PSFC/RR-11-4 DOE/ET-54512-374

Measurements and Modeling of Lower Hybrid Driven Fast Electrons on Alcator C-Mod

Schmidt, A.E.

May 2011

Plasma Science and Fusion Center
Massachusetts Institute of Technology
Cambridge MA 02139 USA

This work was supported by the U.S. Department of Energy, Grant No. DE-FC02-99ER54512.

Measurements and Modeling of Lower Hybrid Driven Fast Electrons on Alcator C-Mod

by

Andréa E. W. Schmidt

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2011

Author .....
Department of Physics
April 21, 2011

Certified by .....
Miklos Porkolab
Professor, Department of Physics
Thesis Supervisor

Certified by .....
Ronald R. Parker
Professor, Department of Nuclear Science and Engineering
Thesis Supervisor

Accepted by .....
Krishna Rajagopal
Associate Department Head for Education

Measurements and Modeling of Lower Hybrid Driven Fast Electrons on Alcator C-Mod

by

Andréa E. W. Schmidt

Submitted to the Department of Physics
on April 21, 2011, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

A Lower Hybrid Current Drive (LHCD) system has been implemented on Alcator C-Mod with successful coupling to the plasma of up to 1 MW of power. Nearly fully non-inductive current drive has been achieved for several current relaxation times at low but ITER-relevant densities. One major advantageous feature of the C-Mod LH system is its phasing flexibility, allowing it to produce spectra with a wide range of peak parallel refractive index ($n_{||}$). Theory predicts that LH power deposition location is strongly dependent on $n_{||}$, as well as on other parameters such as electron temperature, electron density, and plasma current. Several diagnostics exist on Alcator C-Mod which can measure the effects of LHCD power on the plasma. The primary diagnostic for measuring LH-driven fast electrons is a hard x-ray (HXR) camera that measures fast electron Bremsstrahlung emission. This is a horizontally-viewing diagnostic with 32 spatial chords that span the plasma cross-section. Each chord ends at a cadmium zinc telluride (CZT) detector that detects individual x-ray photons as current pulses. The output from these detectors is shaped and digitized. Post-processing of the raw pulse train allows for flexible time and energy binning. Another diagnostic for detecting fast electrons is the Electron Cyclotron Emission (ECE) diagnostic. This diagnostic is designed to measure electron temperature in a Maxwellian plasma that is optically thick in the second harmonic. However, the LH-driven fast electrons have relativistically downshifted electron cyclotron frequencies and contribute to additional emission at frequencies just below the second harmonic, where the plasma is optically thin. A third useful diagnostic for LHCD operation is the Motional Stark Effect (MSE) diagnostic, which measures the magnetic field pitch angle profile and therefore can be used to infer the current profile inside the plasma. This current profile is the sum of ohmic current, LH current, and bootstrap current and has been observed to change when LH power is applied to an ohmic plasma. GENRAY/CQL3D is a ray-tracing/Fokker-Plank code package that solves iteratively for a self-consistent electron distribution function in the presence of LH waves, given a plasma scenario and LH wave spectrum. This code includes synthetic diagnostics that can be compared to experimental HXR and ECE measurements. Al-

though an MSE synthetic diagnostic does not currently exist in CQL3D, the inferred current profile from MSE can be compared with the current profile output by CQL3D. Modeling has been carried out for multiple plasma scenarios to benchmark the code package and to further our understanding of how to interpret the experimental results. An experiment in which LH power is square-wave modulated on a time scale much faster than the current relaxation time does not significantly alter the poloidal magnetic field inside the plasma and thus allows for realistic modeling and consistent plasma conditions for different $n_{||}$ spectra. Inverted hard x-ray profiles show clear changes in LH-driven fast electron location with differing $n_{||}$. Boxcar binning of hard x-rays during LH power modulation allows for 1 ms time resolution, which is sufficient to resolve the build-up, steady-state, and slowing-down of fast electrons. Ray-tracing/Fokker-Planck modeling in combination with a synthetic hard x-ray diagnostic show quantitative agreement with the x-ray data for high $n_{||}$ cases. The time histories of hollow x-ray profiles have been used to measure off-axis fast electron transport in the outer half of the plasma, which is found to be small on a slowing down time scale. This work is supported by the US DOE awards DE-FC02-99ER54512 and DE-AC02-76CH03073.

Thesis Supervisor: Miklos Porkolab
Title: Professor, Department of Physics

Thesis Supervisor: Ronald R. Parker
Title: Professor, Department of Nuclear Science and Engineering

Acknowledgments

Many people in my life have contributed to my love of science and my successful completion of a Ph.D. My parents and sisters have always been loving and have encouraged me to pursue my own interests. Carmen Edgerly, Susan Weidkamp, John Gibbs, Rich Muller, and Joel Fajans have been instrumental in my interest in science and physics in particular.

The graduate women in physics group, particularly Bonna Newman and Qudsia Ejaz, has been a great source of friendship and support. It was a pleasure to get to know the PSFC women's group, particularly Catherine Fiore, Amanda Hubbard, and Jennifer Ellsworth. I also recieved much support from Kathy Simons, Jennifer Recklet, Sylvia Sirignano, and my many neighbors and friends at Eastgate.

The entire Alcator team deserves thanks. They keep the tokamak and all of the diagnostics running. They all work hard and are willing to collaborate. In particular, I would like to thank the lower hybrid team for keeping the lower hybrid system running: Dave Terry, Dave Gwinn, Pat McGibbon, George MacKay, Dave Johnson, and Atma Kanojia. The scientific feedback of Randy Wilson, John Wright, Bob Harvey, Amanda Hubbard, Matt Reinke, and Craig Petty has been invaluable.

Greg Wallace has been a great friend, is always willing to share and explain science, and has helped me with my diagnostic countless times. Paul Bonoli, the "unofficial" member of my thesis committee, has worked many hours to help me with problems big and small, and read my thesis multiple times.

My advisors, Miklos Porkolab and Ron Parker, provided guidance and support and were very accomodating during my years at MIT. I have enjoyed learning from their expertise and in particular, getting multiple viewpoints on every problem. I would also like to thank my thesis readers, Earl Marmar and Rick Temkin, for their scientific guidance and feedback on my thesis.

Thank you to my sons, Nuri and Zakir, for making me get up early, for playing with me in the evenings, and for providing me with extra motivation to finish my dissertation. And finally, thank you to my husband Sal, who has been a loving and

positive influence on my life, who never lets me give up, and who has been a fantastic father to our children.

Contents

List of Figures

Chapter 1

Introduction

Based on the current and predicted future rate of global electricity consumption [1], it is apparent that new energy technologies are needed to meet growing energy needs in a sustainable way. Commercially available energy production methods currently struggle with cost, energy distribution and storage, limited fuel supply, environmental impact, and scalability.

1.1 Magnetic Confinement Fusion

Fusion energy combines the lack of greenhouse gases and scalability of fission with the added benefits of inherent safety, short-lived radioactive waste, and abundant fuel supply. Like fission, fusion is a nuclear process. While fission involves the breaking apart of large, heavy elements, namely uranium and plutonium, fusion involves the fusing of lighter elements into heavier elements. The 26^{th} element, iron, is the element at the lowest nuclear energy state [2]. In general, elements heavier than iron release energy when they break apart during fission, and elements lighter than iron release energy when they fuse.

Fusion is indeed the energy source powering the sun and other stars. It is also the energy source behind Hydrogen bombs, nuclear weapons many times more powerful than fission weapons. There has been a world-wide effort for several decades to reproduce this type of power generation in a controlled fashion on Earth.

Fusion requires the nuclei of two atoms to come in close proximity with each other. Since atomic nuclei are positively charged, they repel each other quite strongly, which is the essence of why fusion is so challenging. It is only after the nuclei come into close enough proximity that they begin to attract each other and release large amounts of energy (several MeV). The energy required to bring two nuclei close enough to fuse is approximately 10-20 keV and therefore fusion necessitates extremely high temperatures, typically millions of degrees. At these high temperatures, molecules no longer exist and atomic electrons have enough energy to escape from their nuclei, resulting in a sea of electrons and ions. Such a sea of charged particles is called a plasma.

Although there are many different fusion reactions involving different elements and isotopes, the lowest temperature reaction (and therefore the easiest to achieve) is a deuterium-tritium fusion reaction. Deuterium is a naturally-occurring isotope of hydrogen, with a single proton and single neutron, that is found in approximately 1 out of every 3000 water molecules. The amount of deuterium on Earth is essentially infinite in terms of fueling world energy needs.

