}{dt} - \langle X_i\rangle\frac{d\langle X_k\rangle}{dt} - \langle X_k\rangle\frac{d\langle X_i\rangle}{dt} \equiv \frac{d\sigma_{ik}}{dt}, \tag{4.15}$$

where $\sigma_{ik} = \langle X_i X_k\rangle - \langle X_i\rangle\langle X_k\rangle$. Taking the limit of Δt goes to zero in Eq. (4.14) we see that

$$\frac{d\sigma_{ik}}{dt} = g_{ij} g_{kj} = g\ {g^T}_{ik} \tag{4.16}$$

To calculate these quantities we follow Carlsson, Hellsten, and Eriksson (1996). Consider the quantity X, then the average value of X over the distribution function is given by

$$\langle X\rangle = 4\pi\int_{-1}^{1}\int_{0}^{\infty}\int d\mathbf{r} X(v,\lambda,\mathbf{r}) f(v,\lambda,\mathbf{r}) v^2 dv d\lambda, \tag{4.17}$$

Take $f(v,\lambda,\mathbf{r}) = \delta(v-v_0)\delta(\lambda-\lambda_0)\delta(\mathbf{r}-\mathbf{r_0})/(4\pi {v_0}^2)$ and take the time derivative of Eq. (4.17), substituting in for the collision operator from Eq. (4.8), and integrating by

parts where appropriate, one can find

$$\frac{d\langle X\rangle}{dt} = \alpha(v)\,\partial_v X + \frac{\beta(v)}{2}\,\partial_v^2 X + \frac{\gamma(v)}{4\,v^2}\,\left(-2\,\lambda\partial_\lambda\,X + \left(1-\lambda^2\right)\,\partial_\lambda^2 X\right). \tag{4.18}$$

Eq. (4.18) will be used to derive the changes in ε, μ, and P_ϕ which are consistent with the collisional evolution of the particle distribution. Consider the quantities $\varepsilon = v^2/2,\ \Lambda \equiv \mu/\varepsilon = (1-\lambda^2)/B,$ and $P_\phi = Fv\lambda/B - \psi$. Note Λ has been substituted for μ here, because it will have no correlation with ε, thereby simplifying the resulting equations. Then, calculating the drift term, we get

$$\frac{d\varepsilon}{dt} = \frac{(2\,v\,\alpha(v) + \beta(v))}{2} \tag{4.19}$$

$$\frac{d\Lambda}{dt} = \frac{(-1+3\,\lambda^2)\,\gamma(v)}{2\,B\,v^2} \tag{4.20}$$

$$\frac{dP_\phi}{dt} = \frac{F\ \lambda\ (2\,v\,\alpha(v) - \gamma(v))}{2\,B\,v}. \tag{4.21}$$

The calculation of the $d\sigma/dt$ is made somewhat easier by two simplifications: one, it is easy to see from Eq. (4.15) that $d\sigma_{ik}/dt = d\sigma_{ki}/dt$, and, two, our choice of Λ. Thus

$$\begin{pmatrix} \dot\sigma_{\varepsilon\varepsilon} \\ \dot\sigma_{\varepsilon P_\phi} \\ \dot\sigma_{\Lambda\Lambda} \\ \dot\sigma_{\Lambda P_\phi} \\ \dot\sigma_{P_\phi P_\phi} \end{pmatrix} = \begin{pmatrix} v^2\beta \\ \frac{F\,v\,\lambda\,\beta}{B} \\ \frac{-2\ \lambda^2\left(-1+\lambda^2\right)\gamma}{B^2\,v^2} \\ \frac{F\ \lambda\left(-1+\lambda^2\right)\gamma}{B^2\,v} \\ \frac{F^2\ \left(\lambda^2\,(2\,\beta-\gamma)+\gamma\right)}{2\,B^2} \end{pmatrix}. \tag{4.22}$$

Note that these expressions are dependent on the position along the particle's orbit, since buried in α, β, and γ are dependencies on the temperature and density of the plasma which depend on ψ. Thus, these expressions need to be orbit averaged to determine the average rates of change. As discussed in Appendix A the orbit average of an expression is

$$\langle f\rangle = \frac{1}{\tau_b}\int_0^{\tau_b} \frac{f(\psi)d\psi}{\dot\psi}. \tag{4.23}$$

Thus $h_i = \langle \dot{X}_i \rangle$. Recall that the matrix $d\sigma/dt$ is a positive definite, symmetric matrix. As such, it has a Cholesky decomposition, g, which will satisfy $d\sigma/dt = gg^T$. This is the same g that we need to complete the Langevin description from Eq. (4.16). It is given by

$$g = \begin{pmatrix} \sqrt{\langle \dot{\sigma}_{\varepsilon\varepsilon} \rangle} & 0 & 0 \\ 0 & \sqrt{\langle \dot{\sigma}_{\Lambda\Lambda} \rangle} & 0 \\ \frac{\langle \dot{\sigma}_{\varepsilon P_\phi} \rangle}{\sqrt{\langle \dot{\sigma}_{\varepsilon\varepsilon} \rangle}} & \frac{\langle \dot{\sigma}_{\Lambda P_\phi} \rangle}{\sqrt{\langle \dot{\sigma}_{\Lambda\Lambda} \rangle}} & \sqrt{\langle \dot{\sigma}_{P_\phi P_\phi} \rangle - \frac{\langle \dot{\sigma}_{\varepsilon P_\phi} \rangle^2}{\langle \dot{\sigma}_{\varepsilon\varepsilon} \rangle} - \frac{\langle \dot{\sigma}_{\Lambda P_\phi} \rangle^2}{\langle \dot{\sigma}_{\Lambda\Lambda} \rangle}} \end{pmatrix} \tag{4.24}$$

So the final equations, which are equivalent to those of Carlsson et al. (1996) except for slight generalization of the magnetic geometry, are:

$$\begin{pmatrix} \Delta\varepsilon \\ \Delta\Lambda \\ \Delta P_\phi \end{pmatrix} = \begin{pmatrix} \frac{1}{2} \langle 2\, v\, \alpha(v) + \beta(v) \rangle \\ \frac{1}{2v^2} \langle (-1 + 3\,\lambda^2)\, B^{-1}\, \gamma(v) \rangle \\ \frac{1}{2v} \langle F\, B^{-1}\, \lambda\, (2\, v\, \alpha(v) - \gamma(v)) \rangle \end{pmatrix} \Delta t +$$
$$\begin{pmatrix} \sqrt{\langle \dot{\sigma}_{\varepsilon\varepsilon} \rangle} & 0 & 0 \\ 0 & \sqrt{\langle \dot{\sigma}_{\Lambda\Lambda} \rangle} & 0 \\ \frac{\langle \dot{\sigma}_{\varepsilon P_\phi} \rangle}{\sqrt{\langle \dot{\sigma}_{\varepsilon\varepsilon} \rangle}} & \frac{\langle \dot{\sigma}_{\Lambda P_\phi} \rangle}{\sqrt{\langle \dot{\sigma}_{\Lambda\Lambda} \rangle}} & \sqrt{\langle \dot{\sigma}_{P_\phi P_\phi} \rangle - \frac{\langle \dot{\sigma}_{\varepsilon P_\phi} \rangle^2}{\langle \dot{\sigma}_{\varepsilon\varepsilon} \rangle} - \frac{\langle \dot{\sigma}_{\Lambda P_\phi} \rangle^2}{\langle \dot{\sigma}_{\Lambda\Lambda} \rangle}} \end{pmatrix} \begin{pmatrix} \zeta_\varepsilon \\ \zeta_\Lambda \\ \zeta_{P_\phi} \end{pmatrix} \sqrt{\Delta t}. \tag{4.25}$$

By separately turning off different terms in Eq. (4.25), the collision terms can be "benchmarked." By "benchmarked" we mean that we can make sure our code is working as intended, which is most easily done by taking limiting cases that can be checked against other codes or by analytical methods.

First, leaving only the ε drift term should give us just energy slowing down. In Fig. 4.2, we plot the energy of an α-particle initially born at 3.5 MeV versus time, in a 20 keV deuterium plasma.

There is also a term which represents diffusion in energy. Leaving only this term on in the code, the variance of the normalized energy for 5000 350 keV α-particles in a 20 keV deuterium plasma is plotted in Fig. 4.3. The solid line is $\dot{\sigma}_{\varepsilon\varepsilon} t$ as calculated above. Note that agreement is good for short times, but significant deviation begins occurring when Var$(\varepsilon/\varepsilon_0) \approx 0.3$ or the standard deviation is on the order of $\varepsilon_0/2$. There are three causes for this: One is that the diffusion coefficient is a function of velocity and thus as

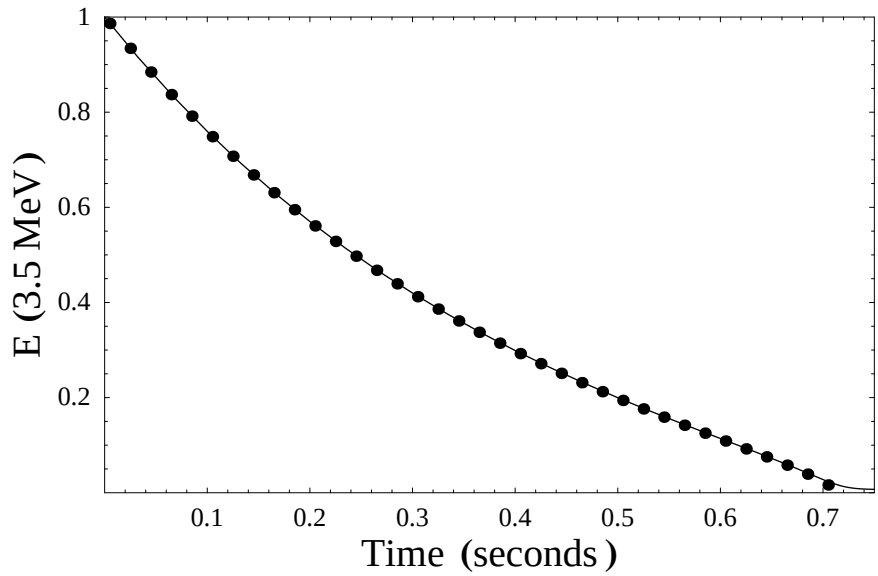

Figure 4.2: The energy of an α-particle, born at 3.5 MeV, versus time in a neutral, 20 keV deuterium plasma. The solid line is determined from Eq. (4.21) using the initial velocity for the evaluation, the dots represent output from the simulation.

particles begin to diffuse significantly they will start to diffuse at different rates, changing the measured value of the variance. More important, and what causes the variance from the simulation to be less than the predicted one, is that boundaries in energy exist both at low and high energy. At high energies particle orbits get large enough to intersect the wall. At low energies a artificial boundary is imposed as we are not interested in the evolution of thermal particles.

As suggested by Carlsson, Hellsten, and Eriksson (1996), the pitch angle scattering can be benchmarked against analytical theory, if we turn off the energy slowing down and energy scattering. We get an equation for f^α which is separable in λ and time. This can be solved in terms of Legendre polynomials. For an initial distribution $f^\alpha(\lambda, v, t = 0) = \delta(\lambda - \lambda_0)$ the solution is

$$f^\alpha(\lambda, v, t) = \sum_{l=0,\infty} (l + \frac{1}{2}) P_l(\lambda_0) P_l(\lambda) e^{-\frac{(l)(l+1)\gamma(v)t}{4v^2}}. \tag{4.26}$$

This expression can be compared to the results of the COM simulation for a high

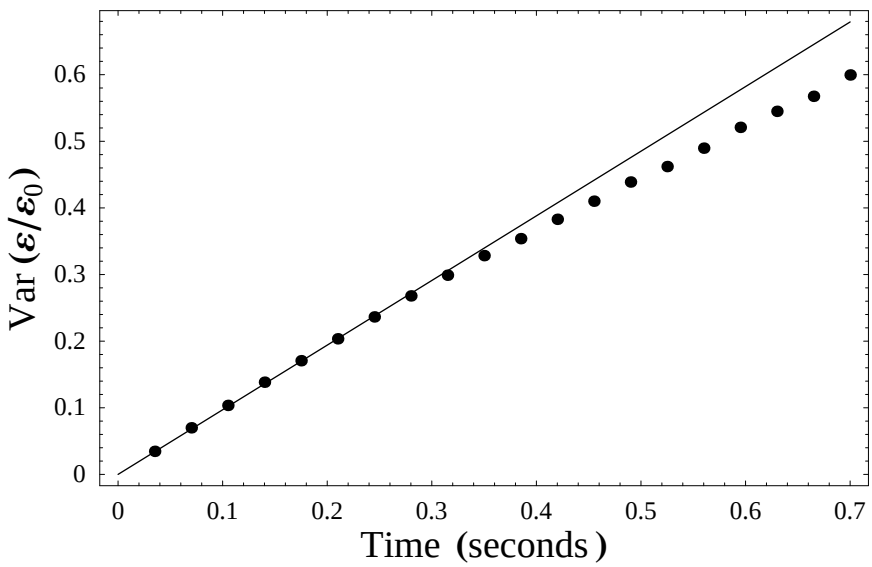

Figure 4.3: The growth in the variance of the normalized energy for a distribution of 5000 α-particles, born at 350 keV in a 20 keV deuterium plasma. The solid line is the analytical rate of growth, dots represent output from the code.

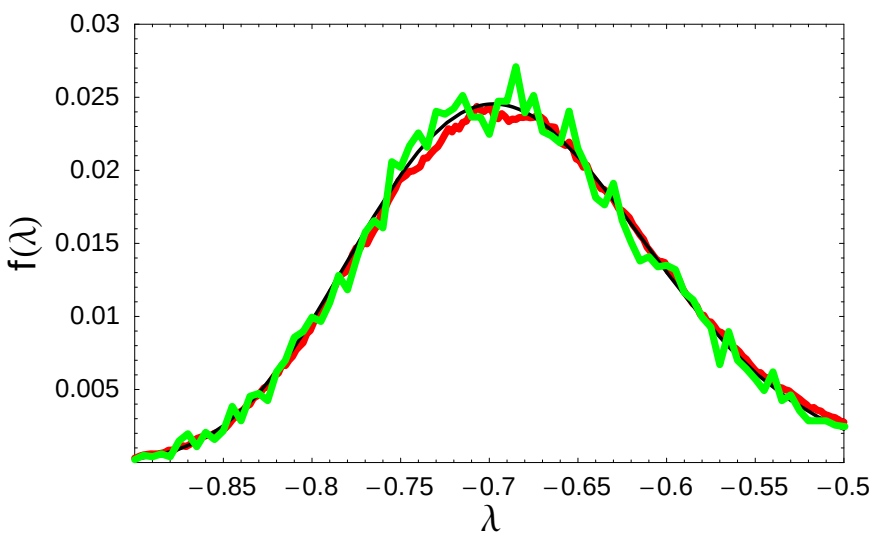

Figure 4.4: Collisional evolution of a δ function in pitch, benchmarked against the ORBIT code and the analytical expression. The smooth black line is the analytical calculation, the jagged black line is the simulation result with 30000 particles, the light line is the result from ORBIT with 50000 particles.

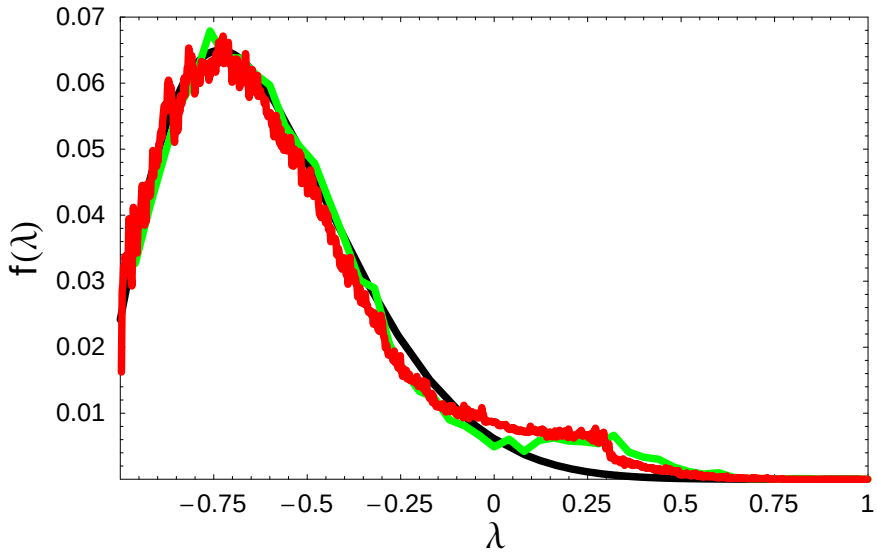

Figure 4.5: Collisional evolution of a δ function in pitch, for longer times. Eventually particles pitch angle scatter far enough that some become trapped. At this point the analytical result for a homogeneous magnetic field and the simulations differ.

aspect ratio equilibrium as shown in Fig. 4.4. The agreement is good. Note that, when the simulation is run for longer, the results disagree with the analytical theory, as shown in Fig 4.5. This is because as time goes on some particles scatter far enough in pitch angle to cross the passing-trapped boundary. When this occurs, the distribution function in the interval $(-\lambda_{pt}, +\lambda_{pt})$ is flattened very rapidly (on the order of a bounce time, which is much less than a pitch angle scattering time). This causes the distribution to be roughly flat between these two points. Both cases also show good agreement with the ORBIT code.

4.3 Stochastic Ripple Diffusion

Stochastic ripple diffusion (SRD) is caused by the distortion from axisymmetry due to the finite number of field coils which create the toroidal field [GOLDSTON, WHITE, and BOOZER 1981]. It can lead to energetic particle losses by breaking the P_ϕ invariant. If the ripple amplitude is sufficient, diffusion in P_ϕ can occur for trapped particles (although the particles μ and ε are conserved under interaction with the ripple). One can see from

an examination of an ε, μ, and P_ϕ space plot, that diffusion in P_ϕ may eventually lead to fast particles hitting the wall. This will be an important consideration for α-channelling, for two reasons: First, it introduces a faster timescale which we must be concerned with into the system, i.e. if we are dealing with trapped particles which are being affected by SRD we must now get them out in a timescale short compared to the SRD time rather than the slowing down time. This timescale can be milliseconds, much faster than the 100's of milliseconds for the slowing down time, and thus presents a concern on the required power level for the waves which are used. Second, in Chapter 5, we will see that α-channelling seems to work best when passing particles are brought out and turned into trapped particles near the edge. This coincides with the region where the ripple is strongest and therefore SRD most likely to be a factor. On the other hand, it may also be the case that interaction with ripple is useful for accomplishing α-channelling, since it provides a means of transporting a select group of particles near to the edge without any energy extraction. In particular this may be a useful aid in extracting deeply trapped particles which do not have enough parallel energy to be pushed out by the TAE.

While the full investigation is yet to be done, we have taken the first step in trying to address SRD within the COM framework. We employ accurate analytical estimates of the stochastic ripple threshold, based on work by White et al. (1996) and more recently White (1998) in the COM code. The basic idea is as follows. Due to the ripple in the magnetic field, a perturbation is added to the guiding center Hamiltonian which varies as $H_1 \propto \delta \sin N\phi$, where δ is the ripple amplitude and N is the number of field coils. This perturbation gives

$$\dot{P}_\phi = \frac{-\partial H_1}{\partial \phi} \tag{4.27}$$

$$\dot{\phi} = \frac{\partial H_0}{\partial P_\phi} + \frac{\partial H_1}{\partial P_\phi}. \tag{4.28}$$

In practice, a simplifying assumption can be made, which is that most of the change in a particle's P_ϕ occurs at its bounce point. When this is done the change in a particles P_ϕ in one half orbit (one pass through a banana tip) can be estimated analytically. Following White et al. (1996) we arrive at a mapping for the evolution of P_ϕ =

-$\psi_{\text{banana tip}}$ due to SRD (we assume up-down symmetry of the equilibrium).

$$P_{\phi_{t+1}} = P_{\phi_t} - \Delta \sin N\phi_t \tag{4.29}$$

$$N\phi_{t+1} = N\phi_t + N\phi_{b,t+1} + N\phi_{p,t+1} \tag{4.30}$$

$$P_{\phi_{t+2}} = P_{\phi_{t+1}} - \Delta \sin N\phi_{t+1} \tag{4.31}$$

$$N\phi_{t+2} = N\phi_{t+1} - N\phi_{b,t+2} + N\phi_{p,t+2}, \tag{4.32}$$

where $\Delta = g\rho\sqrt{\pi N q}/(B\partial_\theta B)^{1/2}$. White (1998) has undertaken a detailed analysis of this mapping to determine the stochastic threshold. His analytical expression has been implemented in the simulation, and a comparison with numerical results is shown in Fig. 4.6. The agreement is good at low energy, as shown here, but not good at higher energies due to finite orbit width corrections to the expressions of White (1998).

For the COM simulation, this routine can be used as in White et al. (1996), that is, assume particles are lost as soon as they cross into the stochastic domain. More ambitiously we could follow the diffusion of the ripple particles in this domain. While in principle the mapping could be followed, this is not desirable. Evaluating the mapping requires keeping track of phase information, which, as discussed earlier, is something to be avoided. Furthermore, the time steps would be limited to 1/2 of a bounce time per step, which is much shorter than is preferred. Instead, the evolution of the distribution function could be modeled by

$$\frac{\partial f}{\partial t} = \frac{\partial}{P_\phi} D(\varepsilon, \mu, P_\phi) \frac{\partial}{P_\phi} f \tag{4.33}$$

where D is zero for regions below stochastic threshold and related to Δ above threshold. Then, applying techniques used in the previous section, an average and random change to P_ϕ due to this operator could be computed and applied to particles in a Monte Carlo fashion.

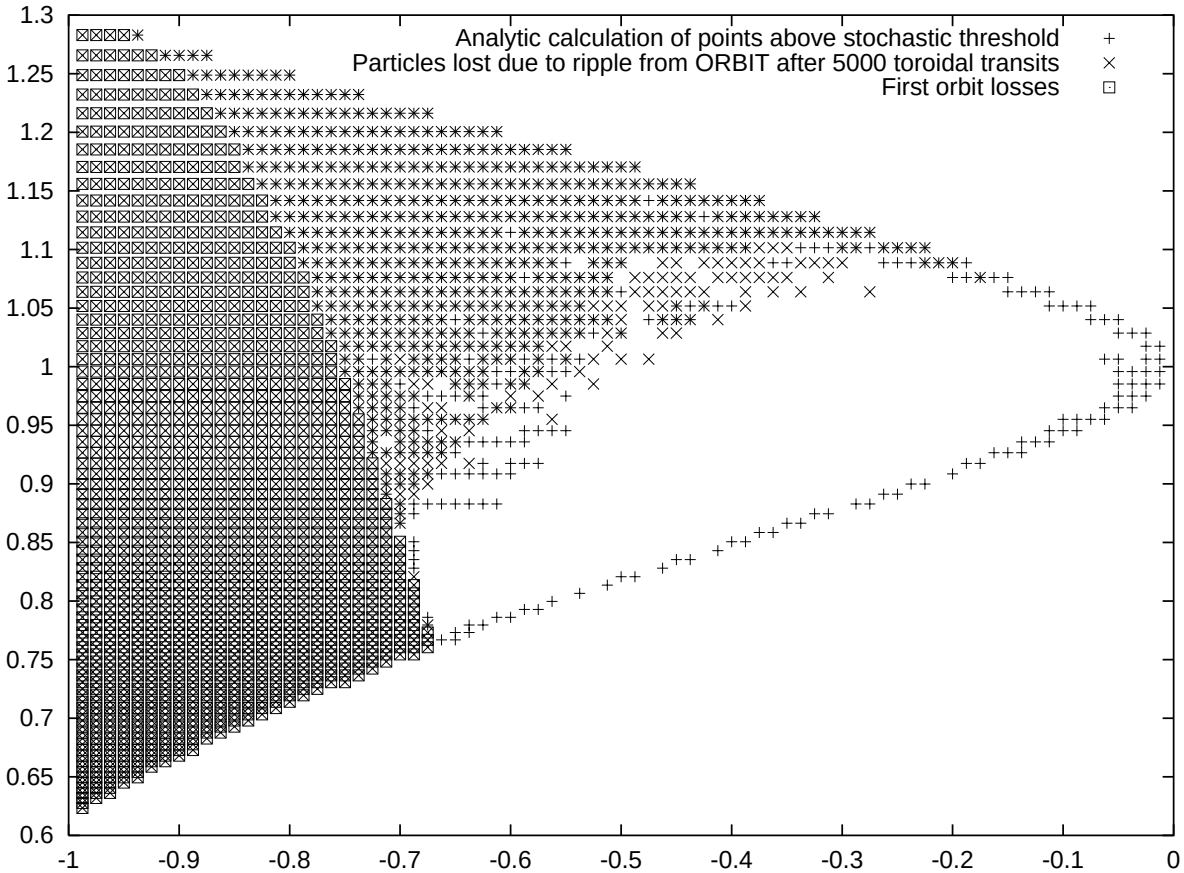

Figure 4.6: Plot of the trapped portion of ε, μ, and P_ϕ space for 700 keV α-particles. The + represents points in μ,P_ϕ space where Δ is above the stochastic threshold. The X represents particles confined for 5000 transits, as determined by the ORBIT code. A * arises when the two overlap.

4.4 Wave Particle Interactions

Let us consider the effect of the MCIBW on the particles. There are two ways to go about calculating the interaction of the MCIBW and the particles. One approach would be to assume that the interaction can be approximated by quasilinear diffusion, then determine the quasilinear diffusion equation, and from that obtain expressions for the change in the COM, much like was done with collisions. An alternate approach would be to derive a mapping from the Hamiltonian of the particle, for the COMs, which would allow us to investigate the approach to stochasticity (or at least estimate the stochastic threshold). In the following section we use the quasilinear formalism to determine the increments to the

COM and we estimate the requirements for the particle to be above stochastic threshold. It will be useful to consider the action angle formulation of the problem, as done by Kaufman (1972). This approach does not require bounce averaging of the quasilinear equation for a homogeneous system, but directly incorporates the full particle dynamics.

4.4.1 Action-angle formulation

We consider a particle's motion in a tokamak. The motion of the α-particles takes place in a 6-dimensional space, however, the unperturbed motion is integrable, and thus, action-angle variables for the motion can be found [KAUFMAN 1972]. The advantage of these variables is that the unperturbed Hamiltonian can be written as

$$H = H_0(\mathbf{J}), \tag{4.34}$$

where the J's are the actions. It can be seen immediately that they are constant in time as the Hamiltonian is independent of the canonical position coordinates, θ. The equations for $\dot{\theta}$ are trivially integrated. In particular the momenta and coordinates are

$$J^1 = \mu \qquad \theta^1 = \Omega^1 t + \theta^1{}_0 \tag{4.35}$$

$$J^2 = \oint d\psi_t \qquad \theta^2 = \Omega^2 t + \theta^2{}_0 \tag{4.36}$$

$$J^3 = P_\phi \qquad \theta^3 = \Omega^3 t + \theta^3{}_0, \tag{4.37}$$

where $\Omega^i = \partial H/\partial J^i$. The Ω^i represent, respectively, the gyrofrequency averaged over a poloidal orbit, the bounce frequency, ω_b, which is 2π divided by the time taken to complete a poloidal transit and the toroidal precession frequency, ω_d, which is the amount of ϕ precessed in one poloidal orbit divided by the bounce time.

Consider the interaction of particles with a perturbation which is of the form

$$H = H_0(\mathbf{J}) + H_1(\mathbf{J}, \theta, t) \tag{4.38}$$

Note that $H_1(\mathbf{J}, \theta, t)$ is deceptively simple. In reality it represents the perturbation felt by the particle as it goes along its unperturbed orbit. Unlike a more standard representation

of the perturbation, i.e. $H_1(\mathbf{r}, t)$, which only depends on the three position variables and time, this representation depends, in general, on all six variables and time.

4.4.2 Quasilinear diffusion in action-angle variables

The approach we take here is to derive the quasilinear diffusion in action angle variables and then transform the diffusion operator to ε, μ, and P_ϕ space. Now, following Kaufman (1972) or Eriksson and Helander (1994), the quasilinear diffusion equation for this system is:

$$\frac{\partial f}{\partial t} = \frac{\partial}{\partial J^i} D^{ij} \frac{\partial f}{\partial J^j} \tag{4.39}$$

$$D^{ij} = \pi \sum_{\mathbf{n},\omega} |H_1(\mathbf{J}, \mathbf{n}, \omega)|^2 \delta(\omega - \mathbf{n} \cdot \Omega) n^i n^j \tag{4.40}$$

and

$$H_1(\mathbf{J}, \mathbf{n}, \omega) = \frac{1}{(2\pi)^4} \int \int H_1(\mathbf{J}, \theta, t) e^{-i\omega t} e^{-i\mathbf{n}\cdot\theta} dt d^3\theta. \tag{4.41}$$

Now let us consider the form of H_1 for the MCIBW. First we write down H_1 in terms of real space variables and then map to $\mathbf{J}, \theta$.

We take the MCIBW to be an electrostatic wave which is localized in a narrow layer between two $|B|$ surfaces. The details of the MCIBW modeling are discussed in Chapter 6, Sec. 6.3. The ray tracing we use for the MCIBW assumes that the MCIBW changes most rapidly in direction of $\boldsymbol{\nabla} B$ which can be taken to be $-\hat{R}$ for the purposes of a 1-d slab geometry. In the slab, k_z is preserved, however, there is a projection of k_x onto $k_\parallel$ due to the poloidal field of the wave, which is important for the wave dynamics and wave particle interaction. In the tokamak the toroidal component of the wave, n_ϕ is conserved. We separate out that part of the wave phase. We take the wave phase, S, to be

$$S = \int^X dX' \cdot k_X(X') + n_\phi \phi - \omega t + \phi_0, \tag{4.42}$$

where X is the $\hat{R}$ direction in the tokamak. Thus, the form of the potential we take is:

$$\Phi = \Phi_0(\mathbf{r})e^{iS}, \tag{4.43}$$

where $\Phi_0(\mathbf{r})$ is both slowly varying and independent of ϕ. Note that we can write $\boldsymbol{\nabla} S$ as

$$\boldsymbol{\nabla} S = k_X(X)\left(\frac{\partial X}{\partial \theta}\boldsymbol{\nabla}\theta + \frac{\partial X}{\partial \psi}\boldsymbol{\nabla}\psi\right) + n_\phi \boldsymbol{\nabla}\phi \tag{4.44}$$

$$\boldsymbol{\nabla} S = k_\perp \mathbf{e_1} + \left(\hat{\mathbf{b}}\cdot\boldsymbol{\nabla} S\right)\hat{\mathbf{b}} \tag{4.45}$$

$$k_\perp^2 = k_X^2 + \frac{n_\phi^2}{R^2} - k_\parallel^2, \tag{4.46}$$

where $\mathbf{e_1}$ is a unit vector perpendicular to the magnetic field and $k_\parallel$ is given by

$$k_\parallel = \hat{\mathbf{b}}\cdot\boldsymbol{\nabla} S = n_\phi \frac{q}{JB} + k_X \frac{1}{JB}\frac{\partial X}{\partial \theta}. \tag{4.47}$$

In the large aspect ratio limit, $k_\parallel$ becomes

$$k_\parallel = n_\phi/R_0 + k_X \frac{\epsilon}{q}\sin\theta. \tag{4.48}$$

If the second term in the above equation is of opposite sign to the first and k_X becomes large enough, $k_\parallel$ can flip (i.e. change) signs from the sign it had when launched. The so called "$k_\parallel$ flip" is very important to the viability of the α-channelling effect by means of the MCIBW as discussed in Sec. 5.5.

The perturbation, H_1 is just $q\Phi$. Let us separate out the fast gyro-motion and the slower guiding center motion. The position of the particle can be written as $\mathbf{X} = \mathbf{X}_{g.c.} + \rho$. Then

$$H_1 \approx \Phi_0(\mathbf{X}_{g.c.})e^{i(S(\mathbf{X}_{g.c.})+\rho\cdot\boldsymbol{\nabla} S)} \tag{4.49}$$

$$= \Phi_0(\mathbf{X}_{g.c.})e^{iS(\mathbf{X}_{g.c.})}e^{ik_\perp\rho\sin\alpha} \tag{4.50}$$

$$= \Phi_0(\mathbf{X}_{g.c.})e^{iS(\mathbf{X}_{g.c.})}\sum_{l=-\infty}^{\infty} J_l(k_\perp\rho)e^{il\alpha} \tag{4.51}$$

where $\alpha = \int^{\mathbf{X}_{g.c.}} \Omega dt$, and Ω is the gyrofrequency at the guiding center of the particle.

We can now calculate $H_1(\mathbf{J}, \mathbf{n}, \omega)$ from Eq. (4.41). Because of the assumptions above, Φ_0 is independent of the angle θ^1, and θ^3. Thus ,

$$H_1(\mathbf{J}, \mathbf{n}, \omega) = \frac{1}{(2\pi)^3} \sum_{l=-\infty}^{\infty} \int \Phi_0(\theta^2, \mathbf{J}, \omega) e^{i(\int^{X_{g.c}} dX' \cdot k_X(X') + n_\phi \phi + \phi_0)} J_l(k_\perp \rho) e^{il\alpha} e^{-i\mathbf{n}\cdot\theta} d^3\theta, \tag{4.52}$$

where the Fourier transform in time has already been performed. The integrations over θ^1 and θ^3 can be written as

$$\begin{aligned} H_1(\mathbf{J}, \mathbf{n}, \omega) = \frac{1}{(2\pi)^3} \sum_{l=-\infty}^{\infty} \int d\theta^2 \Phi_0(\theta^2, \mathbf{J}, \omega) e^{i(\int^{X_{g.c}} dX' \cdot k_X(X') + \phi_0)} J_l(k_\perp \rho) e^{-in^2\theta^2} \\ \times \int d\theta^1 e^{il\alpha - n^1\theta^1} \int d\theta^3 e^{in_\phi \phi - n^3\theta^3} \end{aligned} \tag{4.53}$$

Note that α and ϕ can be written as

$$\alpha = \int^t \Omega dt = \Omega^1 t + \int^t (\Omega - \Omega^1) dt = \theta^1 + f(\theta^2, ...) \tag{4.54}$$

$$\phi = \int^t \dot{\phi} dt = \Omega^3 t + \int^t (\dot{\phi} - \Omega^3) dt = \theta^3 + g(\theta^2, ...), \tag{4.55}$$

so that the integrals over θ^1 and θ^3 can be simplified. This has interesting implications for a wave with one value for n^1, n^3, and ω. It means that $H_1(\mathbf{J}, \theta, t) \propto F(\mathbf{J}, \theta^2) e^{-i(n^1\theta^1 + n^3\theta^3 - \omega t)}$. Returning to the Hamiltonian, by dividing the time derivative of J^i by the time derivative of H, it is now easy to show relationship

$$\frac{dJ^i}{dH} = -\frac{\partial H_1 / \partial \theta^i}{\partial H_1 / \partial t}. \tag{4.56}$$

Because of the simple form of H_1 for the case of a wave with one value for n^1, n^3, and ω, it is easy to show

$$\frac{d\mu}{d\varepsilon} = \frac{n^1}{\omega} \tag{4.57}$$

$$\frac{dP_\phi}{d\varepsilon} = \frac{n^3}{\omega}, \tag{4.58}$$

which can be trivially integrated. Eqs. (4.57) and (4.58) describe how a particle will stay

on a line in ε, μ, and P_ϕ space when interacting with a single wave!

Also, note that the remaining integral over θ^2 can be converted to an integral over time (since $d\theta^2/dt = \omega_b$), leading to an expression for H_1 which is:

$$H_1(\mathbf{J}, n, n^2, n_\phi, \omega) = \frac{1}{\tau_b} \sum_{n=-\infty}^{\infty} \int dt \Phi_0(\theta^2, \mathbf{J}, \omega) J_n(k_\perp \rho) \times e^{i\left(\int^{X_{g.c}} dX' \cdot k_X(X') + n \int^t \Omega dt + n_\phi \int^t \dot{\phi} dt - n\Omega^1 t - n^2\theta^2 - n_\phi \Omega^3 t + \phi_0\right)} \tag{4.59}$$

Note from Eq. (4.40), that D is proportional to $\delta(\omega - \mathbf{n} \cdot \Omega)$, thus everywhere in D we can replace $(n^1\Omega^1 + n^2\Omega^2 + n^3\Omega^3)$ with ω, including in the expression for H_1. We combine all of the terms in the phase to get:

$$H_1(\mathbf{J}, n, n^2, n_\phi, \omega) = \frac{1}{\tau_b} \sum_{n=-\infty}^{\infty} \int_0^{\tau_b} dt \Phi_0(\theta^2, \mathbf{J}, \omega) J_n(k_\perp \rho) e^{i\phi_n(t)} \tag{4.60}$$

$$\frac{d\phi_n}{dt} = k_X(X_{g.c}) v_{X_{g.c.X}} + n\Omega + n_\phi \dot{\phi} - \omega \tag{4.61}$$

$$\frac{d\phi_n}{dt} \approx k_\parallel v_\parallel + n\Omega - \omega \tag{4.62}$$

$$\frac{d^2\phi_n}{dt^2} \approx \mathbf{v}_{g.c.} \cdot \boldsymbol{\nabla}(k_\parallel v_\parallel + n\Omega - \omega) = \frac{v_\parallel}{JB} \frac{\partial}{\partial \theta}(k_\parallel v_\parallel + n\Omega), \tag{4.63}$$

where the drifts are neglected compared to the parallel motion. Note that $|\omega \tau_b H_1|$ is the magnitude of the energy change, $\Delta\varepsilon$, the particle would receive on passing through the wave region (which we will return to later).

4.4.3 Stochasticity

An additional, subtlety involves the summation of Eq. (4.39). This sum has already been collapsed, because we know that $n^1 = n$ the harmonic at which the particle is resonating. Also $n^3 = n_\phi$. What remains is just calculating the sum over n^2 and dealing with the δ function. The δ function arises from the quasilinear theory which assumes that particles are on unperturbed orbits, and therefore a particle only diffuses if it exactly satisfies resonance. There are two distinct ways in which the resonance is broadened

physically. Extrinsic broadening could occur due to collisions or a varying wave state due to a turbulent plasma, or possibly even noise intentionally induced at the antenna. Intrinsic broadening could arise from the nonlinear change in the particles motion due to the kicks themselves. Here we estimate the amplitude of the kicks required to produce a broadening of the resonance. Equivalently, how large must the nonlinear islands around each resonance be in order to overlap each other, and thereby cause the motion of the particles to be stochastic?

Let us determine the wave field necessary for two adjacent resonances to overlap. To begin with, note that points in **J** space can be mapped to ε, μ, and P_ϕ space (really we only need $\varepsilon(\mu, J^2, P_\phi)$ since $J^1 = \mu$ and $J^3 = P_\phi$). Let **I** denote a point in ε, μ, and P_ϕ space. Then, in this space, consider two points **I** and $\mathbf{I} + \Delta\mathbf{I}$, which satisfy resonance, i.e.,

$$\omega = n\Omega^1(\mathbf{I}) + n^2\omega_b(\mathbf{I}) + n_\phi\omega_d(\mathbf{I}) \tag{4.64}$$

$$\omega = n\Omega^1(\mathbf{I} + \Delta\mathbf{I}) + (n^2+1)\omega_b(\mathbf{I} + \Delta\mathbf{I}) + n_\phi\omega_d(\mathbf{I} + \Delta\mathbf{I}). \tag{4.65}$$

Subtracting these two equations, and expanding to first order in $\Delta\mathbf{I}$, we find

$$\Delta\mathbf{I} \cdot \boldsymbol{\nabla}_{\mathbf{I}}(n\Omega^1 + n^2\omega_b + n_\phi\omega_d) = -\omega_b. \tag{4.66}$$

As discussed above the diffusion path for particles interacting with a single wave must occur along a line in ε, μ, and P_ϕ space. Thus, the $\Delta\mathbf{I}$ satisfy:

$$\Delta\mathbf{I} = \begin{pmatrix} 1 \\ \frac{n}{\omega} \\ \frac{n_\phi}{\omega} \end{pmatrix} \Delta\varepsilon, \tag{4.67}$$

so that Eq. (4.66) can be rewritten as

$$\Delta\varepsilon_s = \frac{-\omega_b}{L(n\Omega^1 + n^2\omega_b + n_\phi\omega_d)} \tag{4.68}$$

$$L = \frac{\partial}{\partial\varepsilon} + \frac{n}{\omega}\frac{\partial}{\partial\mu} + \frac{n_\varphi}{\omega}\frac{\partial}{\partial P_\phi}. \tag{4.69}$$

Now the islands will overlap, i.e. the δ function will be broadened if $\Delta\varepsilon > \Delta\varepsilon_s$, where

$\Delta\varepsilon$ is a typical kick given by the MCIBW. For illustrative purposes, consider a passing particle in a high aspect ratio tokamak. Then, converting back to real units for a moment, it is easy to see that $\Omega^1 = \Omega_0$, $\omega_b = |v_\parallel|/(qR_0)$, and $\omega_d = v_\parallel/R_0$. For simplicity, we assume q is constant. Evaluating the resonance condition and the derivatives specified above, and noting that since q is taken constant the derivative with respect to P_ϕ vanishes, we get (again with physical units for clarity):

$$\omega = n\Omega_0 + |v_\parallel|\frac{n^2 + n_\phi \sigma q}{qR_0} \tag{4.70}$$

$$L(n\Omega^1 + n^2\omega_b + n_\phi\omega_d) = \left(1 - \frac{n\Omega_0}{\omega}\right)\frac{n^2 + n_\phi\sigma q}{q|v_\parallel|R_0}, \tag{4.71}$$

where σ is the sign of the parallel velocity. Expressing the resonance condition in a more convenient form,

$$\frac{\omega - n\Omega_0}{{v_\parallel}^2} = \frac{n^2 + n_\phi\sigma q}{q|v_\parallel|R_0}, \tag{4.72}$$

and substituting into Eq. (4.68), we find:

$$\Delta\varepsilon_s = \frac{2\lambda^2\varepsilon\omega_b}{\left(1 - \frac{n\Omega_0}{\omega}\right)^2\omega}, \tag{4.73}$$

where λ is the pitch of the particle. A similar analysis for deeply trapped particles at high aspect ratio (which have $\Omega^1 = \Omega_0$, $\omega_b = \sqrt{\epsilon\varepsilon}/(qR_0)$, and $\omega_d = q\varepsilon/(\epsilon R_0)$) gives

$$\Delta\varepsilon_{s\ \mathrm{trapped}} = \frac{2\varepsilon\omega_b}{\left(1 - \frac{n\Omega_0}{\omega}\right)\omega} = \frac{1}{\lambda^3}\left(1 - \frac{n\Omega_0}{\omega}\right)\sqrt{\frac{\epsilon}{2}}\Delta\varepsilon_{s\ \mathrm{passing}}, \tag{4.74}$$

where here λ is the pitch of the passing particle.

For stochastic threshold we require that kicks given to the particle by the wave be greater than $\Delta\varepsilon_s$. While this is only a crude calculation, it shows that the stochastic threshold is increasing quickly with energy ($\omega_b \propto \sqrt{\varepsilon}$, implying $\Delta\varepsilon_s \propto \varepsilon^{3/2}$ for both passing and trapped particles). Furthermore, the kicks given to the particle (as we shall see later) decrease with increasing particle energy. Thus, for some high enough value of the energy, $\Delta\varepsilon_{\mathrm{MCIBW}} < \Delta\varepsilon_s$, and the motion will no longer be stochastic. Note that

while trapped particles also have a stochastic barrier at high energy (in the high aspect ratio limit), their stochastic threshold is typically much less than the threshold for passing particles and thus they will be stochastic to much higher energies than passing particles.

This barrier at high energy was discussed in the context of particles in a mirror machine by Rosenbluth (1972) and is called superadiabaticity (even though the particles are in the presence of a perturbation with frequency large enough to break their adiabatic invariant, the invariant still holds). Note that the island widths calculated above included only one toroidal mode number and adding more mode numbers should decrease the spacing between resonances, thereby decreasing the threshold for stochastic overlap, conversely increasing the energy at which the particles are superadiabatic. For now, we consider energies low enough that the islands overlap and therefore the sum of n^2 in Eq. (4.39) can be formed yielding just $1/\omega_b$. In practice, some advantage to α-channelling might occur if the particles were superadiabatic at high energies; this is discussed in Chapter 8.

4.4.4 Quasilinear diffusion coefficient

Returning to the diffusion equation, for convenience, we substitute $\Delta\varepsilon/(\omega\tau_b)$ for H_1 . Then, D, the QL diffusion coefficient is

$$D = \frac{\pi}{\omega_b} \sum_{n,n_\phi,\omega} \frac{|\Delta\varepsilon|^2}{\omega^2 \tau_b^2} \begin{pmatrix} n \times n & \frac{n\left(\omega - \Omega^1\, n - n_\phi\, \omega_d\right)}{\omega_b} & n\, n_\phi \\ \frac{n\left(\omega - \Omega^1\, n - n_\phi\, \omega_d\right)}{\omega_b} & \frac{\left(-\omega + \Omega^1\, n + n_\phi\, \omega_d\right)^2}{{\omega_b}^2} & \frac{n_\phi\left(\omega - \Omega^1\, n - n_\phi\, \omega_d\right)}{\omega_b} \\ n\, n_\phi & \frac{n_\phi\left(\omega - \Omega^1\, n - n_\phi\, \omega_d\right)}{\omega_b} & {n_\phi}^2 \end{pmatrix} \tag{4.75}$$

where, as in Eriksson and Helander (1994), the δ function has been used to eliminate n^2 everywhere and the sum over n^2 has been performed. Now that we have D^{ij} for the action-angle variables we must transform it to ε, μ, and P_ϕ space. Following Eriksson and Helander (1994), Eq. (4.39) is transformed to

$$\frac{\partial f}{\partial t} = \frac{1}{\sqrt{g}} \frac{\partial}{\partial I^i} \sqrt{g}\, \bar{D}^{ij} \frac{\partial f}{\partial I^j} \tag{4.76}$$

$$\bar{D}^{ij} = \frac{\partial I^i}{\partial J^k} \frac{\partial I^j}{\partial J^l} D^{kl} = G D G^T, \tag{4.77}$$

where $\sqrt{g}$ is the Jacobian of the mapping from I to J which is $1/\omega_b$ for $I^i = (\varepsilon, \mu, P_\phi)$, and

$$G = \frac{\partial I^i}{\partial J^k} = \begin{pmatrix} \Omega^1 & \omega_b & \omega_d \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{4.78}$$

Note that at this point a significant simplification of the above formula takes place,

$$GDG^T = \begin{pmatrix} \omega^2 & n\,\omega & \omega\, n_\phi \\ n\,\omega & n \times n & n\, n_\phi \\ \omega\, n_\phi & n\, n_\phi & {n_\phi}^2 \end{pmatrix} \tag{4.79}$$

So that, in the end the equation takes on the remarkably simple form:

$$\frac{\partial f}{\partial t} = \frac{1}{\tau_b} \sum_{n, n_\phi, \omega} L \frac{|\Delta\varepsilon|^2}{2} L f, \tag{4.80}$$

where L is defined in Eq. (4.69) above. The fact that only the operator L enters into the equations as above arises from the fact that a single wave diffuses a particle along a line in ε, μ, and P_ϕ space. Equivalently the microscopic kicks satisfy Eqs. (4.57) and (4.58).

Now we can calculate the stochastic and deterministic components (diffusive and drift terms) in much the same way as was done for collisions. We get a drift term and a diffusive term as in Sec. 4.2. For ease of presentation define $\beta^\varepsilon = 1, \beta^\mu = n/\omega, \beta^{P_\phi} = n_\phi/\omega$ then

$$h_i = \sum_{n, n_\phi, \omega} \frac{1}{\tau_b} L \frac{|\Delta\varepsilon|^2}{2} \beta^i \tag{4.81}$$

$$\frac{d\sigma_{ij}}{dt} = \sum_{n, n_\phi, \omega} \frac{|\Delta\varepsilon(n, n_\phi, \omega)|^2}{\tau_b} \beta^i \beta^j \tag{4.82}$$

$$g = \begin{pmatrix} \sqrt{\dot\sigma_{\varepsilon\varepsilon}} & 0 & 0 \\ \frac{\dot\sigma_{\varepsilon\mu}}{\sqrt{\dot\sigma_{\varepsilon\varepsilon}}} & \sqrt{-\frac{\dot\sigma^2_{\varepsilon\mu}}{\dot\sigma_{\varepsilon\varepsilon}} + \dot\sigma_{\mu\mu}} & 0 \\ \frac{\dot\sigma_{\varepsilon P_\phi}}{\sqrt{\dot\sigma_{\varepsilon\varepsilon}}} & \frac{-\frac{\dot\sigma_{\varepsilon P_\phi}\dot\sigma_{\varepsilon\mu}}{\dot\sigma_{\varepsilon\varepsilon}} + \dot\sigma_{\mu P_\phi}}{\sqrt{-\frac{\dot\sigma^2_{\varepsilon\mu}}{\dot\sigma_{\varepsilon\varepsilon}} + \dot\sigma_{\mu\mu}}} & \sqrt{\dot\sigma_{P_\phi P_\phi} + \frac{-2\,\dot\sigma_{\varepsilon P_\phi}\,\dot\sigma_{\varepsilon\mu}\,\dot\sigma_{\mu P_\phi} + \dot\sigma_{\varepsilon\varepsilon}\,\dot\sigma^2_{\mu P_\phi} + \dot\sigma^2_{\varepsilon P_\phi}\,\dot\sigma_{\mu\mu}}{\dot\sigma^2_{\varepsilon\mu} - \dot\sigma_{\varepsilon\varepsilon}\,\dot\sigma_{\mu\mu}}} \end{pmatrix} \tag{4.83}$$

For a single n, n_ϕ, ω the expression for g collapses to just the first column, signifying again diffusion along a line in ε, μ, and P_ϕ space, for multiple n, n_ϕ, or ω diffusion will take place in three dimensions. The expressions for a single wave are:

$$h_i = \frac{1}{\tau_b} L \frac{|\Delta\varepsilon|^2}{2} \beta^i \tag{4.84}$$

$$g = \begin{pmatrix} 1 & 0 & 0 \\ \frac{n}{\omega} & 0 & 0 \\ \frac{n_\phi}{\omega} & 0 & 0 \end{pmatrix} \frac{|\Delta\varepsilon|^2}{\tau_b} \tag{4.85}$$

4.4.5 Calculation of the diffusion coefficient for the MCIBW

We now calculate $\Delta\varepsilon$ for the MCIBW. This can be done using a stationary phase calculation to evaluate Eq. (4.60). For now we consider interaction with only one n, n_ϕ, ω. Assuming only a single n and ω is well justified for the case we are considering. Typically an energetic particle can only resonate with a low frequency wave ($\omega < 2\Omega_0$), at one value of n for an aspect ratio 3 reactor, where the magnetic field varies only by a factor of two across the plasma. For low aspect ratio tokamaks or high harmonic operation or both, ions can resonate with a single frequency wave at many different cyclotron harmonics across the plasma. The single frequency approximation is well satisfied as typically $\delta\omega/\omega < 1\%$. The n_ϕ spectrum in the plasma is determined by the spectrum of the antenna at the wall, and it typically is not very narrow for present day experiments.

The stationary phase calculation will tell us how much of a kick the particle receive when $\omega - k_\parallel v_\parallel - n\Omega = 0$. While in an infinite homogeneous plasma the particle's time in resonance would only be limited by nonlinear effects, in a tokamak there are a number of factors which can cause a resonant particle to dephase. For example, in a tokamak $\Omega \propto B$ and thus as a particle moves across the field, it slips out of phase with the wave because of its changing gyrofrequency. Another dephasing source, which typically is not important, is the variation of $k_\parallel v_\parallel$ due to the dependence of $v_\parallel$ on B. Finally, for the MCIBW, k_x [VALEO and FISCH 1994] is varying rapidly as a function of position at the ion-ion hybrid layer, and, as we have seen (Eq. (4.47)), this has a component parallel to the field. The rapid variation of k_x leads to a rapid variation of $k_\parallel$ [VALEO and FISCH

1995], and this in turn limits the amount of time the particle is resonant. Considering this, the value for $\Delta\varepsilon$ (assuming the particle satisfies the resonance condition) is

$$\Delta\varepsilon = q\omega\Phi_0(\mathbf{r}_{\mathbf{g.c.}})e^{i\varphi_0}J_n(k_\perp\rho)\sqrt{\frac{2\pi}{\frac{v_\parallel}{JB}\frac{\partial}{\partial\theta}(k_\parallel v_\parallel + n\Omega)}}, \tag{4.86}$$

where φ_0 is the phase of the interaction and all quantities are to be evaluated at the location where the resonance condition is satisfied.

4.4.6 MCIBW benchmark

While the formulation of the MCIBW particle interaction appears similar to that of collisions, there are also a number of subtleties which arise in its implementation. It is important to identify and benchmark each simple limit of the complete problem.

In the case of one wave, it is easy to check that the particle moves along the correct diffusion paths, since there are two conserved quantities. When more than one n_ϕ is present, the motion of the particle in (ε, P_ϕ) or (μ, P_ϕ) is no longer along a line, but is a random walk in two dimensions. The COM simulation has been tested for these conditions.

Certain subtleties arise when giving kicks to the particles. As outlined in Sec. 4.1, when a particle is given a kick, it is necessary to ensure that the particle has not been pushed outside the boundaries of ε, μ, and P_ϕ space. When the particle is interacting with the MCIBW it is also necessary to ensure that any kick given to the particle by the MCIBW does not push the particle out of interaction with the wave. That is, the particle must still be able to satisfy resonance with the wave after receiving a kick from the wave. This requirement arises from the microscopic reversibility of the motion (i.e. if time were turned backwards the particle would be able return to its previous orbit). This condition is probably not important for problems like ICRH, which have significant wave fields across much of the plasma. These waves are thus resonant with most parts of phase space, so that the likelihood of the particle getting kicked into a region of nonresonance is small. However, for the MCIBW, which is both spatially localized and requires significant Doppler shift for interaction, it is very important, because particles can be pushed

out of resonance or their orbits can be pushed so that they no longer cross the location of the MCIBW. This condition is implemented by checking that the resonance condition is still satisfied after every wave kick is given to the particle. Of course, in the presence of collisions, this will not necessarily be true, but then it is the collisions which cause the particle to get kicked out of resonance not the wave itself.

4.4.7 Collisionless diffusion

There are both diffusive terms and drift terms acting on the particle due to the wave. By diffusive, we mean terms in the Langevin equation for the change in the particles constants of motion which are multiplied by a stochastic random variable. The drift terms are deterministic. In terms of a one dimensional Fokker Planck equation:

$$\frac{\partial f}{\partial t} = -\frac{\partial}{\partial \varepsilon} A(\varepsilon) f + \frac{1}{2}\frac{\partial^2}{\partial \varepsilon^2} D(\varepsilon) f \tag{4.87}$$

Assuming the principle of dynamic equilibrium or microscopic reversibility of the kicks [ZASLAVSKY 1985], $A = 1/2\ dD/d\varepsilon$, and Eq. (4.87) can be rewritten as

$$\frac{\partial f}{\partial t} = \frac{1}{2}\frac{\partial}{\partial \varepsilon} D(\varepsilon) \frac{\partial f}{\partial \varepsilon} \tag{4.88}$$

Then the change of ε in some time Δt can be written as:

$$\Delta\varepsilon = \frac{\partial D}{\partial \varepsilon}\frac{\Delta t}{2} + \sqrt{D\Delta t}\zeta \tag{4.89}$$

where ζ is a Gaussian random variable with mean zero and variance one. The term proportional to A is a drift term and the term proportional to $\sqrt{D}$ is a diffusive term. From Eq. (4.88) the steady state solution of this equation with reflecting boundary conditions is clearly $f(\varepsilon)$ constant. It is interesting to see how the solution is consistent with a nonzero drift. Since $D(\varepsilon)$ is varying as a function of ε, the drift must be present to ensure that $f(\varepsilon)$ is a constant in steady state. In the absence of the drift term, one would expect that particles would spend the most time where D is small and little time where D is large, and thus the density would not be flat along this line in phase space. The drift term acts to exactly counteract this effect, ensuring that the density of the particles is constant along

the diffusion path. The energy dependence of the diffusion coefficient does not alter the steady state solution, rather the dependence of D on energy changes the time it takes for the distribution to relax to the steady state.

A very useful benchmarking tool is to consider a single initial condition interacting with a wave that has just one (n_ϕ, ω) and take the diffusion coefficient to have an analytical form. Interacting with just one n_ϕ and one ω means that the diffusion path is a line. If this line in ε, μ, and P_ϕ space (the diffusion path of the wave) intersects the wall, then the steady state solution of f along this path will be zero, i.e., eventually all of the particles will leave and hit the wall. If, on the other hand, the diffusion coefficient goes to zero somewhere along this line then the distribution function will flatten in the areas of nonzero D. If the diffusion coefficient takes on a simple form, the diffusion equation can be solved analytically, providing a benchmark for the code. We now formulate this one dimensional problem. Consider Eq. (4.80), recalling the form of L from Eq. (4.69). We want to make this explicitly one dimensional, by a change of coordinates. Consider the coordinate system

$$L = \frac{\partial}{\partial \varepsilon} + \frac{n}{\omega}\frac{\partial}{\partial \mu} + \frac{n_\varphi}{\omega}\frac{\partial}{\partial P_\phi} \tag{4.90}$$

$$u = P_\phi - \frac{n_\phi}{\omega}\varepsilon \tag{4.91}$$

$$v = \mu - \frac{n}{\omega}\varepsilon \tag{4.92}$$

$$w = \varepsilon, \tag{4.93}$$

clearly Lu=Lv= 0, while Lw=1, therefore we can substitute $L = \partial/\partial w$. Since we are considering one initial condition, it is easy to determine the boundary conditions. Below some energy, $\varepsilon_{\text{threshold}}$, the particle will not be able to resonate with the wave, while at some high energy, $\varepsilon_{\text{loss}}$, the particle will hit the wall. Finally, for the simplest comparison, we measure time in units of the bounce time, $dN = dt/\tau_b$. Then Eq. (4.80) can be rewritten

$$\frac{\partial f}{\partial N} = \frac{\partial}{\partial w} D(w) \frac{\partial f}{\partial w} \tag{4.94}$$

$$f(w_{\text{loss}}) = 0, f'(w_{\text{threshold}}) = 0. \tag{4.95}$$

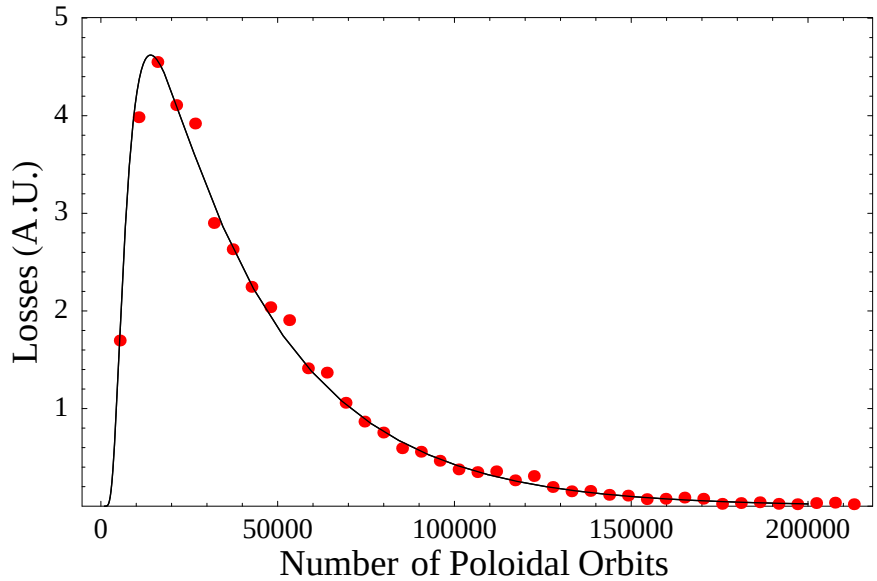

Figure 4.7: Comparison of an analytical calculation (line) and the COM simulation (points) results of the losses versus time for a single particle interacting with a single wave.

This equation is separable in the variables N and w. The second order differential equation in w is a Sturm-Liouville problem that can be expanded in eigenfunctions [RISKEN 1989b]. The steady state solution is $f(w) = 0$, but we are interested in comparing the time behavior of the losses to those of the code. For $D(w)$ constant, the solution is quite simple:

$$f(w,t) = \sum_{n=0}^{\infty} a_n \cos\left(\frac{(2n+1)\pi(w - w_{\text{threshold}})}{2(w_{\text{loss}} - w_{\text{threshold}})}\right) e^{-\frac{(2n+1)^2\pi^2 Dt}{4(w_{\text{loss}} - w_{\text{threshold}})^2}}, \tag{4.96}$$

where the a_n are determine by the initial conditions. Using this form we can calculate the fraction of particles lost by computing the flux through the boundary $w = w_{\text{loss}}$. This can be compared to the predictions of the COM simulation for the same configuration. In Fig. 4.7 the analytical flux is shown as a solid line. The points are obtained by binning the loss times of 10000 instances of a single particle from the COM simulation. The agreement is quite good, confirming that the code is operating correctly with a constant diffusion coefficient.

It is also useful to investigate a case where $D(w)$ is not constant, as such a case

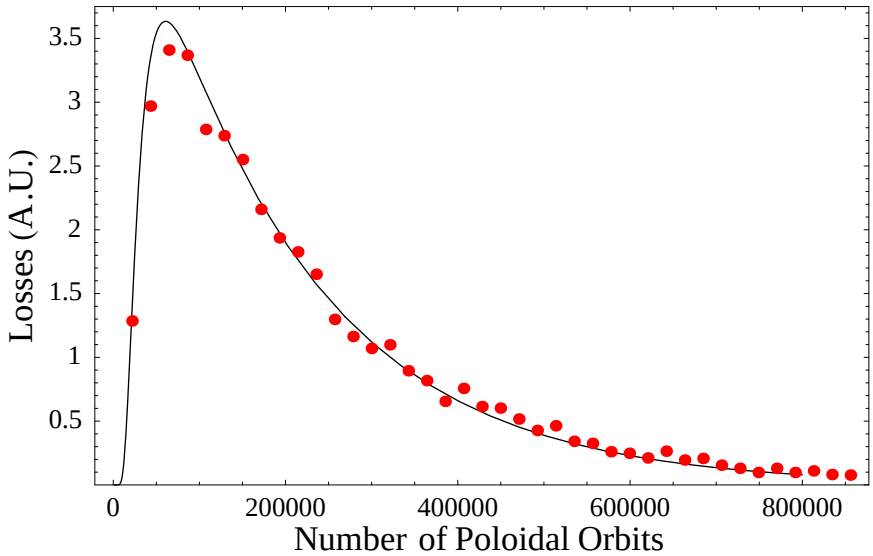

Figure 4.8: Comparison of an analytical calculation (line) and the COM simulation (points) results of the losses versus time for a single particle interacting with a single wave.

will have drift terms as well as diffusive terms. We take $D(w) \propto 1/w$, modeling a diffusion coefficient which might arise naturally if the particles interacted with the wave over a fixed distance, and therefore had a time in resonance inversely proportional to velocity. The eigenfunction expansion is now more complicated (it is in terms of Bessel functions, $J_{-2/3}, J_{2/3}$), but can be done easily enough by Mathematica [WOLFRAM 1996]. In Fig. 4.8 the losses versus time are plotted. Again the agreement is quite good. Note that the drift terms are important. In this case, the drift is in the negative w direction, acting to hold the particles in. (When the simulation is run with the drifts turned off, the mean exit time is 30% shorter than the exit time with the drifts present.)

While this validates that the code is working for simple forms of the diffusion coefficient, there is a more subtle problem which arises when using the calculated diffusion coefficient, Eq. (4.86), in the presence of several waves. This problem arises because particles can fall in or out of resonance with the waves, either by changing their velocity or changing their orbits, so that they no longer intersect the mode conversion layer of one of the waves. When this happens, the diffusion coefficient changes quickly from one time step to the next. In these cases, the drift term may not be treated correctly, unless considerable care is exercised.

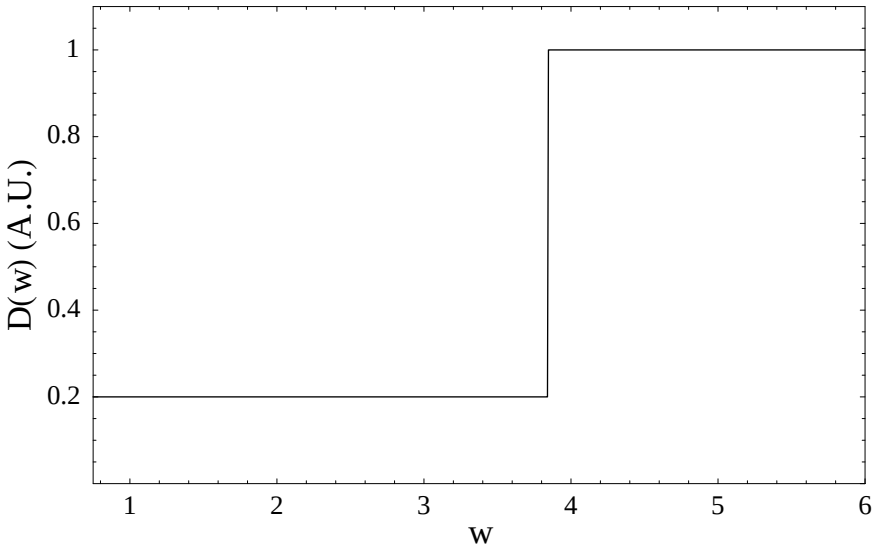

Figure 4.9: A diffusion coefficient with a large jump, as might arise when a particle starts resonating with more than one wave.

For illustration purposes consider a stepwise constant diffusion coefficient case. The diffusion coefficient versus w is shown in Fig. 4.9. To identify the problem such a D causes, it is useful to plot the steady state distribution function, thus we put an artificial barrier at $w = 6.0$ so that the particles do not hit the wall. We then plot the distribution function, which is obtained by sampling many instances of the same particle at different times and then binning the resulting distribution of w's.

First we consider the distribution function for constant D, shown in Fig. 4.10. Note that the distribution is quite flat, in accord with our expectation. The dark line in Fig. 4.11 shows the distribution function in steady state for the diffusion coefficient in Fig. 4.9. Obviously this is not a constant, and therefore not a valid solution. The origin of the discrepancy turns out to be that particles go from one region to the other without feeling the effects of the changing D arising from the sharp jump.

This discrepancy could be mitigated if this jump were smoothed (although this is somewhat complicated for the COM simulation since it is not always easy to determine where jumps will occur), and the time step were chosen small enough so that every particle would "feel" the jump as it stepped across it. The time step required can be determined by ensuring that a typical step is small compared to the spatial scale on which

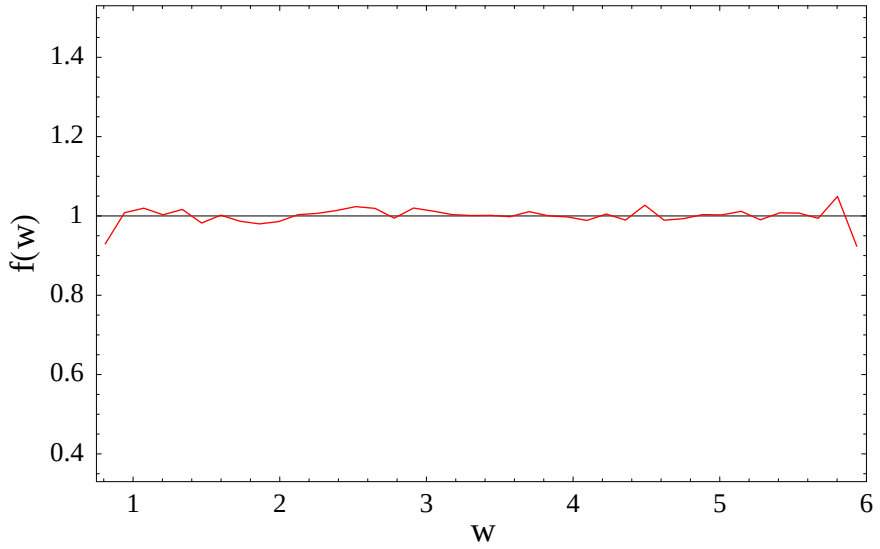

Figure 4.10: The distribution function for the case of constant diffusion coefficient.

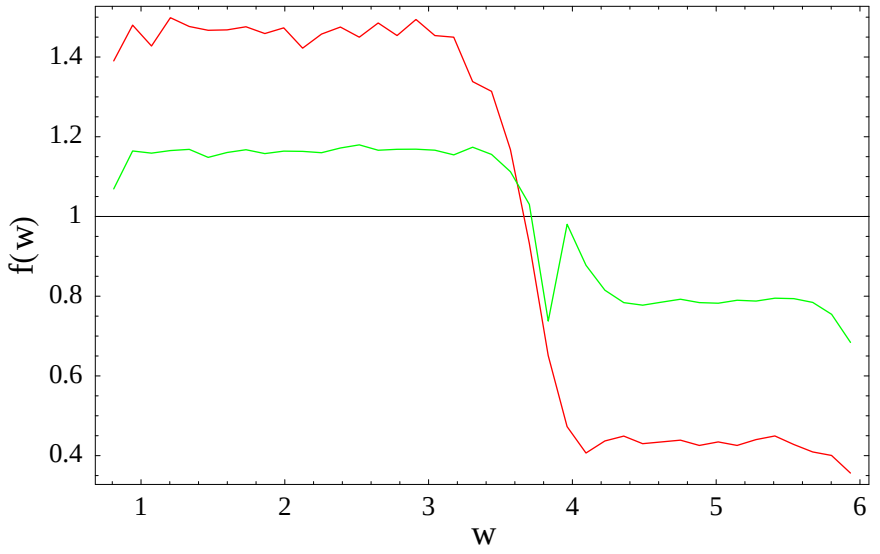

Figure 4.11: The distribution function obtained by the simulation when the diffusion coefficient is as shown in Fig. 4.9. The dark line is the unmodified code, the light line the distribution function when a suitable average of the diffusion coefficient is used.

D changes. Define

$$L_D = \frac{D}{\frac{\partial D}{\partial w}}. \tag{4.97}$$

Then we want

$$\Delta x \approx \sqrt{D\Delta t} < L_D. \tag{4.98}$$

This however can become computationally prohibitive, especially if L_D is small, since, if we are interested in diffusing things across a distance scale L_w (in our case the amount of energy the particle gains before hitting the wall), the number of time steps goes like $(L_w/L_D)^2$. Of course it would be possible to come up with an adaptive method which took small time steps only where D was changing rapidly.

An alternative is an implicit method for determining what kick to give the particle. That is to say, give the particle a trial kick, determine what the diffusion coefficient is where it ends up and then give the particle a kick based on the average of the diffusion coefficients where it starts and where it ends up. This would have to be iterated upon to ensure that the trial location is close to the actual location that the particle jumps to. While these iterations involve more computation, it is possible to take longer time steps than those used in the method above which explicitly uses the current location of the particle for the calculation of D and its derivative.

This method has been implemented approximately (with only one iteration) and benchmarked. The light curve in Fig 4.11 shows the resulting distribution function. Note that the curve is still not flat, but it is much closer than the explicit scheme previously used.

It is also possible to calculate analytically the losses vs. time for the diffusion coefficient shown in Fig. 4.9. The calculation is essentially the same as the calculation outlined earlier, except that the differential equation is solved separately in the different D regions and then the solution is matched at the point where D jumps, giving a condition for the eigenvalues.

Fig. 4.12 shows the result of our comparison between the COM simulation with

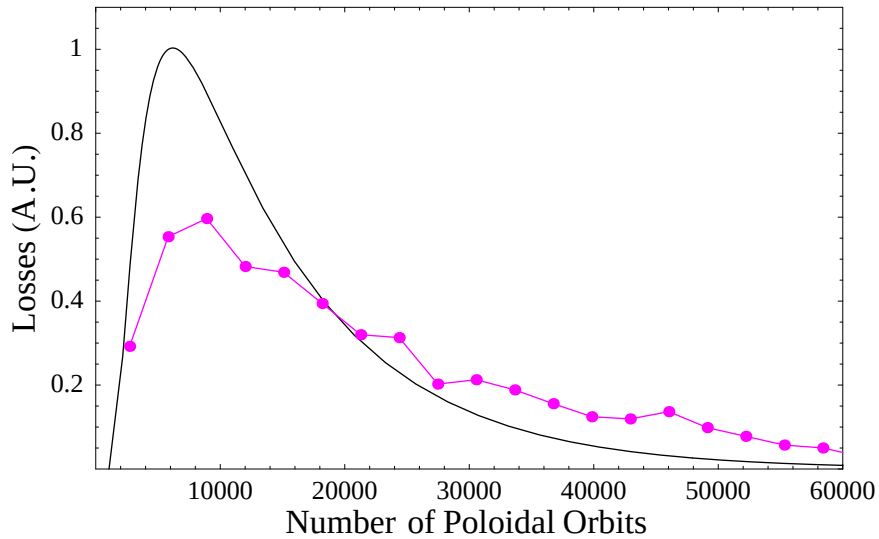

Figure 4.12: Comparison of an analytical calculation (line) and the COM simulation with the explicit scheme (points) for the case of D discontinuous.

the explicit scheme and the analytical calculation. For this case, it is clear that particles are getting "stuck", since the mean time for the particles to leave is almost twice the mean exit time of the analytical calculation. Furthermore there is a long tail on the simulation results making the standard deviation of the loss time about twice that of the analytical calculation. Note that with D being discontinuous, there is no way of taking small enough time steps to satisfy Eq. (4.98).

Fig. 4.13 shows the losses calculated from the approximately implicit method. In contrast, they agree rather well with the analytical calculation, except for the reduction in the peak. The mean and standard deviation of the loss time from the simulation is about 15 % larger than those of the analytical calculation. However, in view of the accuracy of other parts of the calculation such an approximation is within acceptable limits (more accuracy could be had using a fully implicit time step).

4.4.8 Calculation of diffusion due to toroidal Alfvén eigenmodes

The action-angle formalism is useful when investigating the interaction of α-particles with toroidal Alfvén eigenmodes, and their various relatives. For toroidal Alfvén eigen-

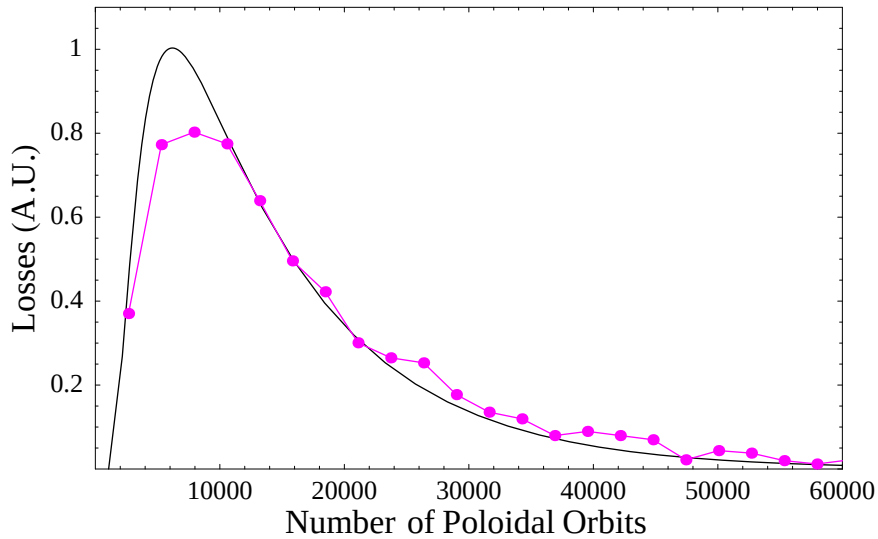

Figure 4.13: Comparison of an analytical calculation (line) and the COM simulation with the implicit scheme (points) for the case of D discontinuous.

modes, ω is much less than the gyrofrequency and therefore μ is conserved. Also, unlike the case above there is a discrete spectrum of waves. As we are interested in describing the motion with a diffusion equation, some care must be exercised since the discrete spectrum suggests that the quasilinear diffusion approach may be invalid.

In fact, this problem is much more than just conceptual. It is a problem for α-channelling. We are interested in using the toroidal Alfvén eigenmodes to transport the particles across much of the plasma, and thus it will be necessary for the islands associated with each mode to overlap throughout phase space. If this is accomplished then the motion in the wave will be diffusive. This requirement will give a power requirement on the toroidal Alfvén eigenmode independent of collisionality of the α-particles, since the island width depends on the mode amplitude and the amplitude depends on the power.

This problem of discrete, well-separated, resonances manifests itself in the formal quasilinear diffusion equation given in Eq. (4.39). Basically the δ function, which can be integrated over for the MCIBW case, must remain in the case of toroidal Alfvén eigenmodes, where there is no resonance overlap. Thus the general form of the quasilinear diffusion equation can be written as (using the form of D_{ql} in ε, μ, and P_ϕ space from

above):

$$\frac{\partial f}{\partial t} = \sum_{n_\phi} \frac{1}{\tau_b} L\pi \sum_{n^2} \omega_b^2 |H_1(\mathbf{J}, n^2, n_\phi, \omega)|^2 \delta(\omega - n^2\omega_b + n_\phi\omega_d) Lf \tag{4.99}$$

$$L = \frac{\partial}{\partial \varepsilon} + \frac{n_\varphi}{\omega}\frac{\partial}{\partial P_\phi} \tag{4.100}$$

This expression gives kicks only to particles which satisfy:

$$\omega - n^2\omega_b - n_\phi\omega_d = 0 \tag{4.101}$$

Unlike the MCIBW case, the steps(in ε, μ, and P_ϕ) between resonances can be large and thus the sum over n^2 (which was used to eliminate the δ function) cannot be performed in the case where stochastic overlap does not occur.

At this point two approaches can be taken, one is to assume that stochastic overlap will be reached via some external mechanism to get rid of the δ function. (This is essentially what was done in an early version of the code.) Alternatively, one can consider the effect of the naturally occurring resonance broadening terms. Berk et al. (1995) did just this in developing a “line-broadened” quasilinear theory. The basic idea is to broaden the resonance condition to some finite width, and allow diffusion to take place in that region. If two resonances overlap than diffusion will occur across those two resonances. Physically the broadening can be due to finite wave amplitude, finite growth rate and finite collisionality. For the case we are interested in, where we assume the mode is not growing, but rather the energy in the wave is being pumped in externally by an antenna, we assume that the first effect is dominant. It is then possible to rewrite the diffusion equation above as in Berk et al. (1996).

4.5 Collisional Evolution of the Beam Ion Distribution

As an example of a problem that would be quite time consuming using a guiding center code consider the collisional evolution of the distribution of neutral beam injected particles. The initial distribution of the particles, i.e. the orbits they ionize on, can be obtained

from a TRANSP run with special flags set [McCune 1997]. This distribution is then used in the COM simulation as the initial distribution of the beam particles. The density and temperature profiles from TRANSP, as well as the magnetic equilibrium based on a splined JSOLVER reconstruction of the TRANSP data is used. For this simulation 50000 particles were used. The runs shown here take about 1.2 hours of CPU time on a DEC Alpha workstation, and have a time step of about 0.5 milliseconds. The total simulation time is 0.1 seconds. Note though that particles are removed from the COM simulation as they drop below a threshold energy (in this case near 15 keV), thus the CPU time per step is reduced as time goes on, since fewer and fewer particles are simulated. Averaged over the length of the run, only 20000 particles were simulated.

A typical run for 0.1 seconds with 1000 deuterons in TFTR would take ORBIT about two to four hours on the CRAY C90, translating to 40 to 80 hours of CPU time for a simulation comparable to the one above. Of course if just collisions are being simulated the pitch angle scattering and slowing down times might be artificially increased in ORBIT. Even in the presence of ripple, if one is only interested in the number of particles lost, it is possible to throw out particles which enter in the stochastic ripple domain and the simulation can be speeded up considerably [White et al. 1996]. Furthermore the more physics added to the COM code the longer it will take, although generally speaking, calculating the bounce time dominates and thus the collisional and wave portions do not add that much time. Also note that the situation will be different for different particles and different machines, as ORBIT must resolve the particles' motion around the torus, and therefore it will take twice as much CPU time to simulate a 400 keV deuteron as it will to simulate a 100 keV deuteron in the same device. On the other hand, the COM simulation is constrained by the rate at which orbits are changing due to the perturbations, for the most part independent of the speed of the particles, but simulating a system with twice the collisionality may take twice as long in order for the COM simulation to be resolved. Finally it is important to note that the COM simulation takes significantly more overhead in its implementation than a guiding center code like ORBIT.

Figures 4.14, 4.15 and 4.16 show the evolution of the number of particles at a given energy, the distribution of particles' λ, and the density of the particles versus ψ at 20 millisecond intervals for TFTR shot 82601 which had only cogoing beam injection.

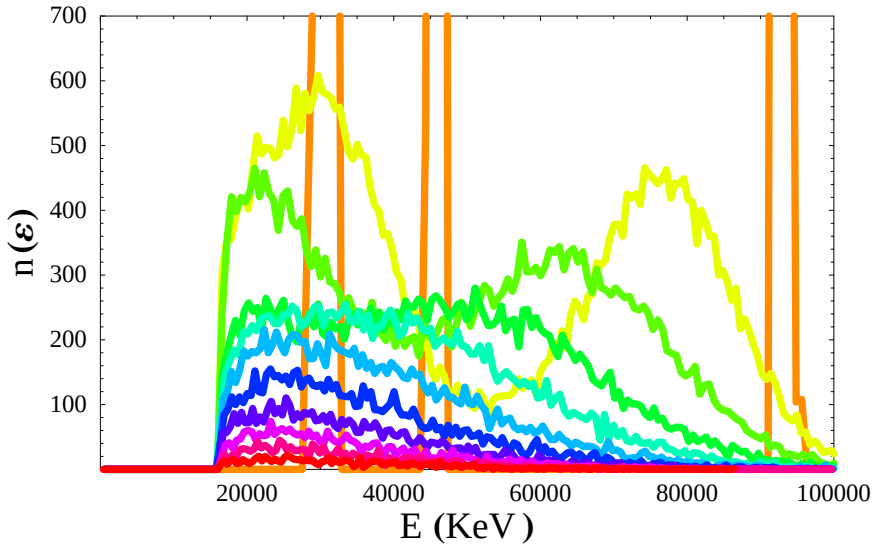

Figure 4.14: The density of particles vs. energy at 20 millisecond intervals. The initial peaks correspond to the full, $1/2$, and $1/3$ energy components of the beam. Time proceeds from light to dark. The total number of particles is decreasing since particles are absorbed when they get to some threshold energy, in this case 15 keV.

4.6 Timing

There are three easy ways to speed up the COM simulation significantly. The first is that the general geometry which is used expresses everything in terms of ψ and θ, however all of our calculations of orbit quantities specify ψ and B. Thus, if we want to know, say the R position of the particle at some ψ, we currently calculate B, (from B_{orb} (ψ)), then do a root finding to find θ which satisfies B (ψ,θ)=B_{orb} (ψ). This typically takes 10 calls to B (ψ,θ). In principle the `spline` code which is used to create the spline representation of the geometry could be modified to give θ as a function of ψ and B. Looking at code timing data this root finding and its associated function calls typically take up 60% of the CPU time, this suggests that the run time could be reduced by a factor of nearly two by implementing the more appropriate coordinate system.

The second thing which could be done is to calculate some quantities on a grid in ε, μ, and P_ϕ space before the run takes place and use look up tables and interpolations to evaluate these quantities, rather than calculate integrals over the orbit at each time step. This was suggested by Eriksson and Helander (1994) and it has been implemented in the

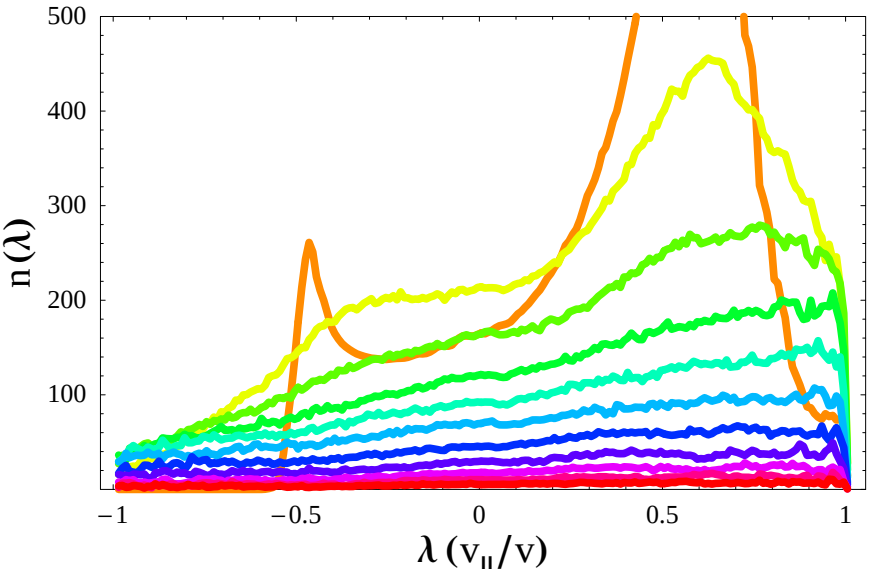

Figure 4.15: The density of particles vs. pitch at 20 millisecond intervals. The initial peaks arises because of coinjection of the beam particles.

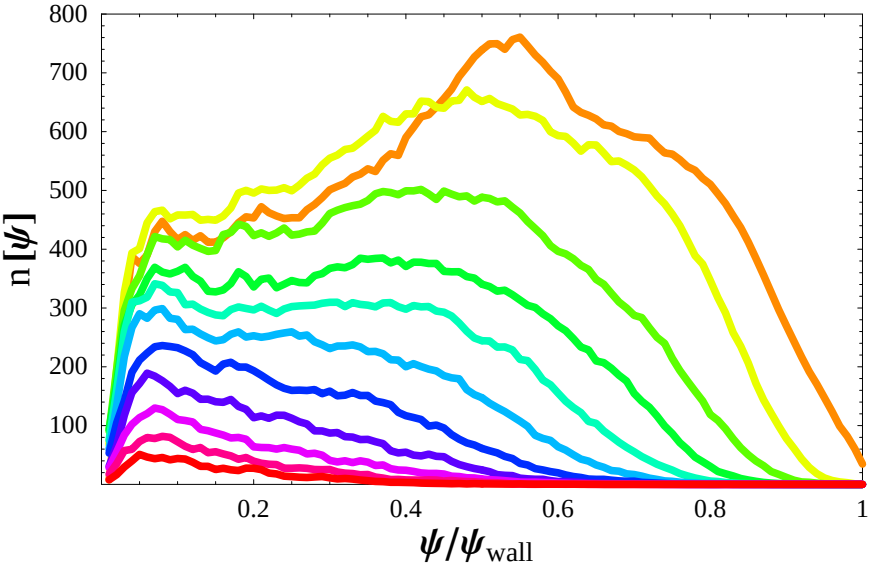

Figure 4.16: The density of particles vs. ψ at 20 millisecond intervals.

FIDO code by [CARLSSON, HELLSTEN, and ERIKSSON 1996; CARLSSON, ERIKSSON, and HELLSTEN 1994], who obtained an order of magnitude speed up.

Finally, this code has only been run on a Sun UltraSparc or Dec Alpha worksta-

tion up until now. The program should be easily parallelizable due to its nature (i.e. test particle codes should parallelize easily). Furthermore the code has been written without any machine specific routines, such as graphics or even numerical libraries, and thus should be easily portable to a massively parallel machine.

4.7 Summary

The general problem of a particle of arbitrary orbit width interacting with waves and collisions in a tokamak has been formulated and implemented in a Monte Carlo simulation. In the next chapter this simulation will be applied to the study of how α-channelling might possibly be implemented in a reactor. In Chapter 6 and Chapter 7 we will show how the simulation can also be used to examine existing experimental data and firmly establish key building blocks for the α-channelling effort.

Chapter 5

Reactor Scenario for α-Channelling with Two Waves

THE FIRST STEP TO ACHIEVING THE CHANNELLING effect is to demonstrate that waves can be amplified as they cool the α-particles, with the α-particles moving to the wall, where they are extracted at lower energy. This chapter addresses just this necessary first step, that is, how waves might control an entire birth distribution of α-particles.

Recall that the channelling effect is like shaking particles out of a bottle through certain holes. The 3D volume here is the energy (ε), the magnetic moment (μ), and the canonical angular momentum (P_ϕ) of the α-particles; the boundary of the bottle corresponds to values of these constants of the motion for orbits intersecting the physical boundary of the tokamak. Waves diffuse particles in this constants-of-motion space (ε-μ-P_ϕ space). The trick is to devise plasma waves that shake most of the α-particles into "holes" in the bottle at low energy. What is described in this chapter is how we used our numerical code to discover promising parameter regimes for α-channelling and the interesting features exhibited by collections of particles in response to the waves.

The chapter is organized as follows: In Sec. 5.1 we discuss the basics of the wave particle interaction. We show that, while a single wave will not be able to accomplish α-channelling, two waves acting in concert can extract significant energy from a single

particle while bringing it to the wall. In Sec. 5.2, we show that two waves can be employed to extract significant energy from an entire distribution of α-particles, not just a single particle. We discuss what configuration of waves is able to accomplish this and the key sensitivities of the energy extraction to variations in this configuration. In Sec. 5.4, we describe the distribution of the exiting particles. Sec. 5.5 highlights the importance of two assumptions that are taken for granted in this chapter. Sec. 5.6 summarizes this chapter.

5.1 Basics

In Chapter 3 we discussed the diffusion paths of α-particles interacting with waves in a simplified geometry. We hypothesized that a single wave alone was unable to accomplish the α-channelling effect in this simplified geometry, because of constraints on the slopes of the waves, and showed that two waves might be able to accomplish the channelling effect. In the toroidal geometry of the tokamak, one may wonder how much of the insight from the simplified geometry is still relevant.

Actually, the insights carry over to a surprising degree. We consider the problem in ε, μ, and P_ϕ space. For a particle interacting with a wave with toroidal mode number n_ϕ and absorbing energy $d\varepsilon$, the particle's P_ϕ changes by

$$dP_\phi/d\varepsilon = n_\phi/\omega. \tag{5.1}$$

For waves with $\omega \sim \Omega_{\alpha 0}$, an α-particle will receive a kick in velocity if $\omega - k_\parallel v_\parallel = n\Omega_\alpha$, where n is the harmonic number and Ω_α is the local α-particle cyclotron frequency, such that

$$d(\mu B_0)/d\varepsilon = n\Omega_{\alpha 0}/\omega \;\Rightarrow\; d\varepsilon_\perp/d\varepsilon = n\Omega_\alpha/\omega. \tag{5.2}$$

Waves with $\omega \ll \Omega_{\alpha 0}$ (e.g. the toroidal Alfvén eigenmode, or TAE), leave μ invariant. Note that throughout this thesis we use TAE as a generic term for any low frequency eigenmode.

While the amplitude of the kick depends on the details of the resonant interaction, the direction in ε-μ-P_ϕ-space is completely determined by Eqs. (5.1) and (5.2). Thus, in a toroidal geometry the problem becomes three dimensional with a complicated boundary, however the diffusion paths remain straight lines in this space, as can be seen from Eq. (5.1)-(5.2).

In analogy to the simplified problem, we can think of $-P_\phi$ as a radial variable (although this is only strictly true in the limit of ε going to zero), with the center being near P_ϕ = 0 and the wall at P_ϕ =$-\psi_{\text{wall}}$. Note that the same constraints on using one wave apply. In order to estimate n_ϕ, consider $\Delta P_\phi \sim -\psi_{\text{wall}}$ is needed to move an α-particle to the wall. Assuming the maximum change in energy $\Delta\varepsilon \sim -\varepsilon_0$, from Eq. (5.1),

$$n_\phi \approx \frac{\omega}{\Omega_{\alpha 0}} \frac{R_0^2}{\rho_0^2} \frac{\psi_{\text{wall}}}{B_0 R_0^2}, \tag{5.3}$$

in order to remove the α-particles from the reactor. For an equilibrium with large aspect ratio A, $\psi_{\text{wall}}/(B_0 R_0^2) \approx 2/(qA^2)$, implying $n_\phi \approx 1000$ to get the α-particles to the edge in a reactor sized tokamak with $R_0 \sim 5$ m, A = 3, $B_0 \sim 6$T and $q = 3$. Such a large n_ϕ is unachievable experimentally. Using only the TAE, which has $n_\phi \Omega_{\alpha 0}/\omega \sim (500-2000) n_\phi$, produces the opposite concern, i.e., α-particles ejected with little energy extracted.

For example, consider extraction by the MCIBW alone. The MCIBW breaks the μ invariant but is unable to drive the α-particle to the periphery if n_ϕ is too small. In Fig. 5.1, we show the orbit of an α-particle interacting with an IBW wave with $n_\phi = 50$ in a simple TFTR like geometry. The outer orbit is the largest orbit the particle can sample and still be resonant with the wave. In this orbit, the α-particle has lost about 2/3 of its initial energy and essentially all of its perpendicular energy. At this point, constrained through Eq.(5.1), the α-particle can lose no more energy and so cannot quite make it to the periphery. The α-particle could be extracted either in a lower current equilibrium (which reduces ψ_{wall}) or in the presence of a higher mode number IBW. However, in practice, lower current equilibria are not reactor relevant and higher mode numbers in a TFTR-size tokamak would be difficult to excite efficiently.

Therefore, our investigations concentrated on using two waves in concert. Al-

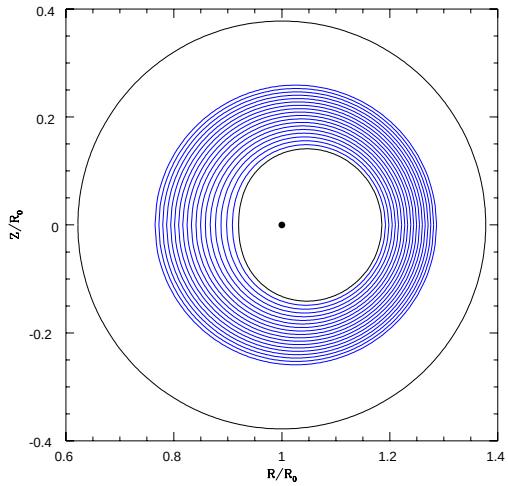

Figure 5.1: Poloidal cross section of orbit of a cogoing α-particle interacting solely with a Mode Converted Ion Bernstein Wave (MCIBW). If the particle were to lose energy under interaction with this wave it would move from the central orbit to the outermost orbit, eventually losing 2.5 MeV.

most all the energy can be extracted from a single α-particle through the use of two waves [FISCH and HERRMANN 1995], one with $\omega \ll \Omega_{\alpha 0}$ (the α-particle cyclotron frequency on axis) such as the toroidal Alfvén eigenmode [CHENG et al. 1985; CHENG and CHANCE 1986] and its relatives, and one with $\omega \sim \Omega_{\alpha 0}$ such as the mode converted ion Bernstein wave (MCIBW) [PARK et al. 1985] which has $1 < \Omega_\alpha/\omega < 3/2$ in deuterium-tritium plasmas. The high frequency wave is able to break the adiabatic invariance of μ, and thereby diffuse the particle in perpendicular energy. The low frequency wave breaks the P_ϕ invariant, diffusing the α-particles radially. Ideally, the combination will diffuse the particle to the tokamak periphery extracting both parallel and perpendicular energy. However, unlike the case of one wave only [FISCH and RAX 1992a], with two waves, there are no constraints on the particle motion, so that some α-particles may be heated while others are cooled.

Let us see what happens to the same particle interacting with a MCIBW as in Fig. 5.1, when a toroidal Alfvén eigenmode is also introduced. The succession of orbits in Fig. 5.1 assumes that between orbits the α-particle has lost energy to the wave ac-

cording to Eqs. (5.1) and (5.2). The result is that the α-particle cannot reach the plasma periphery by interacting only with this wave. Now consider instead the interaction of this α-particle with only a TAE-like mode; in Fig. 5.2 such an interaction is depicted, where, upon repeated interactions, the α-particle orbit can explore the set of orbits shown. The α-particle begins near the plasma center at 3.5 MeV; it can then explore the outer orbits, with the outermost orbit at 2.7 MeV. None of these orbits, however, intersects the plasma boundary, so the α-particle remains bracketed in both position and energy.

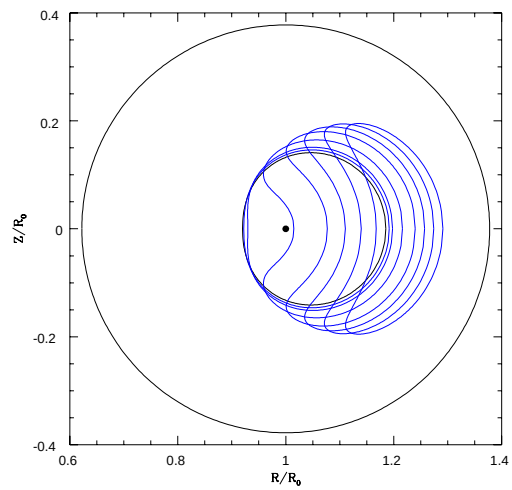

Figure 5.2: Poloidal cross section of cogoing α-particle interacting solely with a TAE like mode. Parameters as in Fig 5.1. If the particle were to lose 0.8 MeV under interaction with this wave it would move from the central orbit to the outermost trapped orbit. The particle cannot proceed further because it would no longer be resonant with the wave.

The simultaneous presence of both of these waves, however, leads to different behavior. In Fig. 5.3, we depict the α-particle succeeding in reaching the plasma periphery, if both waves are present, and if the α-particle always loses energy upon resonating with either wave. In this case, where all kicks are to lower energy, the α-particle loses 2.3 MeV to the waves before exiting at the plasma periphery.

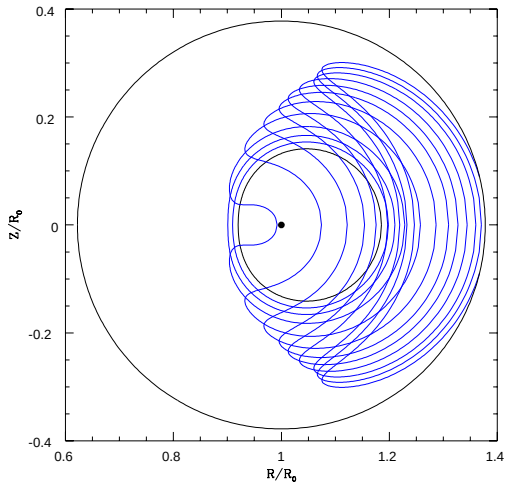

Figure 5.3: Same as Fig. 5.2 except the α-particle interacts with both the MCIBW and a TAE like mode, exiting at the periphery with 1.2 MeV left.

5.2 Cooling Scenario with Two Waves

Of course, a set of waves that cools one particle may be ineffective on another particle, or, worse yet, tend to heat other particles. How can we find an effective set of waves the acts on the complete birth distribution? This search is expedited by first developing a number of insights and and by using our quick turn-around simulations.

To develop these insights, we begin by demonstrating that a good, if not optimal solution exists. Then we will explain the reasoning that guided us to this solution, as well as how better solutions can be found by "perturbing" around this solution.

Consider the following promising case for a reverse shear tokamak reactor with $A = 3$, $R_0 = 5.4$ m, $B_0 = 6$ T, and $I_p = 16.3$ MA (this is a design which was consider by the ARIES-RS team). In this "advanced" reactor, 70% of the energy of the ejected α-particles (73% of those born) is diverted to waves, corresponding to 51% of the α-particle power if we use both MCIBW and toroidal Alfvén eigenmode. The location of the waves and other wave parameters are discussed below. In Fig. 5.4 and Fig. 5.5, the birth locations of 1000 3.5 MeV α-particles are shown in a fixed-energy slice of

constants-of-motion space (see Appendix A for a discussion of this space). Those that eventually reach the tokamak periphery are color-coded to show the total energy exchange with each wave. Particles remaining in the tokamak are shown in solid black. The IBW (Fig. 5.4) extracts the most energy from those particles which have significant amounts of perpendicular energy. In contrast, the TAE (Fig. 5.5), which must conserve μ, extracts the most energy from those particles which have the most parallel energy. In this case, only 0.2% of the α-particles were heated while being extracted.

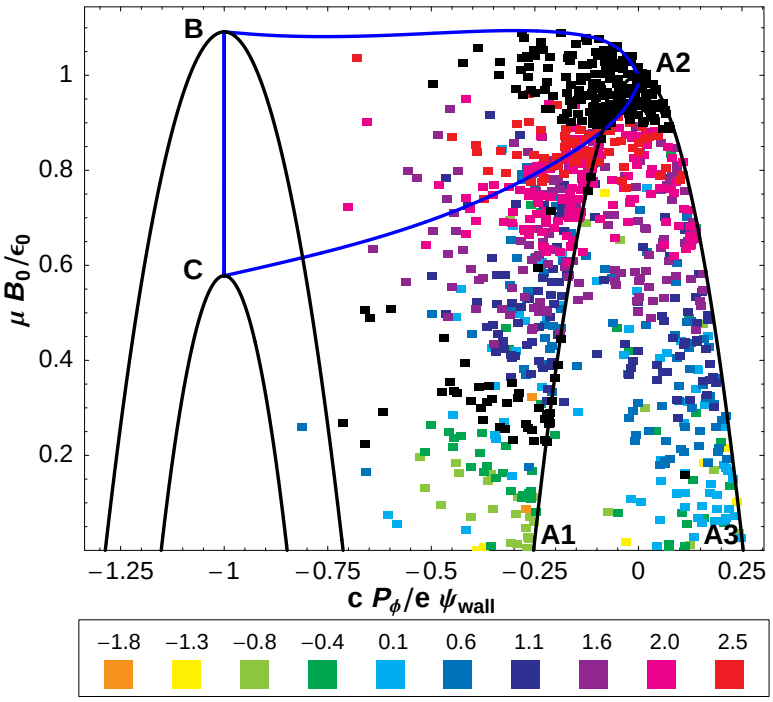

Figure 5.4: Energy extracted (MeV) by the IBW vs. initial location of particle in constants-of-motion space. The IBW extracts an average of 1.14 MeV per ejected α-particle.

Of the 51% of the α-particle power extracted in this example, 28% goes into the TAE and 23% into the IBW. This power flow might, in fact, sustain the wave amplitudes necessary to cool the α-particles in a time short compared to the slowing down time as discussed later

There are a number of idealizations in these simulations. Collisions are not present, and the amplitudes of the kicks the particle receives from the waves is determined as an input parameter, rather than from a diffusion coefficient (these effects, which

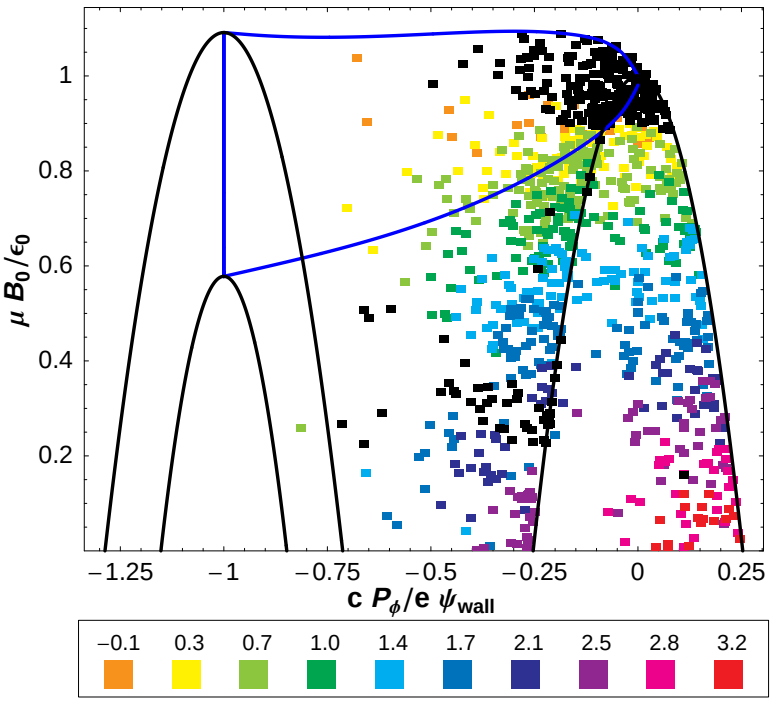

Figure 5.5: Energy extracted (MeV) by the TAE-like mode vs. initial location of particle in constants-of-motion space. The TAE extracts an average of 1.30 MeV per ejected α-particle.

are discussed in Chapter 4, were later added to the code and are included in the simulation of the experiments discussed in Chapter 6 and 7). The $k_\parallel$ spectrum of the MCIBW was taken from $-n_\phi/R_0$ to n_ϕ/R_0, which was based on ray tracing calculations done by Valeo and Fisch (1994). That $k_\parallel$ can be opposite in sign to n_ϕ [VALEO and FISCH 1994], the so-called "$k_\parallel$ flip," is important because cogoing α-particles then satisfy a resonance condition with the IBW that is correctly phased for energy extraction (see Sec. 5.5). For the TAE mode, it is assumed that so many modes are present within a certain radius, that all of the particles within that radius are diffused, eliminating the need for selecting only certain particles through a resonance condition. In addition to specifying where the mode exists, the value of n_ϕ/ω for the mode is specified. Although these assumptions are crude, the diffusion paths used are *precisely* correct for both the MCIBW and the toroidal Alfvén eigenmode. The location of the waves which accomplished this effect are shown in Fig 5.6. Where the light grey is the area over which the AE-like mode is assumed and the dark grey area labeled 1 very near the edge is the region of the MCIBW.

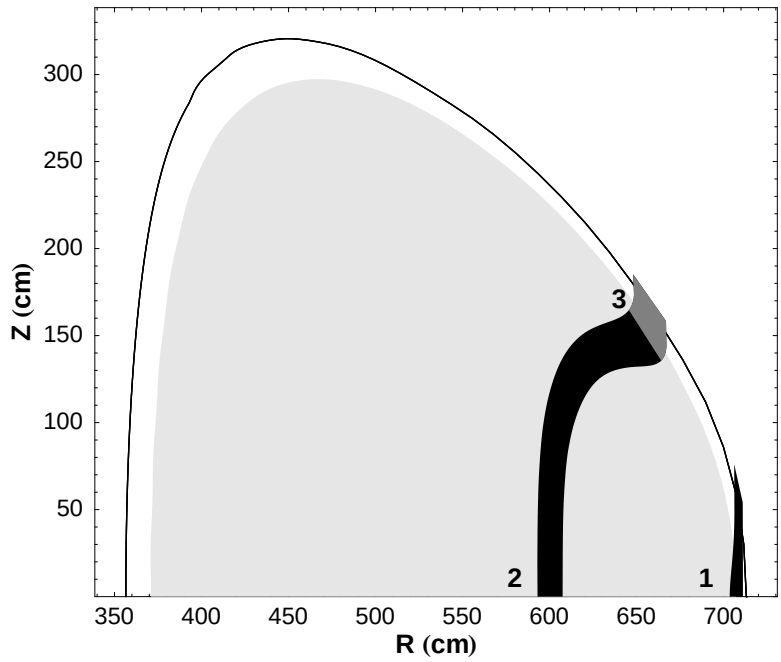

Figure 5.6: Location of waves for the two wave scenarios. The light grey is the area over which the TAE-like mode is assumed to exist. The dark grey regions 1,2,3 represent the location of the MCIBW for three different cases discussed in the text. The strange shape of the MCIBW region 2 arises from the shape of the $|B|$ surface in a reverse shear plasma.

5.3 Other Two Wave Scenarios

Let us now examine how this choice of wave parameters gives a good result and what ways of varying these parameters might even lead to a greater cooling effect.

We find that it is necessary for the MCIBW to be near the edge for significant cooling to occur on average. When the MCIBW is located near the center, for example in region 2 of Fig 5.6, almost all of the α-particles leave the plasma, (95% for this case) in a fixed number of steps, but no net energy is extracted from the wave!

Upon closer examination, we can see that many α-particles are heated by the IBW before being ejected. While the TAE still extracts 20% of the α-particle power, the net effect of the IBW is to heat the α-particles by almost that amount. By looking at Fig 5.7 we see that nearly 50% of the particles leave the plasma after being cooled, while

the other 50% leave the plasma after being heated significantly! Based on many different runs, it seems that locating the MCIBW in the center is a poor strategy.