Tritium is a radioactive isotope of hydrogen with a single proton and two neutrons. Tritium does not occur naturally, but can be "bred" in nuclear reactors by bombarding lithium with neutrons. Estimates of land-based lithium resources indicate that there is enough lithium to meet world energy needs with deuterium-tritium fusion reactions for tens of thousands of years. As currently envisioned, the first generation of commercial fusion reactors will use a deuterium-tritium mix of fuel and that a more advanced generation of fusion reactors will run on the more challenging deuterium-deuterium reaction.

There are two broad approaches to achieving fusion in a controlled fashion on Earth. One method, so-called inertial confinement fusion (ICF) [3], involves the bombardment of a small, cold pellet of fuel by lasers from many directions. The conceptual framework behind ICF is that the energy and momentum from the laser light will heat and compress the pellet enough to fuse faster than the pellet can blow itself apart. This "miniature H-bomb" pellet would release energy in a small

enough increment to be captured and turned into electricity without destroying any equipment and the miniature explosions would be repeated frequently to continuously put power onto the grid. ICF research also has military applications, as it allows scientists to study the dynamics of nuclear weapons without field testing them.

Magnetic confinement fusion (also known as magnetic fusion energy or MFE) is an approach that involves containing a plasma fuel in a magnetic bottle of some sort and then heating it hot enough to fuse [4]. The charged particles that make up the plasma follow magnetic field lines and can therefore be confined in one area using a clever geometry of magnetic field. Given the temperatures involved in fusion, it is not possible to contain a fusible plasma in any physical bottle because there is no material that would remain solid at those temperatures. The Magnetic Fusion Energy Program is purely scientific with no defense applications.

Magnetic fusion energy is being studied in the US, Europe, and Asia, with facilities of various sizes around the globe. In order to achieve a hot enough plasma to fuse, a reactor larger than any previously built will be required. Recently, a group of countries under the auspices of the International Atomic Energy Agency (IAEA) have agreed to collaborate to build a large fusion facility called ITER [5] in Cadarache, France. ITER is currently under construction and experiments are expected to commence in 2019.

1.2 Tokamaks and Current Drive

Many different magnetic geometries for MFE reactors have been proposed and built. One of the oldest and most successful geometry tried thus far is called a tokamak [6]. A tokamak has a comparatively simple magnetic geometry. Like all MFE devices, it is based on the idea that charged particles follow magnetic field lines. Charged particles in a plasma (ions and electrons) undergo cyclotron motion along a magnetic field line and thus are well-confined to that field line. However these particles are free to move along the field line and eventually are lost in magnetic configurations that consist of straight or open field lines. The tokamak takes those straight magnetic

field lines and bends them around so that they meet up with themselves again. This creates a toroidal (donut-like) shape.

However, this shape of magnetic field lines is insufficient on its own to confine the plasma. Due to the curvature of the system, plasma drifts would cause the plasma to be lost to the wall almost immediately in this simple geometry [1]. Another component of the magnetic field must be added in the poloidal direction. The toroidal and poloidal fields combine to make up the magnetic field lines within a tokamak, as shown in Figure 1-1. While the main (toroidal) magnetic field can be created by external current coils, the poloidal magnetic field in an axisymmetric device can only be created by currents that the plasma itself carries [2].

Since plasma is an excellent electrical conductor, toroidal electric fields of only a few tenths of a volt per meter can result in a megaampere (MA) of current. The main method of driving current in all tokamaks built to date is an ohmic transformer. This transformer consists of a solenoid passing through the central hole of the torus. As the current in the solenoid is ramped up slowly, the magnetic flux through the torus changes. This changing magnetic flux induces a loop voltage around the torus, which then drives a current in the plasma. Thus the plasma acts like the secondary winding of a transformer.

While a conventional AC transformer relies on alternating currents for induction, the plasma current must always flow in one direction, rather than change sign as a function of time. This is because the plasma needs a stable magnetic field configuration for confinement. So the solenoid in the center cannot oscillate back and forth; it can only continue to ramp up in current until some physical limit of the solenoid or power supply is reached. Once this physical limit is reached, the plasma eventually loses its current and its confinement.

Because of the inductive current drive used in all present-day tokamaks, there are

[1] So-called Grad B and curvature drifts would cause charge separation, leading to an electric field. The resulting $\mathbf{E} \times \mathbf{B}$ drift would cause the plasma to drift to the wall

[2] In another type of magnetic confinement device, a stellarator, the poloidal field is created entirely by external coils. The lack of a need for a plasma current is the main advantage of a stellarator over a tokamak. However, the coils required to make a stellarator magnetic field are far more complex than those required for a tokamak, which presents construction challenges. The more complex magnetic field of a stellarator is also more difficult to model.

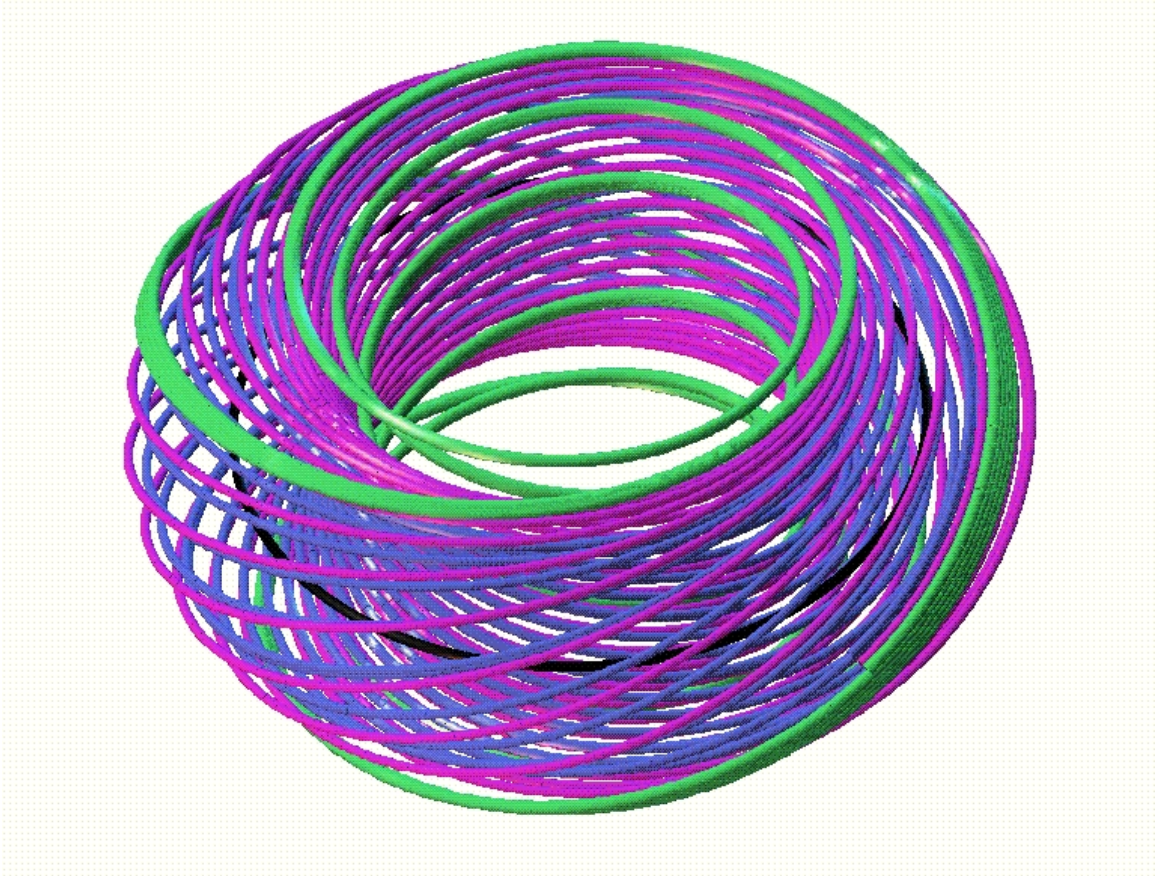

Figure 1-1: Magnetic field lines in a tokamak (toroidal plus poloidal field). Figure courtesy ORNL.

currently no tokamaks that can operate in steady-state, though one tokamak that has since been shut down was able to run in steady-state at low plasma density [7]. An important consequence of running a tokamak in a pulsed mode is that constant mechanical and thermal cycling of the tokamak components would cause them to wear out and need replacement much more frequently than if they were run continuously [8]. Frequent replacement of parts would be very costly and complicate the nuclear waste disposal issue. It is clear that new ways to drive current in steady-state are needed and the development of these new techniques is one of the major problems that needs to be solved before the tokamak design can become part of a commercially viable power source.