What is happening in all of these cases, of course, is that heating and cooling always occur simultaneously; however, in the good case, the heated particles do not hit the wall. Let us think of the web of diffusion paths in ε, μ, and P_ϕ space. When the MCIBW is located near the center, the resonant region in constants-of-motion space and diffusion paths can connect to the boundary at energies significantly above the birth energy. Random walks can easily be formed from the center to the edge which involve heating rather than cooling. One way in which this can happen is for countergoing particles to be heated up significantly. Because the phasing used is $n_\phi > 0$, this means they must move in. However, if they are heated enough, they can cross the passing-trapped boundary, giving them a large radial step outwards. Then they can freely interact with the toroidal Alfvén eigenmode and continue on their way out of the plasma. Placing the MCIBW at the edge, however, means that the particles which come out are cooled and generally trapped by the time they reach the MCIBW layer. If they get heated by the MCIBW here, they tend to move back in; whereas, if they get cooled, they tend to move out closer to the wall. With the MCIBW very close to the edge, this tendency strongly favors paths where particles get out by being cooled.

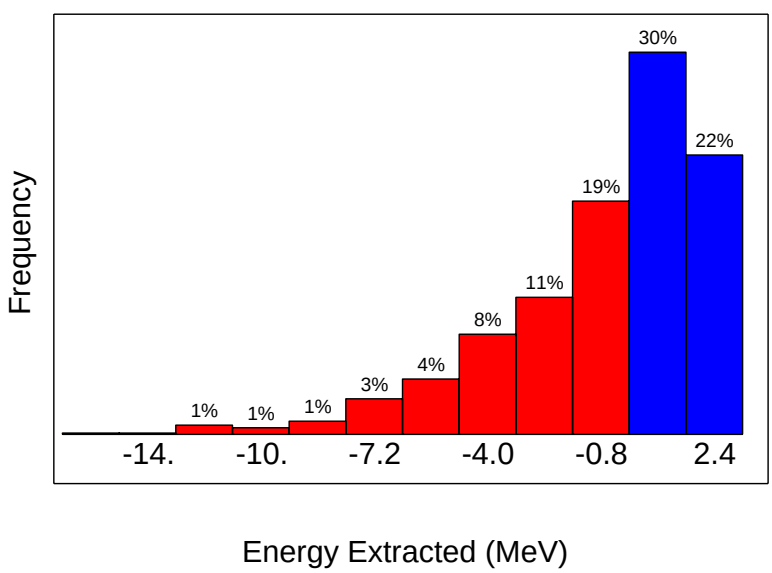

Figure 5.7: Histogram of the energy extracted from particles when the MCIBW is near the center.

This is verified by simulations with the wave location as shown in Fig. 5.6 with the MCIBW located in region 1. Now 64% of the α-particles are extracted, with, on average, 1.8 MeV of energy given to the wave per particle, leading to a total of 34% of the α-particle power given to the wave. It is also of interest to note that if the IBW power is concentrated at point 3, results about the same as those achieved for power concentrated at surface 1 can be obtained.

To gain further insight we depict in Fig. 5.8 the amount of energy extracted from α-particles with different initial ε, μ, and P_ϕ coordinates. The full set of points represents the α-particle birth distribution. We see that while the energy extraction is quite good in some areas of phase space, the deeply trapped particles are not getting out at all. Furthermore, little energy is being extracted from the particles which have low $\mu B_0/\epsilon$ (deeply passing particles). By referring to Fig. 5.8 we can pinpoint the "problem" α-particles, and then alter the wave characteristics so as to address these problems.

Recall that the toroidal Alfvén eigenmode extracts mainly parallel energy from the particles as it moves them out. The amount the particles are moved per unit energy extracted is given by Eq. (5.1). For this 34% extraction case $n_\phi \Omega_{\alpha 0}/\omega = 3500$ for all values of $\mu B_0/\epsilon$; however, this is not optimal. Ideally we would extract more parallel energy by the toroidal Alfvén eigenmode from those particles which have more parallel energy and less from those which have less. This can be accomplished by tailoring the value of $n_\phi \Omega_{\alpha 0}/\omega$ vs. $\mu B_0/\epsilon$. For the case described initially, we used $n_\phi \Omega_{\alpha 0}/\omega = 3500$ for particles with $\mu B_0/\epsilon > .85$, and $n_\phi \Omega_{\alpha 0}/\omega = 3500/3$ for the remaining particles. This choice of waves tends to maximize the energy extraction. The lower resonant frequencies experienced by high $\mu B_0/\epsilon$ particles [BIGLARI et al. 1992] are consistent with this choice. The TAE then extracts almost all the energy from the particles with $\mu \sim 0$ (passing particles), while allowing some particles with $\mu B_0/\epsilon \sim .85$ (trapped particles) to be moved out far enough that they can interact with the IBW, which is located near the edge for this simulation. This increases the amount of energy extracted from 34% to 51%, and increases the number of particles lost from 64% to 73%.

Further improvement can be had if we could tap the energy of the deeply trapped particles that are not leaving the plasma. Fig 5.8 suggests that if these particles were to get out, significant energy could be extracted from them. In order to determine why

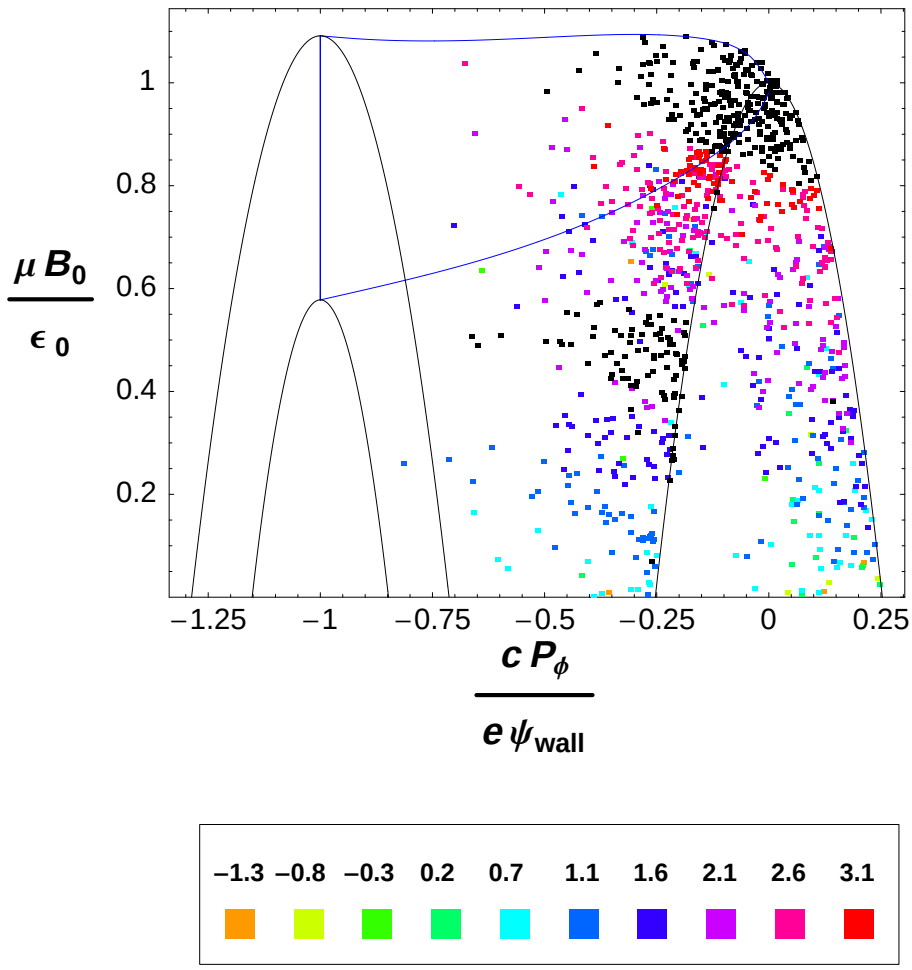

Figure 5.8: Energy extracted (MeV) by the both waves vs. initial location of particle in constants-of-motion space for MCIBW near the edge, but no tailoring of n_ϕ/ω for the TAE. Note that while almost all of the particles energy is being extracted from particles near the passing trapped boundary ($\mu B_0/\epsilon \approx 0.85$), particles which are deeply passing ($\mu B_0/\epsilon$ near zero) have significantly less energy extracted.

these particles are not exiting the plasma, we simulate one of these particles interacting with both waves. To aid in understanding, the particle is given only cooling kicks. The succession of orbits is shown in Fig. 5.9. Note that, as the particle moves out, its banana tips shrink towards the midplane. Recall that while the toroidal Alfvén eigenmode moves particles much greater distances than the MCIBW as it cools them, it still extracts energy as it moves the particles out. However, the toroidal Alfvén eigenmode can only effectively tap the parallel energy of the particle. Once the particle runs out of parallel

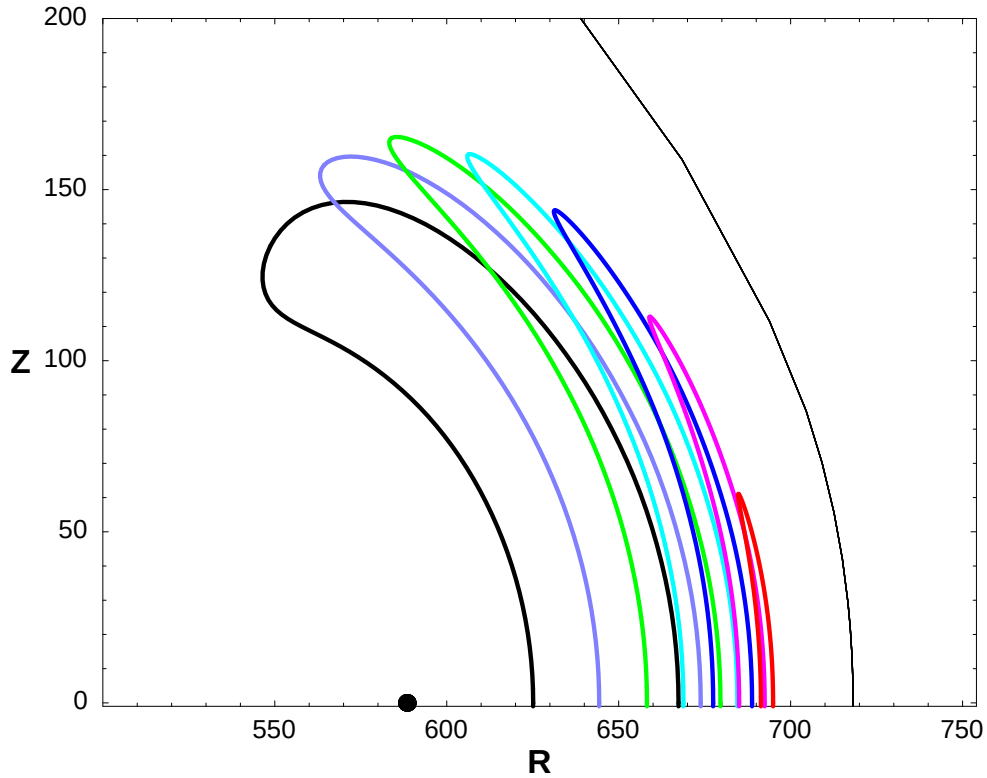

Figure 5.9: A succession of orbits of a deeply trapped particle interacting with the TAE. All of the particle's parallel energy is extracted before it can reach the MCIBW layer, and thus it is stuck in the plasma.

energy, it can no longer be pushed out by the toroidal Alfvén eigenmode. In this case, these deeply trapped α-particles get stuck before they can reach the MCIBW layer and thus are not able to leave the plasma and give their energy to the wave.

However, these deeply trapped particles might be taken out to the IBW layer, and then extracted, through stochastic ripple diffusion [WHITE et al. 1996], which does not extract energy or alter μ as it diffuses particles in P_ϕ. To model this effect, the simulation was modified so that particles with $\mu B_0/\epsilon > 1.0$ were treated as diffusing in P_ϕ with almost no energy extracted. Then, in contrast to 73% of particles extracted 93% of α-particles are now extracted, each cooled on average to $1/3$ of their birth energy, with 61% of the total power going to waves! Figures 5.10 and 5.11 show the extraction vs. ε, μ, and P_ϕ for this case. In arriving at this case, note how important it is to be able to pinpoint exactly where in ε, μ, and P_ϕ space the wave combination works well and where

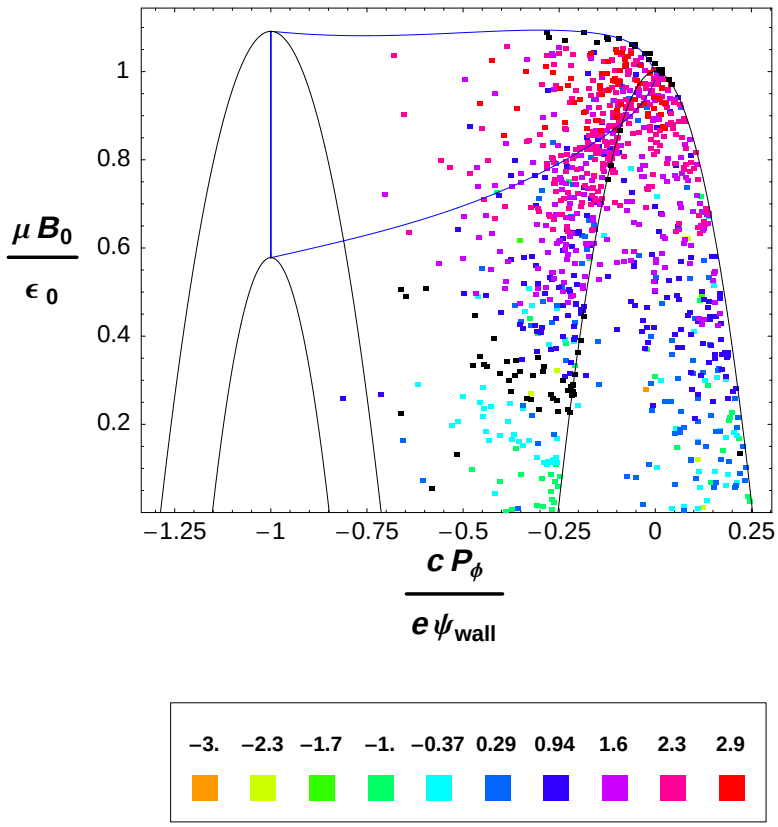

Figure 5.10: Energy extracted by the MCIBW when the n_ϕ/ω for the TAE is tailored.

it did not.

In all of these simulations the TAE is chosen to exist only part of the way to the wall, but, so long as the relative amplitudes of the MCIBW and TAE can be controlled, the result is not significantly changed if the TAE extends to the wall. On the other hand, if the location of the TAE does not overlap with the MCIBW layer, few particles would be lost and little energy extracted.

5.4 Cooled Distribution at the Wall

The distribution function of the exiting α-particles exhibits interesting features. Fig. 5.12 shows the position in velocity space of the α-particles that hit the wall. Note the bunching

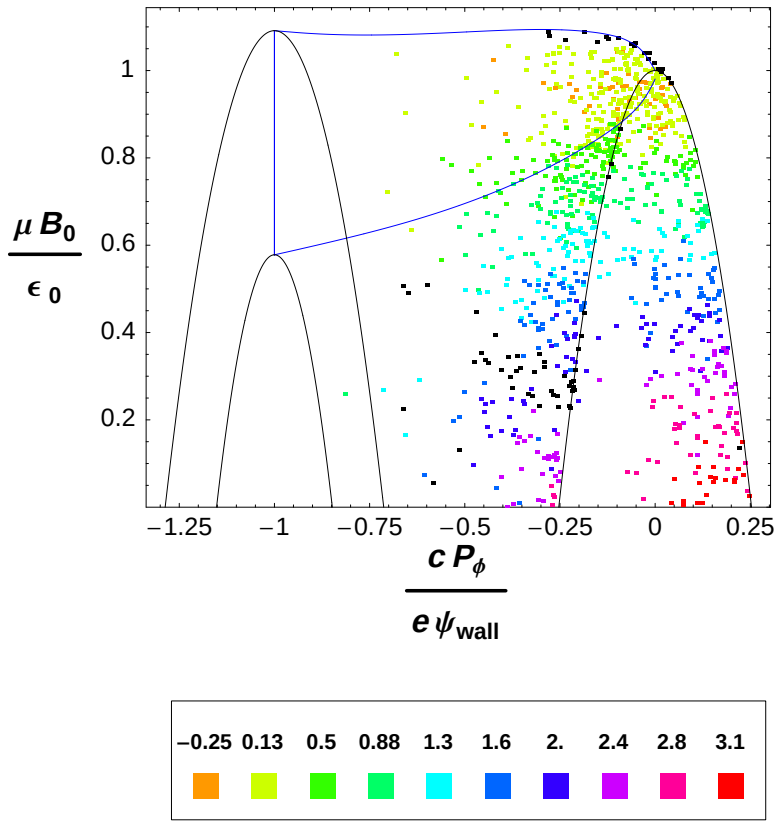

Figure 5.11: Energy extracted by the TAE when n_ϕ/ω for the TAE is tailored.

in $v_\parallel$, with a range of perpendicular velocities.

Because the distribution is almost entirely cogoing or trapped particles, bunching also occurs in the poloidal exit angle. In Fig. 5.13, the distribution of the α-particles on the wall vs. poloidal angle is plotted where $0°$ is at the outer midplane, $180°$ at the inner midplane and whether the loss occurs on the upper or lower half of the tokamak depends on the direction of the ∇B drift. The loss is on the outer midplane, because, if the wall of the tokamak is a flux surface, cogoing ions whose orbits are slowly deformed outward, no matter where they receive a kick from the wave, will eventually scrape off at $0°$. In contrast countergoing ions undergoing a slow deformation of their orbit will strike 180°. By making the size of the last kick the α-particle receives bigger (smaller), this loss can be distributed (localized). One may also significantly alter the distribution of the losses by changing the shape of the wall relative to the shape of the flux surfaces.

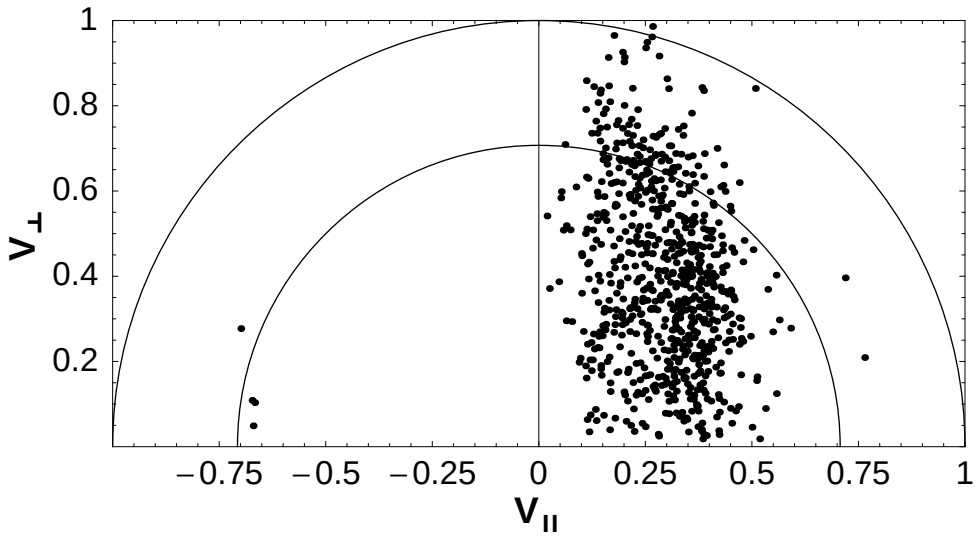

Figure 5.12: Velocity space position for particles leaving the tokamak after having their energy extracted. Semi-circles represent energies of 3.5 MeV and 1.75 MeV.

Note that in these simulations the amount of power flowing into the wall (20–30%) is much larger than the expected tolerance of future reactors (1–5%), if the loss is localized. While the loss might be tolerable if it were not localized, the interesting challenge is to exploit the bunching in phase space for further energy extraction.

5.5 Two Key Questions

In the simulations which demonstrate that significant channelling can be achieved, several assumptions were made. Two key assumptions are that the MCIBW experience a $k_{\|}$ flip and that the collisionless limit be applicable. These appear to be reasonable assumptions. Ray-tracing calculations do predict the $k_{\|}$ flip. At some wave power the collisionless limit is certainly attained. Nonetheless, we do assume that the $k_{\|}$ flip, in fact, occurs and that the collisionless limit is attained at reasonable power levels. In Chapter 7 we show that there is ample data to support these two assumptions.

In Sec. 5.5.1, we explain why the $k_{\|}$ flip is necessary. In Sec. 5.5.2 and Sec. 5.5.3 we estimate the wave power flowing in the plasma that leads to the collisionless limit.

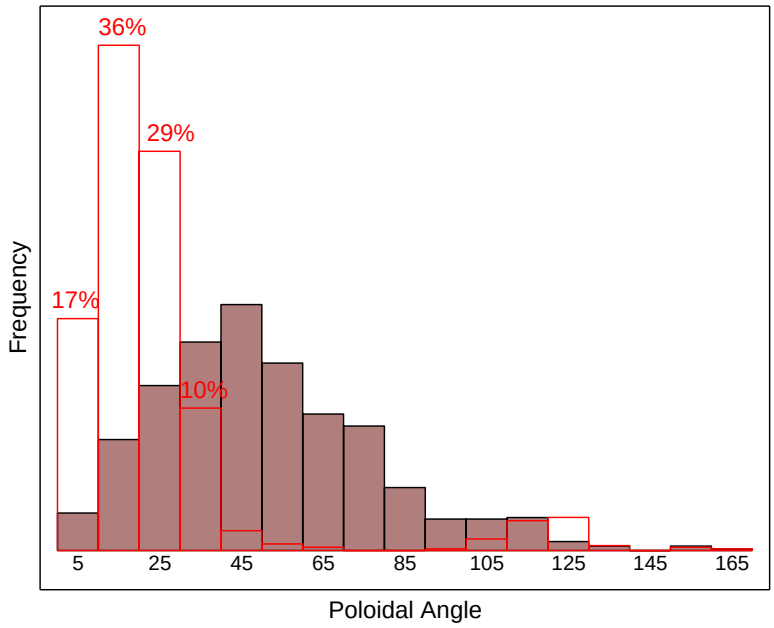

Figure 5.13: Histogram of losses vs. poloidal angle in degrees for α-particles ejected by the two wave scenario. Two cases are considered with the final kick varying by a factor of 8.

5.5.1 $k_{\|}$ flip of the MCIBW

The $k_{\|}$ flip of the MCIBW wave is a major part of what makes the MCIBW very attractive as a candidate wave for α-channelling. This has already been discussed briefly in Sec. 4.4 and will be discussed in more detail in Sec. 6.3.1 and Sec. 7.1. Recall that what we mean by a flip is that the parallel phase velocity of the launched wave is opposite in sign to the parallel phase velocity of the mode converted wave.

That the $k_{\|}$ of the MCIBW could flip sign if the effect of the poloidal field were strong enough was noticed by Valeo and Fisch (1994). This can be seen as follows. As discussed in Sec. 6.3.1, since the dispersion relation for the MCIBW changes most rapidly in the direction of $-\hat{\mathbf{X}}$ after it mode converts, its wavelength in this direction will become short. Recalling Eq. (4.47) from Chapter 4 in the large aspect ratio limit, $k_{\|}$ is just

$$k_{\|} = n_{\phi}/R_0 + k_X \frac{\epsilon}{q} \sin\theta. \tag{5.4}$$

Even though the poloidal field is small compared with the toroidal field, i.e., $\mathcal{O}(\epsilon/q)$, k_X can become quite large for the MCIBW, and the second term in $k_\parallel$ can become larger than n_ϕ/R_0. In the event that n_ϕ/R_0 and $k_X\hat{\mathbf{X}}\cdot\mathbf{B_p}$ are in opposite directions the sign of $k_\parallel$ can "flip".

This peculiarity leads to desirable effects, both in terms of its damping on the background fuel ions as well as its amplification by the fusion byproducts.

As pointed out by Valeo and Fisch (1994), this change in $k_\parallel$ as the wave propagates prohibits the wave from damping on the electrons, since the parallel phase velocity of the wave is very large, in fact infinite at the point of flipping. Eventually, under certain circumstances, the wave reaches large enough $k_\perp$ that it is able to damp on the thermal tritium ions. Thus the wave avoids electron damping while it is amplified convectively, finally damping on tritium. The ability of the waves which are amplified by the α-particle power to heat ions is central to the α-channelling concept.

With respect to the wave amplification, note that α-particles must interact with $n_\phi > 0$ waves in order to move out while being cooled. To see this, note that a particle moving from the center to the periphery of the plasma moves to lower P_ϕ (assuming its drift from its flux surface is small compared to the tokamak minor radius). Then, from Eq. (5.1), α-particles will cool as they leave the plasma only for $n_\phi > 0$. On the other hand, for α-particles to interact with the MCIBW, they must satisfy resonance, which is $v_\parallel = (\omega - \Omega_\alpha)/k_\parallel$. Mode conversion in DT plasmas occurs to the high field side of the α-particle gyroresonance layer, so that $\omega < \Omega_\alpha$. Thus, where the wave particle interaction occurs, $k_\parallel$ must be opposite the sign of $v_\parallel$. In particular, this implies that a cogoing, passing deuteron resonates only with waves such that $k_\parallel$ less than zero. However, the wave that was launched has $k_\parallel$ greater than zero, implying that the cogoing α-particles will not be able to resonate with the wave unless the wave undergoes a $k_\parallel$ flip!

In practice, many of the α-particles are trapped by the time they reach the mode conversion layer, if the MCIBW is located near the edge. This trapping arises from the extraction of parallel energy by the toroidal Alfvén eigenmode and the movement of the particles out so that they feel a larger mirror ratio. In principle, the trapped particles could interact either on their cogoing or countergoing legs, which diminishes (for some α-particles) the need of the $k_\parallel$ flip for particles to satisfy resonance. However, it is still

useful to have the $k_{\|}$ flip even for interactions with trapped particles, as can be seen from Fig. 5.14. As the trapped particles cool upon interaction with the MCIBW, they lose perpendicular energy and gain parallel energy, and thus eventually they cross from trapped to passing. If they do so while interacting on their cogoing leg ($k_{\|}$ flipped), they become cogoing passing particles, while if they are interacting on their countergoing leg, they become countergoing particles one full banana width inside of the cogoing particle, making it that much harder to get the cooled particles to hit the wall. Also, note that countergoing particles will hit the wall on the inner midplane, which may be less desirable than having the α-particles hit the wall on the outer midplane as shown in Sec. 5.4.

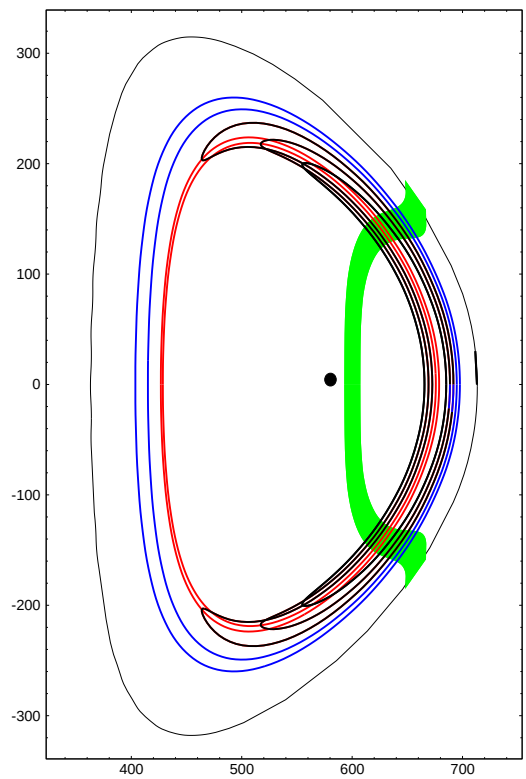

Figure 5.14: Particle interacting with the MCIBW and becoming trapped. If it resonates on the cogoing leg ($k_{\|}$ flipped) it becomes a cogoing particle upon becoming passing. If it resonates on the countergoing leg (no $k_{\|}$ flip) it becomes a countergoing passing particle.

5.5.2 Power for the toroidal Alfvén eigenmode

In order to neglect collisions the energy stored in the TAE must be large enough to diffuse the α-particles across the required region in a time short compared to the α-particle's slowing down time. Thus we need

$$\frac{D_\varepsilon}{\varepsilon_{\alpha 0}^2 \nu_s} \gg 1 \tag{5.5}$$

where D_ε is the energy diffusion coefficient of the α-particles in the TAE, $\varepsilon_{\alpha 0}$ is the birth energy of the α-particles, and ν_s is the energy slowing down rate of the α-particles. One might suppose a further constraint on the stored energy, namely, that the islands in ε, μ, and P_ϕ space associated with the TAE be overlapped so that the particles can diffuse across the affected region. For a fixed number of modes present, this would be a requirement on the energy in the modes. However, this can be eliminated by increasing the number of modes present, reducing the amplitude necessary for stochastic overlap. Hence the inequality Eq. (5.5) gives the requirement in D_ε.

We can compute D_ε by extending the calculation of Wu [WU et al. 1995]. We consider, in the high aspect ratio limit, a passing α-particle in a tokamak interacting with a mode which has **A** =$A_\parallel \hat{\mathbf{b}}$ where $A_\parallel(\psi, \theta, \varphi) = A_m(\psi) e^{i(n\varphi - m\theta - \omega t)}$ and $\Phi(\psi, \theta, \varphi)$ determined by requiring $E_\parallel = 0$. The frequency of the TAE is given by $\omega = \frac{v_A}{2qR}$ where q and v_A are evaluated at the value of ψ (poloidal flux) where $q = \frac{m+1/2}{n}$.

Then the expression for the rate of change of the energy of a particle near resonance (keeping only the slowly varying term) is given by:

$$\frac{d\varepsilon}{dt} = \frac{v_\parallel {\omega_t}^2}{\kappa} \sin(\Theta) \tag{5.6}$$

$$\frac{d^2\Theta}{dt^2} = {\omega_t}^2 \sin(\Theta), \tag{5.7}$$

where $\Theta = n\varphi - (m-1)\theta - \omega t$. The $(m-1)$ term arises due to the wave resonating with the particle drift motion, which has $\sin\theta$ and $\cos\theta$ dependence. The resonance condition

is $d\Theta/dt \approx (n - \frac{m-1}{q})v_\parallel - \omega = 0$. The parameters ω_t and κ are given by

$$\omega_t{}^2 = \omega\kappa v_\parallel \left(1 + \frac{v_\perp^2}{2v_\parallel^2}\right) B^2 \frac{q^2}{2\epsilon(nq-m)} \left(1 + \frac{\epsilon^2 q'}{q^2(nq-m)}\right) \left(\frac{\delta B^\psi}{B}\right) \tag{5.8}$$

$$\kappa = n - \frac{m-1}{q} + \frac{(m-1)q'v_\parallel}{q^2}(1 - \frac{nv_\parallel}{\omega}) \tag{5.9}$$

where q, the safety factor, and ϵ, the local aspect ratio, are functions of ψ. Assuming that the decorrelation time is approximately the trapping time (1/ω_t) we get

$$D_\varepsilon \equiv \frac{(\Delta E)^2}{2\Delta t} \approx \frac{v_\parallel^2 \omega_t^3}{2\kappa^2} = C(n, m, \psi, v_\parallel) \left(\frac{\delta B^\psi}{B}\right)^{\frac{3}{2}}. \tag{5.10}$$

This analysis gives an approximate value of the energy diffusion coefficient for an α-particle which satisfies resonance with a given mode m, n. Note that the resonance condition is only satisfied locally (in fact at one particular value of q). To estimate the minimum amplitude required (so that the global diffusion of α-particles from the center to the edge takes place on a time scale short compared to the energy slowing down time), we assume that many m's and n's are excited. For each m and n there is a q for which a cogoing 3.5 MeV α-particles is resonant. We take the mode amplitude to be nonzero from the resonance point halfway to the next closest resonance point on both sides. We then find the $\delta B^\psi/B$ which ensures that Eq. (5.5) is satisfied across the plasma cross section. For the reverse shear reactor we are considering here we find that at low values of n (1–10), $\delta B^\psi/B \approx 3 \times 10^{-4}$ satisfies the requirement above. At high values of n (20–50) $\delta B^\psi/B \approx 10^{-3}$ is necessary.

The power, P_{TAE}, in the mode that produces a perturbation δB, is given by [BERK et al. 1992],

$$P_{TAE} = 2\gamma_d \int \frac{\overline{(\boldsymbol{\delta}\mathbf{B})^2}}{4\pi} \mathbf{dx}, \tag{5.11}$$

where γ_d is the linear damping rate of the eigenmode. Assuming $\frac{\partial A_m(\psi)}{\partial \psi} = 0$ we find

$\delta B^\theta, \delta B^\varphi \ll \delta B^\psi$. Evaluating the time average we then get

$$P_{TAE} \approx 2\gamma_d \int \frac{\left(g_{\psi\psi}\delta B^{\psi^2}\right)}{8\pi}\mathbf{dx} \approx 2\gamma_d \left(\frac{\delta B^\psi}{B}\right)^2 \int \frac{g_{\psi\psi}B^2}{8\pi} J d\psi d\theta d\varphi. \tag{5.12}$$

We can express this in terms of a ratio of the α-particle power by multiplying by $\frac{P_\alpha \tau_E}{\mathrm{TE}} = 1$ where TE is the thermal stored energy of the plasma. Calling $1/\beta_M = 1/(8\pi\mathrm{TE}) \int g_{\psi\psi}B^2 J d\psi d\theta d\varphi$, we can conveniently write

$$\frac{P_{TAE}}{P_\alpha} \approx 2\frac{\gamma_d}{\omega}\left(\frac{\delta B^\psi}{B}\right)^2 \frac{1}{\beta_M}\omega\tau_E. \tag{5.13}$$

Plugging in typical numbers for these quantities, with $\delta B^\psi/B \approx 3 \times 10^{-4}$, we find that

$$\frac{P_{TAE}}{P_\alpha} \approx 26\frac{\gamma_d}{\omega}. \tag{5.14}$$

The range of γ_d/ω has been experimentally measured in work by Fasoli [FASOLI et al. 1996; FASOLI et al. 1996; FASOLI et al. 1995] to lie between 10^{-4} and 10^{-1}. The first scenario in Sec. 5.2 had 51% of the α-particle power going to waves, with 28% of the total α-particle power diverted to the TAE and the remainder to the IBW. Thus, a value of $\gamma_d/\omega < 10^{-2}$ should be sufficient for achievement of self-consistent steady state operation.

Note that the requirement on D_ε may be even less stringent than Eq. (5.5). In the 51% channelling case, on average 1.3 MeV was extracted from α-particles by the TAE before they left. The perturbation amplitude above assumed 3.5 MeV must be extracted. Thus the required wave amplitude could be reduced by $(3.5/1.3)^{(8/3)}$ reducing the power dissipated by approximately 14. Of course, it is necessary to diffuse some particles the whole 3.5 MeV by the TAE, and these would not exit in a slowing down time if P_{TAE} were 14 times smaller.

Note also that the radial mode extent was taken to be exactly right, so that no power would be “wasted”, supporting the mode in regions where it was not diffusing the α-particles. On the other hand, the diffusion coefficient shows considerable variability

across the plasma minor radius for flat $\delta B^{\psi}/B$ profiles. The level above was chosen so that, at the minimum, the collisionless limit was satisfied. By tailoring the $\delta B^{\psi}/B$ radial profile so that the collisionless limit is never greatly exceeded, the required power might be reduced reduced further from the estimate given in Eq. (5.14).

Note also that we considered only the steady state channelling. We showed that the power flowing directly from the α-particles to the TAE in steady state appears to be sufficient to maintain the necessary wave amplitudes. We did not address the initiation of the channelling effect which requires external power to start. The external power requirements may be significant, although it is possible that ways of bootstrapping the α-channelling effect up from low power levels could be found.

5.5.3 Power for the MCIBW

The power in the MCIBW must also be sufficient to satisfy Eq. (4.60). An estimate for D_{ε} can be obtained from Sec. 4.4.4 and Sec. 4.4.5. Squaring Eq. (4.86) and dividing by the bounce time we find

$$D_{\varepsilon} \approx \frac{q^2 \Phi_0^2 J_n(k_{\perp}\rho)^2}{\tau_b} \frac{2\pi\omega^2}{\frac{v_{\parallel}}{JB}\frac{\partial}{\partial\theta}(k_{\parallel}v_{\parallel} + n\Omega)} \tag{5.15}$$

We now need an estimate for Φ_0. The MCIBW is an electrostatic wave with $k_X \gg k_{\parallel}$. Previous calculations [VALEO and FISCH 1994] have found that (for the correct poloidal phasing of the waves) damping is not significant until after the wave group velocity goes through zero and it the wave reaches its maximum amplitude. For such a convectively damped wave, the relevant field amplitude is determined by balancing the incoming power with the wave energy flux. The cold plasma dispersion relation can be used to estimate these quantities. The power flux P_X is then [STIX 1992]:

$$P_X \approx \frac{\omega}{8\pi}\Phi_0^2 k_X \epsilon_{XX} \tag{5.16}$$

In steady state this power will equal the power flux from the antenna, P_{IBW}/A where A

is the spot size in the plasma. Setting these two equal gives:

$$\Phi_0^2 \approx \frac{P_{IBW}}{A}\frac{8\pi}{\omega k_X \epsilon_{XX}}. \tag{5.17}$$

For the mode converted IBW the cold plasma dispersion tensor term $\epsilon_{XX} \to 0$ at the mode conversion layer, so that Φ_0 blows up. While the derivation of Eq. (5.17) breaks down at the mode conversion layer we can estimate the relevant amplitude by calculating Φ_0 one perpendicular wavelength before the resonance [VALEO and FISCH 1995]. Thus,

$$\Phi_0^2 \approx \frac{P_{IBW}}{A}\frac{4}{\omega \frac{\partial \epsilon_{xx}}{\partial X}}. \tag{5.18}$$

This expression can be plugged into Eq. (5.15) to obtain a diffusion coefficient which in turn can be plugged in to Eq. (5.5) to estimate a power level. Plugging in some typical numbers for the cooling scenario outlined Sec. 5.2 we find a power level of approximately 100 MW would be sufficient. In this scenario 23% of the α-particle power is diverted to the IBW. For a 3 GW thermal reactor, with 600 MW of α-particle power, the MCIBW would be amplified by approximately 140 MW. Thus, in steady state, the diverted power could support the IBW at the required amplitude. However, there are a number of ways in which this calculation might change.

The requirement on D_ε may be less stringent than stated. The MCIBW in the first scenario only extracted 1.1 MeV on average from the particles. If the diffusion coefficient were set at a level which would diffuse only 1.1 MeV in a slowing down time the power could be reduced by $(3.5/1.1)^2 \approx 10$. However this would sacrifice the cooling of those particles which could be cooled by almost 3.5 MeV by the IBW.

An important effect for a self-consistent calculation would be to take into account the distribution of α-particles at the MCIBW layer near the edge. In the 51% channelling case, on average 1.3 MeV was extracted from α-particles by the TAE before they left. Since the particles are first moved out by the TAE and then interact with the IBW, on average, the IBW will be interacting with particles which are slower and tend to have much of their parallel energy extracted. From the expression above for D_ε, such particles would be expected to have higher diffusion coefficients since they tend to spend more

time in resonance.

While the power flowing directly from the α-particles to the IBW in steady state may be sufficient to maintain the necessary wave amplitudes, significant external power will be necessary to get the channelling process started.

While these arguments suggest the collisionless limit is attainable at reasonable power levels, only quite detailed calculations of the wave propagation, damping, and wave-particle interaction can substantiate this assumption.

5.6 Summary

What has been shown here is that low frequency waves and ion Bernstein waves can act in concert to extract upwards of 50% of the α-particle power from a tokamak reactor. The best case reported here was 61% extracted, which was $2/3$ of the energy from 93% of the particles, with the other 7% remaining in the plasma.

This study has been done with a generic advanced tokamak reactor. If the benefits were sufficient, a tokamak designed to optimize α-channelling might be envisioned. Such a reactor may look very different from the reactor used in these simulations, and it may be possible to design one specifically to optimize the α-channelling effect.

The waves utilized in these simulations enjoy substantial experimental documentation. The mode converted ion Bernstein wave has been studied as a means of heating electrons or driving currents [MAJESKI et al. 1994; MAJESKI et al. 1996; MAJESKI et al. 1996] and experiments have documented the interaction of these waves with deuterium beam ions in D-He^3 plasmas [DARROW et al. 1996]. Importantly, for the cooling scenarios presented here, experiments have shown that the layer of mode conversion can be controlled quite precisely by varying ω/Ω_α and the species mix of the plasma.

The TAE, generically used here for any low frequency eigenmode, has many different forms (e.g. TAE, EAE, KTAE, etc.), which have been shown to cause the loss of fast ions [WONG et al. 1991], and recently shown to be driven unstable by energetic

α-particles [NAZIKIAN et al. 1997]. Experiments [FASOLI et al. 1995; FASOLI et al. 1996] indicate that these modes can be launched externally.

For α-channelling it is important that the waves that are amplified at the expense of the α-particle power damp on ions. Theoretical calculations show that certain Alfvén eigenmodes damp on plasma ions [BETTI and FREIDBERG 1992], and mode converted ion Bernstein waves in a moderately deuterium rich reactor damp on tritium ions [VALEO and FISCH 1994]. Recent experiments in DT plasmas show strong ion heating by the MCIBW under certain conditions [WILSON et al. 1998].

Chapter 6

Modeling Fast Particle-MCIBW interactions on TFTR

TESTING THE α-CHANNELLING EFFECT experimentally requires several hurdles to be overcome. Today (and for the foreseeable future), there is only one running experiment, JET, that has the capability to operate in DT and therefore to produce the appropriate birth distribution of α-particles. The α-channelling scenarios discussed in the previous chapter require two very different waves (toroidal Alfvén eigenmodes and MCIBW) to interact with the α-particles in order to accomplish cooling. At present, both of these waves require further experimental investigation for a complete understanding. Even if both waves were present in a DT plasma, diagnosing the occurrence of the α-channelling itself is an experimental project of its own. To convince those working on the only available DT tokamak that this experiment is worth scarce machine time requires not only the hope that, in principle, α-channelling could have a telling effect on the viability of the tokamak reactor concept, but that the experiment itself has a large probability of success. In this respect, the experiments on TFTR provide a critical database.

First, it is possible to validate separately the various elements which make up α-channelling. Both the MCIBW and toroidal Alfvén eigenmode in themselves provide a host of interesting physics problems. Recently, serious efforts have been undertaken

to understand the toroidal Alfvén eigenmodes on JET [FASOLI et al. 1996; FASOLI et al. 1996; FASOLI et al. 1995] and the MCIBW in DT and D^3He plasmas on TFTR [MAJESKI et al. 1996; ROGERS et al. 1996; MAJESKI et al. 1996; WILSON et al. 1998]. Experiments were conducted on TFTR which observed interactions between the MCIBW and fast ions [DARROW et al. 1996; DARROW et al. 1996; FISCH et al. 1996; DARROW et al. 1997].

In this chapter, we review the experimental and theoretical foundations of the MCIBW. We take particular interest in TFTR experiments which demonstrated a strong fast particle MCIBW interaction. In order to fully appreciate the data from these experiments we review the MCIBW, the fast ion species which are present, the lost alpha detectors [ZWEBEN et al. 1990; DARROW et al. 1995], and the ways in which fast particles hit the wall in tokamaks in Sec. 6.1. We then provide an overview of the data from the experiments which exhibited fast ions interacting with the MCIBW in Sec. 6.2. These fast ion losses can be understood in terms of a theoretical framework which is developed in Sec. 6.3. While the theoretical framework is able to account for some of the observations, to include all of the various effects it is necessary to rely on numerical simulations of the losses which is done in Sec. 6.4. We then summarize the results of this chapter.

In Chapter 7 we will pursue analysis of these experimental observations further to show how, from the TFTR experiments, key elements of the building blocks for α-channelling can, in fact, be deduced.

6.1 Overview of Experiments

Many experiments related to mode conversion heating and current drive in DT and D^3He plasmas were carried out on TFTR from December, 1994 to March, 1997, led by R. Majeski and D. Darrow. For the most part, our discussion here will focus on the D^3He discharges carried out with the 43 MHz antenna. These discharges were carried out prior to the summer of 1996, after which most mode conversion experiments used 30 MHz RF on DT plasmas. Figure 6.1 shows a schematic diagram of the experiment. The fast mag-

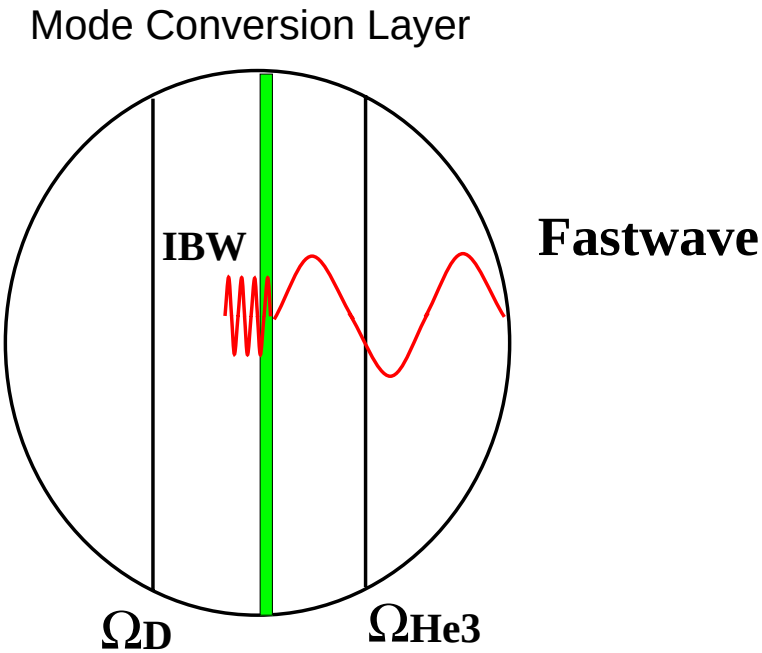

Figure 6.1: Schematic of a D^3He mode conversion experiment on TFTR.

netosonic wave is launched from the low field side into a D^3He plasma and propagates to the ion-ion hybrid layer (whose location depends on the wave frequency, toroidal field, and the density of each ion species). Once the wave reaches the mode conversion layer, it mode converts to a short wavelength electrostatic wave and damps, predominantly on electrons. Typical parameters for these discharges were I_p = 1.4 MA, $n_{^3\mathrm{He}}/n_e$ from 0.1 to 0.25 (the rest was deuterium or deuterium-like, i.e. a charge to mass ratio of one half), and B_T ranging from 4.4 to 5.3 T. For these conditions the mode conversion layer was within 25 cm of the magnetic axis.

In many of these discharges, a strong loss of energetic ions was observed on at least one of the lost alpha scintillator detectors [ZWEBEN et al. 1990; DARROW et al. 1995]. The lost alpha diagnostic detects energetic ions impinging on the wall in four poloidal locations at one point toroidally in TFTR. It has limited resolution in both pitch angle and gyroradius. In principle, these α-particles detectors reveal a wealth of information: the time history of the losses at each poloidal angle, can be had a function of pitch angle and energy (gyroradius). That would be a signature in four dimensions (time, poloidal angle, pitch angle, energy).

These lost alpha signatures make up the bulk of our data set. The characteristics of the losses under varying experimental parameters provide an invaluable database for benchmarking the COM simulation and for uncovering key physics relevant to α-channelling.

6.1.1 MCIBW physics

Mode conversion of the fast magnetosonic wave to the ion Bernstein wave has been experimentally confirmed [PARK et al. 1985]. More recently it has been the subject experiments both on TFTR [MAJESKI et al. 1996; ROGERS et al. 1996; MAJESKI et al. 1996; WILSON et al. 1998] and Alcator C-Mod [BONOLI et al. 1997]. We present here the results of these experiments which are especially relevant for α-channelling and describe the experimental conditions for the fast ion loss experiments.

The MCIBW is launched from the low field side in TFTR and propagates as a fast wave to the ion-ion hybrid layer, whose location is determined by the solution of (assuming $\omega \ll \Omega_e$, and only two species are present),

$$\frac{c^2 k_\parallel^2}{\omega^2} = \frac{{\omega_{p1}}^2}{{\Omega_1}^2 - \omega^2} + \frac{{\omega_{p2}}^2}{{\Omega_2}^2 - \omega^2}, \tag{6.1}$$

where $\omega_{ps} = 4\pi n_s {q_s}^2/m_s$, and $\Omega_s = q_s B/m_s c$. This equation can be solved for the value of B at the mode conversion layer. Note that $\Omega_1 < \omega < \Omega_2$ is satisfied at the mode conversion layer [STIX 1992] (we take the charge to mass ratio of species 2 to be greater than that of species 1).

In Fig. 6.2 the position of the mode conversion layer is plotted versus the toroidal field on axis and the ratio of ^{3}He to electron density for a 43 MHz wave in a TFTR plasma. The mode conversion layer can be swept across the plasma by fixing the ratio of $n_{^3\mathrm{He}}/n_e$ and varying the toroidal field, or by holding the field fixed and changing the ^{3}He concentration. Note that the value of ω/Ω_D or $\omega/\Omega_{^3He}$ at the mode conversion layer depends only on the ratio $n_{^3\mathrm{He}}/n_e$. The mode conversion layer moves to the deuterium cyclotron layer ($\omega = \Omega_D$) as $n_{^3\mathrm{He}}/n_e$ goes to 1/2 (which is the maximum value it can reach given charge neutrality), and moves toward the ^{3}He cyclotron ($\omega = \Omega_{^3He}$) layer as

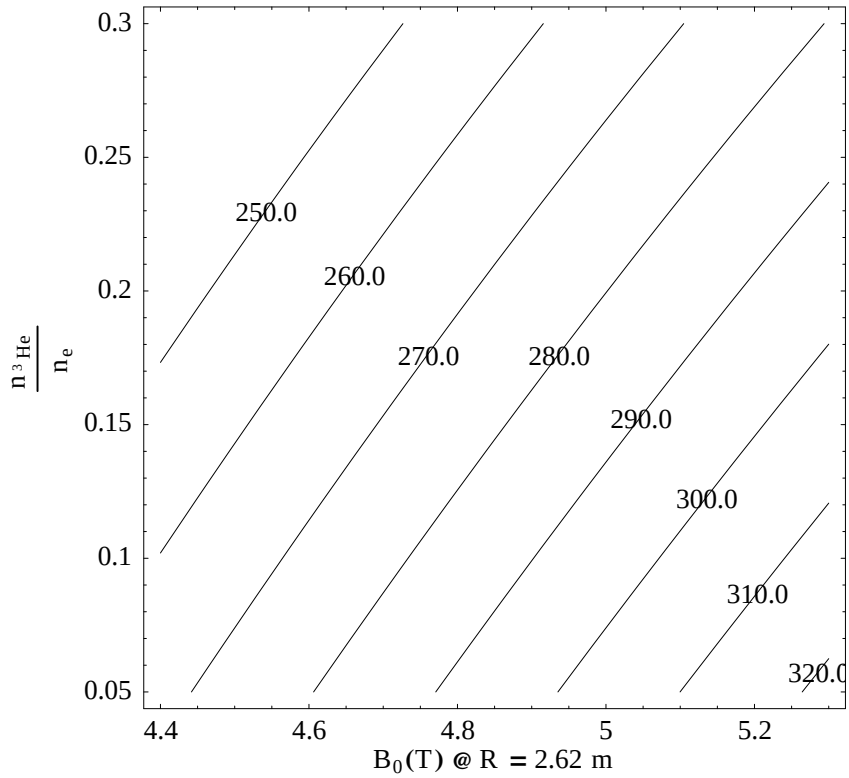

Figure 6.2: The major radius of the mode conversion layer (R_{mc}) for a 43 MHz wave in D^3He TFTR plasmas. B is taken to be B_0R_0/R here, which is a good approximation for TFTR. We have plotted R_{mc} vs. B_0, the toroidal field at the center of the device, and the fraction of ^3He in the plasma.

$n_{^3He}/n_e$ approaches zero.

Experiments were performed to investigate both mode conversion heating and current drive. First, high efficiency mode conversion (greater than 50%) can in fact be achieved with a launch from the low field side of the tokamak [MAJESKI, PHILLIPS, and WILSON 1994]. This theoretical result was confirmed in the TFTR experiments [MAJESKI et al. 1996], using a break-in-slope analysis to analyze changes in electron temperature when the input power was changed rapidly. It was deduced that from 60% to 80% of the coupled RF power was deposited as electron heating at the mode conversion surface. By determining the power deposition profile, it was determined that the power was being damped with a radial full width at half maximum of 15 to 20 cm. This estimate actually overstates the width of the mode conversion layer, because the mode conversion

layer is a vertical strip intersected by many flux surfaces, but, as the heated electrons travel along their flux surfaces, they smear out the power deposition profile. The peak of the electron power deposition profile was found to agree with the location predicted by Eq. (6.1), within the error bars which are on the order of ± 5 cm [MAJESKI et al. 1996], for many different values of the toroidal field, and ^{3}He densities.

The MCIBW current drive experiments demonstrated that currents can be driven both on and off the magnetic axis. More than 100 kA of current were inferred to be driven by 2.2 MW of ICRF power in one case. While the measured currents agree with the predicted current within the errorbars, there are many uncertainties in the prediction, such as the amount of power coupled, the spectrum of the coupled power, the mode conversion efficiency, and how much current was driven in the opposite direction due to the backward lobe of the antenna spectrum. The uncertainty in the theoretically predicted value of the current is about a factor of two [MAJESKI et al. 1996].