Several candidates for non-inductive current drive exist. Neutral beam injection (NBI) [9, 10, 11, 12, 13, 14] is a method of current drive that injects neutral atoms into the plasma in a preferential direction. These atoms can ionize or charge exchange with ions in the system, with the net result that charged ions move preferentially in one direction, creating a current. Electron cyclotron current drive (ECCD) is a method that uses radio frequency (RF) waves at the electron cyclotron resonance frequency in the plasma to create an asymmetry in the electron distribution function which results in a net current [10, 15, 16]. Another RF current drive method, lower hybrid current drive (LHCD), injects RF waves traveling preferentially in one direction. The RF waves impart momentum to the electrons through a wave-particle interaction and create an asymmetric plasma resistivity, which results in a net current [10, 17].

All of the methods of non-inductive current drive have their strengths and weaknesses and in all likelihood, all three may be used in the next generation of tokamaks and eventually in commercial fusion reactors. Lower hybrid current drive has the advantage of having the highest efficiency as well as the ability to drive current near the edge of the plasma for fusion-temperature plasmas. Edge current can be useful in setting up internal transport barriers that can significantly reduce particle and energy losses to the wall [18, 19].

The experiments in this thesis are focused on understanding the physics of lower hybrid current drive. The Alcator C-Mod tokamak at MIT is one of a small handful

of tokamaks in the world that has an LHCD system on it (the only tokamak in the US with LHCD). Lower hybrid current drive experiments on Alcator C-Mod and other present-day tokamaks will guide an LHCD system design and experimental program for the next generation of tokamaks.

1.3 Lower Hybrid Current Drive

The basic principle behind LHCD is that waves are launched into the plasma with a phase velocity a few times that of $v_{te} = \sqrt{2kT_e/m_e}$, the characteristic electron thermal velocity. These waves Landau damp on electrons traveling at or near the phase velocity of the wave, imparting momentum to the electrons. Electrons that are accelerated in the parallel (along B) direction will also pitch angle scatter into the direction perpendicular to the magnetic field. This combination of parallel acceleration and pitch angle scattering results in an asymmetric velocity space distribution function for the electrons, which leads to a net current in the plasma. Since the current-carrying electrons are fast (often relativistic), they thermalize fairly slowly relative to thermal electrons [10]. The resulting distribution function represents a balance between momentum input from the waves, slowing-down from collisions, and pitch-angle scattering from collisions [20].

1.3.1 Cold Plasma Dispersion Relation

To derive the equations that govern lower hybrid physics, we start from the cold plasma dispersion relation for a magnetized plasma. The following derivation of the cold plasma dispersion relation is based on solving the linearized fluid equations for a homogenous plasma in slab geometry. We are assuming that there is an external magnetic field but no external electric field. The plasma is cold, so particle velocity is a perturbed quantity that arises from the oscillating electric field. The various field and plasma quantities can be linearized into 0^{th} order and 1^{st} order quantities, with the order denoted by subscript.

$$B = B_0 + B_1, \tag{1.1}$$

$$E = E_1, \tag{1.2}$$

$$n = n_0 + n_1, \tag{1.3}$$

$$v = v_1. \tag{1.4}$$

Hereafter we shall drop the subscript for perturbed quantities with no 0^{th} order component for simplicity.

We begin with two of Maxwell's Equations:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{1.5}$$

$$\nabla \times \mathbf{B} = \frac{1}{c^2}\frac{\partial \mathbf{E}}{\partial t} + \mu_0 \mathbf{J}. \tag{1.6}$$

Taking the time derivative of Eq. 1.6 and substituting from Eq. 1.5 yields:

$$\nabla \times (\nabla \times \mathbf{E}) = -\frac{1}{c^2}\frac{\partial^2 \mathbf{E}}{\partial t^2} - \mu_0 \dot{\mathbf{J}}. \tag{1.7}$$

We can assume that whatever electromagnetic waves exist in the plasma can be Fourier decomposed in frequency, and thus have an exponential dependence: $e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)}$. Therefore the ∇ operators in Eq. 1.7 can be replaced with $i\mathbf{k}$ and the time derivatives with $-i\omega$:

$$\mathbf{k} \times (\mathbf{k} \times \mathbf{E}) = -\frac{\omega^2}{c^2}\mathbf{E} - i\omega\mu_0 \mathbf{J}. \tag{1.8}$$

We now introduce the vector index of refraction, $\mathbf{n} = \frac{c\mathbf{k}}{\omega}$, and make use of the vector identity $\mathbf{A} \times (\mathbf{B} \times \mathbf{C}) = \mathbf{B}(\mathbf{A} \cdot \mathbf{C}) - \mathbf{C}(\mathbf{A} \cdot \mathbf{B})$ to simplify Eq. 1.8. Bringing all

terms to one side yields:

$$\mathbf{n}(\mathbf{n}\cdot\mathbf{E}) - n^2\mathbf{E} + \mathbf{E} + \frac{i}{\omega\epsilon_0}\mathbf{J} = 0. \tag{1.9}$$

To relate $\mathbf{J}$ and $\mathbf{E}$ in the presence of a magnetic field, we must use the tensor form of ohm's law, $\mathbf{J} = \overline{\overline{\sigma}}\cdot\mathbf{E}$. The elements of $\overline{\overline{\sigma}}$ can be derived by considering the linearized momentum equation (1.10) and the relationship between $\mathbf{J}$ and $\mathbf{v}$ (1.11):

$$\frac{\partial \mathbf{v}}{\partial t} = \frac{q}{m}\mathbf{E} + \mathbf{v}\times\mathbf{B}_0, \tag{1.10}$$

$$\mathbf{J} = \sum_j q_j n_{0j}\mathbf{v}_j. \tag{1.11}$$

In Eq. 1.11, the sum over the index j represents a sum over electron and ion species, including impurities, and q_j is the charge of the species.

As is customary in plasma physics, we can align our coordinate system so that the external magnetic field is in the $\hat{\mathbf{z}}$ direction ($\mathbf{B}_0 = B_0\hat{\mathbf{z}}$). If we replace the time derivative in Eq. 1.10 with $-i\omega$, its three components can then be solved algebraically for v_x, v_y, and v_z. In Eqs. 1.12-1.14, Ω is the species-dependent cyclotron frequency, $\Omega = qB_0/m$ where q is the charge of the species.

$$v_x = \frac{q}{\omega m}\frac{iE_x - (\Omega/\omega)E_y}{1-\Omega^2/\omega^2}, \tag{1.12}$$

$$v_y = \frac{q}{\omega m}\frac{(\Omega/\omega)E_x + iE_y}{1-\Omega^2/\omega^2}, \tag{1.13}$$

$$v_z = \frac{iqE_z}{\omega m}. \tag{1.14}$$

Substituting Eqs. 1.12-1.14 into Eq. 1.11 yields the relation:

$$\begin{bmatrix} J_x \\ J_y \\ J_z \end{bmatrix} = \underbrace{\sum_i \frac{q_j^2 n_{0j}}{\omega m_j} \frac{1}{1-\Omega_j^2/\omega^2} \begin{bmatrix} i & -\Omega_j/\omega & 0 \\ \Omega_j/\omega & i & 0 \\ 0 & 0 & i(1-\Omega_j^2/\omega^2) \end{bmatrix}}_{\overline{\overline{\sigma}}} \begin{bmatrix} E_x \\ E_y \\ E_z \end{bmatrix}. \quad (1.15)$$

We now define the dielectric tensor, $\overline{\overline{K}} = \overline{\overline{I}} + \frac{i}{\omega\epsilon_0}\overline{\overline{\sigma}}$. The coefficients in front of $\overline{\overline{\sigma}}$ can be grouped together to form the plasma frequency, ω_p, so that $\overline{\overline{K}}$ becomes:

$$\overline{\overline{K}} = \overline{\overline{I}} + \sum_j \frac{i\omega_{pj}^2}{\omega^2} \frac{1}{1-\Omega_j^2/\omega^2} \begin{bmatrix} i & -\Omega_j/\omega & 0 \\ \Omega_j/\omega & i & 0 \\ 0 & 0 & i(1-\Omega_j^2/\omega^2) \end{bmatrix}. \quad (1.16)$$

Then Eq. 1.9 becomes:

$$[\mathbf{nn} - n^2 + \overline{\overline{K}}] \cdot \mathbf{E} = 0. \quad (1.17)$$

In order for Eq. 1.17 to have non-trivial solutions of $\mathbf{E}$, it is necessary and sufficient that the determinant of the matrix vanish:

$$det(\mathbf{nn} - n^2 + \overline{\overline{K}}) = 0. \quad (1.18)$$

It is now convenient to introduce the notation used in Stix [21]:

$$S = 1 - \sum_j \frac{\omega_{pj}^2}{\omega^2 - \Omega_j^2}, \quad (1.19)$$

$$D = \sum_j \frac{\Omega_j}{\omega} \frac{\omega_{pj}^2}{\omega^2 - \Omega_j^2}, \quad (1.20)$$

$$P = 1 - \sum_j \frac{\omega_{pj}^2}{\omega^2}. \quad (1.21)$$

We can now rewrite $\overline{\overline{K}}$ as:

$$\overline{\overline{K}} = \begin{bmatrix} S & -iD & 0 \\ iD & S & 0 \\ 0 & 0 & P \end{bmatrix}. \tag{1.22}$$

We have already aligned the z coordinate with the magnetic field. We can now choose our x and y coordinates such that $\mathbf{n}$ lies in the x-z plane. Then Eq. 1.18 becomes:

$$det \begin{bmatrix} S - n_{||}^2 & -iD & n_{||}n_{\perp} \\ iD & S - n_{||}^2 - n_{\perp}^2 & 0 \\ n_{||}n_{\perp} & 0 & P - n_{\perp}^2 \end{bmatrix} = 0. \tag{1.23}$$

In Eq. 1.23, $n_{||}$ is n_z, the component of the index of refraction that is parallel to the magnetic field, and $n_{\perp}$ is n_x, the component of the index of refraction that is perpendicular to the magnetic field. Eqation 1.23 represents the general cold plasma dispersion relationship and describes propagation for all types of electromagnetic waves in homogeneous plasmas, not including finite temperature effects.

1.3.2 Lower Hybrid Wave Dispersion Relation and Accessibility Condition

We now consider the special case where $n_{||}$ is fixed by the geometry of the experiment. This is typically true of lower hybrid current drive experiments on tokamaks, where the spacing of the waveguides in the lower hybrid launcher and the relative phasing between waveguides fixes the launched $n_{||}$ spectrum (it is possible for the parallel wavenumber to evolve as the wave propagates, but we shall ignore that effect for the time being) [22]. It should be noted that any physical experiment will actually launch a spectrum of $n_{||}$ values, but each value of $n_{||}$ can be treated independently and then superimposed.

Evaluating the determinant in Eq. 1.23 yields an equation that is quadratic in $n_{\perp}^2$:

$$an_{\perp}^4 + bn_{\perp}^2 + c = 0, \tag{1.24}$$

$$a = S, \tag{1.25}$$

$$b = -[(S - n_{||}^2)(P + S) - D^2], \tag{1.26}$$

$$c = P[(S - n_{||}^2)^2 - D^2]. \tag{1.27}$$

In simplifying this relationship, it is useful to look at the relative magnitudes of S, D, and P. To do this, let us consider some typical densities and magnetic fields for a tokamak plasma: $n = 2 \times 10^{20}\ m^{-3}$, $B = 5\ T$. Then for a deuterium plasma, the relevant frequencies are:

$$\frac{\Omega_i}{2\pi} = 38\ MHz, \tag{1.28}$$

$$\frac{\Omega_e}{2\pi} = 140\ GHz, \tag{1.29}$$

$$\frac{\omega_{pi}}{2\pi} = 2\ GHz, \tag{1.30}$$

$$\frac{\omega_{pe}}{2\pi} = 127\ GHz. \tag{1.31}$$

Typical lower hybrid system operation frequencies (as we shall see later) are a few GHz, so the relevant frequencies can be ordered as follows:

$$\Omega_i^2 \ll \omega_{pi}^2 \sim \omega^2 \ll \omega_{pe}^2 \sim \Omega_e^2. \tag{1.32}$$

Then S, D, and P can be approximated as follows:

$$S = 1 - \sum_j \frac{\omega_{pj}^2}{\omega^2 - \Omega_j^2} \approx 1 + \frac{\omega_{pe}^2}{\Omega_e^2} - \frac{\omega_{pi}^2}{\omega^2} \sim 1, \tag{1.33}$$

$$D = \sum_j \frac{\Omega_j}{\omega} \frac{\omega_{pj}^2}{\omega^2 - \Omega_j^2} \approx -\frac{\Omega_e}{\omega} \frac{\omega_{pe}^2}{\Omega_e^2} + \frac{\Omega_i}{\omega} \frac{\omega_{pi}^2}{\omega^2} \approx -\frac{\omega_{pe}^2}{\omega \Omega_e} = \frac{\omega_{pi}^2}{\omega \Omega_i} \sim \sqrt{\frac{m_i}{m_e}}, \tag{1.34}$$

$$P = 1 - \sum_j \frac{\omega_{pj}^2}{\omega^2} \approx -\frac{\omega_{pe}^2}{\omega^2} \sim -\frac{m_i}{m_e}. \tag{1.35}$$

Clearly $|S| \ll D \ll |P|$, so a, b, and c can be approximated as:

$$a = S, \tag{1.36}$$

$$b = -[(S - n_{||}^2)(P + S) - D^2] \approx P(n_{||}^2 - S) + D^2, \tag{1.37}$$

$$c = P[(S - n_{||}^2)^2 - D^2] \approx -PD^2. \tag{1.38}$$

Then the two wave roots corresponding to Eq. 1.24 can be written as:

$$n_\perp^2 = \frac{(-b \pm \sqrt{b^2 - 4ac})}{2a} = -\frac{P}{2S}\left[n_{||}^2 - S + \frac{D^2}{P} \pm \sqrt{\left(n_{||}^2 - S + \frac{D^2}{P}\right)^2 + \frac{4SD^2}{P}}\right]. \tag{1.39}$$

The root with the positive sign in Eq. 1.39 is associated with the so-called slow wave, named for its slower phase velocity. The negative root is associated with the fast wave. It is the slow wave that is used for lower hybrid current drive applications.

To derive a simplified slow wave dispersion relation, let us first assume that the two roots are well separated so that $n_\perp^2 \approx -b/a$:

$$n_\perp^2 = -\frac{b}{a} = \frac{(S - n_{||}^2)(P + S) - D^2}{S} \approx \frac{-n_{||}^2 P}{S}, \tag{1.40}$$

based on the ordering established by Eqs. 1.33-1.35.

Replacing $n_\perp^2$ by $n^2 - n_{||}^2$ and rearranging, yields the result:

$$1 = \frac{n_{||}^2}{n^2}\left(1 - \frac{P}{S}\right) \approx -\frac{n_{||}^2}{n^2}\frac{P}{S} = \frac{n_{||}^2}{n^2}\frac{\omega_{pe}^2}{\omega^2 + \omega^2\frac{\omega_{pe}^2}{\Omega_e^2} - \omega_{pi}^2}. \tag{1.41}$$

Solving for ω^2 yields:

$$\omega^2 = \omega_{lh}^2\left(1 + \frac{m_i}{m_e}\frac{n_{||}^2}{n^2}\right), \tag{1.42}$$

where

$$\omega_{lh}^2 \equiv \frac{\omega_{pi}^2}{1 + \frac{\omega_{pe}^2}{\Omega_e^2}} \tag{1.43}$$

is the lower hybrid frequency. This is the appoximate dispersion relation for the slow wave, also known as the lower hybrid wave.

We now note that in order for the wave to propagate freely in the plasma, $n_\perp^2 = (-b \pm \sqrt{b^2 - 4ac})/2a$ must be real and positive. From Eq. 1.35, we know that P is a negative quantity. Therefore $n_\perp^2$ is real and positive if $S > 0$ and $n_{||}^2 - S + \frac{D^2}{P} > 2D\sqrt{\frac{S}{-P}}$. The requirement that $S > 0$ is a constraint on ω:

$$\omega^2 > \frac{\omega_{pi}^2}{1 + \frac{\omega_{pe}^2}{\Omega_e^2}} = \omega_{lh}^2. \tag{1.44}$$

Equation 1.44 states that the frequency chosen for LH operation must be greater than the local lower hybrid frequency everywhere in the plasma.