More recently, experiments investigating mode conversion heating in DT were carried out [ROGERS et al. 1996; WILSON et al. 1998]. Initial results were disappointing, as only 10-20% of the power appeared in the mode converted wave. It was later determined that a small amount of ^{7}Li , (which was used in TFTR for conditioning) was present in these discharges. Since ^{7}Li has a charge to mass ratio between that of tritium and deuterium, it apparently was absorbing almost all of the ICRF power. Presumably small quantities of ^{9}Be or ^{11}B would also absorb large amounts of power. For the purposes of α-channelling it is worth noting that using Li, Be, or B for minority heating would lead to a large fraction of the absorbed power going to the ions. The critical energy at which the energy slowing down rate on ions is equal to that on electrons is [STIX 1992]

$$\varepsilon_c = 14.8 T_e \left(\frac{A^{3/2}}{n_e} \sum_i \frac{n_i {Z_i}^2}{A_i} \right)^{2/3} \approx 9.3 A T_e \tag{6.2}$$

where A is the atomic mass number, and in the latter expression we have taken Z_i/A_i equal to one half (which is an over estimate due to the presence of tritium, but close enough for our purposes). Thus tail ^{7}Li, ^{9}Be, or ^{11}B ions less than $\approx$ 60, 80, or 100 T_e, respectively, should deposit most of their power on the ions.

Experiments carried out without the ^{7}Li present (^{6}Li was substituted) demonstrate mode conversion in DT plasmas [WILSON et al. 1998]. If the ion temperature was high enough, strong localized ion heating, was observed, with up to 75% of the RF power going to the ions. In fact, even though only 1 MW of RF was being injected on top of 18 MW of beam power, RF modulation was found to modulate the ion temperature by several keV at the mode conversion layer.

It is interesting to note that ion heating in DT mode conversion was predicted *a priori* by Valeo and Fisch (1994). Ion heating is crucial for α-channelling as currently envisioned, because the RF power given up by the α-particles to the waves must then flow into ions, so that a hot ion mode can be maintained.

6.1.2 Fast ion species

In the absence of beam heating, the relatively cold plasmas of the MCIBW experiments (T_i around 4 keV) produce almost no charged fusion products (CFPs), so that there are essentially no fast particles present. By comparing experiments with and without neutral beam heating, it is easy to determine that the observed losses during MCIBW experiments must arise from an interaction of the MCIBW with fast particles. However, determining exactly which species of fast particles is being lost is more difficult.

Neutral beams can be injected into TFTR either cogoing or countergoing to the current (or both), creating cogoing or countergoing ions which are predominately at 100 keV (although significant components exist at 1/2 and 1/3 of this energy). The injected species can be either D or T and, in the presence of neutral beam heating, a plethora of energetic charged particles can be present. For D beams on D^{3}He background plasma, one can expect significant populations of 1 MeV tritons, 800 keV ^{3}He ions, and 3 MeV protons from the DD fusion reactions, 3.6 MeV α-particles and 14.7 MeV protons from the D^{3}He fusion reactions, as well as the D beams themselves. The ratio of beam ions to charged fusion products is on the order of 10^4 for these plasmas with D beam injection. For T neutral beam injection into D^{3}He plasmas, the only fast particles in significant numbers will be beam tritons and 3.5 MeV α-particles from the DT reaction.

For the discussion below, note the gyroradius of different particles. Considering particles with all of their energy in the perpendicular direction, with mass m, energy ε, and charge q,

$$\rho = \frac{\sqrt{2\varepsilon/m}}{\Omega} = \frac{\sqrt{m\varepsilon}}{qB}. \tag{6.3}$$

Thus, a 3.5 MeV α-particle, a 1.75 MeV deuteron, and a 1.17 MeV triton all have the same gyroradius if they have the same pitch (equivalently the same fraction of their energy in the parallel direction). Furthermore, the guiding center trajectory of any particle depends only on its pitch, gyroradius, and position. Thus the motion of a 3.5 MeV α-particle and a 1.75 MeV deuteron, with the same pitch and initial location, are indistinguishable on the basis of their guiding center motion in the absence of collisions.

In the presence of collisions, one can distinguish between the α-particle and the deuteron, as the slowing down rate of the α-particle will be twice as fast. Some slowing down times on axis for the plasmas typical of these experiments are given in Tab. 6.1. At $r/a = 0.3$, the slowing down time is shorter by a factor of two than the time on axis, while the pitch angle scattering time remains about the same.

particle	energy slowing down time	pitch angle scattering time
100 keV D	0.076	0.160
1.75 MeV D	0.171	11.3
3.5 MeV α	0.085	10.9
100 keV T	0.078	0.198
1.16 MeV T	0.240	7.66

Table 6.1: Collisional times (the inverse energy slowing down rate and the inverse pitch angle scattering rate in seconds) for various particles for $n_e = 8 \cdot 10^{14}/cm^3$, T_e=T_i=5 keV, $Z_{\rm eff} = 3.5$.

6.1.3 Losses of fast particles and their detection

Good confinement of fast particles is crucial to the future operation of fusion reactors. It is an area that has been heavily studied, both theoretically and experimentally. In TFTR the losses of charged fusion products and ICRF minority tails has been extensively studied [ZWEBEN et al. 1990; ZWEBEN et al. 1994; ZWEBEN et al. 1997; DARROW et al. 1996; HERRMANN 1997]. For our purposes here, it is useful to go over the main loss mechanisms and their detection.

The experiments that were carried could not have been done at all without the lost alpha scintillator detectors mentioned earlier. These detectors are located on the wall of TFTR, 20°, 45°, 60°, and 90° degrees below the outer midplane, at a one toroidal location. Note, because $\mathbf{B} \times \boldsymbol{\nabla} B$ is in the downward direction for TFTR, losses only occur below the midplane. A schematic of the detector is shown in Fig. 6.3. The detector admits particles with gyroradii between 2 cm and 11 cm, and pitch angles between 45° and 83° degrees ($v_{\|}/v$ between 0.71 and 0.12 in the cogoing direction). The 45°, 60°, and 90° have a thin metal film over their aperture, which prevents hydrogenic ions of less than 400 keV and helium like ions with less than 900 keV from penetrating.

These detectors were designed to give information about the confinement of the 3.5 MeV fusion α-particles which were present in large numbers during TFTR's DT campaign, but they can detect other energetic ions as well. The detectors have no capability to distinguish two different species with the same gyroradius. This weakness is somewhat compounded by the fact that the guiding center motion of a particle depends only on its gyroradius and pitch. Thus, the signature from a 3.5 MeV α-particle born with pitch λ_0, at position (ψ_0, θ_0) will be the same as the signature from a 1.75 MeV D ion born with the same pitch at the same location. Of course in most cases this ambiguity is inconsequential since only 3.5 MeV α-particles are present in large numbers. However, as mentioned above, in D^3He plasmas with D beams as many as 5 different types of fast particle are present, many with similar gyroradii.

The detectors have only limited resolution in gyroradius and pitch angle, a consequence of providing the maximum signal, and the widest range of detection of fast particle losses. Also, the absolute calibration of the detectors is very difficult and is un-

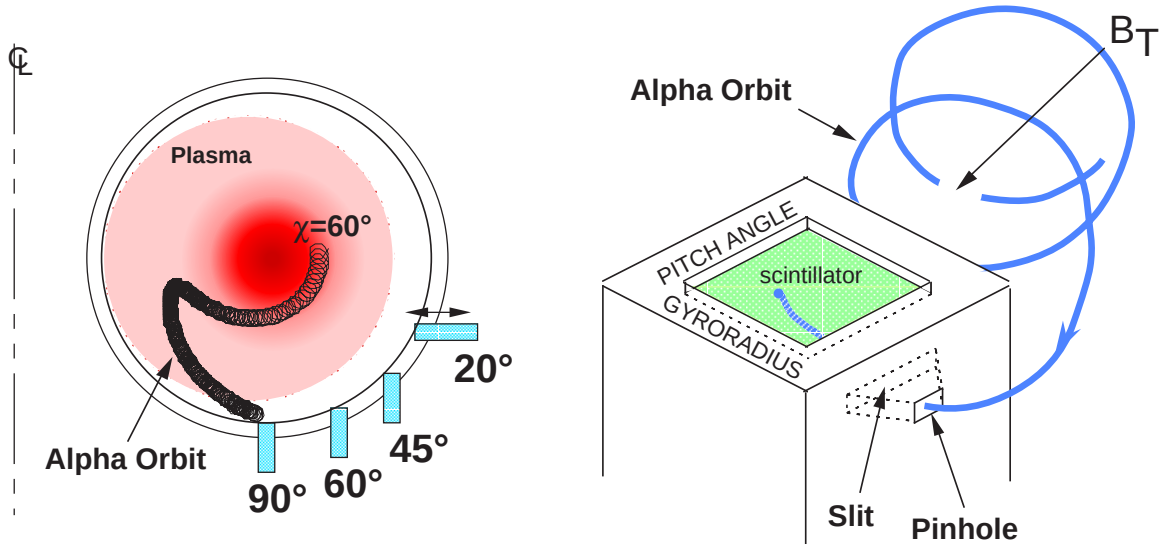

Figure 6.3: The poloidal projection of a 3.5 MeV α-particle born onto a loss orbit, and hitting the lost alpha detector at 90°(left). On the right is a schematic diagram of a lost alpha scintillator detector. Where the escaping particle hits the scintillator is dependent on the particles' pitch angle and gyroradius. Figure from D. Darrow.

certain to within a factor of 2 or 3 [ZWEBEN et al. 1994]. Rather than use an absolute calibration, losses are typically compared with the neutron signal. For first orbit losses of charged fusion products which have associated neutrons (like DT and DD reactions, but not D^3He), if the birth profiles of fusion products and plasma current remain the same throughout the discharge, the first orbit loss rate should be proportional to the neutron signal throughout the discharge, although it can vary considerably from shot to shot as the profiles vary. An estimate of the number of lost particles can be made by estimating the level of first orbit losses and multiplying by the number of neutrons. When the ratio of the lost alpha detector signal to neutron signal changes significantly during the discharge, it is usually indicative of non first orbit losses.

We discuss four different ways in which charged energetic particles can hit the wall in the tokamak. First, we discuss the so-called "first-orbit" loss, which for most TFTR discharges is the dominant loss mechanism. This loss arises from charged fusion products which are born on trajectories that intersect the wall. By far the vast majority of these particles are trapped particles that have banana widths larger than the minor radius. These particles are typically born countergoing near the center, mirror on the high field

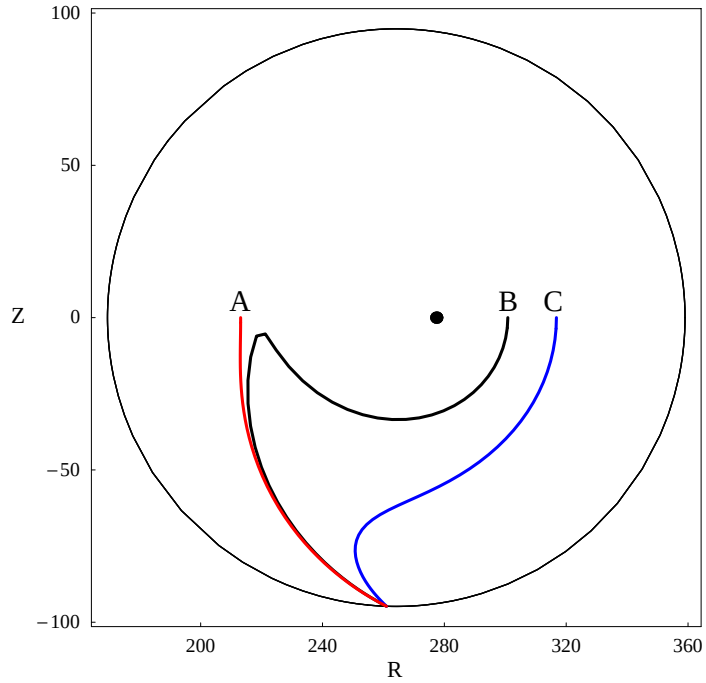

Figure 6.4: The guiding center motion of three 3.5 MeV α-particles which strike the 90° detector. A represents a passing particle which strikes the detector with a pitch angle below the passing trapped boundary. B is a "fattest banana" orbit which strikes the detector with pitch angle at the passing trapped boundary. C is a more deeply trapped particle which strikes the detector above the passing trapped boundary.

side, and intersect the wall before reaching the outer midplane. These losses depend strongly on the birth profile of the fusion products, the total current, and the current profile of the plasma. These losses represent a background level for the experiments we are interested in. Typically, for a 1.4 MA discharge, 10-20% of the charge fusion products might be lost to first orbit losses. In Fig. 6.4 we show the guiding center trajectory of 3 3.5 MeV α-particles born onto loss orbits that hit the wall at the 90° detector.

In Fig. 6.5 we summarize the first orbit losses for the entire plasma for a 1.4 MA discharge. In the top graph, we show the distribution of the losses versus angle along the bottom half of the wall. In the four graphs beneath the large graph, the pitch angle distribution of the losses is shown at the four detectors (20°, 45°, 60°, and 90°). The sharp jump in the pitch angle distribution at each detector signifies the passing trapped boundary. Particles with pitch angles lower than this boundary are passing, and therefore, like particle A in Fig. 6.4, must come from a large radius on the inner midplane. Particles

with pitch angles just greater than the passing trapped boundary are trapped particles, like particle B above. This particle passes much closer to the core of the plasma than does particle A, and since the radial distribution of fusion products is typically taken to be $(1-(r/a)^2)^n$ where n is 8-12, the source rate for particle B is much greater than that for particle A, leading to the much higher level at the detector.

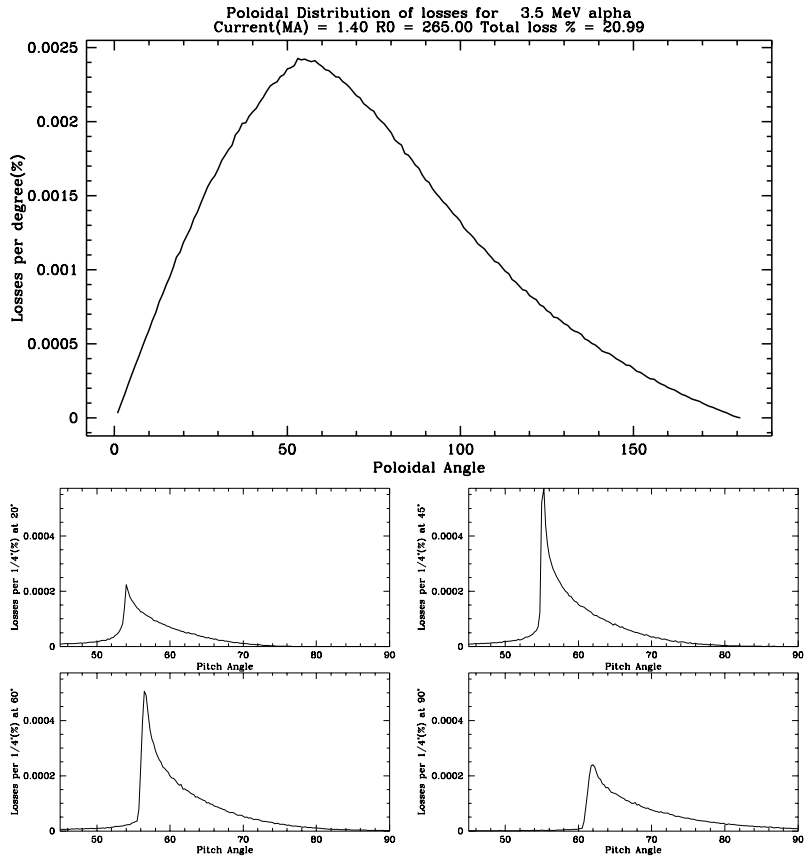

Figure 6.5: First orbit losses of a birth distribution of α-particles. Poloidal distribution of losses(top). Pitch angle distribution of losses at each of the 4 detectors (bottom).

In addition to first orbit losses, there are losses that occur from small changes in the orbits of initially confined particles. In Fig. 6.6, four examples are given of a confined particle which undergoes an orbital change and becomes a lost particle. In both Fig. 6.6A and Fig. 6.6B, the confined particle is countergoing. For case A, the particle crosses from passing to trapped. The resulting trapped particle has a very large banana width (comparable to the minor radius), and thus hits the wall, in this example near the

90° detector. In case B, the particle is given a radial kick moving it onto an orbit which is slightly farther out and hits the wall near the inner midplane. Case C shows the loss of an initially cogoing particle, which also has received a radial kick moving it onto a loss orbit which hits the wall near the outer midplane. Finally, case D shows a trapped particle whose banana tip has moved radially, causing it to intersect the wall near the outer midplane.

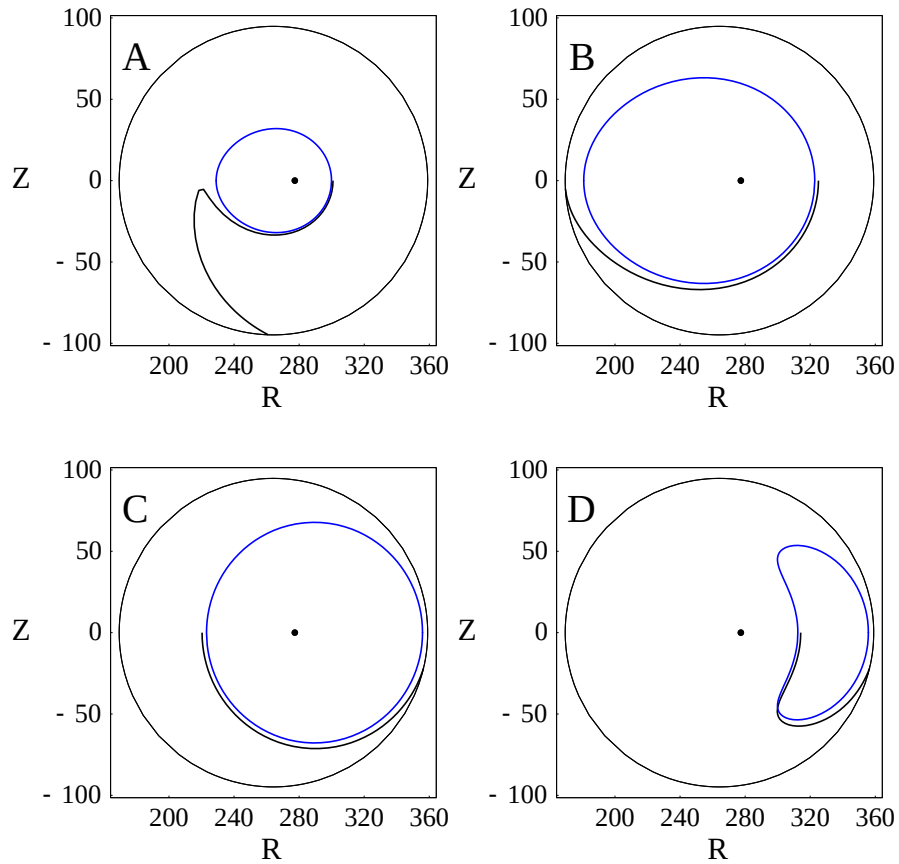

Figure 6.6: Examples of confined particles transitioning to lost orbits.

6.2 Fast Particle MCIBW Interactions on TFTR

Experiments on TFTR have seen large losses of energetic ions to the four lost alpha detectors mentioned above during MCIBW experiments [DARROW et al. 1996]. Data

has been compiled from both piggyback experiments (where losses were detected during mode conversion heating and current drive experiments) and dedicated experiments (where the physics of the losses themselves was being investigated). Many TFTR shots (hundreds) exhibited losses significantly in excess of first orbit losses under varying conditions when MCIBW is present. This represents a vast database which can be used for validating hypotheses about the losses. In this section, we summarize the experimental observations of these losses.

The observations from the lost alpha detector in two mode conversion experiments which are identical except for the ratio of $n_{^3\mathrm{He}}$ to n_e are shown in Fig. 6.7. While these discharges have virtually identical RF power, neutron rate and neutral beam injection power, the loss on the lost alpha detectors at 60° and 90° is as much as ten times higher for $n_{^3\mathrm{He}}$ = 0.20 n_e than the losses for the case corresponding to $n_{^3\mathrm{He}}$ =0.15 n_e, which appears to be at the level of first orbit losses. The loss observed is much larger than the losses which are typically present in rf discharges without mode conversion [DARROW et al. 1996], where as much as 5 MW of RF power only modulated the first orbit losses by 30%.

Discharges without deuterium neutral beams, but otherwise identical to the discharges in the last section, show no sign of enhanced loss, implying that the loss is not due to an energetic ^{3}He tail formed by RF minority heating (which in any event would be more likely for the case with lower $n_{^3\mathrm{He}}$ not higher $n_{^3\mathrm{He}}$). Thus, the loss must be due to either the deuterium beams or one of the charged fusion products that is present interacting with the MCIBW.

Figure 6.8 shows the gyroradius of these losses compared to typical first orbit losses, showing the MCIBW induced loss is at significantly larger gyroradius. This implies significant heating of the lost particles is occurring. The mean energy of the lost particles would be about 1.5 MeV if the loss is DD tritons and 2.25 MeV if the loss is heated beam deuterons. The pitch angle distribution suggests that the loss is strongly peak around the passing trapped boundary. Data at other detectors also shows a loss at the passing trapped boundary; however, the loss at lower poloidal angles is typically at lower gyroradius.

Initially, there was a great deal of uncertainty about which fast ion species was be-

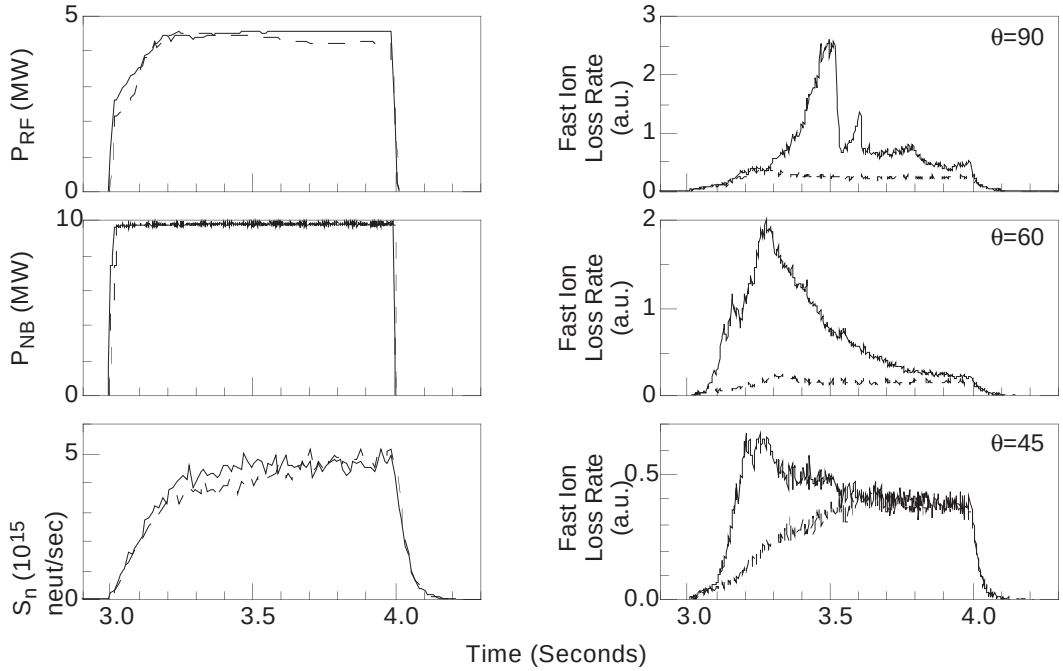

Figure 6.7: Two discharges which are nearly identical except for $n_{^3He}/n_e$ = 0.20 for the solid line and $n_{^3He}/n_e$ = 0.15 for the dashed line. Plotted are the ICRF power, the D neutral beam injection power, the neutron signal which is a measure of the rate of fusion, and hence CFP production, and the signal at the lost alpha detectors at 90°, 60°, and 45°. Figure from D. Darrow.

ing lost. The lost alpha detectors could provide little help in identifying the lost species. In a D^3He plasma with D beams there are 5 different charged fusion products present along with the beams themselves. It was found that, when tritium beams were injected instead of deuterium beams, the loss went away. This tended to rule out energetic α-particles since they would be present in both discharges. Both 3 MeV and 14 MeV protons and 800 keV ^{3}He ions were deemed unlikely to be the lost species since respectively there would be difficulty satisfying wave particle resonance, the sensitivity of the detector was low at very high energies, and the high collisionality left no time for the particle to be accelerated and ejected. This left just the 1 MeV DD tritons and the 100 keV beam deuterons as candidates for the lost species.

Tritons were attractive since the losses would only require about 0.5 MeV of heating. In contrast the D beams would have to be heated on the order of 2 MeV before being lost. Furthermore, typically when a large tail is pulled out on a distribution (such

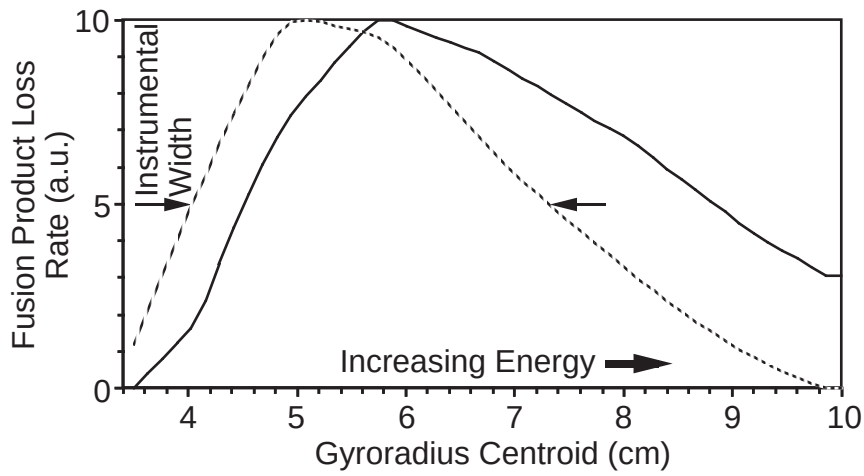

Figure 6.8: The gyroradius distribution of the MCIBW related loss (solid line) and first orbit loss (dashed line) at the 90° detector. The first orbit losses peak at the birth energy of the CFP's, in this case from the DD and D^3He reactions, and the width is representative of the instrumental width. The losses associated with the MCIBW peaks at higher gyroradius and is broader, indicating significant heating of the fast ions that are lost. Figure from D. Darrow.

as the D beams) by RF, magnetics measurements register an anisotropy in the plasma stored energy (i.e. the perpendicular stored energy increases). No such signature was observed. Finally while RF tails can get up to high energies, one would probably expect that if the D beams were being accelerated, the gyroradius distribution would peak at low gyroradius and trail off at higher gyroradius.

On the other hand, D beams could more easily produce the large loss levels seen in Fig. 6.7. If, typically, the charge fusion product losses are around 10% of the number of CFP's, then if the loss is 10 times bigger than first orbit loss almost every CFP would have to be expelled. This argument requires a detailed knowledge of the loss distribution on the wall. If, in fact, the MCIBW losses were localized toroidally or poloidally, this argument would not be a refutation of T loss.

Further evidence on the nature of the loss was provided by "beam blip" experiments, as illustrated in Fig 6.9. These experiments involved blips of neutral beam injection concurrent with the RF. The beam blip was useful for a number of reasons. It provided a distribution of beam particles and fusion products which had a narrow spread in

energy, as opposed to the slowing down distribution which develops after the beams have been on a while. Furthermore it allows the isolation of such things as the time it takes for a particle to get out, which will shed light on the diffusion coefficient. By comparing the lost alpha signal integrated over the discharge to the first orbit losses which were large only early in the discharges, when the CFPs were being produced by the deuteron beams, a more accurate estimate of the ratio of total losses to number of charged fusion products could be obtained. These comparisons show more convincingly that many more particles were coming out than were born as CFPs, again pointing towards beam deuterons as the lost species.

The dilemma regarding which species is lost species is resolved by the data in Fig. 6.9, which shows very different behavior when the D beams are injected cogoing and countergoing to the current, with symmetrically phased RF(both signs of n_ϕ present). Neither fusion product distribution would be affected significantly by the direction of the beams, thus identifying heated beam deuterons, in particular, countergoing beam deuterons, as the dominant lost species.

Besides the strong heating and large losses, there are other signatures of the mode conversion losses. In particular, the losses seem to depend strongly on the location of the mode conversion layer as shown in Fig. 6.10. This plot shows the losses of fast ions to the detector vs. the toroidal field for fixed ratio of $n_{^3\mathrm{He}}$ to n_e. Recall from Fig. 6.2 that as we increase the B field at fixed ^{3}He density the mode conversion layer moves out.

Because of their importance, losses in the presence of phased waves and the dependence of the losses on input power are discussed in Chapter 7.

Finally while the losses for the most part behave in the manner shown here, the losses are not completely reproducible [DARROW 1996]. An exhaustive search of what appear to be the relevant plasma parameters has not come up with an explanation.

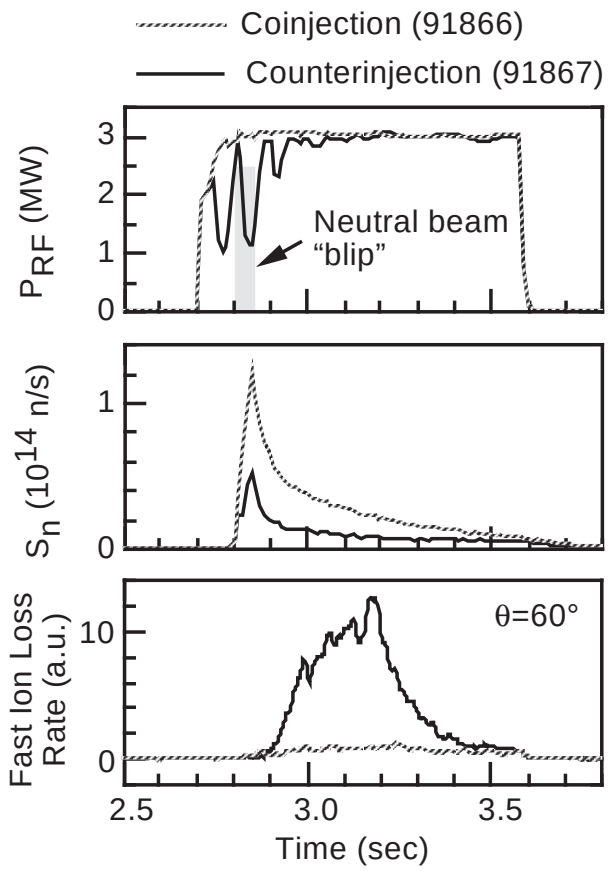

Figure 6.9: Losses of cogoing and countergoing deuterium beam ions in the presence of symmetrically phased RF. Discharge at 4.8 T, 1.4 MA with mode conversion layer on axis. Note that even for the cogoing beam injection, which shows no anomalous loss, the neutron production persists for much longer than the beam slowing down time. Figure from D. Darrow.

6.3 Theoretical Considerations

We now attempt to lay a theoretical basis for analyzing and drawing conclusions from these experiments, by considering some simple elements of this problem, namely the resonance condition, and the path in ε, μ, and P_ϕ space that the particle must follow, and the physics of the MCIBW. Surprisingly, we are able to make some predictions on the basis of these simple calculations. In the next section these simplified models will be verified with simulation results. For now we neglect the effect of collisions, returning to it in Chapter 7.

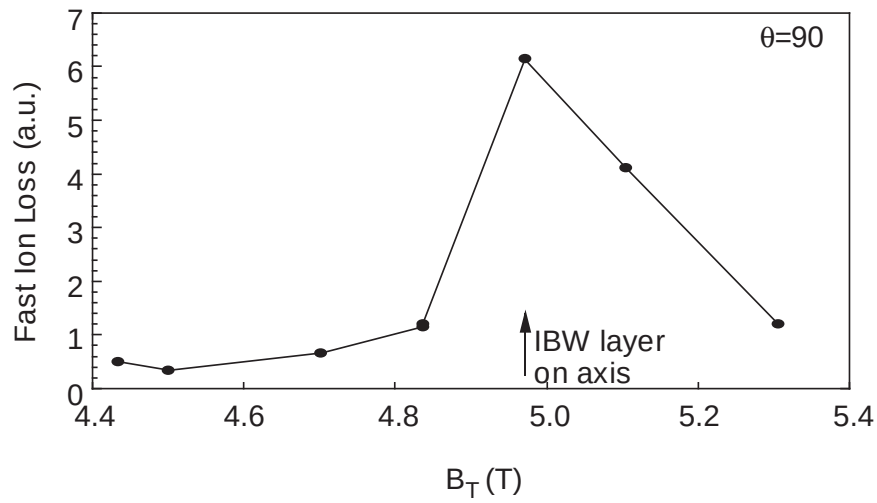

Figure 6.10: Relative losses at the 90° detector for $n_{^3\mathrm{He}}/n_e$ = 0.15 versus the toroidal field. The losses peak when the mode conversion layer is thought to be closest to the axis. Figure from D. Darrow.

6.3.1 IBW wave physics

The details of the MCIBW are important to understanding the fast ion losses. In particular, the MCIBW losses will be sensitive to the location of the mode conversion layer, the extent over which the wave exists, the mode conversion efficiency and coupled RF power, the $k_{\|}$ and $k_{\perp}$ spectrum, the n_{ϕ} spectrum, the spot size, and the electric field amplitude. We begin by giving an overview of how each of the above might be determined, and then explain how the wave is described in the simulation.

The spectrum of modes launched by the antenna can be analyzed by Fourier decomposing in toroidal mode number, since, as the mode propagates, and even under goes mode conversion, the toroidal mode number is conserved. This is true due to the axisymmetry of the tokamak, and remains true even though $k_{\perp}$, k_{θ}, $k_{\|}$ all may change rapidly as the mode propagates through the tokamak. The ICRF antennas used for these experiments in TFTR could launch a fast wave either with symmetric or asymmetric toroidal phasing. For symmetric phasing, which is used in mode conversion heating experiments, and the majority of the MCIBW- fast ion interaction experiments, the spectrum of wave power peaks at about $n_{\phi} \approx \pm 36 - 43$. For asymmetric- phasing, which is used for mode conversion current drive, the majority of the RF power (approximately 75%) can be put

into the lobe in the desired direction with $n_\phi \approx \pm 22 - 25$, however about 25% of the power is directed in the opposite direction with $n_\phi \approx \mp 43$.

The MCIBW is launched into the plasma as the fast wave from antennae on the low field side edge of TFTR. The spectrum of launched $k_\parallel$'s (equivalently, launched n_ϕ) can be estimated by taking the Fourier transform of the current distribution in the antennae and the images currents. The spectrum can be very simple for a single antenna, or very complicated for multiple antennas, especially considering that the relative phasing between the antennas, could vary during the discharges. In Fig. 6.11 the spectrum of $k_\parallel$ for a single antenna with symmetric phasing is shown. In Fig. 6.12 the spectrum for two antennae, symmetrically phased is shown.

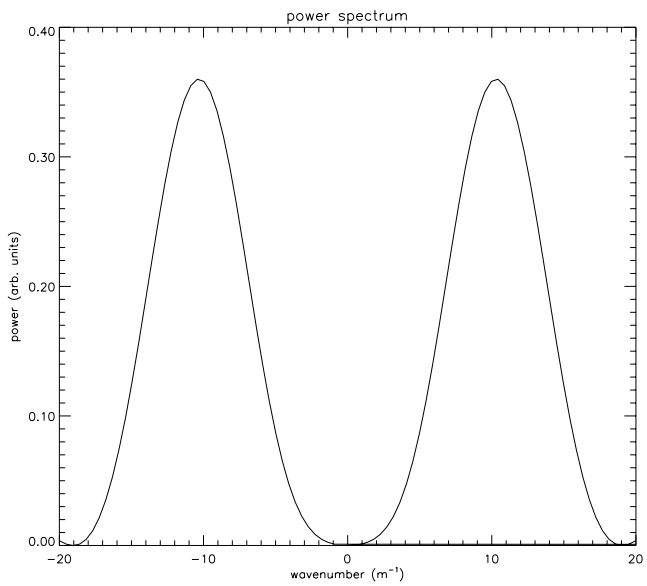

Figure 6.11: Spectrum of coupled power for the TFTR bay M antenna, which was used in the 43 MHz D^3He mode conversion experiments, assuming symmetric phasing.

Once the coupled spectrum has been determined, it is necessary to understand the mode conversion of the fast wave. Majeski, Phillips, and Wilson (1994) showed the effect of the cutoff-resonance-cutoff triplet on the mode conversion efficiency is quite strong. The power mode-converted in these cases is very dependent on parallel wavelength, since it depends on the amplification of the electric field at the mode conversion layer which can occur when there is a standing wave pattern formed due to the effect of the high field

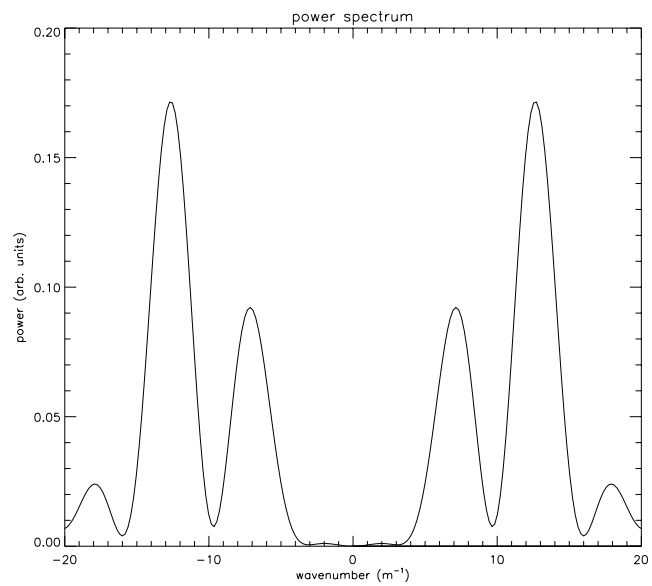

Figure 6.12: Spectrum of coupled power for the TFTR bay L and M antenna, assuming symmetric phasing and no phase difference between the antennas.

side cutoff. Thus they show an example for DT mode conversion which has nearly 80% mode conversion efficiency for $k_{\|}$ equal to 7 or 11 cm^{-1} but almost no mode converted power at 8.5 cm^{-1}.

After the wave mode converts to the MCIBW it is necessary to use ray tracing to follow it. Valeo and Fisch (1994) have done this for the case of DT mode conversion, keeping the important effects of a poloidal field and arbitrary $k_{\perp}\rho_i$. They found that the waves were highly dispersive and typically damp in a narrow region around the mode conversion layer. They also found very interesting behavior depending on the direction of the poloidal field compared to the gradient of the magnetic field. This arises from the shift in $k_{\|}$ which occurs due to the projection of k_X onto the poloidal field (where k_X is the wavenumber in the direction of the magnetic field gradient). For the case where the shift adds to the launched $k_{\|}$, the only effect is for the wave to damp more quickly on the electrons due to the higher $k_{\|}$ and therefore lower values of $\omega/k_{\|}v_{the}$. However in the case where the shift is in the opposite direction of $k_{\|}$, the wave moves to smaller and smaller $k_{\|}$. As it does so, its phase velocity gets very large, and thus the wave cannot damp on electrons. However, as the wave moves to yet higher field, k_X increases further and the wave phase velocity reverses directions, i.e. $k_{\|}$ flips at high enough k_X! After $k_{\|}$

flips, it turns around and starts heading towards lower B. This reversal in space means that the group velocity goes to zero, and therefore power and electric field amplitude builds up at this point in the plasma. Eventually the wave reaches $k_\perp \rho_{thi} \approx 1$ and it damps on ions or $k_\parallel$ turns around enough that it can damp on the electrons again. Which of these occurs depends on the location of the mode conversion layer, and the relative temperatures of the ions and electrons. As pointed out above, significant ion heating has been observed in DT mode conversion experiments, consistent with this picture.

While some of these things are easy to calculate, and have been verified experimentally (such as the mode conversion layer location), others require full wave codes to determine correctly (like the spotsize and the mode conversion efficiency versus $k_\parallel$), and still others are open areas of research (the $k_\parallel$ and $k_\perp$ spectrum). Furthermore the experiments on TFTR were carried under a wide range of experimental conditions. It is the difference in losses between different experiments that one usually wants to understand. Even if the codes existed to calculate all of the above quantities once, doing so for each of the required experimental conditions would be a very large task. Thus in modeling the MCIBW, we have of necessity used a simplified description that captures much, but not all, of the relevant physics. For the physics that is not directly modeled, we have estimated the appropriate parameter and used that in the simulation. Future work should include full 3-D ray tracing of the MCIBW in the magnetic geometry, some accounting for the dependency of the mode conversion efficiency on the $k_\parallel$, and full-wave analysis for determination of the spot size at the mode conversion layer.

For symmetric phasing, we assume that two n_ϕ in each direction are sufficient to account for the antenna spectrum and the filtering effect of the mode conversion itself. Typical n_ϕ's for the ray tracing were $n_\phi = \pm 18, \pm 43$ corresponding to $k_\parallel$ at the edge (R = 360 cm) of .05 cm^{-1} and .12 cm^{-1}. The results of the simulation were found to be very sensitive to the assumed spotsize of the wave, so this parameter was typically varied between 15 and 30 cm.

For the details of the spectrum after mode conversion we have used the ray tracing code of Valeo and Fisch (1995). Here, we summarize the physics included in this code.

Consider the full hot plasma dispersion relation in a uniform magnetic field [STIX 1992]:

$$\begin{vmatrix} \epsilon_{xx} - n_z^2 & \epsilon_{xy} & \epsilon_{xz} + n_x n_z \\ \epsilon_{yx} & \epsilon_{yy} - n_x^2 - n_z^2 & \epsilon_{yz} \\ \epsilon_{zx} + n_z n_x & \epsilon_{zy} & \epsilon_{zz} - n_x^2 \end{vmatrix} = 0 \tag{6.4}$$

where $\mathbf{B} \wedge \hat{\mathbf{z}} = 0$, $\mathbf{n} \cdot \hat{\mathbf{y}} = 0$. Assuming $(k_\perp \rho_i)^2 \sim 1$, $\omega \sim \Omega_i$, $k_\parallel << k_\perp$, $\beta << 1$ and ω not too close to Ω_i, we can justify the neglect of ϵ_{xz} compared to $n_x n_z$, ϵ_{yy} compared to n^2, and of ϵ_{yz}. With these assumptions, we obtain,

$$\Delta \equiv \epsilon_{xx} - \frac{n_z^2 \epsilon_{zz}}{(\epsilon_{zz} - n_x^2)} - \frac{{\epsilon_{xy}}^2}{n^2} = 0 \tag{6.5}$$

The strongest variation of Δ with position is through the variation in B. This is then solved in the ray approximation, including only the inhomogeneity in B and a "poloidal" field: $\mathbf{B} = B_0 R/x(\hat{z} + \theta\hat{x})$ The local parallel wavevector is

$$k_\parallel = k_{\parallel 0} + \theta k_x \tag{6.6}$$

The approach taken was to run the code off line and store the results in a file, which was then processed further and output into a file which was read in by the COM simulation. Important input parameters for this code are the k_z, the plasma temperature and density of both species, the toroidal field strength, and the poloidal field strength. Typical values of the temperature and total density were taken and the code was run for varying ^{3}He concentrations, launched k_z, and toroidal field strength. Furthermore a typical value of the poloidal field strength dotted with the wave vector at the mode conversion layer, θ in the notation above, was chosen. Of course in reality this value varies considerably with height above the midplane where it is zero, and, to further complicate the situation different ions cross the mode conversion layer at different heights, so that each interacts with the MCIBW at a different value of θ. While it would be possible to write the COM simulation, so that it could handle several values of θ, this was not done. Instead a value of θ=0.05 was chosen for most runs. A plot of θ versus height above the midplane for different mode conversion layer positions is shown in Fig. 6.13.

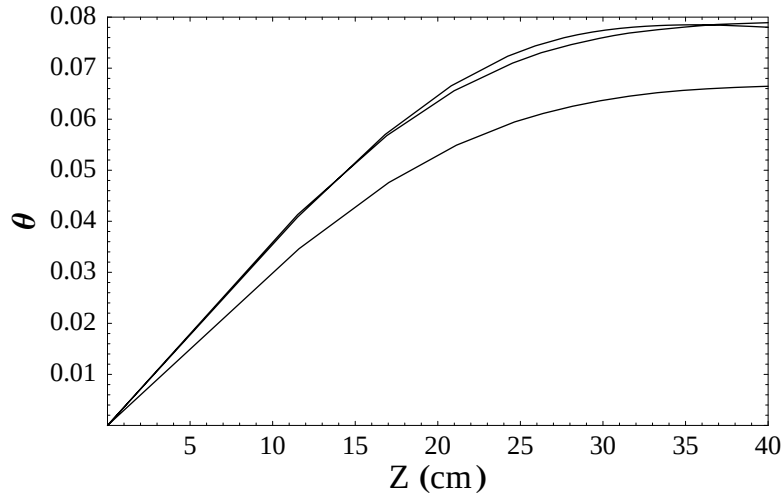

Figure 6.13: Plot of the projection of the poloidal field in the $\hat{R}$ as a function of height along the mode conversion for a typical D^3He mode conversion experiment. The three lines correspond to the mode conversion layer being on axis and 15 cm to either side. The sign of θ would be negative for z< 0.

Typically, four different sets of data were necessary to simulate an experiment if both cogoing and countergoing beams were present. These correspond to $(n_\phi > 0, \theta > 0)$, $(n_\phi > 0, \theta < 0)$, $(n_\phi < 0, \theta > 0)$, $(n_\phi < 0, \theta < 0)$. As discussed below in order for the beam deuterons to interact, different parts of the spectrum must be present. Consider a countergoing deuteron at 100 keV. These deuterons will interact with waves which have $(n_\phi < 0, \theta < 0)$. However, as the deuterons increase their energy they require less Doppler shift to interact, the $k_\|$ they interact with gets closer to zero, and thus they interact with the waves which have $(n_\phi < 0, \theta > 0)$. It is important that the $k_\|$ spectrum of these two sets of waves overlap, otherwise deuterons will reach some maximum energy, and will not be able to resonate with the wave any more, and no deuteron losses would be possible, which is inconsistent with the observations.

In Fig 6.14 a typical ray tracing diagram for the MCIBW is shown, similar to those in Valeo and Fisch (1994). The solid lines represents three different n_ϕ's (-37,-43,-49) ray traced for $\theta > 0$. The dashed lines are n_ϕ's (-17,-22,-27) for $\theta < 0$. While the rays are traced for several initial n_ϕ, this is just to get a feel for the sensitivity of the waves to changing n_ϕ. Typically only one of each group of three would be used in the COM simulation. In Fig 6.15 the k_X of the wave is plotted vs. $k_\|$. The linear

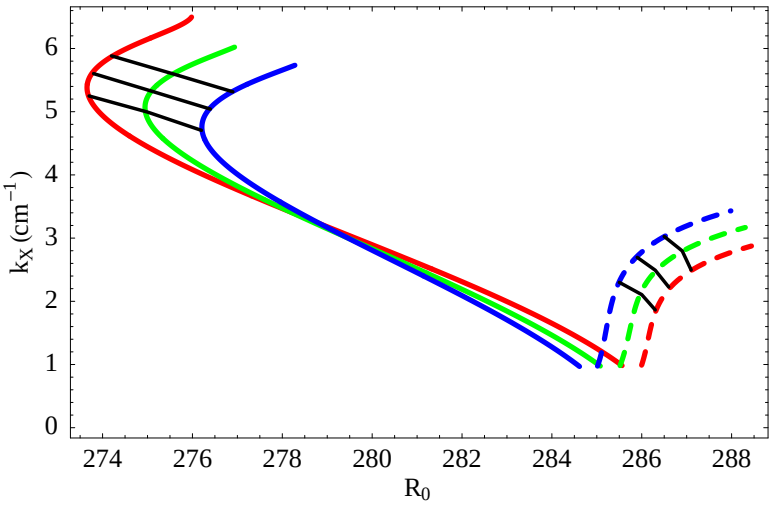

Figure 6.14: k_X vs. R from ray tracing of the MCIBW for $n_{^3He}/n_e$ = 0.15, B = 5.0 T. The two sets of waves shown correspond to n_ϕ = -37,-43,-49 with$\theta = 0.05$(solid) and n_ϕ = -17,-22,-27 for $\theta = -0.05$ (dashed). The lines across the rays show the points along the ray where 75%, 50%, 25% of the initial power is left.

relationship seen comes from Eq. (6.6). Note that because each ray starts with k_X near 1 cm^{-1}, two waves with the same n_ϕ interacting with opposite signs of θ will not overlap, but instead their $k_\parallel$ will be separated by approximately $2\theta k_{X0}$. In light of the importance of spectrum overlap discussed above in ensuring the deuterons can satisfy resonance as they are heated it is necessary to use two groups of n_ϕ's with opposite signs of θ in order to ensure that the spectrum overlaps. Thus the rays shown correspond to two different groups of n_ϕ with opposite signs of θ.

In addition to the spectrum of waves present determining whether or not the beam deuterons can resonate, the details of the wave determine the strength of the deuteron MCIBW interaction. In Fig. 6.17 the electric field amplitude, as determined from the power flux to the mode conversion surface and the group velocity of the waves, versus the $k_\parallel$ of the wave is plotted. The peaks for the rays with $\theta < 0$ arise from the build up in amplitude which occurs when the group velocity of the wave in the R direction goes to zero, and the wave turns around. Eventually, at high enough $k_\parallel$ the waves damp and the electric field amplitude goes to zero.

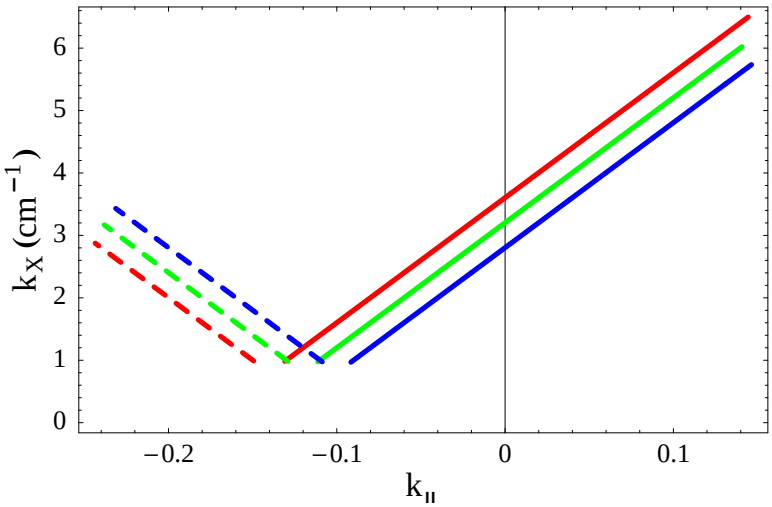

Figure 6.15: k_X vs. $k_{\parallel}$ for the MCIBW for $n_{^3\mathrm{He}}/n_e$ = 0.15, B = 5.0 T.

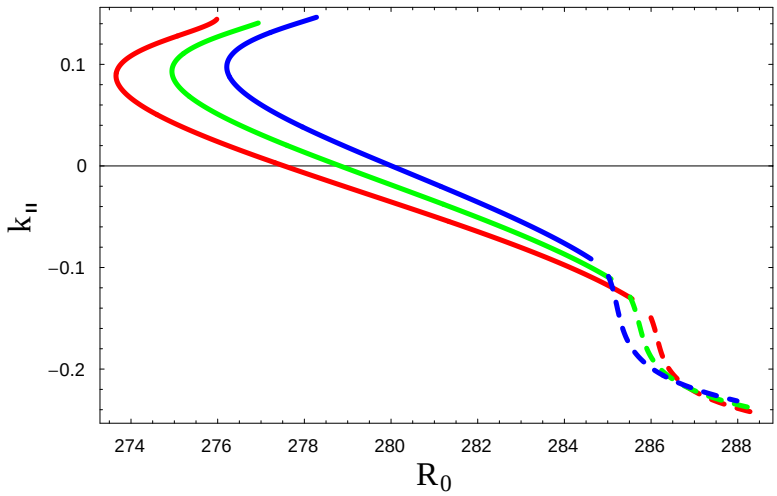

Figure 6.16: $k_{\parallel}$ vs. R for the MCIBW for $n_{^3\mathrm{He}}/n_e$ = 0.15, B = 5.0 T.

6.3.2 Resonance condition

The resonance condition for a particle to interact with the MCIBW is

$$k_{\parallel} v_{\parallel} = \omega - n\Omega(B), \tag{6.7}$$

where n is the harmonic with which the particle is interacting. As discussed above the MCIBW usually damps within a few cm of the mode conversion layer, thus we can take

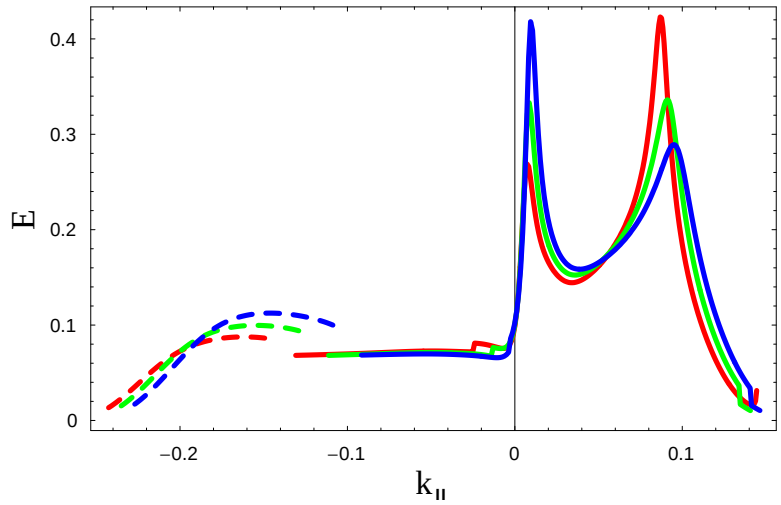

Figure 6.17: Electric field amplitude in esu vs. $k_{\|}$ for the MCIBW for $n_{^3\mathrm{He}}/n_e$ = 0.15, B = 5.0 T, assuming 1 watt/cm^2 incident flux.