The second requirement stemming from Eq. 1.39 is a constraint on $n_{||}$:

$$n_{||}^2 > S - \frac{D^2}{P} + 2D\sqrt{\frac{S}{-P}} = \left(\sqrt{S} + \frac{D}{\sqrt{-P}}\right)^2. \tag{1.45}$$

Using the approximations from Eqs. 1.33, 1.34, and 1.35, Eq. 1.45 becomes:

$$|n_{||}| > n_{||,crit} \equiv \sqrt{1 + \frac{\omega_{pe}^2}{\Omega_e^2} - \frac{\omega_{pi}^2}{\omega^2}} + \frac{\omega_{pe}}{|\Omega_e|}. \tag{1.46}$$

Equation 1.46 is known as the accessibility condition, and defines $n_{||,crit}$, the critical value of $n_{||}$ above which the lower hybrid wave is accessible [23, 24]. The implications of the accessibility condition are a bit subtle. Recall that there are two roots of Eq. 1.39, corresponding to the slow (lower hybrid) wave and the fast wave. When $n_{||} = n_{||,crit}$, the discriminant vanishes and the two roots coalesce. If a wave with fixed $n_{||}$ propagates into a plasma of increasing density and reaches a location where $n_{||} = n_{||,crit}$, the wave will reflect and the slow wave will convert to the fast wave (or vice versa) in a process known as mode conversion.

Recall that Eq. 1.46 must be satisfied everywhere in the plasma, taking into account local densities and magnetic fields. Generally speaking, LH waves are more accessible at high magnetic field and low density. This requirement is often most stringent around the center of the plasma, where the densities (and plasma frequencies) are highest. However, in an improved confinement regime known as H-mode [25], in which plasmas have somewhat flat density profiles, accessiblity can be an issue at the plasma edge on the low-field side.

Maintaining LH wave accessibility by increasing the launched $n_{||}$ can be a problem for current drive because as we shall see in the next section, lower $n_{||}$ corresponds to higher current drive efficiency. Thus the lowest practical value of $n_{||}$ is desired to drive current.

1.3.3 Current Drive Efficiency

Sustaining the entire plasma current in a tokamak with non-inductive current drive requires a non-trivial amount of power. The current drive component of any fusion reactor will be a significant fraction of the cost of the reactor as well as a significant fraction of the power required to run the reactor. Thus the amount of current that can be driven per unit of power has serious implications for the commercial viability of the tokamak as a fusion reactor. It also affects the amount of output power that needs to be generated in order for a power plant to produce more power than it consumes. We derive the current drive efficiency here and similar derivations can be found in [10, 15].

We usually think of current drive efficiency, η, in terms of current driven per unit of input power:

$$\eta = J_{||}/S_d. \tag{1.47}$$

In Eq. 1.47, $J_{||}$ is the parallel current density and S_d is the dissipated power density required to sustain that current density. The most simplified physical picture of current drive, which produces the correct scaling, is the single particle picture. Consider a single fast electron, traveling at a parallel velocity $v_{||}$, that gets a small increase in parallel velocity, Δv (we consider an incremental increase in parallel velocity, so that $\Delta v/v_{||} \ll 1$). The energy required to increase the velocity of the electron is equal to the difference in its final and initial kinetic energy:

$$\Delta E = \frac{1}{2}(v_{||} + \Delta v)^2 - \frac{1}{2}v_{||}^2 \approx m v_{||} \Delta v. \tag{1.48}$$

Then the dissipated power density is given by Eq. 1.49, where n_{fe} is the number density of electrons accelerated by interactions with the wave and $\nu(v_{||})$ is the momentum loss rate due to Coulomb collisions.

$$S_d = n_{fe} m v_{||} \Delta v \nu(v_{||}). \tag{1.49}$$

The current driven by the increase in parallel velocity is:

$$J_{||} = n_{fe} e \Delta v. \tag{1.50}$$

Substituting Eqs. 1.49 and 1.50 into Eq. 1.47, yields:

$$\eta = J_{||}/S_d = \frac{e\Delta v}{m v_{||} \Delta v \nu(v_{||})} = \frac{e}{m v_{||} \nu(v_{||})}. \tag{1.51}$$

The momentum loss rate is proportional to bulk electron density, n_e, and inversely proportional to the cube of the velocity [26]:

$$\nu(v_{||}) = \nu_0 \frac{v_{th}^3}{v_{||}^3}, \tag{1.52}$$

where

$$\nu_0 = \frac{\omega_{pe}^4 log\Lambda}{2\pi n_e v_{th}^3}, \tag{1.53}$$

so that the momentum loss rate has the following dependence on $v_{||}$ and n_e:

$$\nu(v_{||}) \propto \frac{n_e}{v_{||}^3}, \tag{1.54}$$

Then the current drive efficiency from Eq. 1.51 has the following scaling:

$$\eta \propto \frac{v_{||}^2}{n_e}. \tag{1.55}$$

This result shows that even though the energy cost of accelerating electrons goes up with electron velocity, the reduced collisionality at high velocity dominates the efficiency. In other words, it is more efficient to drive current with electrons already moving at high velocity because they take a longer time to thermalize than slow electrons do.

While this simple calculation gets the velocity scaling right, a 2-dimensional calculation gives more insight into the physics of current drive. The 2D derivation was first performed by Fisch [15] and highlights the relative contributions of momentum and energy input to current drive efficiency.

We now consider an electron traveling with velocity v_1 which we shall push to v_2, requiring energy $E_2 - E_1 \propto v_2^2 - v_1^2$. At v_2, the electron has reduced collisionality, such that the current produced by this change in velocity is proportional to:

$$J \propto \frac{v_{||2}}{\nu_2} - \frac{v_{||1}}{\nu_1}, \tag{1.56}$$

where ν_1 and ν_2 are the effective collision frequencies at velocities v_1 and v_2 respectively. Ref. [15] shows that $\nu(v)$ is:

$$\nu(v) = \nu_0 \frac{5 + Z_i}{2u^3} \tag{1.57}$$

where

$$\nu_0 = \frac{\omega_p^4 ln\Lambda}{2\pi n_0 v_{te}^3} \tag{1.58}$$

and

$$u = \frac{v}{v_{te}}. \tag{1.59}$$

Then the current drive efficiency is proportional to:

$$\frac{J}{S_d} \propto \frac{\frac{v_{||2}}{\nu_2} - \frac{v_{||1}}{\nu_1}}{v_2^2 - v_1^2}. \tag{1.60}$$

In the case of LHCD, we give the electron a small increase in parallel velocity (from v_1 to v_2). Thus we can put Eq. 1.60 in differential form:

$$\frac{J}{S_d} \propto \frac{\frac{\partial}{\partial v_{||}}\left(\frac{v_{||}}{\nu(v)}\right)}{\frac{\partial}{\partial v_{||}}\left(v^2\right)} \propto \frac{\frac{\partial}{\partial v_{||}}\left(v_{||}v^3\right)}{\frac{\partial}{\partial v_{||}}\left(v^2\right)} = \frac{\frac{\partial}{\partial v_{||}}\left(v_{||}(v_{||}^2 + v_\perp^2)^{3/2}\right)}{\frac{\partial}{\partial v_{||}}\left(v_{||}^2 + v_\perp^2\right)}. \tag{1.61}$$

Evaluating Eq. 1.61 yields the result:

$$\frac{J}{S_d} \propto v_{||}^{-1}(v_{||}^2 + v_\perp^2)^{3/2} + 3v_{||}(v_{||}^2 + v_\perp^2)^{1/2}. \tag{1.62}$$

Note that the first term in Eq. 1.62 would appear whether or not a perpendicular velocity was included. This is the contribution to the current drive efficiency from the momentum input. However, only by including a perpendicular velocity could we recover the second term, which is approximately 3 times larger and originates from the energy input.

The implication of Eq. 1.62 is that much of the current drive efficiency – roughly 3/4 of it – comes from the asymmetric resistivity created by the energy input into the electrons. In fact, it is possible to drive current using only energy input without any toroidal momentum input, as in the case of ECCD.

It is useful to relate the current drive efficiency scaling to the parameter controlled in the experiment, $n_{||}$. Consider an LH wave launched with a parallel refractive index $n_{||}$ and a parallel phase velocity $v_{||} = c/n_{||}$. If the electrons interacting with the wave have a parallel velocity matching the phase velocity of the wave, then it follows from Eq. 1.55 that a lower hybrid current drive efficiency will scale as follows:

$$\eta \propto \frac{1}{n_e n_{||}^2}. \tag{1.63}$$

Clearly higher current drive efficiency is achieved with lower $n_{||}$ waves. However, there is a lower limit on $n_{||}$ imposed by the accessibility condition described in Section 1.3.2.