B to be B_{mc}, the value of the field at the mode conversion layer and require the particle to be resonant at the mode conversion layer if it is resonant at all. Dividing through by ω, and substituting $\nu = n\Omega_{mc}/\omega$, we get,

$$k_{\|}\lambda = \frac{\omega}{\sqrt{2\varepsilon}}(1-\nu). \tag{6.8}$$

Table 6.2 summarizes the values ν and $k_{\|}\lambda$ which are typical for different fast ion species resonating with the MCIBW. Note that due to the localized nature of the MCIBW and the large distance from the mode conversion layer to the resonance layer, only ions going in one direction will be able to resonate with one $k_{\|}$. From the table we see that deuterons and α-particle interact with $k_{\|}$ of the same sign as their parallel velocity while tritons interact with $k_{\|}$ of opposite sign for mode conversion in D^{3}He.

This is unlike the case of ICRF minority heating in which the wave is able to interact with minority ions going in opposite directions on the low field and high field side of the resonance layer for a fixed $k_{\|}$. For most discharges we are discussing, symmetric phasing was used, and so both signs of $k_{\|}$ are present and thus ions going in both directions can be resonated with. This does not mean cogoing and countergoing ions will behave the same under interaction with the MCIBW. The motion of the particles in ε, μ, and P_ϕ space depends on the n_ϕ that the particle is interacting with, and this creates

an important asymmetry that will be discussed in Sec. 6.3.4. In some discharges with phased RF there were fast ion losses, which will be discussed in Sec. 7.1.

Species	n	$n_{^3He}/n_e$ = 0.20, B = 5 T		$n_{^3He}/n_e$ = 0.15, B = 5 T	
		ν	$k_{\|}\lambda$(cm^{-1})	ν	$k_{\|}\lambda$(cm^{-1})
100 keV D	1	0.83	0.15	0.81	0.164
3.5 MeV α	1	0.83	0.036	0.81	0.04
100 keV T	2	1.11	-0.11	1.07	-0.08

Table 6.2: The $k_{\|}\lambda$ which different fast ion species resonate with is shown for two cases, both with B = 5 T, but one case has 15% ^{3}He and the other 20%. Since the mode conversion layer is closer to the deuterium/α-particle resonance for higher ^{3}He concentration the required $k_{\|}\lambda$ is lower in this case. For tritium, which interacts at the second harmonic, the opposite is the case. Since ν is independent of toroidal field for the same ^{3}He concentrations, the required $k_{\|}\lambda$ is independent of the toroidal field.

While the α-particles and the beam tritons can resonate with waves in the launched spectrum, the $k_{\|}\lambda$ required for deuterons is large compared with the spectrum which is launched into the plasma during symmetric phasing. For a typical beam deuteron pitch of 0.8 the $k_{\|}$ required is $\approx 0.21cm^{-1}$, almost twice the peak of the antenna spectrum. As explained above, the MCIBW does upshift as well as downshift. The combination of this upshift and the presence of some deuterons with λ greater than 0.8 allows a fraction ($\sim$ 5-20%) of the beam ions to resonate with the MCIBW at 15% ^{3}He levels. The typical number of deuterons that interact with the wave is discussed in more detail in Section 6.4 below. The resonance condition is more easily satisfied for the deuterons at higher ^{3}He fractions.

Given the $k_{\|}$ spectrum from the ray tracing, and knowing the details of the beam slowing down, one can estimate the maximum time a beam particle could slow down before it fell completely out of resonance. This estimate can be compared to beam blip experiments which were done to estimate the “threshold energy” of interaction. These experiments varied the amount of time between a neutral beam blip and the start of the RF, thus allowing the deuterons to slowdown before interacting with the wave. The fast ion losses can then be plotted versus the time delay. The experiments suggested a delay

of 20 to 40 ms was sufficient to wipe out almost all the losses, consistent with a threshold energy of 65-80 keV [HEETER et al. 1997]. From the previous section the maximum $k_\parallel$ at which the MCIBW has any significant power is about 0.22 cm^{-1}. If we require the particle to have a pitch of at most 0.9 (few deuterons which are on flux surfaces near the axis have pitch higher than this), the energy threshold can be estimated about 70 keV, giving some confidence in the ray traced MCIBW.

Finally, while one typically thinks of slow particles having a hard time satisfying resonance, for the MCIBW it may also be the case that fast particles cannot satisfy resonance. Since the MCIBW exists only over a narrow layer it is not necessarily true that a fast particle can resonate with a given $k_\parallel$ (for ICRF, a fast particle can satisfy resonance with a given $k_\parallel$ at location given by $n\Omega = \omega - k_\parallel v_\parallel$). Thus when beam deuterons are being heated, it is important that the spectrum be broad enough that both a 100 keV deuteron and a 1.6 MeV deuteron (which is 4 times faster but will more typically have a $v_\parallel$ twice that of 100 keV deuterons since most of the energy gain is in perpendicular energy) can resonate. Fortunately, the up shifts and down shifts of the MCIBW spectrum discussed above are sufficiently broad to cover this range again lending some credibility to the ray tracing.

6.3.3 Trajectory in $v_\perp, v_\parallel$ space

As discussed in Chapter 4 particles interacting with the MCIBW diffuse along a line in ε, μ, and P_ϕ space. Recall from Eqs. (4.57) and (4.58) that

$$\frac{d\mu}{d\varepsilon} = \frac{n\Omega_0}{B_0\omega} \tag{6.9}$$

(where all expressions have been returned to their dimensional form for clarity). This relationship leads to interesting implications for the fast ion MCIBW losses.

For ICRF minority heating Eq. (6.9) is the cause of "resonance localization" [HSU et al. 1984; HAMMETT 1986], which is the tendency for particles interacting with an ICRF wave to have their banana tips approach the resonance layer, i.e. where $\omega = n\Omega$. For ICRF when $k_\parallel$ is small particles will resonate at the resonance layer, and

Eq. (6.9) implies that $\Delta\mu B_{res} = \Delta\varepsilon$, or that particles get their kicks in perpendicular energy only. For finite $k_\parallel$ particles receive their kicks away from the resonance layer, and thus get kicks in parallel energy as well as perpendicular energy. However, resonance localization still occurs, as discussed by Hammett (1986), i.e. the effect of the kick away from the resonance layer is the same as if the particle had gotten a kick in perpendicular energy when it crossed through the resonance layer. This can be seen from the form of Eq. (6.9) which is obviously independent of $k_\parallel$.

As we shall see, considering the case of fast particles interacting with the MCIBW leads to a more complicated view of "resonance localization". Examine a particle satisfying resonance with the MCIBW. Mark its location in $v_\perp$, $v_\parallel$ as passes by a particular mod B surface (choosing the $\omega = \Omega$ resonance layer is not convenient for MCIBW since this is far from where the particle receives its kicks, and the particle may not even pass through this layer).

$$\varepsilon_\perp = \mu B = \mu_0 B + \frac{n\Omega}{\omega}\Delta\varepsilon \tag{6.10}$$

$$\varepsilon_\parallel = (\varepsilon - \mu B) = \left(\varepsilon_0 - \mu_0 B + \Delta\varepsilon\left(1 - \frac{n\Omega}{\omega}\right)\right) \tag{6.11}$$

where $\Delta\varepsilon$ is the total energy exchanged with the wave, Ω is the gyrofrequency at the mod B surface we are interested in. This can be written in terms of the pitch and velocity and using $\nu = n\Omega_{mc}/\omega$ from above:

$$\varepsilon_\perp = \mu_0 B + \frac{\nu B}{B_{mc}}\Delta\varepsilon \tag{6.12}$$

$$\varepsilon_\parallel = \left(\varepsilon_0 - \mu_0 B + \Delta\varepsilon\left(1 - \frac{\nu B}{B_{mc}}\right)\right) \tag{6.13}$$

$$\lambda = \sqrt{1 - \frac{\nu B}{B_{mc}} - \frac{v_0^2}{v^2}\left(1 - \lambda_0^2 - \frac{\nu B}{B_{mc}}\right)} \tag{6.14}$$

$$v = \sqrt{2(\varepsilon_0 + \Delta\varepsilon)} \tag{6.15}$$

Note that $v_\perp/v = \sqrt{1-\lambda^2}$, $v_\parallel/v = \lambda$. λ will be real so long as the particle reaches the mod B surface we are discussing. These equations are valid for ICRF if B_{mc} is taken to be B_{res} and ν is determined at the resonance layer.

Consider Eq. (6.12), and take $B = B_{mc}$. From Table 6.2, for T interacting with the MCIBW in D^3He (at the second harmonic), ν is greater than 1. For $\Delta\varepsilon$ positive, we see that the gain in the triton perpendicular energy is greater than $\Delta\varepsilon$, i.e. on gaining energy from the MCIBW the triton converts some of its parallel energy to perpendicular energy. For the deuteron, ν is less than one and therefore as it gains energy from the wave it gains both perpendicular and parallel energy. This implies that a triton interacting with the MCIBW will cross the passing trapped boundary quickly while a deuteron will take longer to cross the passing trapped boundary. Recall that a major loss mechanism for fast ions, and the only one which can readily explain losses at the 90° detector, is crossing the passing-trapped boundary.

Consider the two separate cases of a countergoing 100 keV beam deuteron interacting with the MCIBW in D^3He, with the mode conversion layer on axis and the same deuteron resonating with the fast wave (i.e. no mode conversion) with the deuterium resonance layer on-axis (note that these two examples must have different RF frequencies). The solid line in Fig 6.18 shows the deuteron's $v_\perp$ and $v_\parallel$ (normalized to its injection energy) when it crosses $B = B_0$, as determined by Eq. (6.14,6.15). The passing trapped boundary is marked by the dotted line, which turns into a solid line (marked $\theta(\varepsilon)$) at high enough energy, representing particles which are energetic enough to hit the wall when they become trapped. Additionally, trapped particles which are energetic enough to scrape off on the outer midplane are marked by the line 0°, (see the caption for more details). The particle interacting with the MCIBW, which in this case has $\nu B/B_{mc}$ of 0.8 heads to higher $|v_\parallel|$ as it gains energy, eventually crossing the passing trapped boundary after gaining 1.5 MeV and hitting the wall near the 90° detector. The deuteron heated by the fast wave, which has $\nu B/B_0$ of 1 crosses the passing trapped boundary after gaining just 300 keV of energy. Eventually it too gains more than a MeV and is scrapped off at the outer midplane. In Fig. 6.19 the poloidal projection of the deuteron orbit interacting with the MCIBW is shown.

Now let's consider the case of a 100 keV countergoing beam triton interacting with the MCIBW in D^3He, shown in Fig. 6.20. The MCIBW, which has $\nu B/B_{mc}$ of 1.07 for this case, pushes the triton across the passing trapped boundary before it reaches 400 keV, which is below the threshold energy at which the particle would hit the wall. This explains the absence of losses at the 90° detector when beam tritons are injected instead

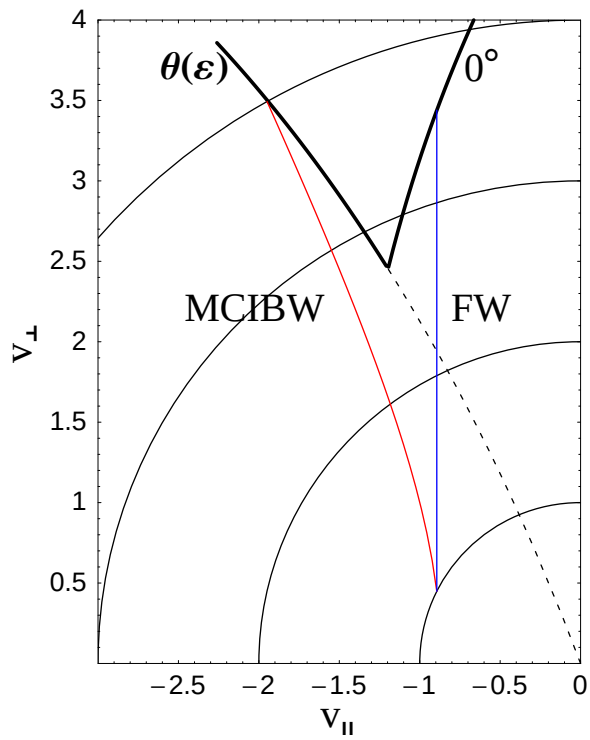

Figure 6.18: The trajectory of a 100 keV deuteron gaining energy from either the MCIBW or the fast wave (FW). The dotted line represents the passing trapped boundary at the point the plasma where $B = B_0$ and $\psi = 0.3\psi_{\text{wall}}$ (particles below are passing while particles above are trapped). At sufficiently high energy if a passing particle becomes trapped it will hit the wall. This is noted by the dark solid line marked $\theta(\varepsilon)$. The angle at which it hits the wall will depend on its energy, with lower energies corresponding to the outer midplane while higher energies correspond to the inner midplane. The solid line at an angle to the passing trapped boundary (marked 0°) denotes trapped particles whose energy is so large their orbits scrape off on the outer midplane (see Fig. 6.6C). If the boundary were plotted for a location with larger ψ, it would be shifted to higher pitch, i.e. for the same $v_\parallel$ it would be at lower $v_\perp$, and conversely if the we move the location closer to the center.

of beam deuterons. Unlike the fast wave heated deuteron, the triton does not even scrape off on the outer midplane (although it is close and a different initial condition might have scraped off and hit the wall). Eventually (when $\lambda \to 0$) the triton stops reaching the $B = B_0$ surface. The poloidal projection of the tritons orbit is shown in Fig 6.21. Interestingly, tritons have a maximum amount of energy they can gain from the wave, which can be seen from Eq. (6.15). Taking $B = B_{mc}$ we see that as the particle gains energy upon passing through the mode conversion layer, $\lambda \to \sqrt{1-\nu}$. But ν is greater

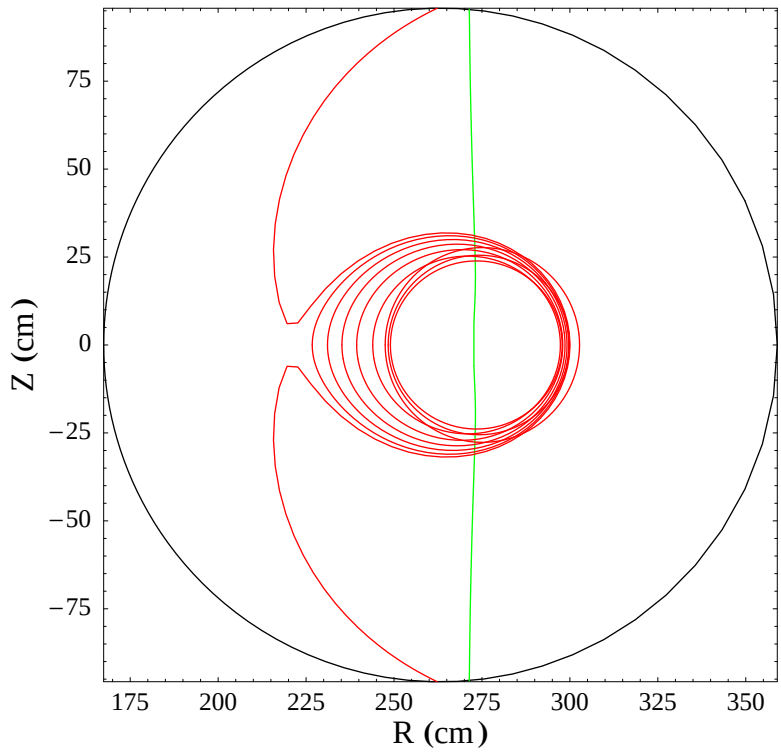

Figure 6.19: Poloidal projection of the orbit of a D interacting with the MCIBW.

than 1. Physically what will happen is that the triton will keep slowing down as it passes through resonance. But the triton *needs* a Doppler shift to resonate with the MCIBW. Thus as it slows down it will resonate with higher and higher $k_{||}$, which will have smaller and smaller amplitudes. Eventually the diffusion coefficient will be zero, effectively barring the triton from gaining any more energy. Another way of seeing this is that the triton is trying to get its banana tips to the resonance layer, but the resonance layer is on the low field side of the mode conversion layer, and the particle must pass through the mode conversion layer to interact with the wave. In contrast, the deuteron interacting with MCIBW is able to continue interacting so long as $k_{||}$'s are present which are lower (in absolute value) than the $k_{||}$ it interacted with initially. This is because the deuteron gains parallel energy as it gains perpendicular energy.

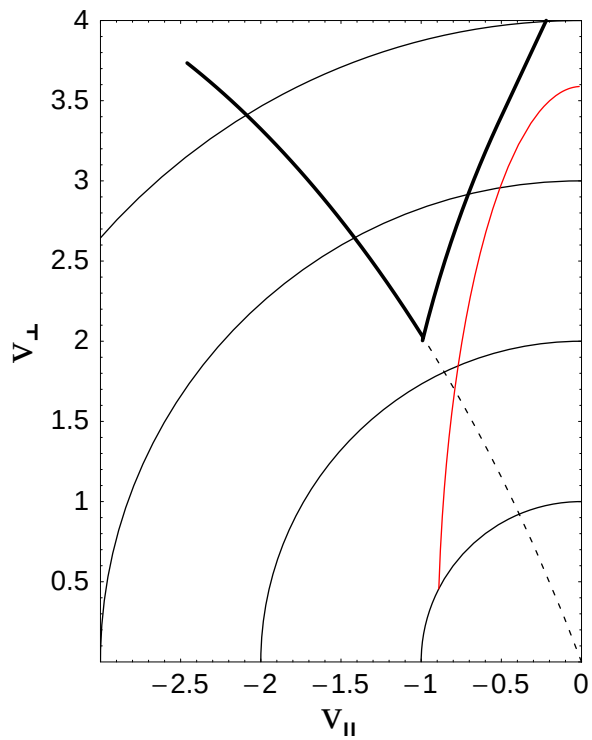

Figure 6.20: Passing/trapped boundary and trajectory of a 100 keV T interacting with MCIBW vs. $v_{\|}$ and $v_{\perp}$. See Fig. 6.18 for details.

6.3.4 Effect of n_ϕ on the particle trajectories

While the movement in $v_{\|}$ and $v_{\perp}$ plays a very important role in the particle losses, it does not tell the whole story. In particular, we have not addressed the interaction of cogoing particles with the wave, nor the importance of radial transport. The equation relevant to these issues is

$$\frac{dP_\phi}{d\varepsilon} = \frac{n_\phi}{\omega} \tag{6.16}$$

Since $P_\phi = Fv_{\|}/B - \psi$, and for low energies the $v_{\|}$ term is small, -P_ϕ can be treated somewhat like a radial variable. On this basis, interacting with a wave with $n_\phi/\omega > 0$ means that as the particle gains energy, P_ϕ increases, -P_ϕ decreases, hence, the particle moves in. Conversely, with $n_\phi/\omega < 0$, as the particle gains energy, P_ϕ decreases, -P_ϕ increases, the particle moves out. This notion is not precise due to the parallel velocity term in P_ϕ.

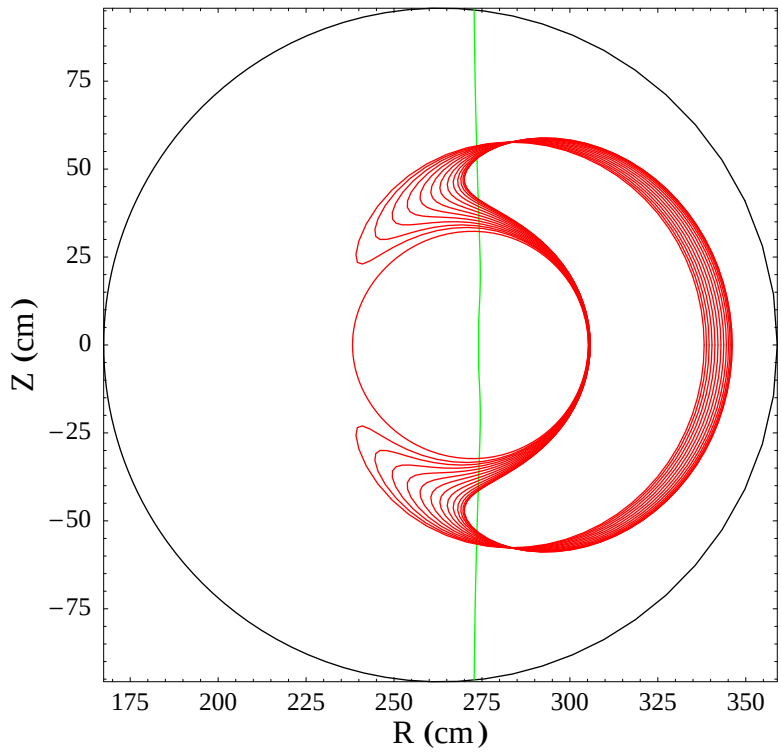

Figure 6.21: Poloidal projection of the orbit of a T interacting with the MCIBW.

This radial transport is important in the MCIBW particle dynamics for both co going and countergoing particles. While the motion in $v_{\|}$ and $v_{\perp}$ is specified, as discussed above, the location of the passing trapped boundary depends strongly on the position of the particle, (the farther out the particle is the larger the inverse aspect ratio leading to larger regions of trapped space). Thus if the particle is moved out while gaining energy it is likely to cross the passing trapped boundary sooner, conversely if it is moved in as it gains energy it will require a larger energy gain before the particle crosses the passing trapped boundary. In fact, once n_{ϕ} is above some value the particle will continue to move in as it gains energy, never crossing the passing trapped boundary. Figure 6.22 illustrates a countergoing deuteron interacting with the MCIBW assuming different values of n_{ϕ}. The amount of energy the particle gains before crossing the passing trapped boundary (or even if the particle crosses the passing trapped boundary) depends strongly on n_{ϕ}.

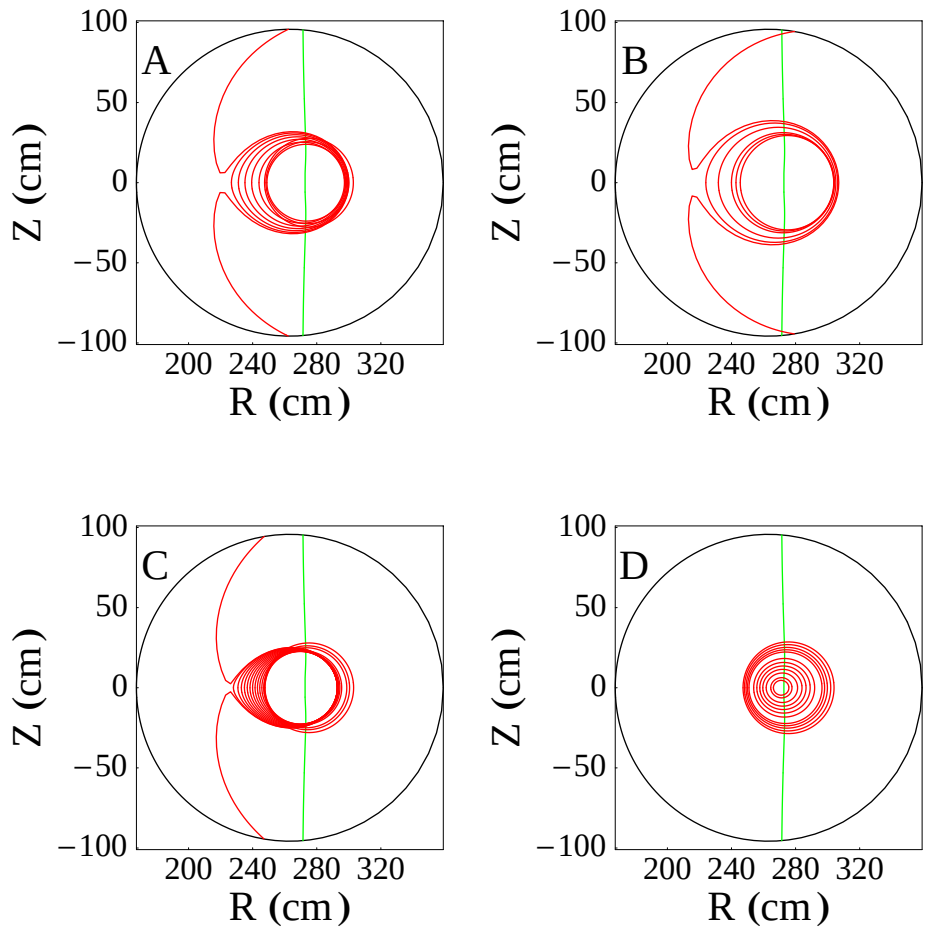

Figure 6.22: A countergoing deuteron interacting with the MCIBW for various values of $n_\phi < 0$. Graph A shows a particle interacting with two waves with $n_\phi = -18, -43$, which correspond to the values used in the ray tracing. It gains nearly 2 MeV before crossing the passing trapped boundary and hitting the wall near 90°. B shows the same particle interacting with two waves, both of which have $n_\phi = -60$. This particle gains 1.2 MeV and then crosses the passing trapped boundary hitting the wall to the low field side of the 90°. C shows the same particle interacting with two waves with $n_\phi = -23$. This particle gains 3.1 MeV and then crosses the passing trapped boundary hitting the wall to the high field side of the 90°. Finally, D shows a particle interacting with two waves with $n_\phi = -10$. This particle gains 2.5 MeV and can gain no more, because if it did it would stop crossing the MCIBW layer, and the wave cannot push the particle to a place where it can no longer interact with it.

Figure 6.23 illustrates a cogoing deuteron interacting with the MCIBW assuming different n_ϕ. Recall that because of the Doppler shift required for deuterons to resonate they interact predominantly with waves which have $n_\phi > 0$

The reason that particle D in Fig. 6.22 (particle B from Fig. 6.23) moves in (out) as it is heated, even though $n_\phi < 0$ ($n_\phi > 0$) can be understood as follows. Consider a particle interacting with a wave, with n_ϕ = 0 and therefore P_ϕ remains constant. As this particle is heated its orbit will move in if it is countergoing and out if it is cogoing and move out on the cogoing leg and in on the countergoing leg if it is trapped. This can be seen by consider the curve of the particle in B,ψ space as discussed in Ap. A in particular Fig. A.2. As the energy goes up the width of this the B_{orb} curve spreads out for μ/ε fixed (which these particles tend to have since μ/ε approaches a limiting value after enough energy is gained). This spreading makes the intersections with the inner and outer midplane move closer to the center for countergoing ions, closer to the wall for cogoing ions and spreading out the intercepts for trapped particles. For large enough $|n_\phi|$ this effect is overcome by the change in P_ϕ, and the expected behavior is recovered.

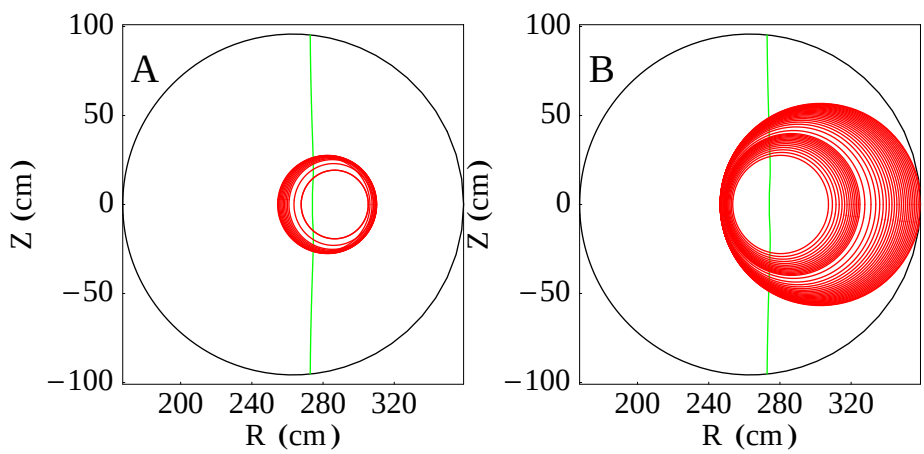

Figure 6.23: A cogoing deuteron interacting with MCIBW for two values of $n_\phi > 0$. Graph A shows a particle interacting with two waves with $n_\phi = 18, 43$, which correspond to the values used in the ray tracing. It gains nearly 2.3 MeV before it can proceed no further. B shows the same particle interacting with two waves with $n_\phi = 10$. This particle gains nearly 5 MeV and then hits the wall on the low field side.

6.3.5 Theoretical summary

By examining the theoretical issues involved in the MCIBW beam deuteron interaction we have been able to get some idea of the important issues involved in the experiments. On the basis of these insights, predictions are able to be made about the nature of the losses.

The way the heated particles move in velocity space leads to several observations consistent with the experiments. For example, the absence of magnetics signature of a large tail formation (which initially led to speculation that it could not be beam losses), could be due to the deuterons gaining parallel and perpendicular energy upon interacting with the wave, unlike the case of fast wave heating where the particles gain mostly perpendicular energy. Of course it can be difficult to see a tail via excess perpendicular energy when beams are injected, since the beam particles tend to have excess parallel energy. Also, the absence of a loss at lower gyro radii is seen is a consequence of the parallel heating of the deuterons, which significantly increases the energy they must obtain before they can cross the passing trapped boundary. At the same time the absence of significant losses during triton beam injection is explainable by the tritons movement in velocity space upon interaction with the MCIBW.

The toroidal mode number the particles interact with can also explain some of the observations. In particular, the absence of deuteron loss under co injection, even though the extremely long tail in the neutron production shown in Fig. 6.9 for coinjection suggests that the deuterons are being significantly heated. This can be seen to be consistent with the fact that cogoing deuterons interact with the $n_\phi > 0$ waves which tend to move the particles in as they heat them. This heating tends to move the hot cogoing deuterons to regions of high wave power (near the midplane) and high density, unlike the case for counter injection. Together with the natural tendency for cogoing beam particles to be better confined than countergoing ones, this is a possible explanation for the significantly longer neutron tail for coinjection than counter.

Finally we have seen that information about the waves may gleaned by the presence of the deuteron interaction, and its details. In the next section will expand upon these hypothesis by doing simulations of many particles under more realistic conditions.

6.4 Simulation Results

The previous section has given us a feel for what to expect when simulating the beam ion MCIBW interaction, and given a plausible explanation for some of the observed phenomena. However, to really be convinced that the model we have is describing some of what is actually going it is necessary to carry out full fledge simulation. These simulations significantly increase the number of effects modeled. For instance we use a distribution of beam ions, taken from the TRANSP code prediction of where incident neutral beam particles are ionized, to see the effects of a full fledged distribution of beam particles on the losses. We now include the effect of the spot height of the MCIBW which, we will see, will severely limit which particles can be ejected. Furthermore scans over various wave parameters will tell us how the losses change as the toroidal field is varied or the ratio of $n_{^3\mathrm{He}}/n_e$ changes.

However, we still do not have a complete description of the problem. To give one example, stochastic ripple diffusion is not included in these simulation. Its effect will not be important for the losses of counter going particles, since they remain passing until they receive their last kick and then hit the wall. Since they are not trapped while they are confined SRD (which only effects trapped particles) is not relevant. On the other hand some cogoing particles interacting with the wave may become trapped, and upon doing so they do not hit the wall. Thus some of them will be affected by SRD. The effect of SRD on these particles will be to diffuse their banana tips, possibly causing them to hit the wall somewhere on the outer midplane. An investigation by Boivin, Zweben, and White (1993) suggests that most of these losses would be to the 45° and 20 ° degree detectors.

The description of the wave physics is also not complete as discussed above. Many details of the waves enter into the strength and nature of the wave particle interaction, and thus we cannot expect a complete description of the problem given the limited amount of wave physics which is included. In particular, the results are sensitive to some of the assumed parameters, which a more detailed modeling of the wave would provide but is dealt with here by scanning over the parameters.

Finally, all of these simulations are done collisionlessly. Collisional slowing

down, which is the dominant collisional process for these ions, will strongly affect the number of deuterons which get out. Collisions will be discussed in the section on the diffusion coefficient in Chapter 7.

In this section, the interest lies in the poloidal and energy distribution of the losses, which the collisions do not strongly affect. For now the amplitude of the wave is set to some level and the particles are allowed to diffuse in the wave for a given amount of time. In comparing shots with different parameters the fraction of beam particles lost will be used as one gauge of the effect of the MCIBW on the discharge. This number should not be taken as an absolute number, as collisions and the wave amplitude will strongly affect the fraction of particles that are lost. However for relative comparison between two shots, it is a very useful number, since, presumably, collisions would affect the two discharges in similar ways.

Note, however, that even collisionlessly it is important to let the particles diffuse (as opposed to just heating the particle until it leaves the plasma). It is found that diffusive kicks in energy to the particle creates a different distribution of the losses than heating only kicks. This is not surprising, since in general the particles are interacting with more than one n_ϕ, and thus their diffusion is no longer along a line, but now along a surface in two dimensions. By always giving heating kicks the full two dimensional surface cannot be explored and the loss distribution is different than the diffusive case.

For these simulations, the beam distribution was taken from a TRANSP simulation of a TFTR discharge which included both co going and countergoing beam deuterons, with a current, toroidal field, and plasma density which was typical for the discharges with MCIBW. TRANSP has a post processor which will create a Monte Carlo distribution of beam orbits consistent with the beam deposition profile [McCune 1997]. Each orbit is specified by the square root of the normalized toroidal flux on the outer midplane, the pitch, and the energy. Typically 10000 particles were used. The magnetic geometry was obtained as described in Appendix A with the TRANSP characterization of the equilibria at a given time as an input. The geometry used had B = 4.8 T, I = 1.4 MA with a 12 cm Shafranov shift. Note that for one of the scans done below the toroidal field was varied from 4.4 to 5.3 T. In the actual experiment this scan was performed at constant current, however, for the simulations the geometry used was the same in each

discharge, but scaled by the toroidal field. This means that the simulation at 5.3 T is really a simulation of a B = 5.3 T, I = 1.54 MA discharge. The effect of this discrepancy is discussed below. The waves were ray traced using the code of Valeo and Fisch (1994) as described above.

6.4.1 Toroidal field scan

Before proceeding to the simulation results, we consider one of the implications of the preceding section's discussion on the losses as the toroidal field is varied. In Fig 6.24 the trajectories of a deuteron in $v_{\|}, v_{\perp}$ space at B = B_0 is shown. The three different lines correspond to the deuteron interacting with the MCIBW at three different toroidal fields, but with $n_{^3\mathrm{He}}/n_e$ = 0.15 in all three cases. At the lowest field, 4.7 T, the mode conversion layer is on the high field side of the magnetic axis, while at 5.3 T the mode conversion layer is on the low field side of the axis. Examining Eq. (6.15), note that $\nu = n\Omega_{mc}/\omega$ depends only on the fraction of ^{3}He, however, the term that matters is $\nu B/B_{mc}$. This leads to the deuterons crossing the passing trapped boundary much sooner with the mode conversion layer on the low field side of the axis (5.3 T), than with the mode conversion layer on the high field side (4.7 T). For the case shown here, the deuteron at 5.3 T crosses the passing trapped boundary and does not hit the wall. At 5.0 T the deuteron hits the wall near the outer midplane, after gaining more than 1.0 MeV. The 4.7 T case hits the wall near 90° after gaining almost 2.0 MeV. It is possible that some of this behavior is responsible for the localization of the losses as the toroidal field is scanned as shown in Fig. 6.10.

The insight from Fig. 6.24 is borne out in simulation results shown in Fig. 6.25. In this figure, the loss energy of a distribution of beam deuterons is plotted versus the poloidal angle at which they exit. In all cases the ratio of $n_{^3\mathrm{He}}/n_e$ = 0.15, however the toroidal field is scanned from B = 4.4 T, 4.7 T, 5.0 T and 5.3 T, moving the mode conversion layer from the high field side to the low field side. Note that the losses move to lower energy and lower poloidal angle as the toroidal field is raised, just what would be expected from the preceding discussion. The losses near the 90° are at energies close to 2 MeV consistent with the experimental data shown in Fig. 6.8. Furthermore at lower poloidal angles, the losses tend to be at lower energy, consistent with the observations.

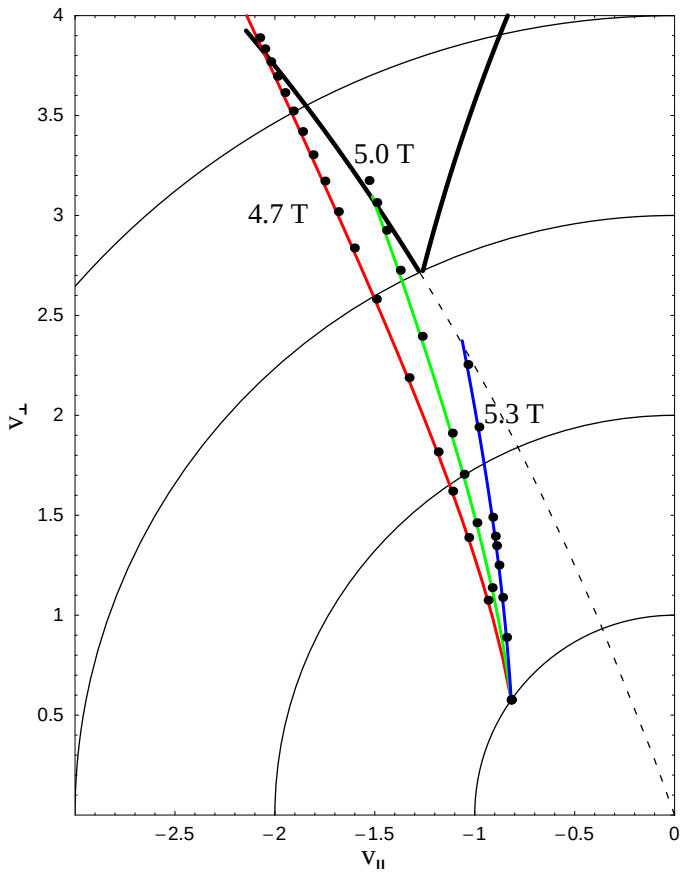

Figure 6.24: The movement of an initially 100 keV deuteron in $v_\parallel$,$v_\perp$ space due to interaction with the MCIBW for $n_{^3\mathrm{He}}/n_e$ = 0.15, but varying toroidal field. The passing trapped boundary is shown as a dotted and solid line. The solid lines represent the analytical calculation of the particles path as well as some points corresponding to the simulation of these particles. Note that the lower the field the higher the energy at which the particle crosses the passing trapped boundary.

In the case above, the fraction of the beam distribution which could resonate with the wave was typically 13% - 21.5%, with the peak coming at 4.86 T when the mode conversion layer was closest to the axis. Since the Doppler shift required to resonate with the MCIBW depends only on the ^{3}He fraction, these variations arise from geometric factors, like whether or not the particles pass through the MCIBW layer, and what their $v_\parallel$ is when they do. The number of particles kicked out is about 0.5% for the 4.4 T, 4.7 T and 5.3 T cases, and almost 2% for the 4.86 T and 5.0 T cases. Note that this strong peaking of the number of lost particles, as well as their localization near 90° when the

mode conversion layer is nearest the magnetic axis (4.86 T, 5.0 T) is strikingly consistent with the observations shown in Fig. 6.10.

However, data for the detectors other than the 90° detector also peak when the mode conversion layer is on axis, unlike the somewhat localized losses for 4.86 T and 5 T shown in the simulation results. Furthermore the experimental data shows that the losses at the 90° detector decay more slowly as the field gets higher and whereas the losses at the 60° detector decay more slowly as the field gets lower, seemingly contradicting the trend shown above. Two possible explanations lie in stochastic ripple diffusion spreading out the losses especially to the 20° detector and 45° detectors, some wave phenomenon not accounted for. Also as noted above the simulations all use the same magnetic geometry, based on a TFTR shot with 4.8 T, 1.4 MA, but with a scaled magnetic field, which also results in a scaled current. In the experiments the current was held fixed as the magnetic field was varied. This means that the simulations at 5.3 T had somewhat higher currents than the experiment, while at 4.4 T the current was somewhat lower than experiments. Using the correct equilibria for each run would tend to shift the losses poloidally a few degrees for both the higher and lower field cases.

The pitch angle distribution of the lost particles is shown in Fig. 6.26. This distribution is consistent with losses at the passing trapped boundary, with the slight variation in pitch angle at a given poloidal angle due to the spread in energy of the particles hitting the wall at that location.

One parameter which strongly affected the losses was the assumed spot size of the MCIBW. For these simulations the spot size was assumed to be constant. In practice though, the wave power is focused as it travels towards the magnetic axis, and then defocused on the other side. Thus when the mode conversion layer is on the high or low field side of the axis the spotsize will probably be larger than assumed here. The electric field was taken to fall off in a Gaussian manner to half its value on the midplane at about 24 cm, giving a power deposition with a full width at half maximum (FWHM) of about 16 cm. The results of a 50% increase in the spot size are shown in Fig. 6.27 and 6.28. In the first case, which has B = 4.4 T, the number of particles lost goes up by a factor of 4, including particles lost near the outer (inner) midplane, which previously could not be lost. These particles are cogoing (countergoing) particles which cross through the mode

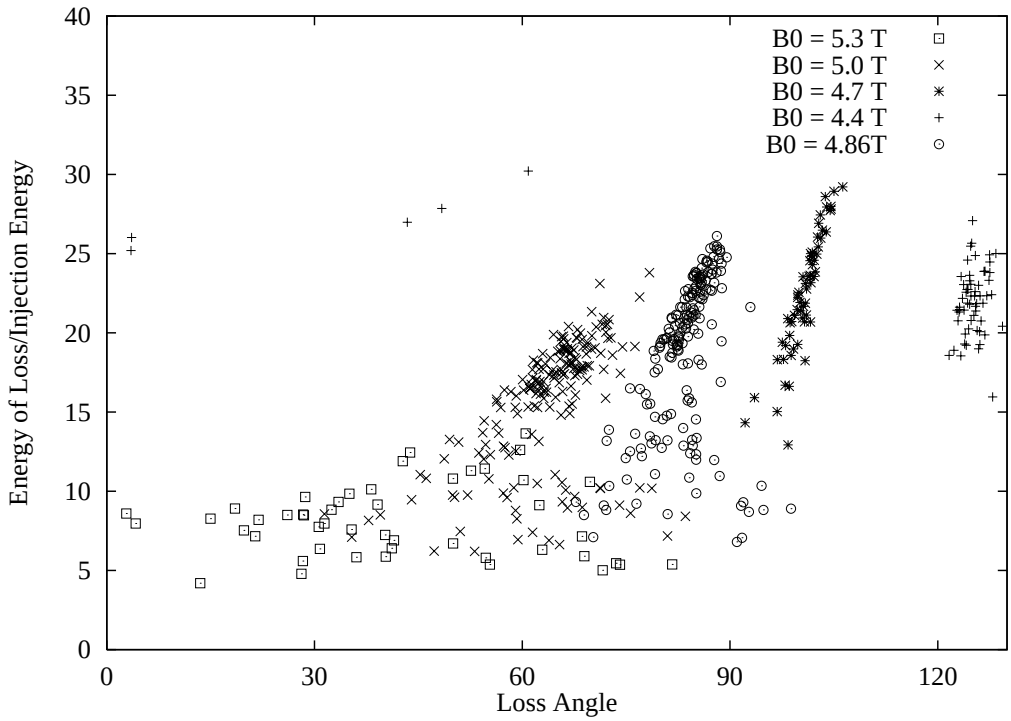

Figure 6.25: Energy of the lost deuterons versus poloidal angle where they hit the wall for B = 4.4 T, 4.7 T, 4.86 T, 5.0 T, 5.3 T.

conversion layer significantly above or below the midplane. With the expanded spot size these particles can move out much more rapidly than before, and the wave amplitude is large enough that they can receive a "last kick" to push them into the wall. In the second case, which has B = 5.0 T the number of particles lost doubles.

6.4.2 $\mathbf{n}_{^3\mathrm{He}}/\mathbf{n}_e$ scan

Recall Fig. 6.7, which shows the strong dependence of the losses on the density of ^{3}He. To see if this effect was present in the simulation, a scan was performed of the ratio of $n_{^3\mathrm{He}}/n_e$ at fixed field. The effect of this scan is to move the mode conversion layer from the outer midplane to the inner midplane. At low $n_{^3\mathrm{He}}/n_e$, the mode conversion layer comes very close to the ^{3}He resonance layer, and is correspondingly far from the D layer. Thus the Doppler shifts required to resonate with the wave become prohibitively

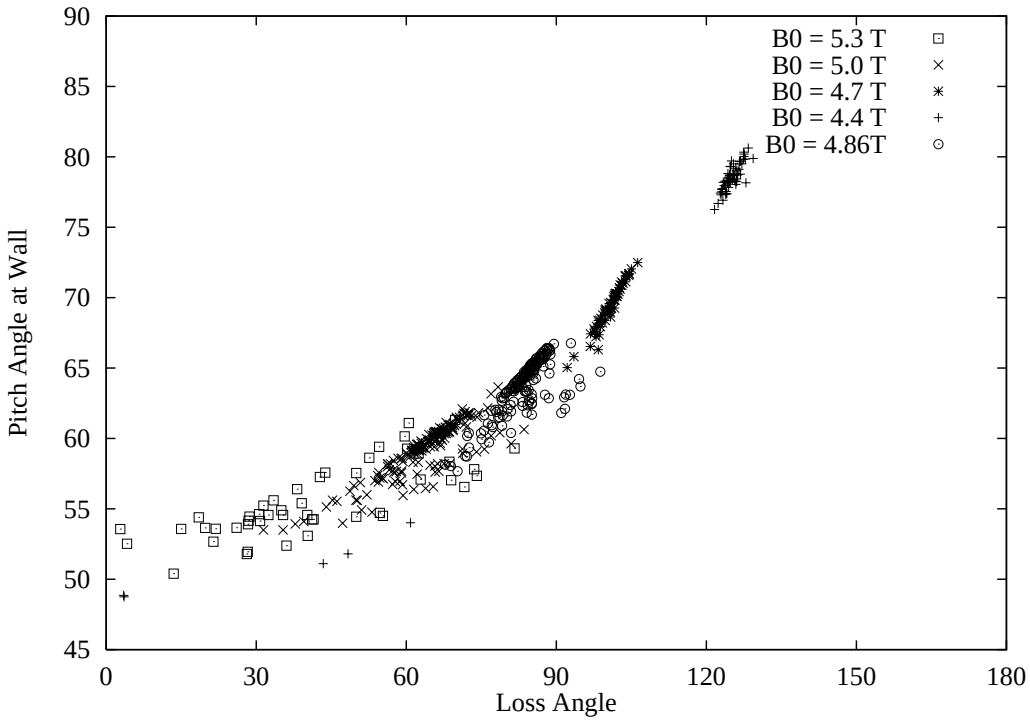

Figure 6.26: Pitch angle (degrees) of the lost deuterons versus poloidal angle where they hit the wall for B = 4.4 T, 4.7 T, 4.86 T, 5.0 T, 5.3 T.

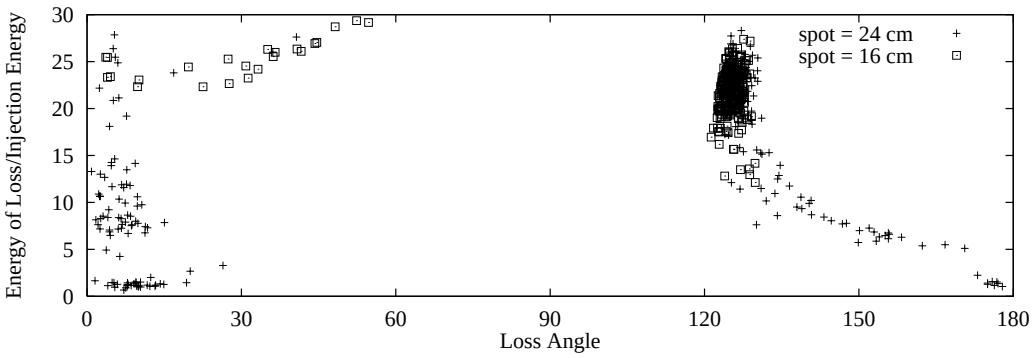

Figure 6.27: Loss Energy vs. poloidal angle for two different spot sizes. The toroidal field is 4.4 T.

large, and no particles are able to resonate.

The results of a scan of $n_{^3He}/n_e$ from 0.1 to 0.25 at 5 T are shown in Fig. 6.29 and 6.30. The equilibrium for this case has a current of 1.7 MA, as opposed to the previous

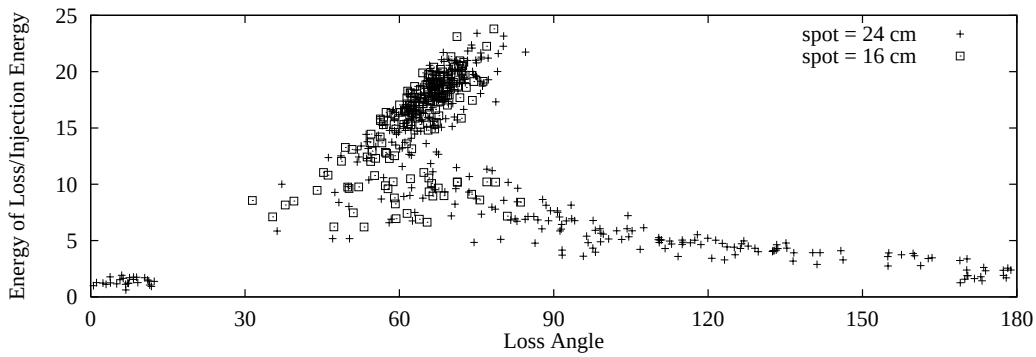

Figure 6.28: Loss Energy vs. poloidal angle for two different spot sizes. The toroidal field is 5.0 T.

results in which the current was 1.4 MA. The change in current is to more accurately reflect the conditions of Fig. 6.7. The fraction of deuterons resonant with the waves goes from 8% at $n_{^3\mathrm{He}}/n_e$ = 0.1 to 42% at $n_{^3\mathrm{He}}/n_e$ = 0.25. As can be seen from the figure, the fraction of particles lost increases substantially going from 0.05% at $n_{^3\mathrm{He}}/n_e$ = 0.1, 2.1% at $n_{^3\mathrm{He}}/n_e$ = 0.15, 5.5% at $n_{^3\mathrm{He}}/n_e$ = 0.2, 6.5% at $n_{^3\mathrm{He}}/n_e$ = 0.25. The distribution of the losses also changes markedly. Note that the large change in losses seen in Fig. 6.7 when $n_{^3\mathrm{He}}/n_e$ is changed could indeed be consistent with the simulation results shown here, although the largest change in the simulation results is seen when $n_{^3\mathrm{He}}/n_e$ goes from 0.1 to 0.15 whereas experimentally it was thought that $n_{^3\mathrm{He}}/n_e$ was changing from 0.15 to 0.2 in Fig. 6.7. One possible explanation lies in the determination of $n_{^3\mathrm{He}}/n_e$. Changing $n_{^3\mathrm{He}}/n_e$ by 0.05 corresponds to movement of the mode conversion layer by only about 8 cm at 5 T. The experimental determination of $n_{^3\mathrm{He}}/n_e$ is uncertain by an amount similar to this [MAJESKI et al. 1996]. Alternatively, there are uncertainties in the modeling of the mode converted wave which could shift the mode conversion layer by a few cm. For instance only nominal values of temperature and density were used in the ray tracing as opposed to determining the actual density and temperature at the mode conversion layer for each discharge. Furthermore, the simulation assumes that the deuterons interact with the mode converted wave at one value of $|B|$. In practice, the mode converted wave may exist over several cm.

Another result of this strong sensitivity to $n_{^3\mathrm{He}}/n_e$ may be an explanation for the sometimes erratic time history or the losses. In many discharges the loss frequently

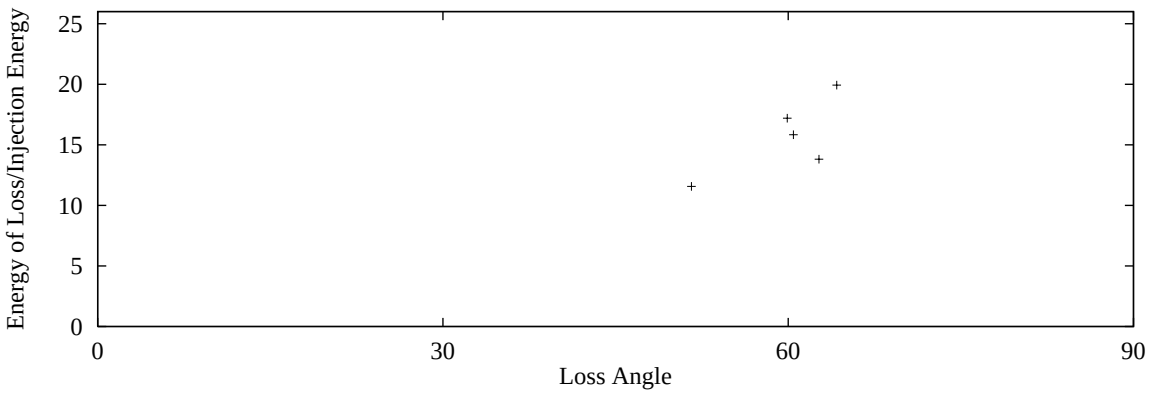

(a) $n_{^3\mathrm{He}}/n_e$ = 0.10

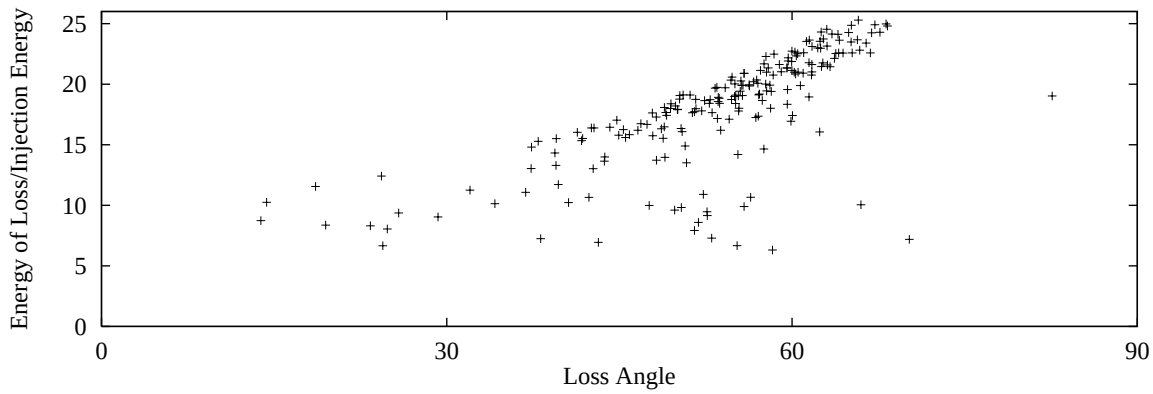

(b) $n_{^3\mathrm{He}}/n_e$ = 0.15

Figure 6.29: Losses of beam deuterons for $n_{^3\mathrm{He}}/n_e$ = 0.1 (a), 0.15 (b) and B = 5.0 T

peaks early in the discharge and then decays away with time. This could be due to small, undetectable shifts in the mode conversion layer location due to a decreasing ratio of $n_{^3\mathrm{He}}/n_e$ as the shots evolves (arising, for instance, from beam fueling of deuterium).