Now we can relate the local driven current density and dissipated power to total current and total power. For a tokamak with annular cross-section area A and volume $2\pi RA$, the total power dissipated is $P_d = 2\pi RA\langle S_d \rangle$ and the total current driven is $I = A\langle J \rangle$, where $\langle x \rangle$ denotes the volume average of the quantity x. Then:

$$\frac{I}{P} = \frac{\eta}{2\pi R} \propto \frac{1}{R n_e n_{||}^2}. \tag{1.64}$$

Then for a fixed value of $n_{||}$:

$$\frac{n_e I R}{P} \approx const. \tag{1.65}$$

The practical consequence of Eq. 1.65 is that the amount of power needed to drive a given amount of current goes up linearly with both density and major radius. This scaling is universal to all non-inductive current drive schemes. The $1/n_e$ scaling of the efficiency comes from the momentum loss rate, which is also common to all non-inductive current drive schemes. Also a fixed amount of current being carried in a ring requires a quantity of current carriers proportional to the circumference of the ring, which leads to the $1/R$ scaling of efficiency.

These realities of current drive efficiency scaling make it difficult to drive current non-inductively at the high densities needed for fusion, especially on tokamaks of

large major radius, like ITER. The scalings in Eq. 1.65 have been confirmed both experimentally and theoretically [27, 28].

1.3.4 Quasi-Linear Theory

The evolution of the electron distribution function is a competition between RF wave-particle interactions and collisions between fast electrons and the bulk plasma that serve to return the distribution function to a Maxwellian. These two effects are described mathematically in the Fokker-Plank equation [10]:

$$\frac{\partial f}{\partial t} = \frac{\partial}{\partial \mathbf{v}} \mathbf{D}_{QL} \cdot \frac{\partial}{\partial \mathbf{v}} f + \left(\frac{\partial f}{\partial t} \right)_c . \tag{1.66}$$

In Eq. 1.66, $\mathbf{D}_{QL}$ is the quasi-linear diffusion coefficient [29], which describes the diffusion of electrons in velocity space due to the presence of the RF wave electric field. The quantity $(\frac{\partial f}{\partial t})_c$ is the Fokker-Plank collision operator. For lower hybrid waves, the dominant component of $\mathbf{D}_{QL}$ is the $v_{||}v_{||}$ term (the transfer of parallel momentum by the parallel wave electric field). Thus Eq. 1.66 can be simplified as:

$$\frac{\partial f}{\partial t} = \frac{\partial}{\partial v_{||}} D_{QL}(v_{||}) \frac{\partial}{\partial v_{||}} f + \left(\frac{\partial f}{\partial t} \right)_c . \tag{1.67}$$

A simplified derivation of the quasilinear diffusion coefficient [29] can be derived [30] by starting from the 1-D Vlasov equation:

$$\frac{\partial f}{\partial t} + v_{||} \frac{\partial f}{\partial z} - \frac{e}{m} E_{||} \frac{\partial f}{\partial v_{||}} = 0. \tag{1.68}$$

We let the electric field vary in space and time as $e^{i(k_{||}z - \omega t)}$. We now represent the electron distribution function as a spatially constant 0^{th} order component and a perturbed 1^{st} order component that also varies as $e^{i(k_{||}z - \omega t)}$:

$$f(z, v_{||}, t) = f_0(v_{||}, t) + f_1(z, v_{||}, t), \tag{1.69}$$

where $f_0(v_{||}, t)$ is the spatial average of $f(z, v_{||}, t)$, defined as:

$$f_0(v_{||},t) = \lim_{L\to\infty} \frac{1}{L}\int_{-L}^{L} dz\ f(z,v_{||},t). \tag{1.70}$$

Substituting Eq. 1.69 into Eq. 1.68 yields:

$$\frac{\partial f_0}{\partial t} - i\omega f_1 + ik_{||}v_{||}f_1 - \frac{e}{m}E_{||}\frac{\partial f_0}{\partial v_{||}} - \frac{e}{m}E_{||}\frac{\partial f_1}{\partial v_{||}} = 0. \tag{1.71}$$

Taking the spatial average of Eq. 1.71 eliminates several terms:

$$\frac{\partial f_0}{\partial t} = \frac{e}{m}\frac{\partial}{\partial v_{||}}\langle E_{||}f_1\rangle. \tag{1.72}$$

In Eq. 1.72, $\langle x\rangle$ denotes the volume average of the quantity x. Gathering all terms from Eq. 1.71 with a $e^{i(k_{||}z-\omega t)}$ component and solving for the perturbed distribution function to first order yields:

$$f_1(v_{||}) = \frac{\frac{ie}{m}E_{||}\frac{\partial f_0}{\partial v_{||}}}{\omega - kv_{||}}. \tag{1.73}$$

Substituting Eq. 1.73 into Eq. 1.72 gives the evolution of the 0^{th} order distribution function:

$$\frac{\partial f_0}{dt} = \frac{ie^2}{m^2}\frac{\partial}{\partial v_{||}}\left(\frac{\langle E_{||}^2\rangle}{\omega - kv_{||}}\frac{\partial f_0}{\partial v_{||}}\right). \tag{1.74}$$

It is clear from the form of Eq. 1.74 that D_{QL} in Eq. 1.67 is:

$$D_{QL} = \frac{ie^2}{m^2}\frac{\langle E_{||}^2\rangle}{\omega - kv_{||}}. \tag{1.75}$$

Thus changes in the electron distribution function are second order in the perturbed RF wave field amplitude.

1.4 Hard X-Ray Bremsstrahlung Emission

All charged particles emit radiation when they are accelerated by their interactions with other nearby charged particles [31]. Such radiation is called Bremsstrahlung

radiation. For the relativistic fast electrons generated by LHCD, Bremsstrahlung radiation is typically in the hard x-ray range of energies, and originates from the fast electrons' interactions with both the bulk electrons and the bulk ions. The photons emitted from a fast electron can be of any energy up to the fast electron kinetic energy [32]. There is an angular anisotropy in the amount of Bremsstrahlung emission originating from a fast electron, with the dominant emission radiating in the direction of the velocity of the electron [31]. The lack of one-to-one correspondence between fast electron energy and photon energy as well as the angular anisotropy of the radiation leads to inherent ambiguities in the interpretation of fast electron Bremsstrahlung emission. More specifically, it is not possible to uniquely map a spatially- and energy-resolved Bremsstrahlung emission spectrum to an electron distribution function. Extensive measurements of fast electron bremsstrahlung emission during lower hybrid operation, including reconstructions of approximate electron distribution functions, are presented in Refs. [33, 34, 35, 36, 37, 38, 39].

The emission resulting from both the electron-electron and electron-ion interactions can be characterized by bremsstrahlung differential cross-section functions, $\frac{d\sigma_{ee}}{dk\,d\Omega}$ and $\frac{d\sigma_{ei}}{dk\,d\Omega}$ [40]. Each differential cross-section is a function of photon energy, k, electron momentum, $\mathbf{p}$, and the direction of photon emission with respect to the electron momentum, $\hat{\mathbf{k}} \cdot \hat{\mathbf{p}}$. Additionally, the electron-ion Bremsstrahlung cross-section is a function of ion charge, Z. Cross-sections are expressed in area per unit photon energy, per solid angle (Ω).

Electron-electron and electron-ion cross-sections are well known and appear in various forms for different mathematical limits [41]. Sample cross-sections are shown in Figures 1-2 and 1-3. Figure 1-2 shows the Bremsstrahlung cross-sections for 60 and 90 keV photons, for various electron energies. This is useful for thinking about what energy electrons may have contributed to emission at a specific photon energy. In contrast, Figure 1-3 shows the Bremsstrahlung cross-sections for 300 and 500 keV electrons, at various photon energies. This is useful for thinking about what spectrum is emitted from an electron of a single energy.