One way in which the Fig. 6.29 and 6.30 do not replicate the losses in Fig. 6.7 is the absence of a loss at 90°. Actually those figures were for B = 5 T, however the actual experiment was at B = 4.86 T. The losses for this value of B, with $n_{^3\mathrm{He}}/n_e$ = 0.15, 0.20 are shown in Fig. 6.31 where the losses are seen to be near 90°.

As a final note on the simulations, when tritium beams are substituted for deu-

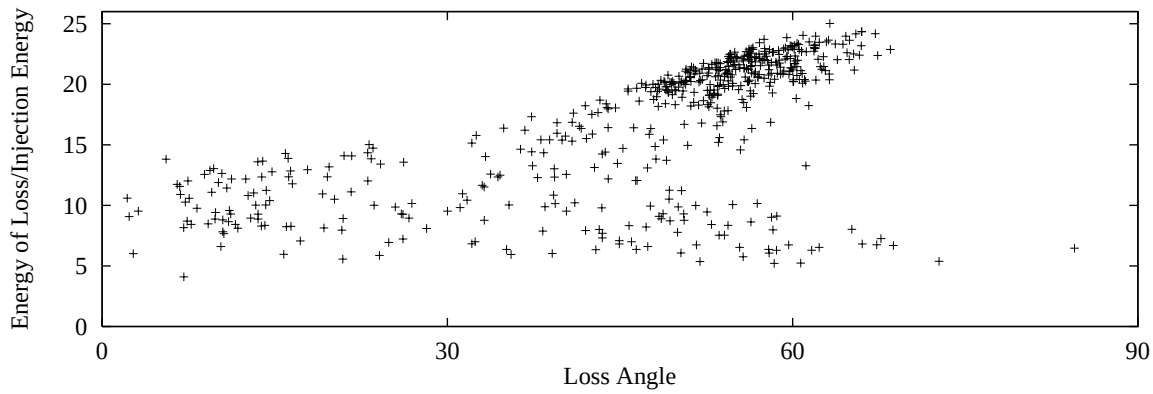

(a) $n_{^3\mathrm{He}}/n_e$ = 0.20

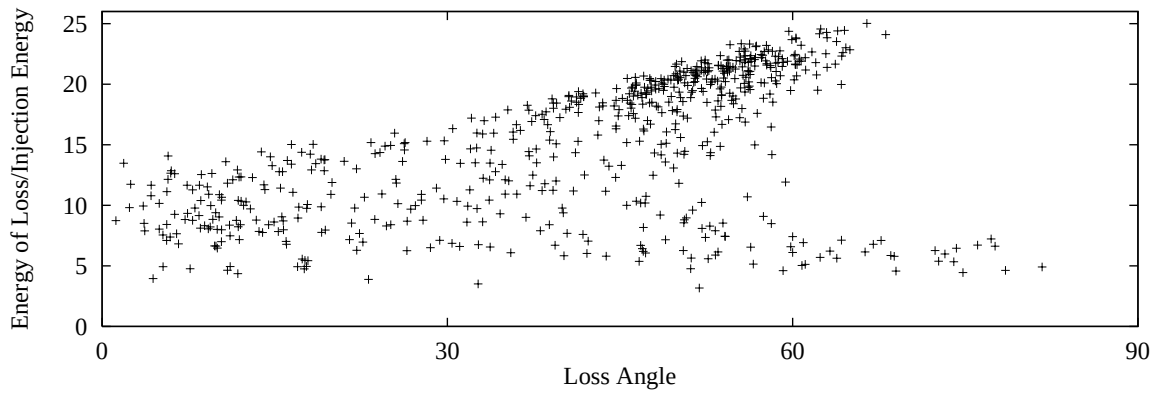

(b) $n_{^3\mathrm{He}}/n_e$ = 0.25

Figure 6.30: Losses of beam deuterons for $n_{^3\mathrm{He}}/n_e$ = 0.2 (a), 0.25 (b) and B = 5.0 T

terium beams almost 70% of the particles are resonant for B = 4.8 T, $n_{^3\mathrm{He}}/n_e$ = 0.15 (the mode conversion layer near the axis). However only 0.3% of the tritons are lost, and all of them are below the 400 keV threshold for the lost alpha detectors (although a few particles seem to hit the 20° detector, which does not have the foil prohibiting losses below 400 keV). Note that in addition to the particles crossing the passing trapped boundary at low energy, examination of the orbit shown in Fig. 6.21 shows that as the triton gains energy its intersection with the mode conversion layer tends to move to higher Z, thus the 20 cm spot height severely limits the number of tritons which can be lost. This is consistent with our expectation from the previous section, as well with the experimental

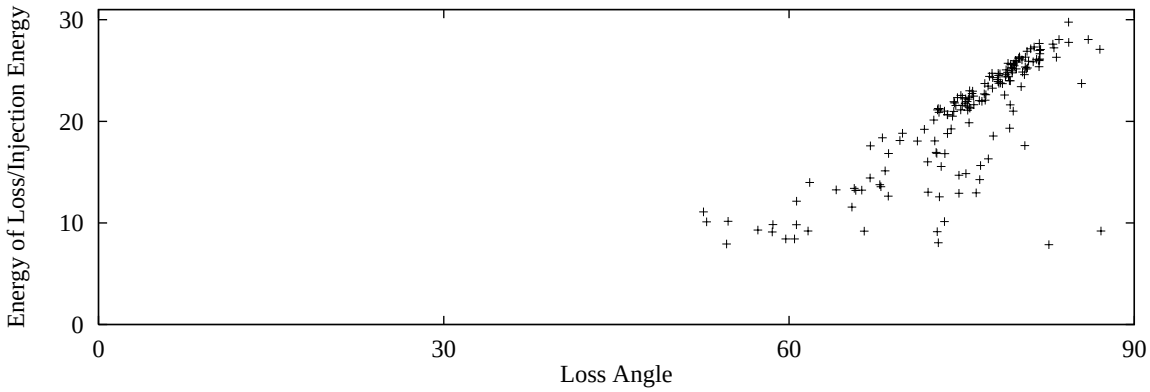

(a) $n_{^3\mathrm{He}}/n_e$ = 0.15, B = 4.86 T

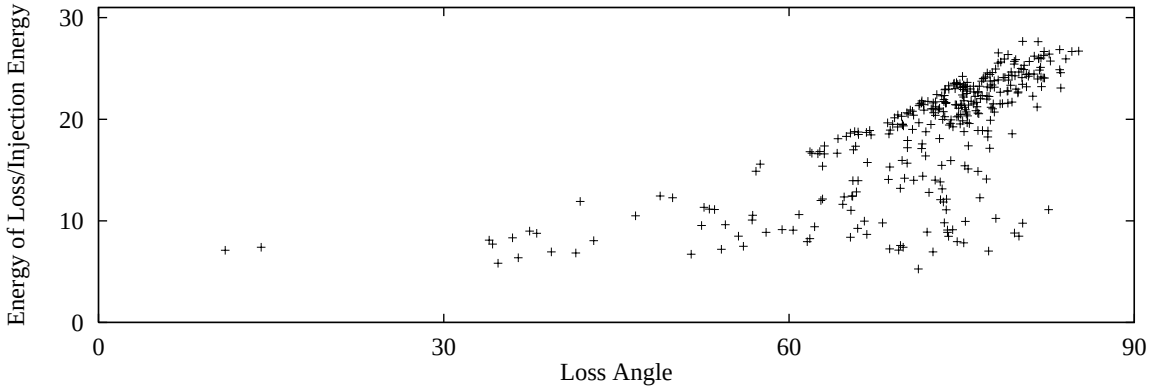

(b) $n_{^3\mathrm{He}}/n_e$ = 0.2, B = 4.86 T

Figure 6.31: Losses of beam deuterons for $n_{^3\mathrm{He}}/n_e$ = 0.15 (a), 0.2(b) and B = 4.86 T

observations.

Chapter 7

Key Implications of the TFTR experiments

THE EXPERIMENTS CONDUCTED ON TFTR are extremely important to α-channelling for two reasons. First, they provide us with invaluable experimental data on the interaction of the MCIBW and fast ions, and thereby give data which can be used to benchmark the simulation. Second, they allow us to infer key physics results which we then make use of to accomplish α-channelling in a reactor.

In particular, this data is so valuable because it provides an integrated test of the various facets of α-channelling, albeit with heated particles rather than cooled ones, i.e. it includes wave physics, fast particle physics, and the wave-particle interaction. Even if we were able to test each component separately, we could not be very confident about our prediction without an integrated test. If the experimental data can be used to validate the simulation, than we will be much more confident about the predictions of the code when extrapolating to full α-channelling scenarios. This is the use to which the data was put in the previous chapter.

In this chapter, we put the data to a different use, namely, the determination of key physics results in support of α-channelling. The two key results discussed here are the $k_{\|}$ flip of the MCIBW in Sec. 7.1 and the diffusion coefficient for the deuterons interacting

with the MCIBW in Sec. 7.2

7.1 The $k_{\|}$–flip

One of the most interesting predictions of the work of Valeo and Fisch (1994), which has already been discussed in some detail in this thesis, is the flip of the $k_{\|}$ of the wave, due to the effects of the poloidal field. This wave has a parallel phase velocity in the opposite direction of the launched wave phase velocity. As we have already pointed out in Sec. 5.5 this effect is very useful for α-channelling. It also is important for the current drive problem, as a wave with $k_{\|}$ flipped would drive power in a direction opposite to the launched direction, which would presumably be undesirable. Is there evidence of such a $k_{\|}$ flip in the TFTR data? If so, how would one detect it ?

One way of detecting such a wave would be to infer that a set of ions interacted with that wave by means of a resonant interaction. Recall that for $n_\phi > 0$ the $k_{\|}$ flipped wave has $k_{\|} < 0$ and for $n_\phi < 0$, $k_{\|} > 0$.

Since $k_{\|}v_{\|} = \omega - \Omega_D > 0$ cogoing ions must interact with waves which have $k_{\|} > 0$ and countergoing ions resonate with waves which have $k_{\|} < 0$. Thus losses of cogoing ions when only $n_\phi < 0$ waves were launched, or losses of countergoing ions when only $n_\phi > 0$ waves were launched would provide unambiguous evidence of a $k_{\|}$ flip.

Note though, that normally both co and countergoing ions are present, hence any detected interaction could have been of the MCIBW with a particle going in the direction that did not require a $k_{\|}$ flip for it to interact. Another problem is that, in general, for symmetric phasing, both $n_\phi < 0$ and $n_\phi > 0$ waves are launched by the antenna. This also makes the detection of the $k_{\|}$ flip ambiguous, since a cogoing (countergoing) deuteron, which must interact with a wave which has $k_{\|} > 0$ ($k_{\|} < 0$) could be interacting with the $n_\phi > 0$ ($n_\phi < 0$) wave or the $k_{\|}$ flipped $n_\phi < 0$ ($n_\phi > 0$) wave (see Sec. 6.3.2). It could be the case that comparing the simulation with and without the $k_{\|}$ flipped wave would give a significant enough discrepancy with the experimental result that one could argue that the $k_{\|}$ flipped wave must be present, but, given the ambiguity in the modeling

of the wave physics, it is unlikely that such an argument would be convincing.

Both of these problems can be overcome. TFTR has the capability of injecting neutral beams in either the cogoing or countergoing directions, and the RF can be phased to give predominantly n_ϕ greater or less than zero. Thus in order to directly test the $k_{\|}$ flip hypothesis, a few shots were carried out in July of 1996. These experiments used both toroidally phased RF and cogoing and countergoing beams, in all four configurations. Due to lack of run time time and difficulty with RF phasing, there was only time for one or two shots in each configuration, all at low RF power levels (around 1.5 MW). The only loss observed was for countergoing injection of the beams with $n_\phi < 0$, showing no signs of a $k_{\|}$ flip. However, a $k_{\|}$ flip cannot be excluded on the basis of these experiments. The absence of losses may well be due to low power levels or diffusion paths which do not reach the wall, rather than the absence of the $k_{\|}$ flip.

Fortuitously, in looking over old data in preparation of Fisch et al. (1996), we stumbled across some runs which were used to investigate mode conversion current drive [MAJESKI et al. 1996; MAJESKI et al. 1996]. These runs had asymmetric phasing of the RF, and experiments were carried out with $n_\phi < 0$ (these waves would damp on countergoing electrons pulling a tail of electrons in the counter direction which would drive current in the cogoing direction), and $n_\phi > 0$ (driving counter current). Fortunately many of these experiments had coinjection only of the beam deuterons. In Fig. 7.1 we show a comparison between two of these experiments. By convention +90° phasing is the phasing which drives co current and therefore has $n_\phi < 0$. Thus we have evidence of loss of cogoing injected beam deuterons with $n_\phi < 0$ demonstrating the $k_{\|}$ flip!

While this picture is consistent there are a number of subtleties which cloud the picture. First, the power launched from the antenna is not entirely directed in the direction of the phasing. Rather a significant portion (about a quarter) goes in the opposite direction with a peak n_ϕ about twice that of the dominant n_ϕ. If this backward propagating lobe were responsible for the losses there would be no need to invoke the $k_{\|}$ flip. Second, these experiments show significant losses of cogoing injected beam deuterons, which is in contrast with later experiments showing no significant losses when only coinjection occurs. Third, it is conceivable that cogoing injected beam deuterons have pitch-angle scattered into countergoing or trapped particles and then are lost, again, removing the

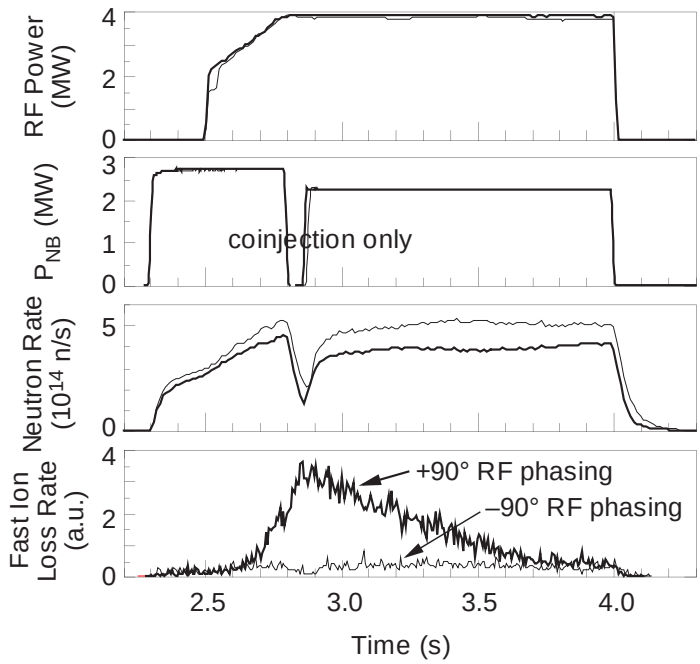

Figure 7.1: Losses of cogoing deuterium beams in the presence of phased rf. +90° phasing (counter to the current) results in significantly enhanced losses, whereas -90° phasing (co to the current) shows first orbit losses of fusion products only. Figure from D. Darrow.

need for the $k_{\|}$ flip to explain the data.

These questions have been addressed with several simulations which include both the backward lobe of the RF and the effect of collisions, and show agreement with the experimental results. In particular a simulation of the losses shows that for the case of $n_\phi < 0$ there were significant losses, however, when $n_\phi > 0$ was chosen losses were absent which is consistent with Fig. 7.1. Simulations with only the backward lobe (i.e. no $k_{\|}$ flipped wave) show no heated beam losses, ruling out the backward lobe as a possible complication. Collisionless simulations as well as collisional simulations show the losses, suggesting that cogoing deuterons do get out without pitch angle scattering from cogoing to countergoing.

The reason cogoing deuterons are lost in these experiments while their loss is typically absent in other experiments has to do with the location of the mode conversion layer. The experiments investigating current drive effect typically had a ratio $\mathrm{n}_{^3\mathrm{He}}/\mathrm{n}_e$ of

0.25 placing the mode conversion layer far on the high field side, as close as any of the TFTR experiments got to putting the mc layer near the D resonance layer. Because of this the wave power, which focuses down to a spot size of $\pm$ 15 - 20 cm when the the mode conversion layer is on axis tends to defocus considerably as it goes to the high field side. This defocusing spreads the power out considerably, increasing the spot size. This increase in spot size allows cogoing deuterons to keep resonating with the wave as they move farther and farther out, giving them a diffusion path to the wall. When the mode conversion layer is on axis it is very difficult for cogoing deuterons to be lost.

Let us consider the last source of ambiguity concerning the $k_{\|}$ flip. Conceivably some deuterons pitch angle scatter from cogoing or trapped to counter going and then start resonating with the wave. A $k_{\|}$ flip would not be necessary to explain the interaction of these countergoing particles with the wave. But the number of particles which can do this is small (see Fig. 4.15, which shows a cogoing injected beam pitch angle scattering and slowing down), since the pitch angle scattering time for theses plasmas is a factor of two to three longer than the slowing down time, and the particles will stop resonating with the wave after slowing down below some threshold energy. Simulations with just the part of the spectrum before the $k_{\|}$ flip do indeed show losses due to pitch angle scattering of cogoing particles to countergoing and then those countergoing particles resonating with the wave. However, these losses are typically at low energies (300-400 keV), below the detection threshold of the detectors (about 400 keV). Because so few pitch-angle scattered particles could be made resonant, and we see from simulations that those that are would be lost only at undetectably low energy, the scattering of cogoing particles cannot explain the very large losses (relative to background) seen in Fig. 7.1. The $k_{\|}$ flip remains the only explanation which is consistent with all of the data, according to the COM simulation.

7.2 Diffusion Coefficient

Understanding the diffusion coefficient, D, of the beam deuterons interacting with the MCIBW is very important for filling in the α-channelling puzzle, as it directly relates to the power level required for accomplishing the channelling effect in a reactor.

How can the diffusion coefficient be determined? It would be nice if we could measure the in situ distribution function of the beam deuterons in all parts of phase space, however such a measurement is not possible (pellet charge exchange can give a limited amount of information about the distribution function at a single time in a small volume of phase space). We do have information like that contained in Fig. 7.2, which shows the neutral beam injection, the neutron rate, and the losses to the 60° lost alpha detector versus time from a typical MCIBW beam blip experiment conducted on TFTR with 3.2 MW of RF power. As mentioned in Chapter 6, these experiments blipped the beams early in the shot, providing a "delta" function source of beam deuterons to interact with the wave.

This data contains a wealth of information about the strength of the MCIBW beam particle interaction. Note the long decay of the neutron production rate. In contrast, consider Fig. 7.3 which has only 1 MW of RF power, and a much more rapid neutron decay. Shots which have rf only show neutron production two orders of magnitude lower than those with beams. Together this is indicative of a significant tail of RF heated beam deuterons being created by the MCIBW. Furthermore the losses appear to have characteristic onset, peak and decay times, giving data points that can be compared to theory or simulation. Finally, by comparing the overall level of losses from different shots with different power levels an estimate of the diffusion coefficient can be obtained. In this section we outline how one can estimate the strength of the wave particle interaction using the data and simulation results.

A cursory glance at Fig. 7.2 shows that that we cannot treat the beam losses as collisionless. For these plasmas a typical slowing down time for a 100 keV beam deuteron is on the order of 100 milliseconds, and yet losses persist for several hundred milliseconds after the blip. If the losses were truly collisionless it would be necessary for the particles to exit in a time short compared with the slowing down time. Furthermore, the level of losses varies strongly with the applied RF power. Collisionlessly one would expect that all particles which can hit the wall (i.e. have a diffusion path intersecting the wall) will, and thus the losses shouldn't change strongly with power. Collisions, however, change this; by dragging particles down in energy they increase the time for particles to leave the plasma. They also can pull particles below the minimum energy required for interacting with the wave. Once a particle falls below this energy it is most

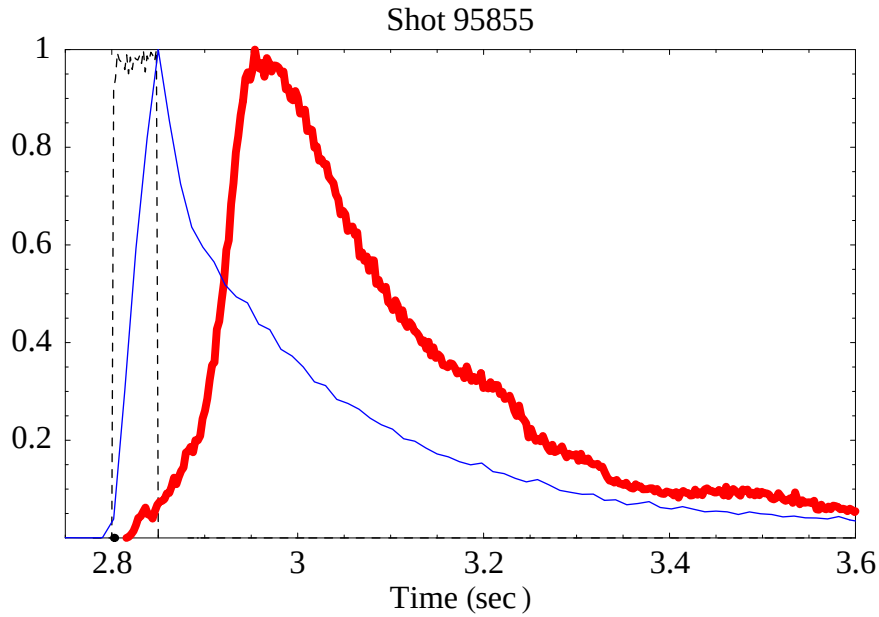

Figure 7.2: The results from a typical beam blip experiment. The dashed line shows the neutral beam injection, the thin solid line shows the neutron measurement and the dark black line shows the loss signal at the 60° detector. The RF power (3.2 MW) is on for the whole time.

likely permanently removed from interacting with the wave.

Thus in order to get a an accurate representation of the experiments it is necessary to include the effect of collisions in the COM simulation as done in Chapter 4. Now, unlike the collisionless case with a single wave where the diffusion path was one dimensional, or even the case of one ω and more than one n_ϕ which is two dimensional , the diffusion will be fully three dimensional in ε, μ, and P_ϕ space.

Before undertaking a simulation of particles interacting with both the waves and collisions it is worthwhile to consider a modified version of the one dimensional problem discussed in Sec. 4.4.7. Heeter et al. (1997) have used a one dimensional model which includes diffusion in velocity due to the waves and collisional effects in order to determine how the losses might scale as the power is increased. For the correctly chosen D and using experimental values for the threshold energy and estimates for the loss energy the one dimensional model is able to nicely replicate the time history of the losses.

Before proceeding to the simulation and experimental results we get a simple

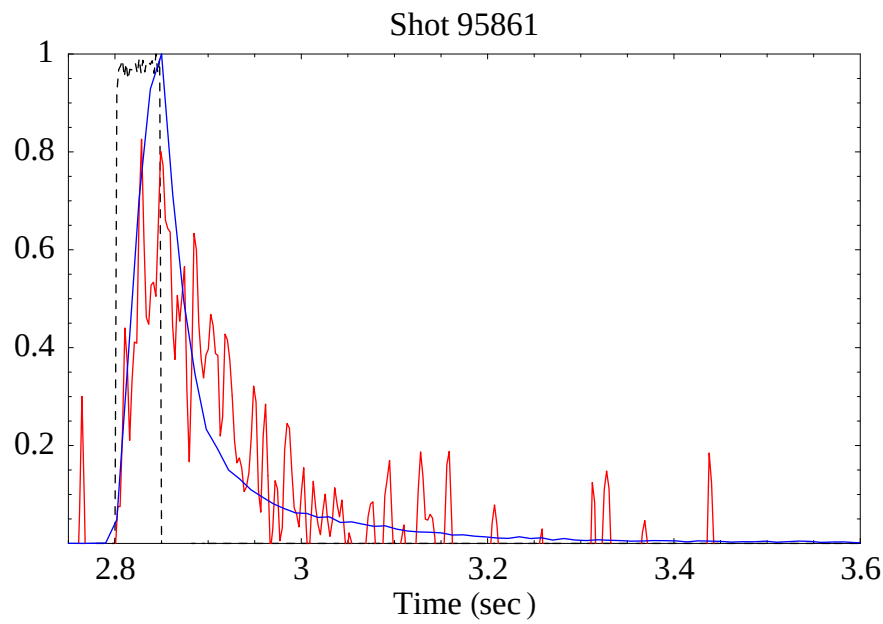

Figure 7.3: The results from a beam blip experiment with only 1 MW of RF power. The dashed line shows the neutral beam injection, the thin solid line shows the neutron measurement. The signal at the 60° detector is not shown, since the losses show no enhancement over first orbit losses.

estimate of the diffusion coefficient, and estimate the timescale for the losses to occur.

7.2.1 Simple estimate

We reproduce here Eq. (4.86) which gives the change in the particle energy upon one pass through wave region.

$$\Delta\varepsilon = q\omega\Phi_0(\mathbf{r}_{\mathbf{g.c.}})e^{i\varphi_0}J_n(k_\perp\rho)\sqrt{\frac{2\pi}{\left|\frac{v_\parallel}{JB}\frac{\partial}{\partial\theta}(k_\parallel v_\parallel + n\Omega)\right|}} \tag{7.1}$$

We want a representative value for this quantity, keeping in mind that in practice this expression can vary by factors of 10 or more for a single particle at different points in ε, μ, and P_ϕ space, and even more from particle to particle. From Fig. 6.15-6.17 we choose , $k_\perp$ = 1.5 cm^{-1},$k_\parallel$ = -0.15 cm^{-1}, Φ_0 can be estimated as $E/k_\perp$. Note that Fig 6.17 shows the electric field in statvolts/cm for an incident flux of 1 Watt/cm^2 on

the mode conversion layer. In a typical MCIBW experiment 2 MW of RF power was launched and mode converted at R = 277 cm, with a spot height of about 35 cm ($\pm$17 cm), giving a typical intensity of 32 Watts/cm^2. The electric field scales as the square root of the intensity giving a value of Φ_0 of about 0.3 statvolts. We choose a countergoing beam deuteron at about 250 keV which has $\rho = 1.9cm$, and $\lambda = 0.63$ at the mode conversion layer. Note that the expression for $\Delta\varepsilon$ above goes to zero at the zeroes of the Bessel function. In practice, many $k_\perp$ are present at the MCIBW layer so that the effect of these zeros is averaged over. The large argument expansion for the Bessel function is squared and averaged over $k_\perp\rho$, giving a term which goes like $1/(\pi k_\perp \rho)$, For calculating the kick we use the square root of this term, which is about 0.3 for the case at hand. The term under the square root above is a resonance time, and ω times this term gives the number of wave periods that the particle is resonant. The denominator of this term can be rewritten

$$\frac{v_\parallel}{JB}\frac{\partial}{\partial\theta}(k_\parallel v_\parallel + n\Omega) \approx \frac{v_\parallel}{JB}\frac{\partial B}{\partial\theta}v_\parallel\frac{\partial k_\parallel}{\partial B} \tag{7.2}$$

where we have neglected the variation of $v_\parallel$ and B with θ compared with the variation of $k_\parallel$ which is quite large. This gives a resonance time, $\omega t_{\rm res} \approx 20$. Multiplying these terms together we get $\Delta\varepsilon$ = 1.8 statvolts. Finally converting from statvolts to volts gives a factor of 300 yielding a "typical" $\Delta\varepsilon$ on the order of 0.5 keV per bounce. The bounce time for one of the beam ions interacting with the wave is on the order of 10^{-5} seconds.

It is useful to estimate a collisionless loss time. A typical particle that leaves gains on the order of 1 MeV before leaving. If the deuteron were to receive only heating kicks it could leave after 2000 bounces or about 20 milliseconds. Diffusively, however, the number of bounces required goes like the square, implying that the bulk of the resonant particles would leave after 40 seconds, much longer than the time of the discharge! In fact given this diffusion coefficient, the fraction of particles leaving in less than a second would be vanishingly small (roughly this is the probability that out of 10000 kicks (the number of bounces in one second), there would be 2000 more plus kicks than minus kicks, but given equally probable kicks, this is a more than 6 σ event.) It would be necessary to increase the kick size by a factor of 3 - 6 in order to have a non negligible fraction of the beam deuterons extracted. We will see from the simulations that the diffusion coefficient predicted from first principles is too small (by a large factor) to

explain the experimental losses.

7.2.2 Simulation results

While the above ball park estimate is useful, numerous effects included in the simulation could change these estimates substantially. Thus an extensive scan of the simulated losses vs. electric field amplitude was performed. The simulation of beam deuterons under the effect of multiple ion Bernstein waves and collisions is the most computationally challenging project in this thesis, and it puts to use almost every aspect of the COM simulation. Typically 10,000 to 50,000 particles are simulated for the duration of a shot (on the order of 1 second or 10^5 bounce times). Each run can take anywhere from 20 minutes to several hours on a Dec Alpha workstation. In order to have good statistics, at the lowest electric field amplitudes, where few particles are lost, several runs with different random number seeds have been combined.

Figure 7.4 shows the fraction of resonant beam particles which hit the wall in the course of 1 second (note that only about 10% of the beam particles are resonant). This scan shows both the collisional(lower) and collisionless losses. Note that some particles, though they resonate with the wave, do not have access to diffusion paths which hit the wall, and thus are stuck in the plasma. Furthermore the small spotsize (which is taken to have Gaussian dependence on height above the midplane), means that some particles which resonate feel extremely small amplitudes. This in turn means that the time scale for these particles to leave the plasma is very, very long (essentially causing these particles to be stuck as well). These two effects account for the fact that the collisionless fraction lost only approaches 0.6 asymptotically in the plot shown here. The effect of collisions is to reduce the losses by a factor of less than two at the highest fields shown and over a hundred at the lowest.

Using the data from TFTR shots 95854-95862 [DARROW 1996], during which a scan in RF power was performed, it is possible to make a comparison between the theoretical and observed value of the diffusion coefficient. In order to do this however it is necessary to make an estimate for the absolute level of losses in these experiments. This is an exercise fraught with uncertainties. Recall that the absolute calibration of the lost

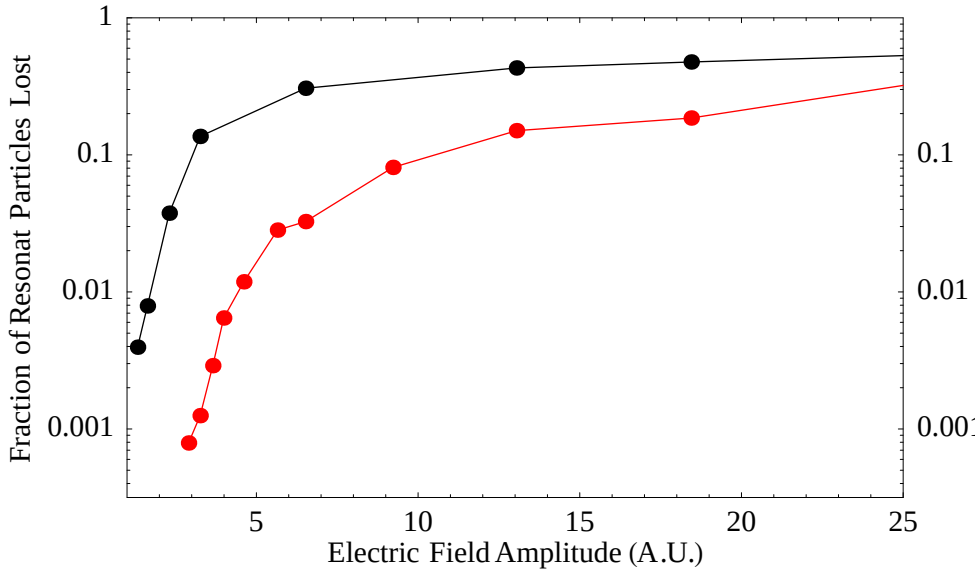

Figure 7.4: Collisionless and collisional loss of resonant ions vs. electric field amplitude

alpha detectors can be off by as much as a factor of two. Furthermore in computing a global loss fraction we must make an assumption about the poloidal and toroidal distribution of the loss, even though the detectors measure the loss rate at just one toroidal and only four poloidal locations. In the absence of any evidence to the contrary we assume that the loss is toroidally symmetric, but distributed poloidally on the outer bottom half of the wall. For simplicity we take the loss level at the 60° detector to be representative and calculate the total flux of energy to the wall over the course of the discharge. For a beam blip it is natural to compare this energy to the total energy injected into the plasma during the neutral beam injection. The results are shown in Fig. 7.5.

We would now like to compare the data to the simulation results. The simulation can also be used to estimate the fraction of the energy of the beam blip which impinges on the wall, and it can be plotted versus the inferred power. Upon doing this, though we find that the simulation predicts no loss at the power levels which were present in the TFTR experiments. This is not surprising. The simple estimate above suggested that "typical" kicks were a factor of 6 to low for there to be any significant losses. This factor of 6 in the kick size would translate to a factor of 36 in the power! Thus rather than directly comparing the simulation results to the experimental losses, we attempt to infer how far off the theoretical estimate of the diffusion coefficient is from the experimental results

by lining up the experimental curve and the simulation results. The result is shown in Fig. 7.5. When the simulation data has been shifted the best estimate for the enhancement in the diffusion coefficient is somewhere between 30-70! Note that if the experimental data were shifted up (down) by a factor of 10 the enhancement would decrease (increase) by a factor of 2. From this enhancement, we can infer from Fig. 7.4 that the collisionless limit could have been attained in TFTR at power levels from 12 to 27 MW.

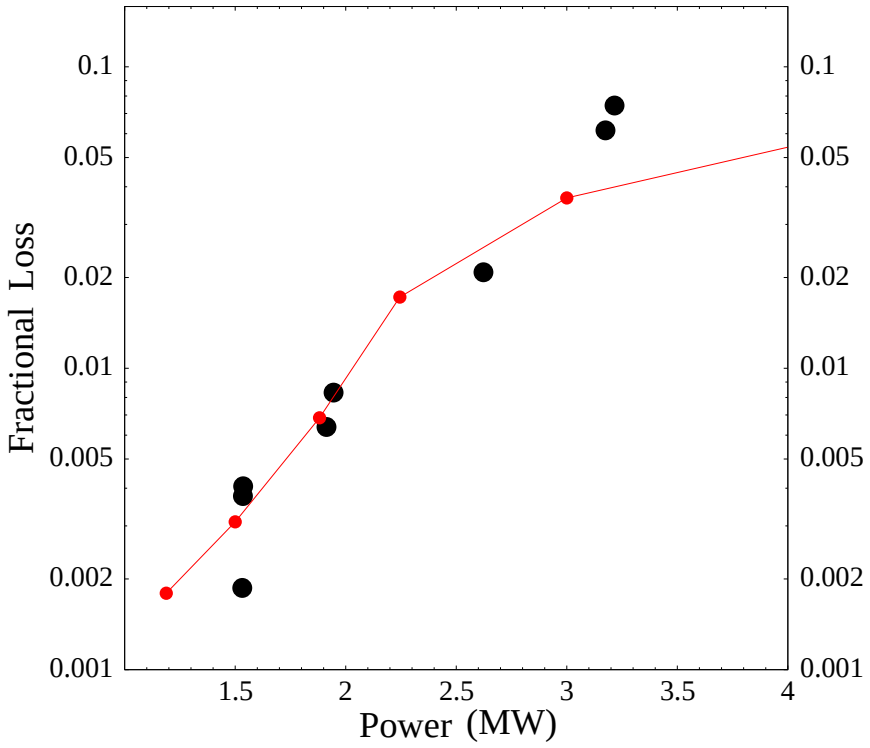

Figure 7.5: Fraction of the injected energy impinging on the wall due to acceleration of the beam ions by the MCIBW. The solid dots are the experimental data, and the line connects the simulation results. Note that the error bars on both curves are quite large, as there is large uncertainty about the experimental losses and the simulation results (especially at low power) are based on just a few particles.

Several consistency checks have been done, to ensure that the COM simulation is operating as intended. In particular, as pointed out in Sec. 4.4.7, the code has benchmarked against the analytical prediction for simple cases.

Furthermore the results from the simulation appear to be consistent with the experiments, in terms of the poloidal and energy distribution of the lost particles. It has not been possible to compare the losses vs. time to those of the experimental data since too few particles are lost from the simulations to provide meaningful histograms of the losses vs. time. However, we can make a comparison between the mean exit time of the simulations and the mean exit time for the experiments. The data is presented in Fig. 7.6 and shows good agreement (when the shift in power mentioned earlier is performed).

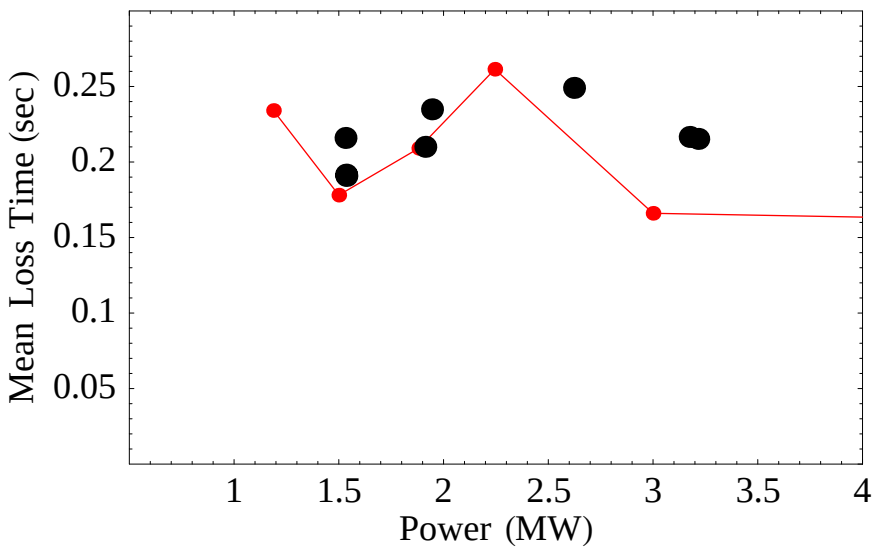

Figure 7.6: Mean loss time of the exiting particles vs. power level. The solid dots are the experimental data, and the line is the simulation results.

The simulation results appear to be consistent with the experiments when the diffusion coefficient is enhanced as above. How might we explain this large enhancement in the diffusion coefficient? We think it unlikely that the anomaly arises from a bug in the COM simulation, given the benchmarking against analytical results, and given the simple estimate above using a random walk argument in one dimension which also suggested that the analytical D was much less than needed to explain the experimental result.

The most likely explanation is that our modeling of the wave from the ray tracing discussed in Sec. 6.3.1 is too simple. The diffusion coefficient depends strongly on the wave parameters which are determined from the ray tracing, in particular, the electric field amplitude and the rate at which $k_{||}$ is changing along the particle's orbit. An in-

crease by a factor of 5 to 9 of the electric field amplitude, a similar increase in the time the particle spends in resonance (due too a less rapid spatial variation of $k_{\|}$), or some combination of the two, could account for the discrepancy.

How could the ray tracing be off by so much? One possibility is that the simplified slab geometry of the ray tracing code we used may miss some large effect associated with the toroidal geometry of the experiment. This could be rectified by a more complete ray tracing code. It is also possible that the physics missing from the ray tracing approach is important enough to cause such a large discrepancy. In this case we would have to resort to a full wave code such as PENN [JAUN, HELLSTEN, and CHIU 1998]. The physics effects that come to mind are correlated wave-particle interactions or the presence of cavity modes.

In any event, the power estimates for α-channelling in a reactor (Sec. 5.5.3), which are based on the analytically calculated diffusion coefficient (assuming uncorrelated kicks and geometrical optics), should be reconsidered in light of the apparently much larger experimentally deduced diffusion coefficient. While at the present we do not know if the enhancement demonstrated here will, in fact, scale to a reactor, if it did, the power levels required for the MCIBW would be significantly lower.

Note that, because we cannot be sure that we are modeling correctly the detailed wave physics, nor that we can ascribe the discrepancy in TFTR data to specific physical effects (such as internal modes or correlated kicks) that might or might not scale to a reactor, we cannot draw the conclusion that the required 100 MW estimate that we gave in Sec. 5.5.3 for collisionless α-channelling in a reactor should be reduced by exactly the factor of 30-70. However the data does motivate intense investigation of just this possibility.

Chapter 8

Directions for Future Work and Conclusions

COOLING α-PARTICLES WITH WAVES in a tokamak is a fascinating challenge. In this thesis, we combined theory, numerical simulations, and analysis of experimental data to establish a framework for approaching this problem. Significant advances have been made on all three efforts.

The α-channelling scenarios presented here are impressive, but the work we have done suggests more impressive scenarios might yet be found. Of course, there are a number of caveats and limitations which must be addressed, and experimental validation of each assumption is crucial. Future work might well address both the caveats affixed to this work, as well as the opportunities for extending and improving the concept.

8.1 Improvements to α-Channelling Scenarios

The α-channelling results for an advanced tokamak reactor using two waves presented in Chapter 5 are promising. It is remarkable that the same set of waves was able to affect the entire distribution of α-particles in such a favorable way. There are several important extensions of this work.

The level of realism in the simulations should be improved. In particular the TAE (or one of its variants) should be properly modeled. This will require considerable work, since up to now this wave has been treated as a single wave with one n_ϕ and ω, existing across much of the plasma. In reality, many modes will be necessary to get most of the α-particles to diffuse from the center to the edge. Each mode has a different n_ϕ, ω, and location. While this will undoubtedly make α-channelling more complicated, the greater degree of control (we are presuming that these modes will be excited externally) may allow improved α-channelling scenarios.

Attention should also be paid to other ways of improving α-channelling, in particular increasing the amount of power channelled or the fraction of power extracted from those particles which leave. For example, can the 32% of the α-particle power hitting the wall (for the case with maximum extraction) be reduced? If it cannot be reduced, is there some way in which that power might be put to good use by taking advantage of the localized nature of the losses?

Along these lines there are several ideas which might be tried. One idea is to modify the shape of the first wall of the tokamak, either to focus the losses on a single point (one could imagine a small, easily replaced rod which is slightly inside the first wall on the outer midplane to scrape off the cooled α-particles before they hit the wall itself), or spread them out across as wide an area as possible to minimize the wall loading.

Another idea would be to try to move the MCIBW into the center. As discussed in Chapter 5, this leads to the heating and ejection of many of the α-particles. This heating might be avoided if there were a superadiabatic barrier to particle heating above a certain energy. This barrier would function as a reflecting boundary in the constants of motion space of the α-particles at high energy, and may allow a cooling scenario to be developed with the MCIBW in the center. This in turn might significantly reduce the α-particle power flowing to the wall.

Note also that the α-channelling scenarios in this thesis considered only an aspect ratio 3, reverse shear tokamak. Other possibilities include exploring various tokamak geometries at different aspect ratios and with different currents, or even optimizing the geometry and plasma parameters for the α-channelling effect.

8.2 Further Experimental Validation

Since it is unlikely that an integrated test of two wave α-channelling scenarios will be performed any time soon, it is vital that each aspect of α-channelling be independently tested where possible. If we have confidence in each aspect on its own then we have a greater degree of confidence in the simulation which combines these effects to simulate a two wave α-channelling scenario.

In this thesis we were able to explore and test two such aspects for the MCIBW, the $k_{\|}$ flip and the diffusion coefficient. We verified that the $k_{\|}$ flip of the MCIBW, which is important for α-channelling, occurs. Our analysis of the diffusion coefficient suggests a large discrepancy between experimental deductions and simple applications of the theory. An obvious area of future work is to resolve this discrepancy. More refined calculations might point to the existence within the TFTR data of important and new phenomena, such as internal modes or correlated kicks.

An area of future work certainly relates to establishing similar levels of confidence in the relevant physics of the TAE. Heeter (1998) is exploring experiments investigating this mode on JET and its implications for α-channelling.

One experimental observation which would significantly enhance our confidence in the two wave scenarios would be the observation of cooling of energetic particles by the MCIBW. This was not observed on TFTR, instead beam deuterons were heated and ejected, although this is consistent with the predictions of the simulation. The basic idea behind α-channelling is that a correctly chosen wave can diffuse particles along a path such that the particles are cooled on average and the wave is amplified. In general, it is very difficult to construct the circumstances under which the MCIBW alone would lead to cooled particles hitting the wall in a large tokamak, as discussed in Sec. 5.1. However, simulations suggest that in a reverse shear discharge on TFTR, 3.5 MeV α-particles might be cooled with the correct phasing of the RF.

The simulation results of such a scenario are shown in Fig. 8.1. Although an experiment was proposed for the observation of cooling of α-particles in a reverse shear plasma on TFTR, unfortunately, it was not possible to perform this experiment in the lim-

ited run time available to TFTR in its final year of operation. The prospects of observing

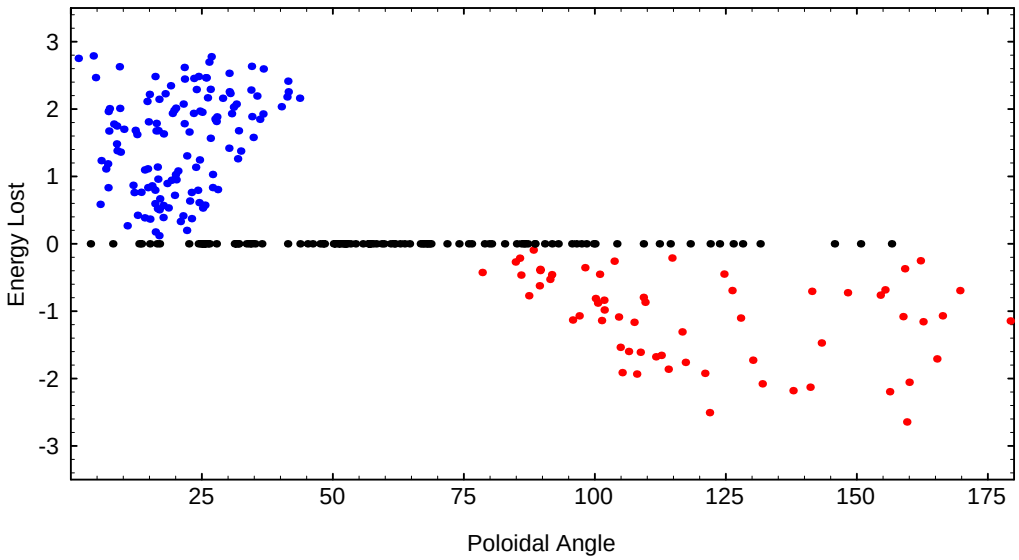

Figure 8.1: Energy lost to the wave (MeV) and poloidal exit angle of α-particles in simulated reverse shear TFTR discharge with B = 5.3 T, I = 1.85 MA (scaled shot #84011). 1000 particles are simulated. Particles exiting with zero energy lost correspond to first orbit losses (11%). 13.7% of the α-particles are cooled, exiting near the outer midplane. 3.8% are heated, exiting between the inner midplane and the bottom.

α-particles cooled by MCIBW in the near future vanished with the shutdown of TFTR. However, recently a proposal has been made to investigate fast particle interactions with MCIBW on a modification of the Current Drive Experiment-Upgrade (CDX-U), a low aspect ratio tokamak at Princeton, to be named the Channelling Physics Experiment (CPX) [MAJESKI 1998]. One of the goals of this experiment will be to diagnose the interaction of the MCIBW with fast ions injected via a neutral beam.

8.3 Applications to Alternate Concepts

This thesis has focused on α-channelling in tokamaks. It is likely that alternate concepts will also benefit from α-channelling, although the degree to which they benefit, and the ease of implementing α-channelling will surely vary from device to device. Many of the ideas and approaches here will be useful in alternate concepts. Below, we outline

how these ideas might be transferred to other devices, and suggest some ways in which α-channelling might take advantage of some devices' unique capabilities.

Both the benefits that accrue from α-channelling, and the challenges of implementing it, shift when moving from the tokamak to other devices. For example, the large losses of energetic α-particles from devices which lack symmetry (such as conventional stellarators), suggest that the win from α-channelling might even be greater in these devices. On the other hand, the higher density associated with devices which have high β means that the α-particles will slow down rapidly, possibly increasing the power required in the waves to accomplish α-channelling.

For the purposes of this consideration, we separate alternate concepts into two different categories, those which have a symmetry or quasisymmetry (and therefore have three constants of the guiding center motion, like the tokamak), and those that do not (like a conventional stellarator).

For devices with a symmetry or quasisymmetry, such as the reversed field pinch or the quasisymmetric stellarator, the framework of the α-channelling problem outlined in this thesis can be mapped over straightforwardly. For instance, for these devices there exist three dimensional constants-of-motion spaces similar to the ε, μ, and P_ϕ space used here. Presumably this space can be separated into different regions of various orbit topologies as done in Appendix A. Waves will diffuse particles along diffusion paths in this constants-of-motion space in much the same way as was the case for the tokamak. Significant differences from the approach taken here may arise due to the different orbit topologies, orbit width to system size ratios, or different wave physics.

For devices which lack a symmetry, there may exist an opportunity for significantly simplified α-channelling. As an example, consider a conventional stellarator which has a significant loss cone for α-particles. As pointed out by Ho and Kulsrud (1986), if an energetic α-particle is born in or scattered into the loss cone it will be immediately lost from the plasma, while lower energy particles will remain confined due to the effect of the radial electric field. Pitch angle scattering will try to fill in this hole in phase space. However, energetic α-particles are much faster than thermal ions and will tend not to pitch angle scatter until they slow down to a few times the thermal velocity of the background. This means that most of the energy of the energetic α-particles will be

deposited in the plasma (being transferred to the electrons through slowing down), but there will not be significant ash accumulation because α-particles at intermediate energies will be scattered into the loss cone and leave the plasma. The neoclassical transport of the thermal plasma will not be prohibitive due to the confining influence of the radial electric field.

The loss cone, which has usually been considered a serious drawback in stellarators, may allow for simplified implementations of α-channelling relative to tokamaks. In a tokamak, it was necessary to move the cooled particles to the periphery to find a sink; in a stellarator, the loss cone provides the sink. Rather than relying on collisions with electrons to slow down the α-particles, one would look for a wave-particle interaction which moves particles in energy and either radius or pitch angle so that if the α-particle takes a step toward lower energy it also takes a step closer to the loss cone and if the particle is heated, it becomes better confined. On average this wave would extract energy from the α-particle-distribution, and cause the α-particles to leave the device. If this same wave were then to damp on ions it would be able to accomplish α-channelling in stellarators in a significantly easier way than is possible in tokamaks.

If a scenario like the one described above appeared promising, it would be desirable to choose a stellarator equilibrium which has been optimized for α-channelling. One way this might be done would be by taking advantage of a code being developed by Reiman (1997) which searches for stellarator equilibria which maximize a given objective function.

8.4 Summary

The substantial benefit in α-channelling (Chapter 2) exhibited by the zero dimensional reactor simulations, motivated the search for ways in which α-channelling might be implemented. The maximum increase in fusion power was found to occur when the ion and electron heat confinement times were not equal. In particular, it was found that in the presence of α-channelling, it was desirable to have good ion heat confinement but poor electron heat confinement.

To gain insight into this effect, α-channelling was considered in a two dimensional phase space associated with a simple slab geometry (Chapter 3). The diffusion path of α-particles interacting with waves in this space is a straight line. Unfortunately, it appears that waves with the optimal slope for energy extraction are difficult to excite. However, we showed that using two waves, with very different slopes, it was possible to diffuse α-particles to the edge while cooling them. The three dimensional constants of the motion space associated with particle orbits in tokamaks was found to be extremely useful for considering α-channelling in tokamaks. In particular, the diffusion paths due to waves in this space are also straight lines, and much of the insight developed in the simple geometry was found to apply to the more complicated situation of α-particles in a tokamak.

A rapid particle simulation was developed in ε, μ, and P_ϕ space. The full particle dynamics are calculated from the particle's constants of the motion (see Appendix A) and the details of the orbits are used to calculate the increments to each particle's ε, μ, and P_ϕ due to waves and collisions (See Chapter 4). This approach is equivalent to a Monte Carlo solution of the orbit-averaged Fokker Planck equation. For the problems where this approach is applicable, a significant speed up is attained over guiding center codes.

The wave characteristics necessary to significantly cool an α-particle in a tokamak were found to be available in a combination of the MCIBW, which is high frequency and can break the μ invariant, and the TAE, which is low frequency and accomplishes the transport of the α-particles (Chapter 5). Remarkably, configurations of these two waves were found which did a very good job of cooling an entire birth distribution of α-particles. By examining how particles at various points in ε, μ, and P_ϕ space were responding to different combinations of waves, we were able to iterate on these cooling scenarios to eventually obtain one which demonstrated how two waves can be combined in a reverse shear tokamak reactor to absorb 2/3 of the energy from the 93% of the α-particles ejected!

By testing the simulations against TFTR experiments, which showed large losses of fast ions when the MCIBW was present (Chapter 6), we were able to gain confidence in the simulation and therefore the two wave α-channelling scenarios. We were also

able to use the simulations in conjunction with the experiments to infer two aspects of the MCIBW which are critical to the implementation of α-channelling: the existence of the $k_{\|}$-flip of the MCIBW and the diffusion coefficient of particles interacting with the MCIBW (Chapter 7). The MCIBW diffusion coefficient was found to significantly exceed that which is predicted by geometrical-optics estimates. While possible mechanisms for this discrepancy were suggested, the precise reason is an open question, with possibly important ramifications for a reactor implementation of α-channelling.

Taken together, the advances in this thesis show how experiments to date give us a measure of confidence in both the simulations themselves, the underlying physical assumptions, and ultimately the reasonableness of the application of these ideas to α-channelling in a tokamak reactor. In the process of focusing on the problem of α-channelling we built a useful numerical tool and used it to determine certain fundamental characteristics of MCIBW, such as the $k_{\|}$ flip and the effective diffusion coefficient.

Appendix A

Energetic Particle Orbits in Tokamaks

Following the detailed motion of an energetic charged particle orbit in the complicated magnetic field of a tokamak, interacting with various perturbations, represents a daunting computational problem. Ideally one would like to follow the particles for their slowing down time ($\sim$ 0.1 second), however, the particle orbits have spatial variations of order $\rho \sim$ a few cm, and vary on a time scale of $1/\Omega_0 \sim 10^{-9}$ seconds. Simulations following hundreds of thousands of particles for a slowing down time are beyond present computational capabilities. Fortunately one can usually follow, to a very good approximation of the actual orbit, the guiding center of the particle, eliminating the short time scales associated with the gyromotion.

Starting from the particle Hamiltonian, Littlejohn (1981) derived the guiding center Hamiltonian, which has the conservation properties one desires (energy, phase volume, angular momentum), by averaging over the gyromotion. This Hamiltonian formulation can be used to write down guiding center equations of motion, which are particularly convenient if represented in terms of magnetic coordinates [WHITE, BOOZER, and HAY 1982; WHITE and CHANCE 1984]. In the absence of perturbations, and assuming axisymmetry of the tokamak, the guiding center equations of motion are integrable. Thus there exist three constants of the motion (COM) which will completely characterize the motion. Following Rome and Peng (1979) we use three constants of the unperturbed motion to solve for the particle's guiding center orbit. The three constants we choose

are:

$$\varepsilon = \mu B + \frac{1}{2} m v_{\parallel}^2 \tag{A.1}$$

$$\mu = \frac{m v_{\perp}^2}{2B} + \epsilon f(\mathbf{r}, \mathbf{v}, \mathbf{B}) \tag{A.2}$$

$$P_{\phi} = \frac{m R B_{\phi} v_{\parallel}}{B} - \frac{e}{c}\psi, \tag{A.3}$$

where ψ is the poloidal flux enclosed by the flux surface the particle is on divided by 2 π. These expressions are shown to first order in ϵ which is the ratio of gyroradius to the scale length of variations in the magnetic field. As shown by Littlejohn (1983) these constants of the guiding center motion are equivalent to the constants for the actual motion to all orders in ϵ, ensuring that for ϵ small the guiding center approximation is quite good. Since the COM approach conserves the same quantities as the guiding center Hamiltonian the orbits found are equivalent to integrating the guiding center equations of motion, such as those put forth by White and Chance (1984).

A.1 Constants of Motion Approach

Throughout the thesis we normalize length, time, and mass to R_0, $1/\Omega_0$, and m. It is useful to get a sense of the terms which appear here, and with this choice of normalization $\varepsilon, \mu \sim \mathcal{O}(\rho^2/R_0^2)$, $(P_{\phi} + \psi) = F(\psi) v_{\parallel}/B \sim \mathcal{O}(\rho/R_0)$, $F(\psi), B \sim \mathcal{O}(1)$.

Now substituting $v_{\parallel} = \sigma \sqrt{2(\varepsilon - \mu B)}$ gives,

$$P_{\phi} = \frac{F(\psi)\sigma\sqrt{2(\varepsilon - \mu B)}}{B} - \psi, \tag{A.4}$$

where $F = R B_{\phi}$ and σ is the sign of $v_{\parallel}$. This equation can be solved for B to give:

$$B_{orb}(\psi) = \frac{-\mu F(\psi)^2 + \sqrt{(\mu F(\psi)^2)^2 + 2F(\psi)^2 \varepsilon (P_{\phi} + \psi)^2}}{(P_{\phi} + \psi)^2}. \tag{A.5}$$

Eq. (A.5) gives us the value of B_{orb}, as a function of ψ, which the particle must be at in order to conserve all three quantities ε, μ, and P_{ϕ}. In the equilibria we consider, there is

a one to one mapping from a pair (ψ,B) to a (R,Z) in the upper half plane (this is violated in bean shaped equilibria and in some low aspect ratio equilibria at extremely high β). Thus if we know how a particle orbit is mapped out in (ψ,B) (from Eq. (A.5)), then we can draw the real space orbit of that particle. Figure A.1 illustrates the correspondence between the two spaces.

Equation (A.5) can be rewritten as

$$B_{orb}(\psi) = \frac{\varepsilon\left(-1+\sqrt{1+2y(\psi)^2}\right)}{\mu y(\psi)^2} \tag{A.6}$$

$$y(\psi) = \frac{\sqrt{\varepsilon}(P_\phi+\psi)}{\mu F(\psi)} = \frac{\sqrt{2}\lambda}{(1-\lambda^2)}, \tag{A.7}$$

where $\lambda \equiv v_\parallel/v$. We can gain some insight from expanding Eq. (A.6) for large and small values of y.

$$B_{orb}(\psi) = \frac{\varepsilon}{\mu}\left(1-\frac{y^2}{2}+\frac{y^4}{2}+\mathcal{O}(y^6)\right) \qquad y \ll 1 \tag{A.8}$$

$$B_{orb}(\psi) = \frac{\varepsilon}{\mu}\frac{\sqrt{2}}{|y|} \qquad |y| \gg 1 \tag{A.9}$$

We see that $B(\psi)$, for small y, is an inverted parabola which peaks at a value of ε/μ when $y = 0$. Note that when $y = 0$ ψ = -P_ϕ, and, from Eq. (A.4),λ = 0. For large $|y|$, $B_{orb}(\psi)$ falls off like $1/|y|$. Also note that $y > 0 \quad (\psi > -P_\phi) \Rightarrow \lambda > 0$ and $y < 0 \quad (\psi < -P_\phi) \Rightarrow \lambda < 0$.

Thus for any ε, μ, and P_ϕ, each curve $B_{orb}(\psi)$ satisfying Eq. (A.5) has the generic properties shown in Fig A.2. Note, however, that each point on the curve does not necessarily correspond to a point along a particle's orbit. Only the sections of the curve which lie within the physically realizable portion of ψ,B space, as in Fig. A.1(b) are meaningful.

As an example consider, Figure A.3 which shows some $B_{orb}(\psi)$ curves and their orbits vs. R,Z for two values of ε, μ, and P_ϕ, at relatively low energies (the width of the orbit in ψ is small compared to the change in poloidal flux from the center to the wall). One curve has a peak (which is where $v_\parallel$ = 0, ψ = P_ϕ) inside the realizable region,

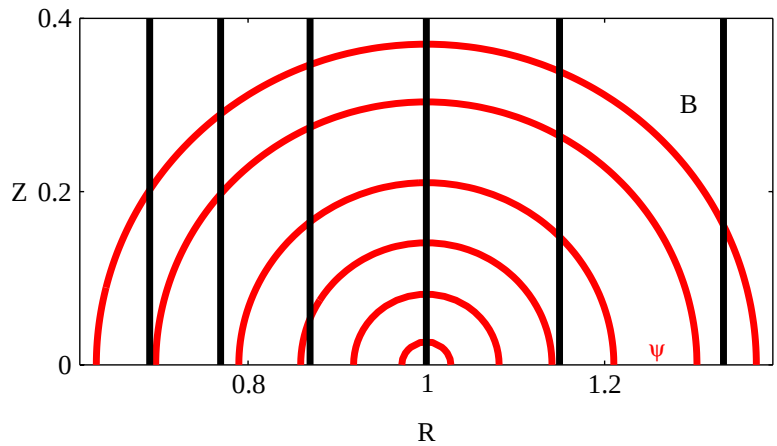

(a) Contours of constant B and ψ, plotted vs. R, Z

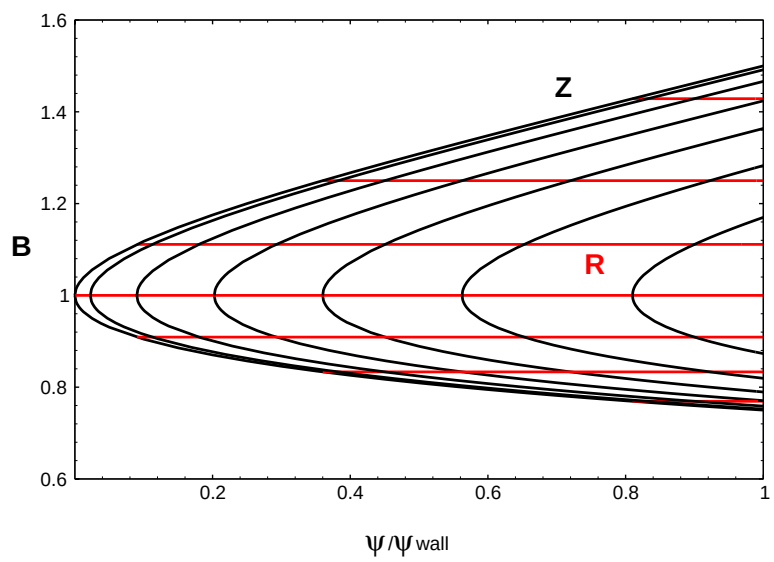

(b) Contours of constant R and Z plotted in ψ, B coordinates

Figure A.1: Alternate representations of a poloidal cross section. Note from (a) that all R, Z (within the walls) correspond to a value of ψ,B, however, all ψ,B do not correspond to a R,Z. The limiting curves are the inner, $B(\psi,\theta = \pi)$ (outer, $B(\psi,\theta = 0)$) midplane corresponding to the maximum (minimum) value of B for a given ψ.

and this corresponds to a trapped particle. The other curve has its peak outside of the realizable region, and corresponds to two distinct passing orbits with the same ε, μ, and P_ϕ, but different signs of $v_\parallel$.

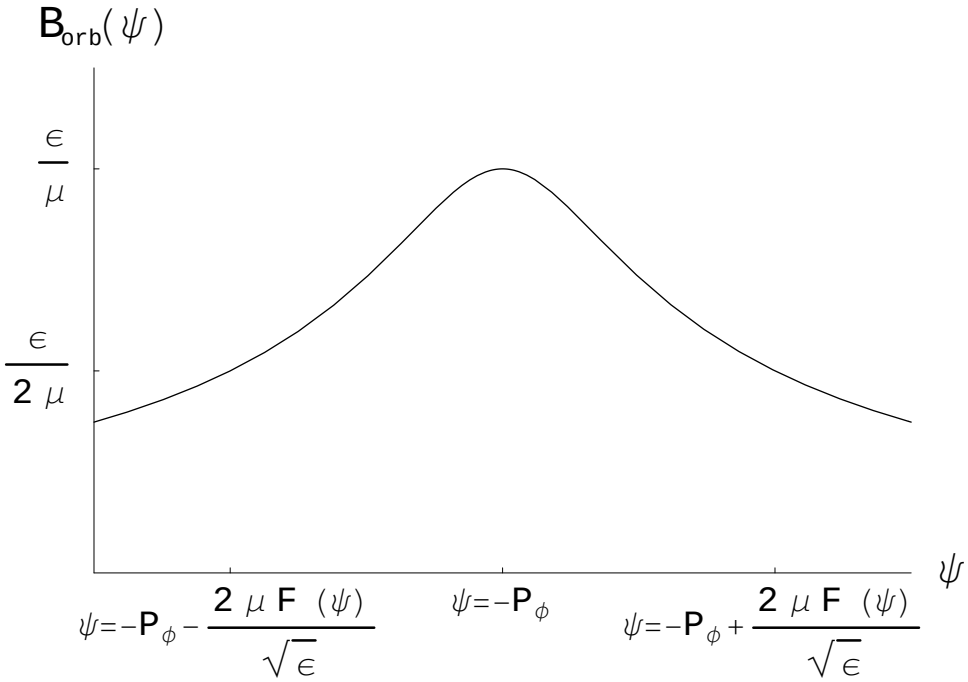

Figure A.2: A generic plot of B_{orb} versus ψ along a particles orbit.

Note from Figure A.2 that the full width at half maximum is about $4\sqrt{\varepsilon}\mu/\varepsilon$. As this width becomes comparable to the total poloidal flux in the plasma new types of orbits appear, created by the different ways in which the generic curve of Fig. A.2 intersects the realizable region. This is discussed in Section A.3.

A.2 General Geometry

Before discussing the orbit topology at high energy we divert for a moment to discuss the magnetic geometries used for the calculations carried out in this thesis. Throughout the thesis we use the formulation of magnetic coordinates used by White (1989). These coordinates are straight-field line coordinates, with the added requirement that the Jacobian,

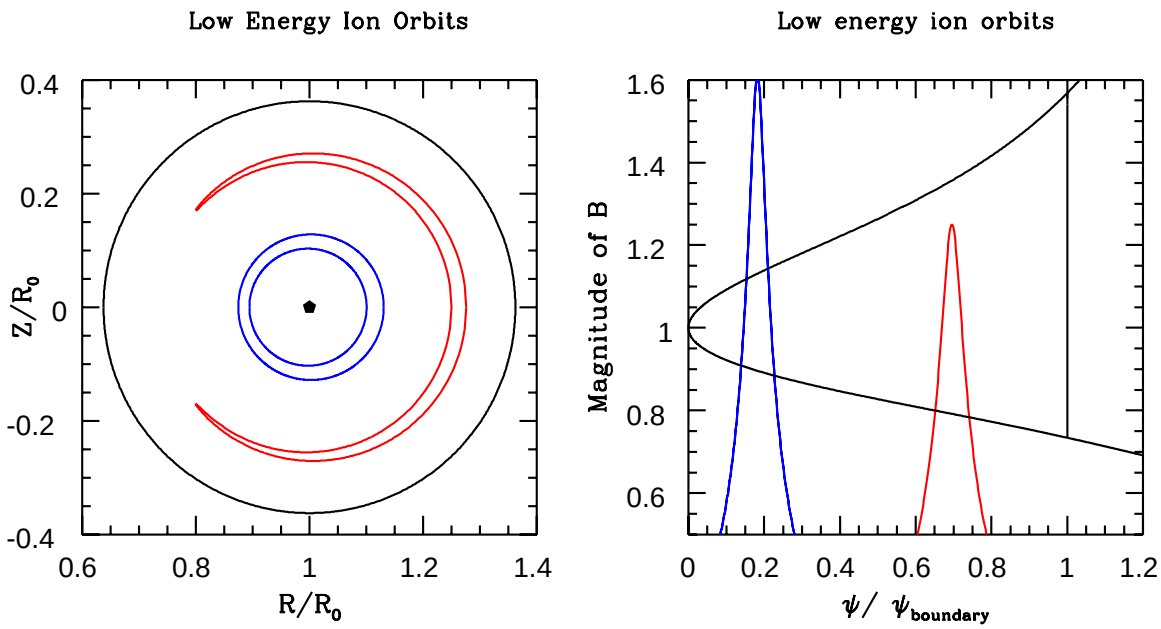

Figure A.3: Low energy orbits in both R,Z space and B, ψ space. Trapped and the two branches of passing particles are depicted.

$J(\psi,\theta) \propto f(\psi)/B^2$. Then the definitions of $\mathbf{B}$ and J are [WHITE 1989]:

$$\mathbf{B} = F(\psi)\boldsymbol{\nabla}\zeta + I(\psi)\boldsymbol{\nabla}\theta + \delta(\psi,\theta)\boldsymbol{\nabla}\psi \tag{A.10}$$

$$= \boldsymbol{\nabla}\zeta \times \boldsymbol{\nabla}\psi + q(\psi)\boldsymbol{\nabla}\psi \times \boldsymbol{\nabla}\theta \tag{A.11}$$

$$J(\psi,\theta) = \frac{1}{\boldsymbol{\nabla}\psi \cdot (\boldsymbol{\nabla}\theta \times \boldsymbol{\nabla}\zeta)} \tag{A.12}$$

$$= \frac{I(\psi) + F(\psi)q(\psi)}{B(\psi,\theta)^2} \tag{A.13}$$

For all these discussion we consider only magnetic equilibrium which satisfy the following:

$$\frac{\partial B}{\partial \theta}|_{\psi=\text{constant}} > 0 \qquad 0 < \theta < \pi \tag{A.14}$$

, that is the value of $|B|$ is monotonically increasing from the outer midplane to the inner one. While one can redefine θ in many different ways, if Eq. (A.14) is true for an equilibria in one coordinated system it will be true in all coordinates systems which have θ = 0 on the outer midplane and $\theta = \pi$ on the inner midplane. When this is true there

is a one to one mapping between (ψ,B) to (R,Z) as discussed earlier. For simplicity we consider equilibria which are up down symmetric, although, in principle this would not have to be the case. For these equilibria, $\partial B/\partial\theta$ at fixed ψ vanishes at $\theta = 0$ and π.

The actual equilibria used are obtained in a series of steps. First an magnetic equilibria is created using `J-Solver` [DeLucia, Jardin, and Todd 1980], which solves the Grad-Shafranov equation. If it is to be based on a TFTR experiment the parameters for this equilibrium come from the output of a `TRANSP` run. Otherwise they are specified directly to `J-Solver`. The output from `J-Solver` is mapped to a straight field coordinate system using the `mapmc` code. Finally, the equilibrium is represented in terms of a biquadratic spline using the `spline` code the basis of which is described in White and Boozer (1995). Figure A.4 show contours of $|B|$ and ψ vs. R and Z, and the converse for a reverse shear aspect ratio 3, reactor with nominal parameters of $A = 3$, $R_0 = 5.4$ m, $B_0 = 6$ T, and $I_p = 16.3$ MA. Note that there is a minimum in $|B|$ on the outer midplane, however, on any flux surface, $|B|$ is monotonically increasing from the outside to the inside and so Eq. (A.14) is satisfied.

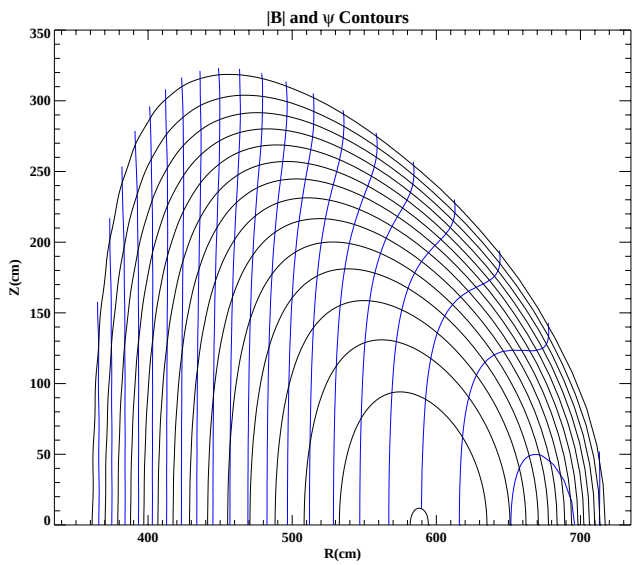

(a) Contours of constant B and ψ, plotted vs. R, Z

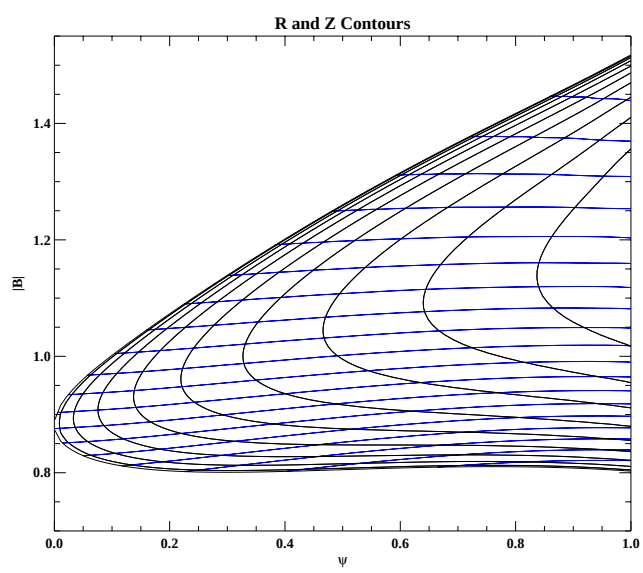

(b) Contours of constant R and Z plotted in ψ, B coordinates

Figure A.4: Alternate representations of a poloidal cross section, as in Fig A.1

A.3 Solving for Orbits

We now discuss, in detail, how one gets from Eq. (A.5) to the poloidal projection of the orbits in a computationally efficient way. Much of the work here has already been discussed by Rome and Peng (1979) and Hsu and Sigmar (1992), for completeness, and to discuss the detailed implementation we expand on these works here.

To begin with, it is useful to consider how the particle moves in B,ψ as it traverses its orbit. Starting with the guiding center velocity

$$\mathbf{v}_{g.c.} = v_{\parallel}\hat{\mathbf{b}} + \mathbf{v}_{\mathrm{d}} \tag{A.15}$$

$$\mathbf{v}_{\mathrm{d}} = K\mathbf{B} \times \boldsymbol{\nabla} B + L\mathbf{B} \times \boldsymbol{\nabla} p \tag{A.16}$$

$$K = \frac{2\varepsilon - \mu B}{B^3} \tag{A.17}$$

$$L = \frac{\beta_0(\varepsilon - \mu B)}{B^4} \tag{A.18}$$

$\beta_0 = 8\pi p_0/B_0^2$, p is normalized to p_0, and all other quantities have been made dimensionless as above. Using the definitions of $\mathbf{B}$ and J Eq. (A.13), we can calculate:

$$\frac{d\psi}{dt} = \mathbf{v}_{\mathrm{d}} \cdot \boldsymbol{\nabla}\psi = K\boldsymbol{\nabla} B \cdot (\boldsymbol{\nabla}\psi \times \mathbf{B}) \tag{A.19}$$

$$= K\left(\frac{\partial B}{\partial \theta}\boldsymbol{\nabla}\theta\right) \cdot (\boldsymbol{\nabla}\psi \times \mathbf{B}) \tag{A.20}$$

$$= \frac{-KF(\psi)}{J(\psi,\theta)}\frac{\partial B}{\partial \theta} \tag{A.21}$$

where several terms vanish because $(\boldsymbol{\nabla}\psi \times \mathbf{B}) \cdot \boldsymbol{\nabla}\psi$ = 0 and $\hat{\mathbf{b}} \cdot \boldsymbol{\nabla}\psi = 0$. The net result is that along the orbit (in the upper half midplane (UHP)) $\dot{\psi} < 0$ except at θ = 0 and π, where $\dot{\psi}$ vanishes. Thus the orbits maximum and minimum values of ψ occur on the midplane, and further in the UHP the particle travels from its maximum ψ_{max} value to the minimum ψ_{min} as it traces out its orbit (since $\dot{\psi}$ is negative). When Eq. (A.14) is violated $\dot{\psi}$ can change signs in the UHP and it is possible for particles to have orbits which are stuck above or below the midplane, bouncing between the points where $\partial B/\partial\theta$ is changing sign, for an example see Mikkelsen et al. (1997).

For reference we also list the equations for $\dot{\theta}$ and $\dot{\zeta}$:

$$\frac{d\theta}{dt} = \frac{v_\parallel}{JB} + \left(K\frac{\partial B}{\partial \psi} + L\frac{dp}{d\psi}\right)\frac{F}{J} \tag{A.22}$$

$$\frac{d\zeta}{dt} = \frac{v_\parallel q}{JB} - \left(K\frac{\partial B}{\partial \psi} + L\frac{dp}{d\psi}\right)\frac{I}{J} \tag{A.23}$$

A term of order δ in $\dot{\zeta}$ has been neglected. Note that in many cases, especially for particles with $\varepsilon \approx \mu B$ or in plasmas with low β_0, $Ldp/d\psi \ll KdB/d\psi$ and this term can be neglected. Unless specifically mentioned later we will neglect this term.

A.3.1 Orbit types

Returning to the COM approach, there are 8 different ways in which the realizable region can be intersected by the $B_{orb}(\psi)$ curves. For the sake of completeness, examples of each possible intersections of $B(\psi)$is shown below along with a real space orbit. This categorization is quite useful, as it will be used extensively in the COM simulation which we have developed. Furthermore, there is a correspondence between these orbit types and ε, μ, and P_ϕ space which will be discussed in the next section. Finally, understanding each of these orbit types is extremely helpful when searching for the intersections of $B_{orb}(\psi)$ with the boundaries of the realizable space.

Some definitions will be useful in the following discussion. An orbit is passing if $|v_\parallel| > 0$ everywhere on its orbit. A particle is cogoing (countergoing) if its parallel velocity is in the same (opposite) direction as the plasma current. Conversely, a particle is trapped if $v_\parallel = 0$ somewhere on its orbit. An orbit is consider encircling if it encircles the magnetic axis, or equivalently, $\dot{\theta}$ is nowhere equal to zero. These two concepts can also be defined in terms of B_{orb} curves, a particle is trapped if B(-P_ϕ,0) $< B_{orb}$(-P_ϕ) = $\varepsilon/\mu < B$(-P_ϕ,π), (hereafter we refer to $B(\psi$,0) and $B(\psi$,π) as B_{out} and B_{in}). An orbit is encircling if it intersects both B_{out} and B_{in}). At low energies, almost all particles are either trapped and non-encircling or passing and encircling.

Type 1 (Fig. A.5) orbits are passing particles which are non-encircling. These orbits exists near the outer midplane and at low energy (with the same pitch, and position)

they would be trapped. However, as seen from Eq. (A.22), there is a competition between the drift term and the parallel motion. For these orbits the parallel motion is insufficient to overcome the drifts and $\dot{\theta}$ goes to zero before the orbit can reach large enough B to mirror. From Eq. (A.22) this competition can only occur when $v_{\|}$ has the opposite sign as $\partial B_{out}/\partial\psi$. Note, from the B,$\psi$ figure that both intersections are on the outer midplane (the lower curve) and lie to the right of ψ = -P_ϕ for this example since $\partial B_{out}/\partial\psi < 0 \Rightarrow v_{\|} > 0 \Rightarrow \psi > -P_\phi$.

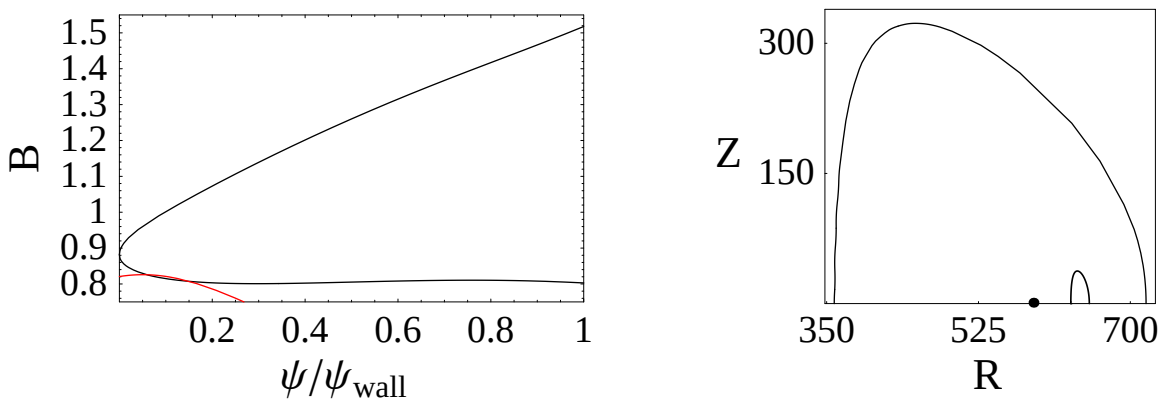

Figure A.5: Orbit type 1. The graph on the left is a plot of $B_{orb}(\psi)$ with $B_{in}(\psi)$(upper) and $B_{out}(\psi)$(lower) providing the boundaries of the realizable space for the reactor described in Sec. A.2. On the right is a poloidal cross section of the upper half plane, with the orbit and the location of the wall shown. The magnetic axis is represented by the dot.

Type 2 (Fig. A.6) orbits are trapped. For these orbits $\dot{\theta}$ is dominated by the parallel motion and thus changes sign near the turning point, where $v_{\|}$ changes sign. Thus these orbits are non encircling. $B_{orb}(\psi)$ has two roots on the outer midplane, one on each side of -P_ϕ.

Type 3 (Fig. A.7) orbits are trapped but encircling. They have been dubbed potato orbits. These orbits have their turning point (where $v_{\|}$ = 0) near the inner midplane, where, typically, the sign of the $\partial B_{out}/\partial\psi$ has changed from it value on the outer midplane. For these orbits the drifts are sufficient to over come the reversed parallel velocity and pull the orbit across the inner midplane. These orbits intersect B_{out} with $\psi >$ -P_ϕ and B_{in} for $\psi >$ -P_ϕ.

Type 4 (Fig. A.8) orbits are cogoing passing orbits which are encircling. These

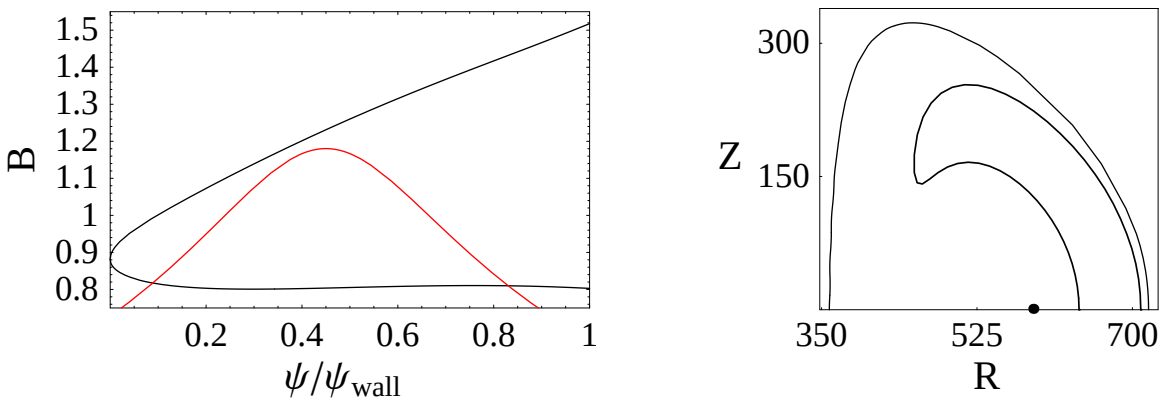

Figure A.6: Orbit type 2.

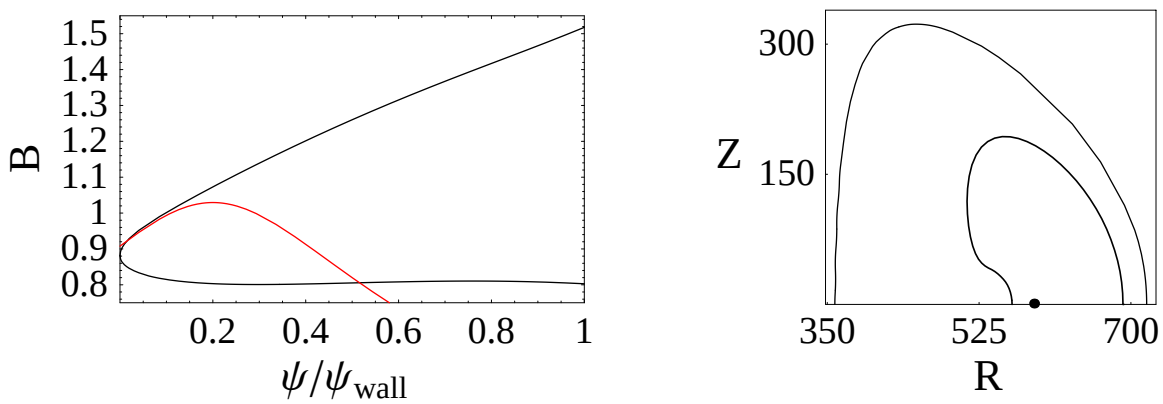

Figure A.7: Orbit type 3.

are similar to cogoing type 7 orbits, however the values of ε, μ, and P_ϕ are such that there is no countergoing complement to these orbits, unlike for type 7. These orbits intersect both B_{in} and B_{out} with $\psi >$ -P_ϕ.

For orbit types 1 through 4 there was only one possible orbit for each ε, μ, and P_ϕ, however for orbit types 5 through 8, there exists two possible orbits for each value of ε, μ, and P_ϕ. One orbit is always a countergoing passing orbit, though it may be encircling or nonencircling, while the other curve is encircling but it may be trapped or cogoing passing.

Type 5 B_{orb} curves (Fig. A.9) have two possible orbits. One is a cogoing passing orbit which is encircling, just like type 4 in terms of intersections with B_{in} and B_{out}.

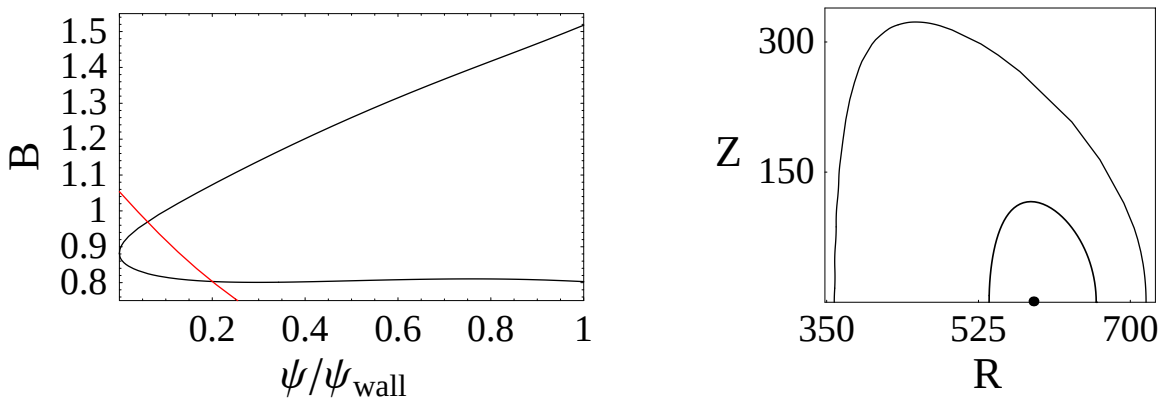

Figure A.8: Orbit type 4.

The other orbit however is quite different, it is a countergoing passing, non encircling orbit, which exists on the inner midplane. In fact, it is the equivalent of a type 1 orbit, on the inner midplane, i.e. it arises from a balancing of drifts and parallel motion. The countergoing portion of this orbit intersects B_{in}, twice, with $\psi <$ -P_ϕ.

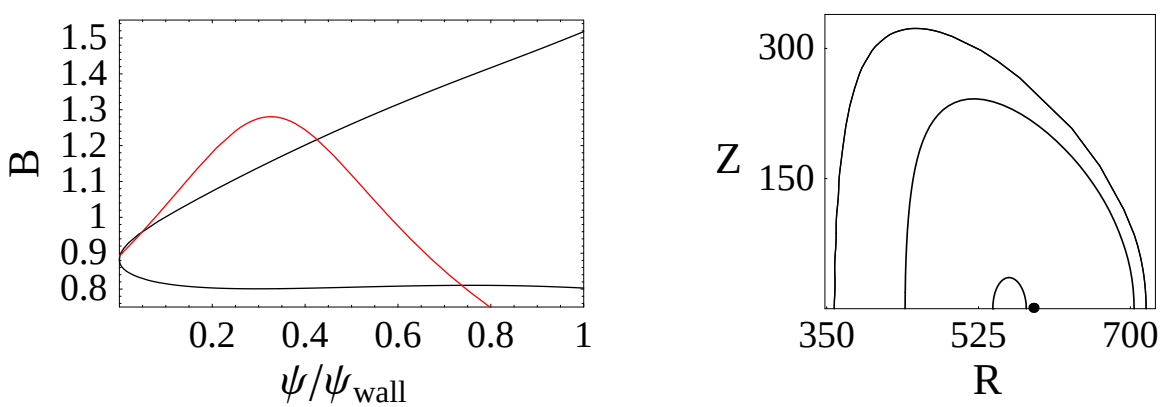

Figure A.9: Orbit type 5. Two orbits are shown(right), representing the two possible tokamak orbits with the specified values of ε, μ, and P_ϕ and the two separate intersections of B_{orb} with the realizable region (left).

Type 6 B_{orb} curves (Fig. A.10) also have two possible orbits. One is an encircling trapped orbit, like type 3, while the other is a countergoing, encircling, passing orbit, describe in type 7 below.

Type 7 B_{orb} curves (Fig. A.11) correspond to one countergoing encircling passing

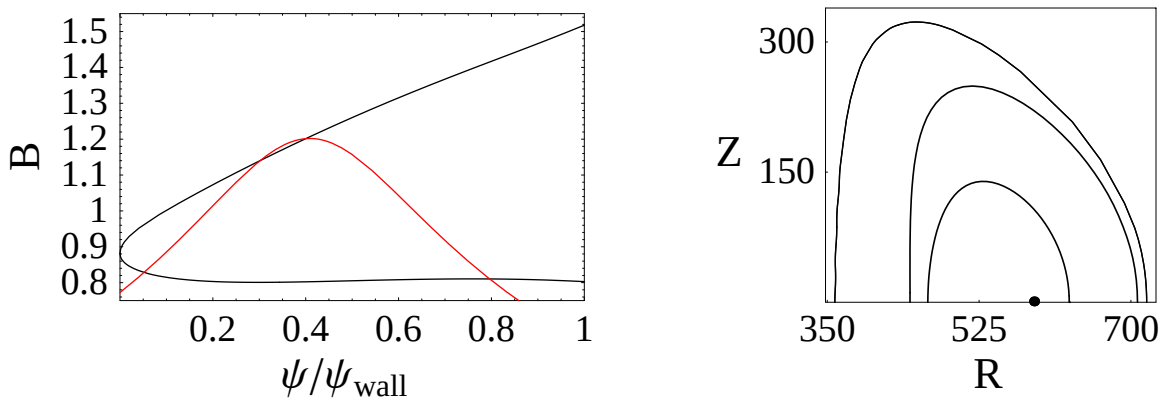

Figure A.10: Orbit type 6.

orbit and a cogoing encircling passing orbit. Each orbit intersects B_{in} and B_{out}.

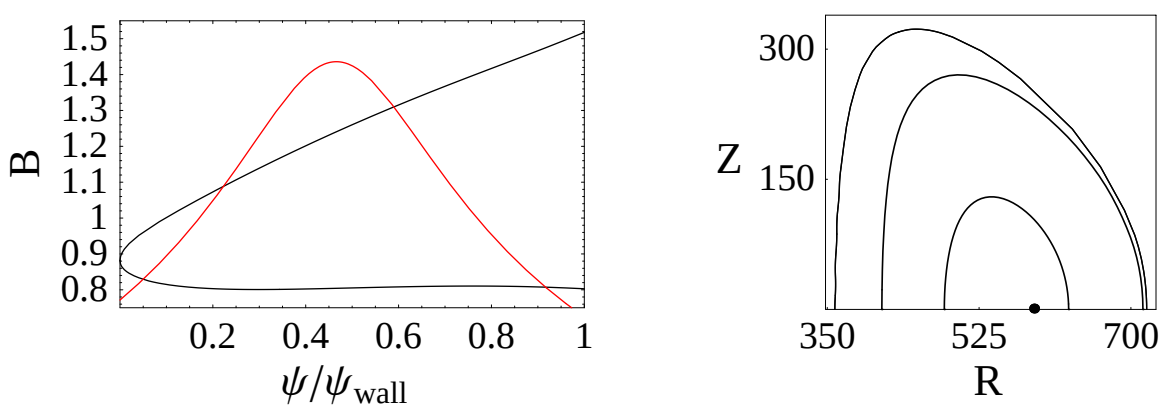

Figure A.11: Orbit type 7.

Type 8 B_{orb} curves (Fig. A.11) correspond to one countergoing non encircling passing orbit (equivalent to the type 5 countergoing orbit) and a trapped encircling orbit (like type 3).

A.3.2 Bounds on orbits

When simulating the energetic α-particles interacting with waves, or colliding with the background plasma (see Chapter 4)it is not necessary to solve for whole poloidal pro-

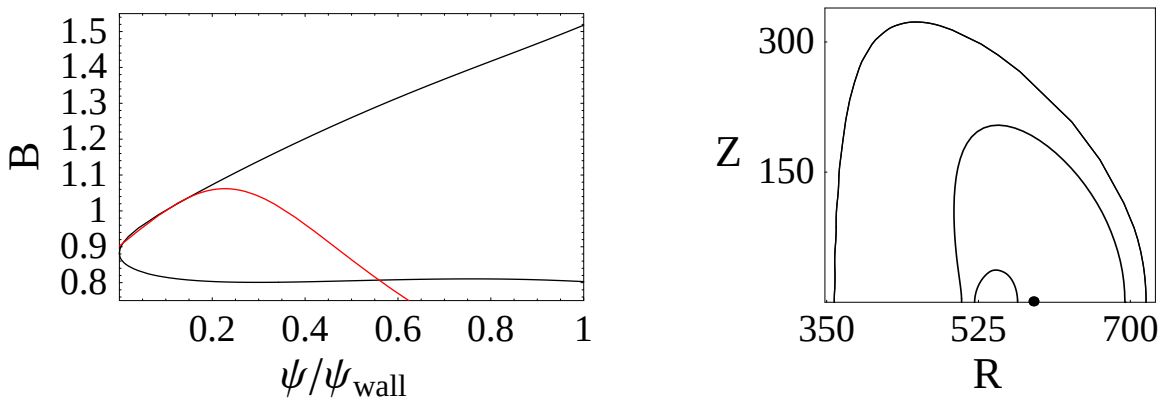

Figure A.12: Orbit type 8.

jection of the orbit for each particle at each time step. However, it is useful to know the extent of the orbit in real space, or to be able quickly determine points on the orbit, if necessary.

Each of these different orbit types occupies a distinct region in ε, μ, and P_ϕ space which is discussed in Section A.4. Thus given ε, μ, and P_ϕ for the particle, we will be able to determine the orbit type, and which branch the particle is on. It is then important to be able to determine ψ_{max} and ψ_{min}. For instance, once ψ_{max} and ψ_{min} are known, the real space orbit can be traced out as in Alg. 1.

Algorithm 1 Calculate poloidal projection of orbit from COM

```
for ψ in (ψ_min,ψ_max) do
    compute B_orb(ψ)
    compute R(ψ,B),Z(ψ,B)
end for
```

If ψ were taken to be outside the boundaries, the computation of R,Z would fail, or worse yet, two branches of the B_{orb} curve might be lumped as one producing real space orbits with discontinuities because of jumps between branches. Furthermore ψ_{min},ψ_{max} are used in computation of the bounce averages (Sec.A.5) which are used when considering collisions.

Depending on the orbit type and branch, ψ_{max} and ψ_{min} are roots of the functions:

$$B_{in}(\psi) - B_{orb}(\psi) \quad \text{or} \quad B_{orb}(\psi) - B_{out}(\psi) \tag{A.24}$$

There can be as many as three and as few as no roots to either of these equations in typical equilibria. Given an orbit type and a branch, we are only interested in the roots which are appropriate for that branch. The root search is most efficiently done with a detailed knowledge of the orbit types. Below, assuming the orbit type is known, as well as, for those types with two orbits, which branch, we detail how the roots are found. The computational approach to the root finding adopted, as advocated by Press et al. (1995a), is to always bracket the roots, that is find two values of ψ such that the functions above change sign in the interval.

For example orbit type 2 is easy. It has two roots with B_{out}, ψ_{min} between $(0, -P_\phi)$ and ψ_{max} between $(-P_\phi, \infty)$. Of course the particle will hit the wall when $\psi = \psi_{\text{wall}}$ and in the simulation when the orbit type and branch are determined, it is also determined if the particle hits the wall, if so a flag is set and ψ_{max} is taken to be ψ_{wall}. For the rest of this discussion ψ_{max} is assumed to be less ψ_{wall} and the issue of particles hitting the wall is returned to in Sec. A.4.

Some orbit types prove to be more difficult than orbit type 2 in that the roots are not easily bracketed. In particular, orbit types 1, 5, 6, and 8 are challenging. Our first approach valid for orbit type 6,8 is to try to analytically determine find the minimum of the curve B_{in}-B_{orb}. Knowing this point will give another point to use for bracketing the roots. We know that this minimum must occur some place where the derivative with respect to ψ of this is equal to zero, i.e. $B'_{in}(\psi) - B'_{orb}(\psi) = 0$. From looking at Fig A.10 and A.12 we see that this point is near the maximum of B_{orb}, i.e. ψ = -P_ϕ. Thus we expand both B'_{in} and B'_{orb} in a Taylor series near this point and solve for the ψ which sets the derivative to zero. The result is

$$\psi_{xi} = -P_\phi - \frac{\mu^3 F(-P_\phi)^2 B'_{in}(-P_\phi)}{\varepsilon^2 + \mu^3 F(-P_\phi)^2 B''_{in}(-P_\phi)} \tag{A.25}$$

We can do the exact same thing for otype 1, the only difference being that we find the

maximum of $B'_{orb}(\psi) - B_{out}(\psi)$ near -P_ϕ, giving

$$\psi_{xo} = -P_\phi - \frac{\mu^3 F(-P_\phi)^2 B'_{out}(-P_\phi)}{\varepsilon^2 + \mu^3 F(-P_\phi)^2 B''_{out}(-P_\phi)} \tag{A.26}$$

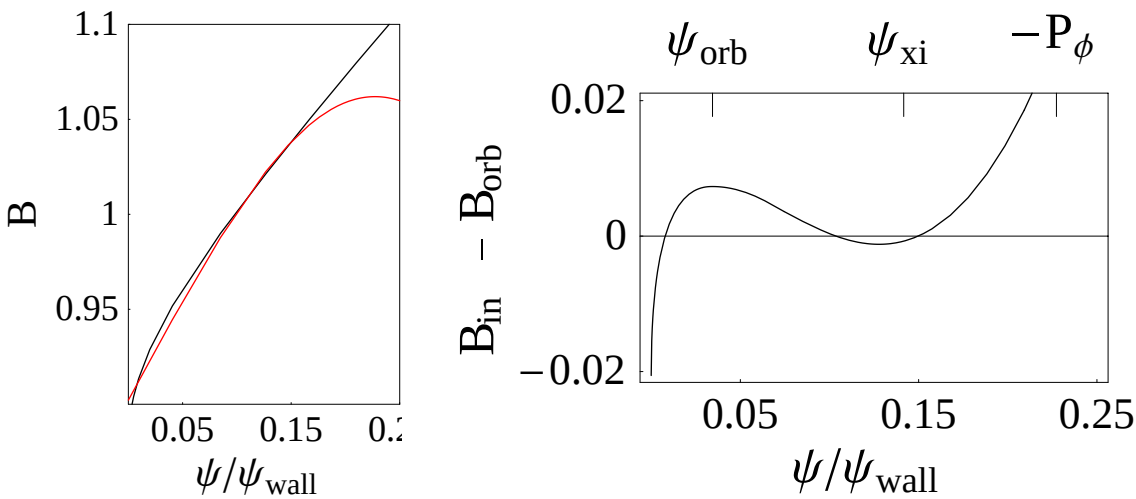

(a) (b)

Figure A.13: (a) Expanded view of Fig. A.12. (b) Graph of B_{in}-B_{orb} versus ψ. This function is positive when B_{orb} is inside the realizable region. ψ = ψ_{orb},ψ_{xi} and -P_ϕ are labeled.

Finally for the case of the countergoing portion of otype 5 and 8 it is necessary to actually search for a point on the orbit, that is a point where $B_{in}(\psi) - B_{orb}(\psi) > 0$. This is implemented by searching for a maximum of $B_{in}(\psi) - B_{orb}(\psi)$ between the points ψ = 0, ψ = ψ_{xi} where the function is negative. We label the ψ associated with this maximum ψ_{orb}. An example, which illustrates these concepts for orbit type is given Fig. A.13, which is based on Fig. A.12.

In Table A.1 the types of orbits, their properties and brackets for ψ_{min},ψ_{max} are given.

Orbit Type	Trapped/ Passing	Encircling	ψ_{min} Brackets & B_{orb} intersect	ψ_{max} Brackets & B_{orb} intersect
1	co	N	$(0, \psi_{xo})$,B_{out}	$(\psi_{xo}, \psi_{\text{wall}})$,$B_{out}$
2	trapped	N	$(0, -P_\phi)$,B_{out}	$(-P_\phi, \psi_{\text{wall}})$,$B_{out}$
3	trapped	Y	$(0, -P_\phi)$,B_{in}	$(-P_\phi, \psi_{\text{wall}})$,$B_{out}$
4	co	Y	$(0, \psi_{\text{wall}})$,B_{in}	$(0, \psi_{\text{wall}})$,B_{out}
5	counter	N	$(0, \psi_{orb})$, B_{in}	$(\psi_{orb}, -P_\phi)$, B_{in}
	co	Y	$(-P_\phi, \psi_{\text{wall}})$,$B_{in}$	$(-P_\phi, \psi_{\text{wall}})$,$B_{out}$
6	counter	Y	$(0, -P_\phi)$,B_{out}	$(0, \psi_{xi})$,B_{in}
	trapped	Y	$(\psi_{xi}, -P_\phi)$,B_{in}	$(-P_\phi, \psi_{\text{wall}})$,$B_{out}$
7	counter	Y	$(0, -P_\phi)$,B_{out}	$(0, -P_\phi)$,B_{in}
	co	Y	$(-P_\phi, \psi_{\text{wall}})$,$B_{in}$	$(-P_\phi, \psi_{\text{wall}})$,$B_{out}$
8	counter	N	$(0, \psi_{orb})$, B_{in}	(ψ_{orb}, ψ_{xi}), B_{in}
	trapped	Y	$(\psi_{xi}, -P_\phi)$,B_{in}	$(-P_\phi, \psi_{\text{wall}})$,$B_{out}$

Table A.1: Description of all the orbit types and their branches. Values of ψ which bracket each of the roots (ψ_{min},ψ_{max}) of Eq. (A.24) are given. ψ_{orb}, ψ_{xi}, ψ_{xo} are defined in the text.

A.3.3 Comparison with ORBIT code

The simulation which we have developed implements the preceding discussion, allowing ψ_{min},ψ_{max} to be rapidly found, and then for example using Alg. 1, we can compute the real space orbits. To benchmark that this is working we also computed the orbits using the ORBIT guiding center code [WHITE and CHANCE 1984]. The results are presented in Fig. A.14 showing excellent agreement.

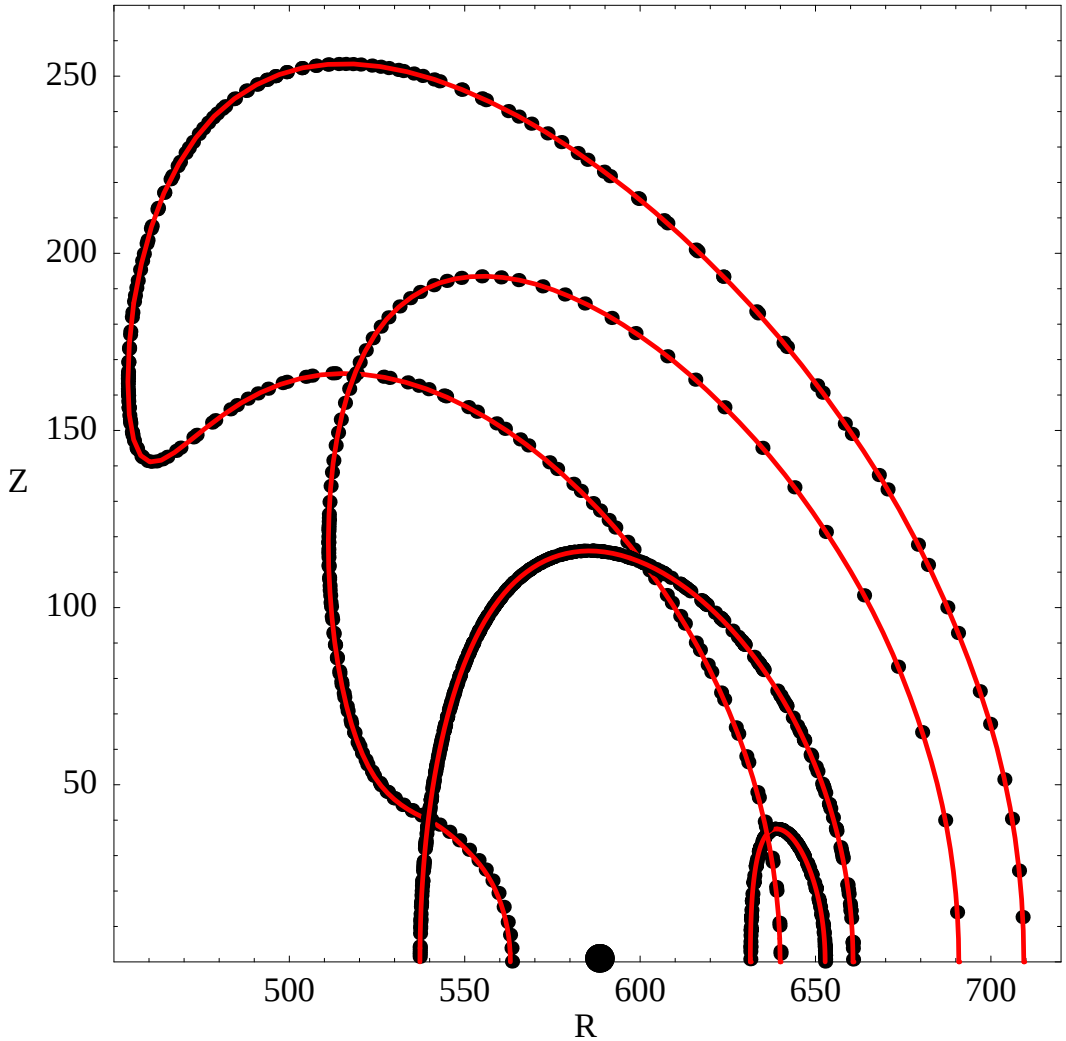

Figure A.14: Comparison of the COM approach(line) and the ORBIT code(points) for determining the poloidal projection of an energetic particle's orbit. The equilibrium used is described in Sec. A.2.

A.4 ε, μ, and P_ϕ space

We have seen that, given the orbit type, a branch choice and three constants of the motion we can determine the orbit of a particle. For the COM simulation, though, we do not want to carry around all the information about the orbit. By understanding the structure of ε, μ, and P_ϕ space we can determine relevant information about the particles orbit without resorting to actually computing the orbit.