The cross-section plots show several important trends:

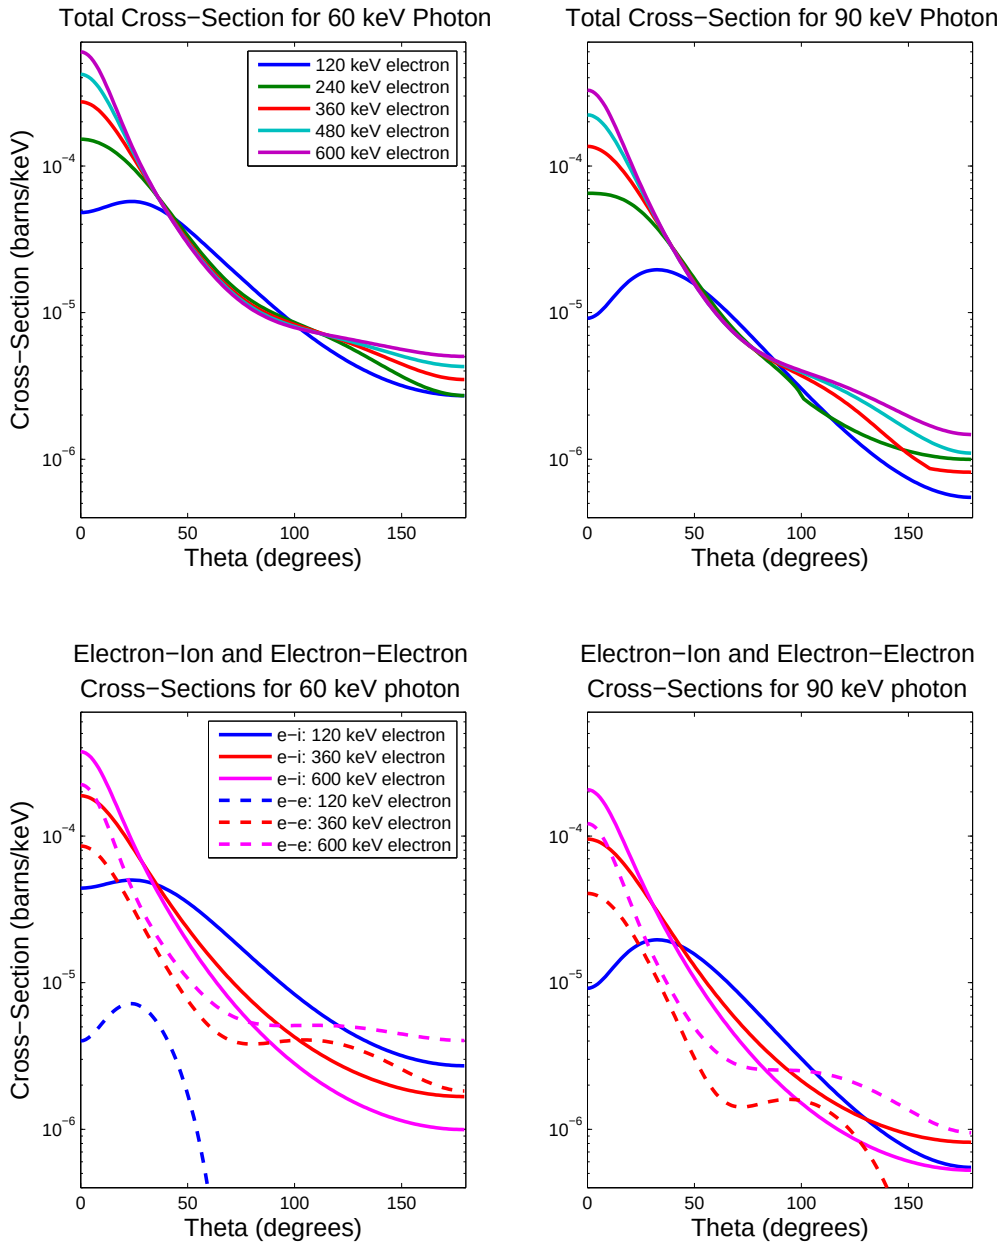

Figure 1-2: Bremsstrahlung cross-sections at fixed photon energies, for various electron energies. Theta is the angle between the fast electron velocity and the viewing chord. Electron-ion cross sections are calculated using the Bethe-Heitler model [42] including the Elwert factor [43]. Electron-electron cross sections are calculated using the Haug model, including the Coulomb correction factor. [44] The 90 keV photon e-e cross-sections for a 120 keV electron are extremely small and do not fall within the bounds of the lower right plot.

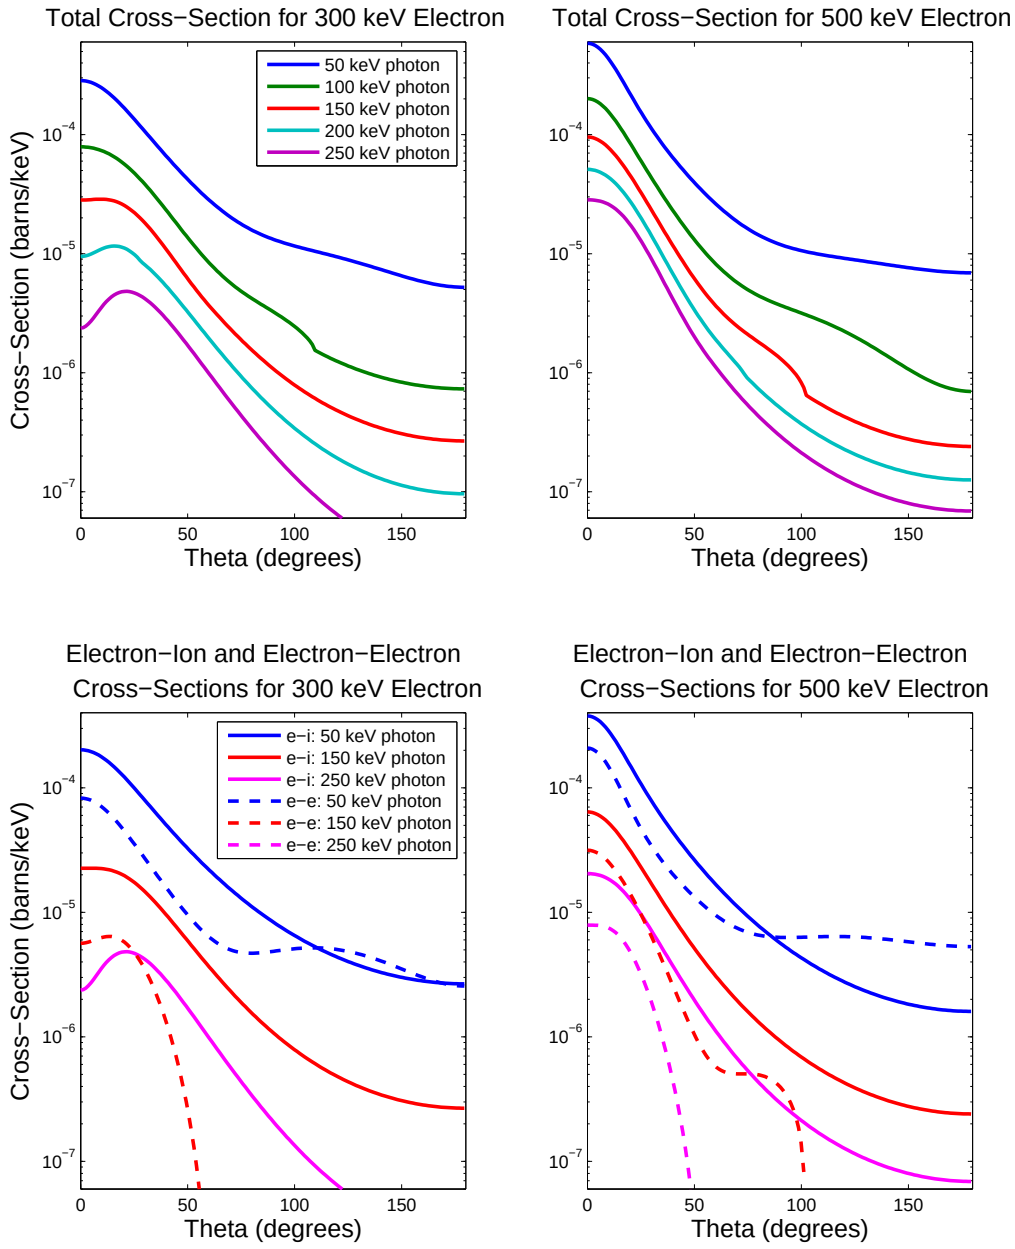

Figure 1-3: Bremsstrahlung cross-sections for fixed electron energies, at various photon energies. Theta is the angle between the fast electron velocity and the viewing chord. Electron-ion cross sections are calculated using the Bethe-Heitler model [42] including the Elwert factor [43]. Electron-electron cross sections are calculated using the Haug model, including the Coulomb correction factor. [44] The 250 keV photon e-e cross-sections for a 300 keV electron are extremely small and do not fall within the bounds of the lower left plot.