A.4.1 Basics

Thus given a particle's initial position, pitch and energy we would like to determine its orbit type. Furthermore, we would like to know if the particle is confined or lost. Understanding ε, μ, and P_ϕ space helps with these determinations. In particular, slices of ε, μ, and P_ϕ space at fixed energy are instructive. In these cross sections we can use Equations (A.1)-(A.4) to solve for μ in terms of P_ϕ, ε, ψ, and B.

$$\mu = \frac{\varepsilon}{B} - \frac{B}{2F^2}(P_\phi + \psi)^2 \tag{A.27}$$

P_ϕ is expressed in terms of μ, ε, ψ, B, and σ in Eq. (A.4). Then an orbit which passes through position, (ψ,B), with energy, ε, must have μ and P_ϕ somewhere on the curve traced out by Eq. (A.27). This curve has the shape of an inverted parabola, when plotted in a P_ϕ, μ coordinate system. In fact the orbit of a particle could be determined by finding (ψ,B) for all of the parabolas which intersect its value of μ and P_ϕ. In Figure A.15 we plot a constant energy slice of ε, μ, and P_ϕ space, with μ/ε plotted vs. $P_\phi/\psi_{\text{wall}}$. We use Eq. (A.27) to plot μ versus P_ϕ curves for the magnetic axis (which is labeled A in the figure), and the outer(C) and inner(B) midplane at the wall. Also plotted, representing a curve opening to the left, is $1/B_{in}(-P_\phi)$ (E) and $1/B_{out}(-P_\phi)$ (D). Particles with orbits which pass through the magnetic axis lie on the curve marked by A1, A2 and A3. A1 is an orbit with $\lambda = -1$ at the axis. A2 and A3 are orbits with $\lambda = 0$ and $\lambda = 1$ respectively. Orbits passing through the outer (inner) midplane at the wall, would lie on curve C (B). Trapped particles must lie between curves D and E, i.e. they must have high μ /ε between $1/B_{in}(-P_\phi)$ and $1/B_{out}(-P_\phi)$. A cogoing or trapped orbit will become lost (hit the wall) if it moves across the right half of curve C from right to left. A countergoing orbit is lost when it crosses the left half of curve B, or if it crosses the curve connecting B and C (i.e., if it crosses the passing trapped boundary). The curves for different energies can be obtained by scaling the width of curves A,B,C in proportion to $\sqrt{\varepsilon/\varepsilon_0}$ and, at the same time, scaling the height of all of the curves by $\varepsilon/\varepsilon_0$.

With the addition of two more curves ε, μ, and P_ϕ space can be broken up into 8 region which correspond to the 8 orbit types discussed in Sec. A.3. This is shown in Fig. A.16, note that curves B and C from the Fig A.15 are missing, as these just provide a reference for the location of the wall, and do not affect the orbit classifications. The

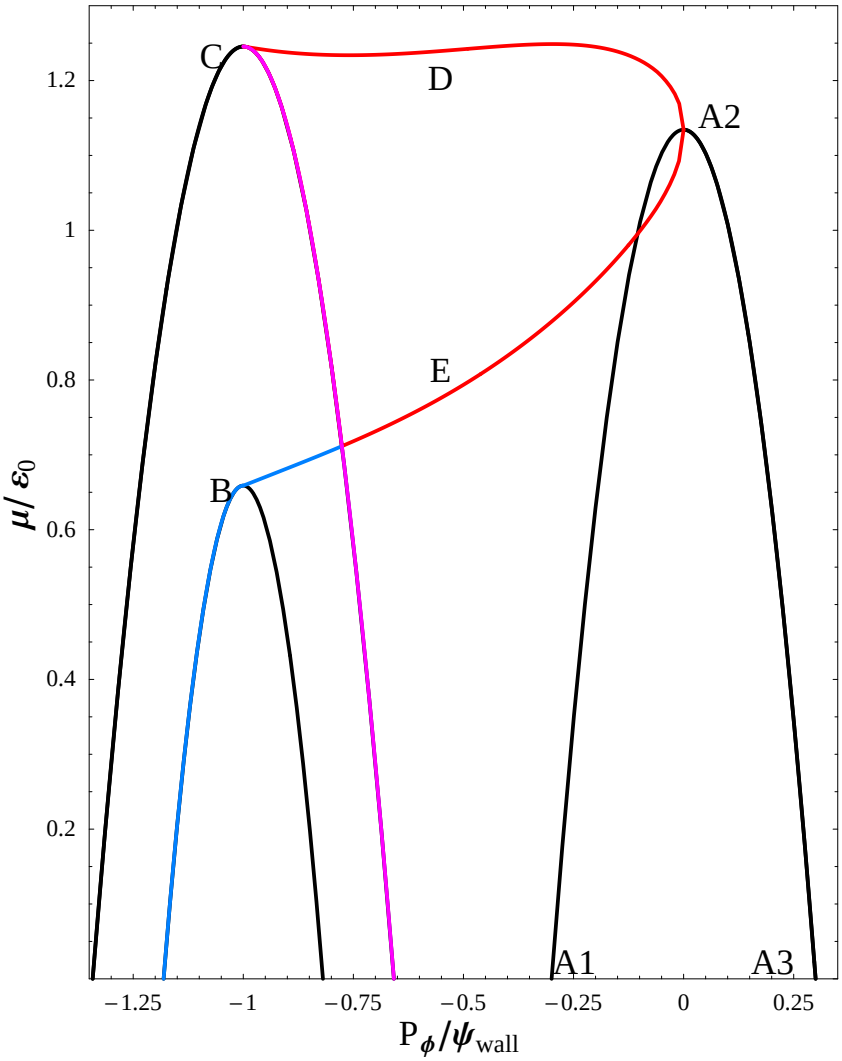

Figure A.15: A constant energy slice of ε, μ, and P_ϕ space. The various curves are described in the text.

added two curves are labeled F and G, and they play an important role in the makeup of ε, μ, and P_ϕ space. Both curves are curves of stagnation point orbits, that is guiding center orbits which are just a point in the poloidal cross section. Above curve F no orbits exist, that is no particle with energy ε could have values of P_ϕ and μ above curve F. Between curve F and curve G, each point corresponds to exactly one orbit. Inside the region outlined by curve G two orbits exist for every point. Note that Fig A.16 has 9 distinct regions. One of these corresponds to unphysical orbits, leaving eight regions. The details of the stagnation points and how to draw curves F and G is discussed below.

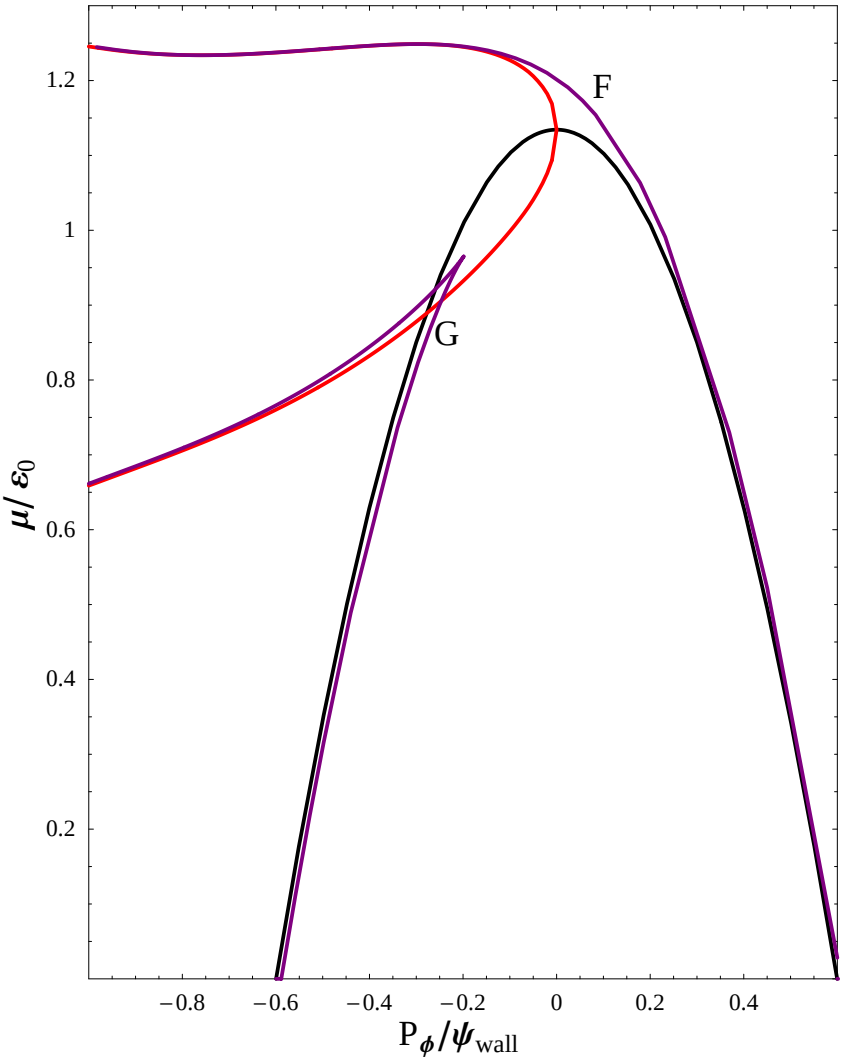

Figure A.16: Boundaries in ε, μ, and P_ϕ space. Curves F and G denote stagnation point orbits.

A.4.2 Stagnation point orbits

These orbits represent a boundary between various regions in ε, μ, and P_ϕ space and the limiting extension of orbit types 1,5,6,8. Stagnation point orbits are guiding center orbits which are just a point. They arise from the exact balancing of the parallel motion and the drifts on the midplane. The two requirements for stagnation orbit are $\dot{\psi} = 0$ and $\dot{\theta} = 0$. Recall, $\dot{\psi}$ is zero at $\theta = 0$ and π. Setting Eq. (A.22) to zero, one can find the stagnation points, in particular we would like to know the pitch, at a given energy and location, of

the orbit which is stagnated. Below B_{mid} refers to either B_{in} or B_{out}.

$$\frac{\sqrt{2\varepsilon}\lambda}{B_{mid}} = -\left(\frac{\varepsilon\left(1+\lambda^2\right)F}{B_{mid}^3}\frac{\partial B_{mid}}{\partial\psi} + \frac{\beta_0\varepsilon\lambda^2 F}{B_{mid}^4}\frac{dp}{d\psi}\right) \tag{A.28}$$

$$2\lambda = -\left(\frac{\sqrt{2\varepsilon}F}{B_{mid}^2}\frac{\partial B_{mid}}{\partial\psi}\left(1+\lambda^2\right) + \frac{\beta_0\sqrt{2\varepsilon}F}{B_{mid}^3}\frac{dp}{d\psi}\lambda^2\right) \tag{A.29}$$

, defining α and η as ,

$$\alpha = \frac{\sqrt{2\varepsilon}F}{B_{mid}^2}\frac{\partial B_{mid}}{\partial\psi} \tag{A.30}$$

$$\eta = \frac{\beta_0\sqrt{2\varepsilon}F}{B_{mid}^3}\frac{dp}{d\psi} \tag{A.31}$$

we see that

$$\lambda_{stag} = \frac{-1+\sqrt{1-\alpha^2-\alpha\eta}}{\alpha-\eta} \tag{A.32}$$

where we have chosen the root which is physical in the $\alpha \to 0, \eta \to 0$ limit. Thus in order to draw curves F and G from Fig. A.16, one would plot parametrically plot

$$\mu_{stag} = \frac{\varepsilon(1-\lambda_{stag}(\psi)^2)}{B_{mid}(\psi)} \tag{A.33}$$

$$P_{\phi stag} = \frac{\sqrt{2\varepsilon}F(\psi)\lambda_{stag}(\psi)}{B_{mid}(\psi)} - \psi \tag{A.34}$$

over ψ from 0 to ψ_{wall}, for $B_{mid} = B_{in}$, corresponding to curve G and $B_{mid} = B_{out}$, corresponding to curve F. In the next subsection these results will be used to determine how to classify an orbit given its ε, μ, and P_ϕ values.

A.4.3 Classifying orbits

In Fig. A.17 we show a close up of Fig. A.16, focusing on the area near μ /ε = 1,P_ϕ =0. Each region has been labeled according to the orbit type which occupies that region, and a point has been placed on the figure corresponding to the orbits shown in

Figures A.5 through A.12. Once the overall structure of this space is known the graph-

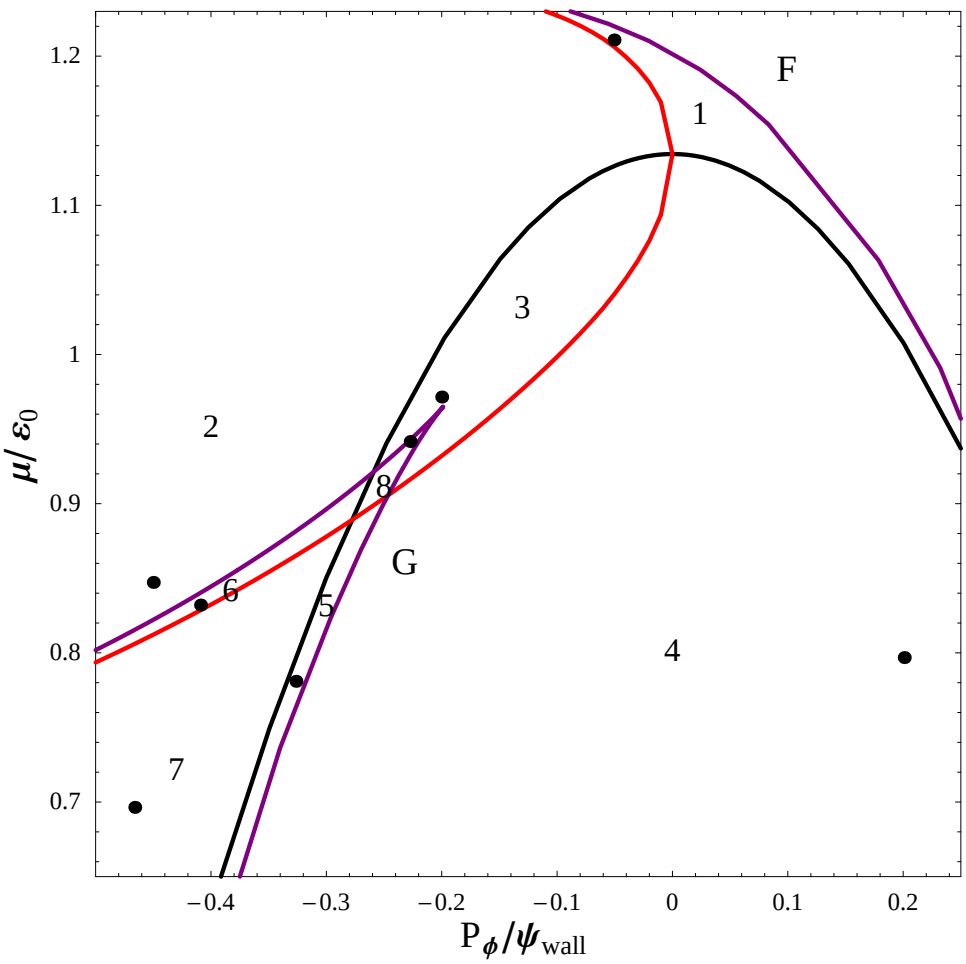

Figure A.17: ε, μ, and P_ϕ blowup

ical determination of an orbit type from just ε, μ, and P_ϕ is quite easy, and for most orbit types the rapid numerical determination is equally simple, However for some orbit types it is more difficult. For example it is quite easy to tell if a ε, μ, and P_ϕ correspond to a type 7 orbit. The two requirements are just 1) $\mu/\varepsilon < 1/B_{in}(-P_\phi)$ and $P_\phi < -F(0)\sqrt{2(\varepsilon - \mu B(0,0))}/B(0,0)$. These requirements can be determined from the behavior of B_{orb} outlined in Sec. A.3.

However to determine if a ε, μ, and P_ϕ correspond to a type 1 orbit is more complicated. Type 1 orbits are points which lie above curve A if $P_\phi > 0$ or above curve D if $P_\phi < 0$. They also must lie below curve F. The first condition is easily determined, as

there are simple analytic expressions for curves A and D, but the latter is more involved, since we do not have a closed form expression for curve F, we only have a parametric expression for it in terms of ψ (Eq. (A.34)). The approach we take is to find the ψ which satisfies $P_{\phi stag}$ (ψ)=P_ϕ, then calculate μ_{stag} from this ψ, then if $\mu < \mu_{stag}$ the point lies below curve F, and therefore, assuming the first condition is satisfied, this corresponds to a type 1 orbit. The only difficulty of this approach is that a root finding routine is necessary to find the ψ.

A.5 Averages Along the Particle Motion

In many instances it is necessary to integrate a quantity over the orbit of the particle, for example when considering the effect of collisions on the particle orbits, or when computing bounce or precession frequencies. In this section, the method used in the COM simulation is explained.

As an example let us consider the computation of the bounce frequency, ω_b, which is the frequency with which a particle completes its poloidal orbit. Typically this is done by

$$\omega_b = 2\int_0^{\theta_{max}} \frac{d\theta}{\dot{\theta}} \tag{A.35}$$

where θ_{max} is π for encircling particles and the point where $\dot{\theta}$ vanishes for non-encircling particles. For our formulation this approach is not convenient, since the orbits are not naturally expressed in terms of θ, rather they are expressed in terms of ψ. Thus an especially convenient way of doing this integral would be

$$\omega_b = 2\int_{\psi_{max}}^{\psi_{min}} \frac{d\psi}{\dot{\psi}} \tag{A.36}$$

However, $\dot{\psi}$ vanishes at ψ_{min} and ψ_{max}, and thus the integral is singular. This singularity is integrable, though, since the equivalent calculation in terms of θ has no singularity at the end points.

How can we quickly evaluate this integral while avoiding problems caused by the endpoint singularities? Consider rewriting the integral in the following way:

$$u = 2\frac{\psi - (\psi_{max} + \psi_{min})/2}{\psi_{max} - \psi_{min}} \tag{A.37}$$

$$\omega_b = (\psi_{max} - \psi_{min}) \int_{-1}^{1} \frac{\sqrt{1-u^2}}{\sqrt{1-u^2}} \frac{du}{\dot{\psi}(u)} \tag{A.38}$$

$$= (\psi_{max} - \psi_{min}) \int_{-1}^{1} \frac{f(u)du}{\sqrt{1-u^2}} \tag{A.39}$$

where $f(u) = \sqrt{1-u^2}/\dot{\psi}(u)$. The integral of Eq. (A.39) can be done with the Gauss-Chebyshev quadrature:

$$\int_{-1}^{1} \frac{f(u)}{\sqrt{1-u^2}} du \approx \frac{\pi}{n} \sum_{k=1}^{n} f\left(\cos\left(\frac{2\,k-1}{2\,n}\,\pi\right)\right) \tag{A.40}$$

Gaussian quadratures have well known advantages [PRESS et al. 1995b], e.g. their order is twice that of other routines for the same number of function calls. In fact, the answer will be exact for a function $f(u)$ which is a polynomial of degree less than $2n$. In Fig A.18 we plot $1/\dot{\psi}$ versus u as well as $f(u)$ for the orbit shown in Fig. A.6. We see that $f(u)$ is a smooth, slowly varying function, and thus expect the Gauss-Chebyshev quadrature would work quite well.

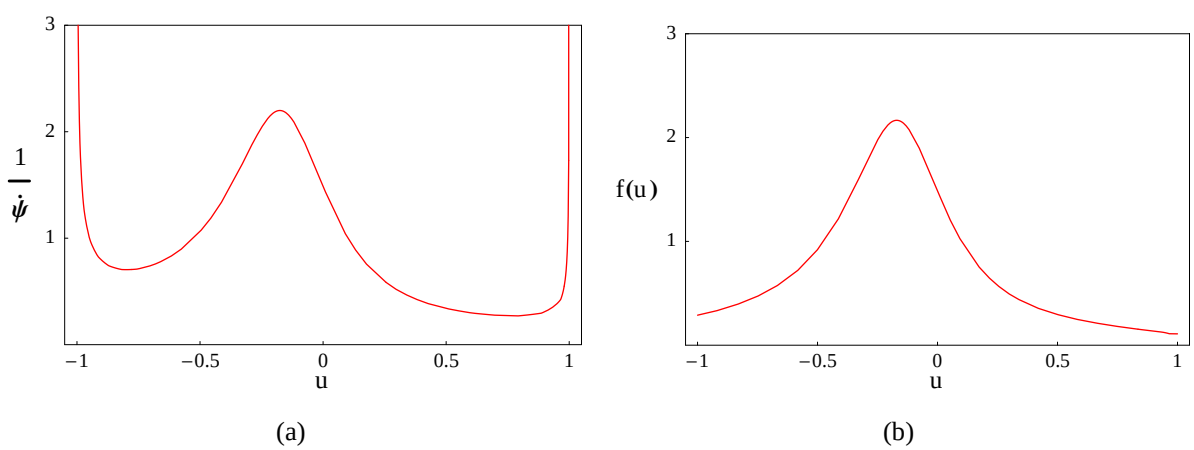

Figure A.18: (a)$1/\dot{\psi}$ along a particle's orbit.(b) $f(u) = \sqrt{1-u^2}/\dot{\psi}$ along a particle's orbit. Both are in arbitrary units.

There is one problem in using Gaussian quadrature, however. For our problem it is necessary to do adaptive integration because there are many different orbit types, which converge to the answer at very different rates. To set the number of points high enough so that all integrals were done to within 1% would be expensive computationally. An adaptive iteration scheme allows one to get around this problem, by increasing the number of points until the answer stops changing within the desired tolerance. In using a Gaussian quadrature, however, the freedom to pick the abscissas where desired is forfeited. Furthermore, the abscissas for the n point integration, do not include, as a subset, the abscissas for the m point integration for any value of $m < n$. Thus the problem with the Gaussian quadrature approach is that when one wants to increase the number of points, the previous function calls must be thrown out.

There do exist, however so called extended Gaussian quadratures rules for certain weights and certain n's which allow the overlap of the old abscissas with the new ones. Such an extension exists for the Gauss-Chebyshev quadrature [MONEGATO 1982]. It is

$$\int_{-1}^{1} \frac{f(u)}{\sqrt{1-u^2}} du \approx \frac{\pi}{2n} \left(\frac{1}{2} f(-1) + \sum_{k=1}^{2n-1} f\left(\cos\left(\frac{\pi k}{2n}\right)\right) + \frac{1}{2} f(1) \right) \tag{A.41}$$

This is the method which is used in the simulation to carry out all of the orbit averages. Since f cannot be evaluated exactly at the end points $(0/0)$ we take it very close (we use $u = \pm .999$ which works fine). In Fig. A.19, A.20 we show a comparison between the (COM) calculation of ω_b and the value determined by ORBIT.

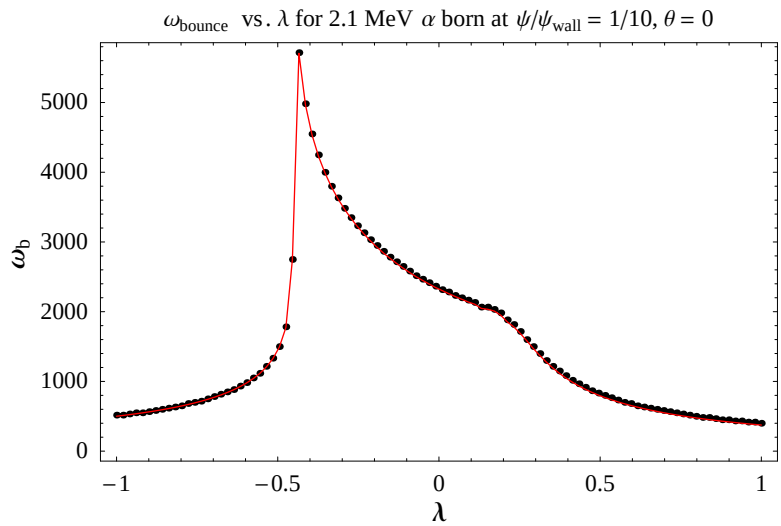

Figure A.19: ω_b vs. λ at high energy. The points are from ORBIT, the solid line from the COM simulation.

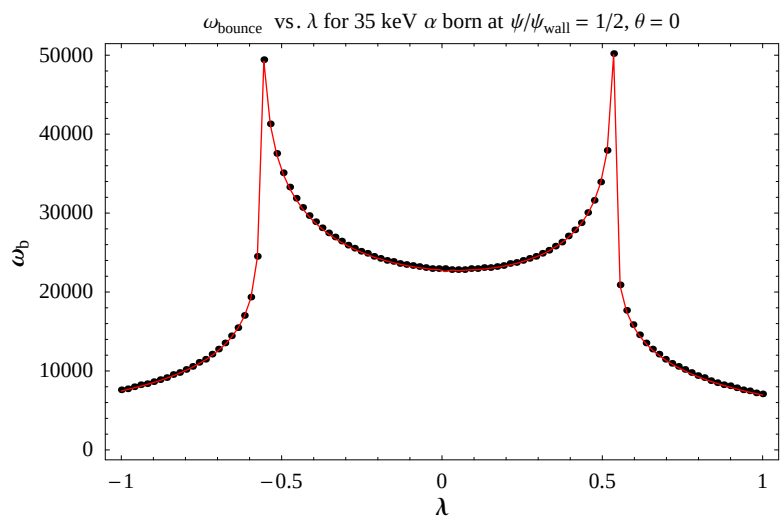

Figure A.20: ω_b vs. λ at low energy. The points are from ORBIT, the solid line from the COM simulation.

Bibliography

BATHKE, C. G., K. A. WERLEY, R. L. MILLER, R. A. KRAKOWSKI, and J. F. SANTARIUS (1992). The ARIES-III D-He^3 tokamak reactor: Design-point determination and parametric studies. In *Proc. IEEE 14th Symp. on Fusion Eng., San Diego, CA (Oct. 1991)*, Piscataway, NJ, pp. 219. IEEE.

BERK, H. L., B. N. BREIZMAN, J. FITZPATRICK, M. S. PEKKER, and H. V. WONG (1996). Nonlinear response of driven systems in weak turbulence theory. *Phys. Plasmas* **3**, 1827.

BERK, H. L., B. N. BREIZMAN, J. FITZPATRICK, and H. V. WONG (1995). Line-broadened quasilinear burst model. *Nucl. Fusion* **35**, 1661.

BERK, H. L., B. N. BREIZMAN, and H. YE (1992). Scenarios for the nonlinear evolution of α-particle-induced Alfvén wave instability. *Phys. Rev. Lett.* **68**, 3563.

BETTI, R. and J. P. FREIDBERG (1992). Stability of Alfvén gap modes in burning plasmas. *Phys. Fluids B* **4**, 1465.

BIGLARI, H., F. ZONCA, and L. CHEN (1992). On resonant destabilization of toroidal Alfvén eigenmodes by circulating and trapped energetic ions/α particles in tokamaks. *Phys. Fluids B* **4**, 2385.

BOIVIN, R. L., S. J. ZWEBEN, and R. B. WHITE (1993). Study of stochastic toroidal field ripple losses of charged fusion products at the midplane of TFTR. *Nucl. Fusion* **33**, 449.

BONOLI, P. T., P. O'SHEA, M. BRAMBILLA, S. N. GOLOVATO, A. E. HUBBARD, M. PORKOLAB, Y. TAKASE, R. L. BOIVIN, F. BOMBARDA, C. CHRISTENSEN, C. L. FIORE, D. GARNIER, J. GOETZ, R. GRANETZ, M. GREENWALD, S. F. HORNE, I. H. HUTCHINSON, J. IRBY, D. JABLONSKI, B. LABOMBARD, B. LIPSCHULTZ, E. MARMAR, M. MAY, A. MAZURENKO, G. MCCRACKEN,

R. NACHTRIEB, A. NIEMCZEWSKI, H. OHKAWA, D. A. PAPPAS, J. REARDON, J. RICE, C. ROST, J. SCHACHTER, J. A. SNIPES, P. STEK, K. TAKASE, J. TERRY, Y. WANG, R. L. WATTERSON, B. WELCH, and S. M. WOLFE (1997). Electron heating via mode converted ion Bernstein waves in the Alcator C-Mod tokamak. *Phys. Plasmas* **4**, 1774.

CARLSSON, J., L.-G. ERIKSSON, and T. HELLSTEN (1994). FIDO a code for calculating the velocity distribution function of a toroidal plasma during ICRH. In E. Sindoni, F. Troyon, and J.Vaclavik (Eds.), *Theory of Fusion Plasmas*, pp. 351.

CARLSSON, J., L.-G. ERIKSSON, and T. HELLSTEN (1997). The effect of the toroidal wave-number spectrum on the high-energy tail formation during ICRH. *Nucl. Fusion* **37**, 719.

CARLSSON, J., T. HELLSTEN, and L.-G. ERIKSSON (1996). FIDO, a code for computing the resonant-ion distribution function during ICRH. Technical Report TRITA-ALF-1996-04, Alfvén Laboratory.

CHEN, K. R. (1994). Anomalous thermalization of fast ions in magnetized plasma. *Phys. Rev. Lett.* **72**, 3534.

CHENG, C. Z. and M. S. CHANCE (1986). Low-n shear Alfvén spectra in axisymmetric toroidal plasmas. *Phys. Fluids* **29**, 3695.

CHENG, C. Z., L. CHEN, and M. S. CHANCE (1985). High-n ideal and resistive shear Alfvén waves in tokamaks. *Ann. Phys. NY* **161**, 21.

CLARKE, J. F. (1980). Hot-ion-mode ignition in a tokamak reactor. *Nucl. Fusion* **20**, 563.

CONN, R., J. HOLDREN, S. SHARAFAT, D. STEINER, D. EHST, W. J. HOGAN, R. KRAKOWSKI, R. L. MILLER, F. NAJMABADI, and K. SCHULTZ (1990). Economic, safety and environmental prospects of fusion reactors. *Nucl. Fusion* **30**, 1919.

DARROW, D. (1996). Private Communication.

DARROW, D. S., H. W. HERRMANN, D. W. JOHNSON, R. J. MARSALA, R. W. PALLADINO, and S. J. ZWEBEN (1995). Measurement of loss of DT fusion products using scintillator detectors in TFTR. *Rev. Sci. Instrum.* **66**, 476.

DARROW, D. S., R. MAJESKI, H. H. DUONG, N. J. FISCH, R. F. HEETER, H. W. HERRMANN, M. C. HERRMANN, S. S. MEDLEY, M. PETROV, J. ROGERS, G. SCHILLING, J. R. WILSON, and S. J. ZWEBEN (1997). Fast ion losses arising

from mode-converted ion Bernstein waves in TFTR. In *5th IAEA TCM on Alpha Particles in Fusion Research*, JET, Abingdon, UK.

DARROW, D. S., R. MAJESKI, N. J. FISCH, R. F. HEETER, H. W. HERRMANN, M. C. HERRMANN, M. C. ZARNSTORFF, and S. J. ZWEBEN (1996). Enhanced loss of fast ions during mode conversion ion Bernstein wave heating in TFTR. *Nucl. Fusion* **36**, 509.

DARROW, D. S., S. J. ZWEBEN, S. BATHA, R. V. BUDNY, C. E. BUSH, Z. CHANG, C. Z. CHENG, H. H. DUONG, J. FANG, N. J. FISCH, E. D. F. R. FISCHER, G. Y. FU, R. F. HEETER, W. W. HEIDBRINK, H. W. HERRMANN, M. C. HERRMANN, K. HILL, E. F. JAEGER, R. JAMES, R. MAJESKI, S. S. MEDLEY, M. MURAKAMI, M. PETROV, C. K. PHILLIPS, M. H. REDI, E. RUSKOV, D. A. SPONG, E. J. STRAIT, G. TAYLOR, J. R. W. R. B. WHITE, K.-L. WONG, and M. C. ZARNSTORFF (1996). Alpha particle losses from Tokamak Fusion Test Reactor deuterium tritium plasmas. *Phys. Plasmas* **3**, 1875.

DELUCIA, J., S. C. JARDIN, and A. M. M. TODD (1980). An iterative metric method for solving the inverse tokamak equilibrium problem. *J. Comp. Phys.* **37**, 183.

DORLAND, W., M. KOTSCHENREUTHER, M. A. BEER, G. W. HAMMETT, R. E. WALTZ, R. R. DOMINGUEZ, P. M. VALANJU, W. H. MINER, J. Q. DONG, W. HORTON, F. L. WAELBROECK, T. TAJIMA, and M. J. LEBRUN (1994). Comparisons of nonlinear toroidal turbulence simulations with experiment. In *Plasma Physics and Controlled Nuclear Fusion Research, 1994*, Volume III, Vienna, pp. 463. International Atomic Energy Agency. paper IAEA-CN-60/D-P-I-6.

EMMERT, G. A., L. A. ELGUEBALY, G. L. KULCINSKI, J. F. SANTARIUS, I. N. SVIATOSLAVSKY, and D. M. MEADE (1994). Improvement in fusion-reactor performance due to ion channeling. *Fusion Technology* **26**, 1158.

ERIKSSON, L.-G. and P. HELANDER (1994). Monte Carlo operators for orbit-averaged Fokker-Planck equations. *Phys. Plasmas* **1**, 308.

FASOLI, A., D. BORBA, G. BOSIA, D. J. CAMPBELL, J. A. DOBBING, C. GORMEZANO, J. JACQUINOT, P. LAVANCHY, J. B. LISTER, P. MARMILLOD, J.-M. MORET, A. SANTAGIUSTINA, and S. SHARAPOV (1995). Direct measurement of the damping of toroidicity-induced Alfvén eigenmodes. *Phys. Rev. Lett.* **75**, 645.

FASOLI, A., J. A. DOBBING, C. GORMEZANO, J. JACQUINOT, J. B. LISTER, S. E.

SHARAPOV, and A. SIBLEY (1996). Alfvén eigenmode excitation by ICRH beat-waves. *Nucl. Fusion* **36**, 258.

FASOLI, A., J. B. LISTER, S. SHARAPOV, D. BORBA, N. DELIYANAKIS, C. GORMEZANO, J. JACQUINOT, A. JAUN, H. A. HOLTIES, G. T. A. HUYSMANS, W. KERNER, J.-M. MORET, and L. VILLARD (1996). Observation of multiple kinetic Alfvén eigenmodes. *Phys. Rev. Lett.* **76**, 1067.

FISCH, N. J. (1987). Theory of current drive in plasmas. *Rev. Mod. Phys.* **59**, 175.

FISCH, N. J. (1997). Uses of intense microwaves in tokamaks. In R. A. Cairns and A. D. R. Phelps (Eds.), *Generation and application of high power microwaves*, pp. 219. Scottish Universities Summer School in Physics: Institute of Physics Publishing.

FISCH, N. J., A. FRUCHTMAN, C. F. F. KARNEY, M. C. HERRMANN, and E. J. VALEO (1995). α channelling using ion Bernstein waves. *Phys. Plasmas* **2**, 2375.

FISCH, N. J. and M. C. HERRMANN (1994). Utility of extracting α particle energy by waves. *Nucl. Fusion* **34**, 1541.

FISCH, N. J. and M. C. HERRMANN (1995). Channelling of α particle power with two waves. *Nucl. Fusion* **35**, 1753.

FISCH, N. J., M. C. HERRMANN, D. DARROW, H. P. FURTH, R. HEETER, H. HERRMANN, J. C. HOSEA, R. MAJESKI, C. K. PHILLIPS, J. ROGERS, G. SCHILLING, and S. J. ZWEBEN (1996). Prospects for α channelling: Initial results from TFTR. In *Fusion Energy 1996: Proceedings of the Sixteenth International Conference on Fusion Energy*, Volume 1, International Atomic Energy Agency, Vienna, pp. 271.

FISCH, N. J. and J.-M. RAX (1992a). Interaction of energetic α-particles with intense lower hybrid waves. *Phys. Rev. Lett.* **69**, 612.

FISCH, N. J. and J.-M. RAX (1992b). Interaction of energetic α-particles with intense lower hybrid waves. In *International Conference on Plasma Physics and Controlled Nuclear Fusion Research*, Volume I, International Atomic Energy Agency, Vienna, pp. 769.

GALAMBOS, J., L. PERKINS, S. HANEY, and J. MANDREKAS (1995). Commercial tokamak reactor potential with advanced tokamak operation. *Nucl. Fusion* **35**, 551.

GIBSON, A. and the JET Team (1997). D-T plasmas in JET: Behaviour and implications. *Bull. Am. Soc.* **42**, 2063.

GOLDSTON, R. J., R. B. WHITE, and A. H. BOOZER (1981). Confinement of high-energy trapped particles in tokamaks. *Phys. Rev. Lett.* **47**, 647.

HAMMETT, G. W. (1986). *Fast ion studies of ion cyclotron heating in the PLT tokamak*. Ph. D. thesis, Princeton University.

HEETER, R. F. (1998). *Studies of Alfvén Eigenmodes and Ion Bernstein Waves in the Tokamak Fusion Test Reactor and the Joint European Torus, and Implications for Alpha Particle Control*. Ph. D. thesis, Princeton University. In preparation.

HEETER, R. F., N. J. FISCH, D. S. DARROW, M. C. HERRMANN, and R. MAJESKI (1997). Modelling of ion Bernstein wave-driven deuterium beam ion losses in TFTR. In *5th IAEA TCM on Alpha Particles in Fusion Research*, JET, Abingdon, UK, pp. 173.

HERRMANN, H. (1997). *Anomalous Loss of DT Alpha Particles in the Tokamak Fusion Test Reactor*. Ph. D. thesis, Princeton University.

HERRMANN, M., N. FISCH, and P. SNYDER (1994). Reactor implications of diverting α-power to fuel ions. *Bull. Am. Soc.* **39**, 1759.

HERRMANN, M. C. and N. J. FISCH (1997). Cooling energetic α-particles in a tokamak with waves. *Phys. Rev. Lett.* **79**, 1495.

HILL, K. W., S. D. SCOTT, M. BELL, et al. (1998). Tests of local transport theory and reduced wall impurity influx with highly radiative plasmas in the Tokamak Fusion Test Reactor. To be published.

HO, D. D.-M. and R. M. KULSRUD (1986). α-particle confinement and helium ash accumulation in stellarator reactors. *Plasma Phys. Controlled Fusion* **28**, 781.

HOLDREN, J. P. (1997). Federal energy research and development for the challenges of the twenty first century. Technical report, Executive Office of the President of the United States, President's Committee of Advisors on Science and Technology, Energy Research and Development Panel.

HSU, C. T. and D. J. SIGMAR (1992). α-particle losses from toroidicity-induced Alfvén eigenmodes. part I: Phase-space topology of energetic particle orbits in tokamak plasma. *Phys. Fluids B* **4**, 1492.

HSU, J. Y., V. S. CHAN, R. W. HARVEY, R. PRATER, and S. K. WONG (1984). Resonance localization and poloidal electric field due to cyclotron wave heating in tokamak plasmas. *Phys. Rev. Lett.* **53**, 564.

JAUN, A., T. HELLSTEN, and S. CHIU (1998). Toroidal mode conversion in the ICRF.

Nucl. Fusion **38**, 153.

KAUFMAN, A. N. (1972). Quasilinear diffusion of an axisymmetric toroidal plasma. *Phys. Fluids* **15**, 1063.

LEVINTON, F. M., M. C. ZARNSTORFF, S. H. BATHA, M. BELL, R. E. BELL, R. V. BUDNY, C. BUSH, Z. CHANG, E. FREDRICKSON, A. JANOS, J. MANICKAM, A. RAMSEY, S. A. SABBAGH, G. L. SCHMIDT, E. J. SYNAKOWSKI, and G. TAYLOR (1995). Improved confinement with reversed magnetic shear in TFTR. *Phys. Rev. Lett.* **75**, 4417.

LITTLEJOHN, R. G. (1981). Hamiltonian formulation of guiding center motion. *Phys. Fluids* **24**, 1730.

LITTLEJOHN, R. G. (1983). Variational principles of guiding centre motion. *J. Plasma Phys.* **29**, 111.

MAHLMAN, J. (1997). Uncertainties in projections of human-caused climate warming. *Science* **278**, 1416.

MAJESKI, R. (1998). Private Communication.

MAJESKI, R., C. K. PHILLIPS, and J. R. WILSON (1994). Electron heating and current drive by mode converted slow waves. *Phys. Rev. Lett.* **73**, 204.

MAJESKI, R., J. H. ROGERS, S. H. BATHA, A. BERS, R. BUDNY, D. DARROW, H. H. DUONG, R. K. FISHER, C. B. FOREST, E. FREDRICKSON, B. GREK, K. HILL, J. C. HOSEA, D. IGNAT, B. LEBLANC, F. LEVINTON, S. S. MEDLEY, M. MURAKAMI, M. P. PETROV, C. K. PHILLIPS, A. RAM, A. T. RAMSEY, G. SCHILLING, G. TAYLOR, J. R. WILSON, and M. C. ZARNSTORFF (1996). Ion cyclotron range of frequency experiments in the Tokamak Fusion Test Reactor with fast waves and mode converted ion Bernstein waves. *Phys. Plasmas* **3**, 2006.

MAJESKI, R., J. H. ROGERS, S. H. BATHA, R. BUDNY, E. FREDRICKSON, B. GREK, K. HILL, J. C. HOSEA, B. LEBLANC, F. LEVINTON, M. MURAKAMI, C. K. PHILLIPS, A. T. RAMSEY, G. SCHILLING, G. TAYLOR, J. R. WILSON, and M. C. ZARNSTORFF (1996). Mode conversion heating and current drive experiments in TFTR. *Phys. Rev. Lett.* **76**, 764.

MCCUNE, D. (1997). Private Communication.

MCGUIRE, K. et al. (1995). Review of deuterium-tritium results from the Tokamak Fusion Test Reactor. *Phys. Plasmas* **2**, 2176.

MCNALLY, J. R. (1982). Physics of fusion fuel cycles. *Nuclear Technology/Fusion* **2**,

9.

MIKKELSEN, D., R. B. WHITE, R. AKERS, S. KAYE, D. MCCUNE, and J. MENARD (1997). Energetic particle orbits in the National Spherical Tokamak Experiment. *Phys. Plasmas* **4**, 3667.

MONEGATO, G. (1982). Stieltjes polynomials and related quadrature rules. *SIAM Rev.* **24**, 137.

NAJMABADI, F., C. G. BATHKE, M. C. BILLONE, J. P. BLANCHARD, L. BROMBERG, E. CHIN, F. R. COLE, J. A. CROWELL, D. A. EHST, L. A. EL-GUEBALY, J. S. HERRING, T. Q. HUA, S. C. JARDIN, C. E. KESSEL, H. KHATER, V. D. LEE, S. MALANG, T.-K. MAU, R. L. MILLER, E. A. MOGAHED, T. W. PETRIE, E. E. REIS, J. SCHULTZ, M. SIDOROV, D. STEINER, I. N. SVIATOSLAVSKY, D.-K. SZE, R. THAYER, M. S. TILLACK, P. TITUS, L. M. WAGNER, X. WANG, and C. P. C. WONG (1997). Overview of the ARIES-RS reversed shear tokamak power plant study. *Fusion Eng. Design* **38**, 3.

NAZIKIAN, R., G. Y. FU, S. H. BATHA, M. G. BELL, R. E. BELL, R. V. BUDNY, C. E. BUSH, Z. CHANG, Y. CHEN, C. Z. CHENG, D. S. DARROW, P. C. EFTHIMION, E. D. FREDRICKSON, N. N. GORELENKOV, B. LEBLANC, F. M. LEVINTON, R. MAJESKI, E. MAZZUCATO, S. S. MEDLEY, H. K. PARK, M. P. PETROV, D. A. SPONG, J. D. STRACHAN, E. J. SYNAKOWSKI, G. TAYLOR, S. VON GOELER, R. B. WHITE, K. L. WONG, and S. J. ZWEBEN (1997). Alpha-particle-driven toroidal Alfvén eigenmodes in the Tokamak Fusion Test Reactor. *Phys. Rev. Lett.* **78**, 2976.

ONO, M., R. BELL, S. BERNABEI, G. GETTELFINGER, R. HATCHER, R. KAITA, S. M. KAYE, H. KUGEL, B. LEBLANC, J. MANICKAM, M. OKABAYASHI, S. F. PAUL, P. RONEY, N. SAUTHOFF, S. SESNIC, H. TAKAHASHI, W. TIGHE, J. R. TIMBERLAKE, S. VON GOELER, J. DUNLAP, A. C. ENGLAND, J. H. HARRIS, R. ISLER, A. POST-ZWICKER, L. BLUSH, R. W. CONN, R. DOERNER, Y. HIROOKA, L. SCHMITZ, G. R. TYNAN, E. DE LA LUNA, I. GARCIA, J. SANCHEZ, S. H. BATHA, F. M. LEVINTON, S. LUCKHARDT, F. PAOLETTI, P. WOSKOV, A. ZOLFAGHARI, T. MUTOH, T. SEKI, N. ASAKURA, A. CARDINALI, R. CESARIO, G. PETRAVICH, and M. TALVARD (1994). Formation of core transport barrier and CH-mode by ion Bernstein wave heating in PBX-M. In *Plasma Physics and Controlled Nuclear Fusion Research, 1994*, Volume 1, International Atomic

Energy Agency, Vienna, pp. 469.

PARK, H., P. S. LEE, W. A. PEEBLES, and N. C. LUHMANN (1985). Tokamak ion temperature determination via cw far-infrared laser scattering from externally excited ion Bernstein waves. *Nucl. Fusion* **25**, 1399.

PERKINS, L. J., J. H. HAMMER, and R. P. DRAKE (1995). The prospects for alternative fusion concepts. Presented at the Advanced Approaches for Economical Fusion Power Workshop, Monterey, California, Sept.11-14, 1995.

POST, D. E., R. V. JENSEN, C. B. TARTER, W. H. GRASBERGER, and W. A. LOKKE (1977). Steady-state radiative cooling rates for low-density high-temperature plasmas. *Atomic Data and Nuclear Data Tables* **20**, 397.

PRESS, W. H., S. A. TEUKOLSKY, W. T. VETTERLING, and B. P. FLANNERY (1995a). *Numerical Recipes in FORTRAN: the Art of Scientific Computing* (Second ed.)., Chapter 9, pp. 340–376. In Press, Teukolsky, Vetterling, and Flannery (1995c).

PRESS, W. H., S. A. TEUKOLSKY, W. T. VETTERLING, and B. P. FLANNERY (1995b). *Numerical Recipes in FORTRAN: the Art of Scientific Computing* (Second ed.)., Chapter 4, pp. 140–155. In Press, Teukolsky, Vetterling, and Flannery (1995c).

PRESS, W. H., S. A. TEUKOLSKY, W. T. VETTERLING, and B. P. FLANNERY (1995c). *Numerical Recipes in FORTRAN: the Art of Scientific Computing* (Second ed.). Cambridge University Press.

REIMAN, A. (1997). Private Communication.

RISKEN, H. (1989a). *The Fokker-Planck Equation: Methods of Solution and Applications* (Second ed.)., Chapter 3, pp. 32–62. In Risken (1989c).

RISKEN, H. (1989b). *The Fokker-Planck Equation: Methods of Solution and Applications* (Second ed.)., Chapter 5, pp. 96–125. In Risken (1989c).

RISKEN, H. (1989c). *The Fokker-Planck Equation: Methods of Solution and Applications* (Second ed.). New York: Springer-Verlag.

ROGERS, J., J. HOSEA, B. LEBLANC, R.MAJESKI, C. PHILLIPS, G.SCHILLING, J. WILSON, et al. (1996). ICRF heating and current drive experiments on TFTR. In *Fusion Energy 1996: Proceedings of the Sixteenth International Conference on Fusion Energy*, Volume 3, International Atomic Energy Agency, Vienna, pp. 317.

ROME, J. A. and Y.-K. M. PENG (1979). The topology of tokamak orbits. *Nucl. Fu-*

sion **19**, 1193.

ROSENBLUTH, M. N. (1972). Superadiabicity in mirror machines. *Phys. Rev. Lett.* **29**, 408.

SCHULTZ, J., L. BROMBERG, and D. COHN (1980). Ignition and thermal stability characteristics of advanced-fuel tokamak plasmas with empirical scaling. *Nucl. Fusion* **20**, 703.

SCOTT, S., C. W. BARNES, D. ERNST, J. SCHIVELL, et al. (1994). Parametric variations of ion transport in TFTR. In W. Horton, M. Wakatani, and A. Wootton (Eds.), *Ion Temperature Gradient Driven Turbulent Transport*. New York: American Institute of Physics.

SIGMAR, D. J. (1979). Velocity space instabilities of alpha particles in tokamak reactors. In *Proc. Physics of Plasmas Close to Thermonuclear Conditions*, Varenna, Italy, pp. 271. Commission of the European Communities.

SNYDER, P. B., M. C. HERRMANN, and N. J. FISCH (1994). Optimization of non-thermal fusion power consistent with channelling of charged fusion product energy. *J. Fusion Energy* **13**, 281.

STIX, T. H. (1992). *Waves in Plasmas*. New York: American Institute of Physics.

SUTTON, W. R., D. J. SIGMAR, and G. H. MILEY (1985). Alpha-driven fast magnetosonic wave heating in tokamak plasmas. *Fusion Technology* **7**, 374.

SYNAKOWSKI, E. J., R. E. BELL, R. V. BUDNY, C. E. BUSH, P. C. EFTHIMION, et al. (1995). Measurements of the production and transport of helium ash in the TFTR tokamak. *Phys. Rev. Lett.* **75**, 3289.

The ARIES Team (1996). ARIES-RS final system code results. `http://aries.ucsd.edu/PUBLIC/ARIES-RS/RS6/output.html#plasmaes`.

VALEO, E. J. and N. J. FISCH (1994). Excitation of large-k_θ ion-Bernstein waves in tokamaks. *Phys. Rev. Lett.* **73**, 3536.

VALEO, E. J. and N. J. FISCH (1995). Private Communication.

WHITE, R. B. (1989). *Theory of Tokamak Plasmas*. North Holland.

WHITE, R. B. (1998). Chaos in trapped particle orbits. *Phys. Rev. E*. To be published.

WHITE, R. B. and A. H. BOOZER (1995). Rapid guiding center calculations. *Phys. Plasmas* **2**, 2915.

WHITE, R. B., A. H. BOOZER, and R. HAY (1982). Drift Hamiltonian in magnetic coordinates. *Phys. Fluids* **25**, 575.

WHITE, R. B. and M. S. CHANCE (1984). Hamiltonian guiding center drift orbit calculation for plasmas of arbitrary cross section. *Phys. Fluids* **27**, 2455.

WHITE, R. B., R. J. GOLDSTON, M. H. REDI, and R. V. BUDNY (1996). Ripple-induced energetic particle loss in tokamaks. *Phys. Plasmas* **3**, 3043.

WILSON, J. R., R. E. BELL, S. BERNABEI, K. HILL, J. C. HOSEA, B. LEBLANC, R. MAJESKI, R. NAZIKIAN, M. ONO, C. K. PHILLIPS, G. SCHILLING, S. VON GOELER, C. E. BUSH, and G. R. HANSON (1998). Ion cyclotron range of frequencies heating and flow generation in deuterium-tritium plasmas. *Phys. Plasmas* **5**, 1721.

WOLFRAM, S. (1996). *The Mathematica Book* (Third ed.). Cambridge University Press.

WONG, K. and M. ONO (1984). Effects of ion cylotron harmonic damping on current drive in the lower hybrid frequency range. *Nucl. Fusion* **24**, 615.

WONG, K. L., R. J. FONCK, S. F. PAUL, D. R. ROBERTS, E. D. FREDRICKSON, R. NAZIKIAN, H. K. PARK, M. BELL, N. L. BRETZ, R. BUDNY, S. COHEN, G. W. HAMMETT, F. C. JOBES, D. M. MEADE, S. S. MEDLEY, D. MUELLER, Y. NAGAYAMA, D. K. OWENS, and E. J. SYNAKOWSKI (1991). Excitation of toroidal Alfvén eigenmodes in TFTR. *Phys. Rev. Lett.* **66**, 1874.

WU, Y., R. B. WHITE, Y.CHEN, and M. N. ROSENBLUTH (1995). Nonlinear evolution of the α-particle-driven toroidicity-induced Alfvén eigenmode. *Phys. Plasmas* **2(12)**, 4555.

ZASLAVSKY, G. M. (1985). *Chaos in Dynamic Systems*, Chapter 3, pp. 79–86. Harwood Academic Publishers.

ZWEBEN, S., V. ARUNASALAM, S. BATHA, R. BUDNY, C. BUSH, S. CAUFFMAN, C. CHANG, Z. CHANG, C. CHENG, D. DARROW, R. DENDY, H. DUONG, N. FISCH, E. FREDRICKSON, R. FISHER, R. FONCK, G. FU, V. GOLOBORODKO, N. GORELENKOV, R. HAWRYLUK, R. HEETER, W. HEIDBRINK, H. HERRMANN, M. HERRMANN, D. JOHNSON, J. MACHUZAK, R. MAJESKI, K. MCGUIRE, G. MCKEE, S. MEDLEY, H. MYNICK, R. NAZIKIAN, M. PETROV, M. REDI, S. REZNIK, J. ROGERS, G. SCHILLING, D. SPONG, J. STRACHAN, B. STRATTON, E. SYNAKOWSKI, G. TAYLOR, S. WANG, R. WHITE, K. WONG, and V. YAVORSKI (1997). Alpha-particle physics in the Tokamak Fusion Test Reactor DT experiment. *Plasma Phys. Controlled Fu-*

sion **39**, 275.

ZWEBEN, S. J., R. L. BOIVIN, M. DIESSO, S. HAYES, H. W. HENDEL, H. PARK., and J. D. STRACHAN (1990). Loss of α-like MeV fusion products from TFTR. *Nucl. Fusion* **30**, 1551.

ZWEBEN, S. J., C. E. BUSH, C. S. CHANG, Z. CHANG, D. S. DARROW, E. D. FREDRICKSON, H. W. HERRMANN, H. E. MYNICK, J. SCHIVELL, M. BELL, R. BOIVIN, R. V. BUDNY, C. Z. CHENG, D. ERNST, G. HAMMETT, L. C. JOHNSON, D. MCCUNE, M. MURAKAMI, D. K. OWENS, J. PARK, C.-K. PHILLIPS, M. H. REDI, S. SCOTT, J. D. STRACHAN, G. TAYLOR, M. TUSZEWSKI, R. B. WHITE, J. R. WILSON, , and M. ZARNSTORFF (1994). Anomalous losses of deuterium-deuterium fusion products in the Tokamak Fusion Test Reactor. *Phys. Plasmas* **1**, 1469.