1. At most angles, a higher energy electron will emit more at any given photon energy than a lower energy electron.

2. For a fixed electron energy, the cross-section goes up as photon energy goes down.

3. Emission is dominant in the direction of the fast electron velocity ($\theta = 0$), sometimes as much as 100 times greater than in the opposite direction ($\theta = 180$).

4. Electron-ion emission often dominates over electron-electron emission.

The electron-electron and electron-ion contributions to the photon density can be expressed as integrals over the corresponding cross-sections [40]:

$$\frac{dn_k^{ee}(t,k,\mathbf{x},\hat{\mathbf{b}}\cdot\hat{\mathbf{r}})}{dt\,dk\,d\Omega} = n_e(t,\mathbf{x})\int d^3\mathbf{p}\frac{d\sigma_{ee}(k,p,\hat{\mathbf{k}}\cdot\hat{\mathbf{p}})}{dk\,d\Omega}vf(t,\mathbf{x},\mathbf{p}), \tag{1.76}$$

$$\frac{dn_k^{ei(j)}(t,k,\mathbf{x},\hat{\mathbf{b}}\cdot\hat{\mathbf{r}},Z_j)}{dt\,dk\,d\Omega} = n_{i(j)}(t,\mathbf{x})\int d^3\mathbf{p}\frac{d\sigma_{ei}(k,p,\hat{\mathbf{k}}\cdot\hat{\mathbf{p}}),Z_j}{dk\,d\Omega}vf(t,\mathbf{x},\mathbf{p}). \tag{1.77}$$

In Eqs. 1.76 and 1.77, t is time, $\mathbf{x}$ is the location of emission, $\hat{\mathbf{b}}$ is the direction of the magnetic field at the emission location, $\hat{\mathbf{r}}$ is the direction of the line of site from the emission location to where we observe, and Z_j is the charge of the j^{th} ion species. The terms in front of the integrals, n_e and $n_{i(j)}$, are the densities of the bulk electrons and j^{th} ion species. The velocity v is a test particle velocity and f is the electron distribution function.

The total photon density is a sum over interactions with bulk electrons and each species of ions:

$$\frac{dn_k(t,k,\mathbf{x},\hat{\mathbf{b}}\cdot\hat{\mathbf{r}})}{dt\,dk\,d\Omega} = \frac{dn_k^{ee}(t,k,\mathbf{x},\hat{\mathbf{b}}\cdot\hat{\mathbf{r}})}{dt\,dk\,d\Omega} + \sum_j \frac{dn_k^{ei}(t,k,\mathbf{x},\hat{\mathbf{b}}\cdot\hat{\mathbf{r}},Z_j)}{dt\,dk\,d\Omega}. \tag{1.78}$$

The total photon density, $\frac{dn_k}{dt\,dk\,d\Omega}$, is the number of photons being emitted by interactions with bulk electrons and ions per unit volume per unit time per unit

photon energy per solid angle. The number of photons reaching the detector in a given energy range can be expressed as a volume and solid angle integral over the photon density, taking into account a geometrical factor that characterizes the detector and its solid angle exposure.

1.5 Electron Cyclotron Emission Diagnostic

The electron cyclotron emission (ECE) diagnostic is a commonly used plasma diagnostic, found on most tokamaks [45]. This diagnostic measures the frequency-dependent amplitude of electron cyclotron emission, from which the electron temperature of a thermal (Maxwellian) plasma can be determined [32]. In the presence of non-thermal features of the electron distribution function, such as the fast electron tail produced by lower hybrid RF power, the ECE signals can be affected significantly and the resulting data may yield some information about the non-thermal electrons.

In a thermal plasma, the ECE diagnostic can measure the electron temperature at optically thick frequencies. The optical depth, τ, is defined as:

$$\tau(\omega) = \int_{path} \alpha(\mathbf{x}, \omega) dl. \tag{1.79}$$

In Eq. 1.79, $\mathbf{x}$ is the plasma position, α is the fractional absorption per unit length at a frequency ω, and the integral is performed over the path taken by radiation of frequency ω through the plasma and into the diagnostic detector. Note that at relevant ECE frequencies, refraction may play a role and these paths may not be straight lines.

A plasma is considered optically thick in a range of frequencies if $\tau(\omega) \gg 1$ for that range of frequencies. In this case, the electron plasma is a good absorber and emitter whose emission can be approximated by the black body equation:

$$I(\omega) = \frac{\hbar\omega^3}{8\pi^3 c^2} \frac{1}{e^{\hbar\omega/T_e} - 1}. \tag{1.80}$$

Note that ω is the electron cyclotron frequency (or a harmonic of it), which is

a function of the radius-dependent tokamak magnetic field. For a typical magnetic field of 5 Tesla, $\hbar\omega \approx 6 \times 10^{-4} eV$, while typical electron temperatures are a few keV. Thus the exponential in Eq. 1.80 can be expanded to yield:

$$I(\omega) \approx \frac{\omega^2 T_e}{8\pi^3 c^2}. \tag{1.81}$$

Since the local electron cyclotron frequency is proportional to the magnetic field and the magnetic field decreases monotonically with increasing major radius, the radial temperature profile can be easily obtained by re-arranging Eq. 1.81:

$$T_e(R) = \frac{8\pi^3 c^2 I(\omega)}{\omega^2}, \tag{1.82}$$

$$\omega = \omega_{ec}(R). \tag{1.83}$$

The electron cyclotron frequency for relativistic electrons, such as those generated by lower hybrid operation, is downshifted by the relativistic factor γ due to the effective mass increase:

$$\omega_{ce} = \frac{eB}{m_e} \rightarrow \omega_{ce} = \frac{eB}{\gamma m_e}. \tag{1.84}$$

This leads to an ambiguity in the location of emission because the emission frequency now depends on both the local magnetic field and the γ of the emitting electron. Furthermore, the relativistic electrons on the outboard side of the tokamak, where the magnetic field (and thus the electron cyclotron frequency) is already low, may emit at low frequencies which are not seen in the ECE spectrum of a thermal plasma. Because the density of relativistic electrons is rather low compared to the density of bulk electrons, the plasma is not necessarily optically thick at these lower frequencies. Thus the blackbody approximation breaks down and the resulting emission can be more difficult to interpret. Nonthermal ECE spectra can be compared with synthetic diagnostics that employ nonthermal electron distributions simulated by RF modeling codes, which will be discussed further in Chapter 4.

1.6 Contributions of the Author

The author used a ray tracing/Fokker-Planck simulation package, Genray-CQL3D, to simulate LHCD experiments on a tokamak at MIT, Alcator C-Mod. Predictions from the x-ray and electron cyclotron emission synthetic diagnostics were compared with experimental measurements. Predicted current profiles were compared with current profiles inferred from motional stark effect measurements. This was the first such benchmarking of Genray-CQL3D against experimental data from a LHCD experiment. Importantly, it provided one of the first benchmarks of a 3D Fokker-Planck code against an LHCD experiment[3].

The use of fast electron radial diffusion in these models predicted that if diffusion coefficients were above approximately 0.05 m^2/s, most of the driven current would be lost (an effect not seen in experiment). The author then designed and carried out a set of LH power modulation experiments, similar to those performed on Alcator C [36] and PBX-M [46], in order to quantify fast electron transport. A model was developed to use time-dependent x-ray profiles to calculate fast electron diffusion and convection for off-axis peaked inverted x-ray profiles. The measured fast electron transport was found to be very small on a slowing down time scale, and measured diffusivities were found to be consistent with values typically used in simulations.

A phasing scan from the LH power modulation experiments was also simulated. It was found that Genray-CQL3D predicted x-ray profiles agreed in magnitude and shape for high $n_{||}$ cases, while the predictions disagreed with experimental observation at low $n_{||}$. An interpretation of the disagreement at low $n_{||}$ is presented. An analysis of our current understanding of the regimes of validity of Genray/CQL3D calculations with respect to Alcator C-Mod LHCD simulations, based on modeling performed by the author as well as by other Genray/CQL3D users, is presented in the final chapter.

[3]Early comparisons of measurements with predictions of another ray tracing/3D Fokker-Planck simulation package, DKE [40] (now known as LUKE), began around the same time.

Chapter 2

Alcator C-Mod

Alcator C-Mod is a tokamak experiment located at the Plasma Science and Fusion Center at MIT, in Cambridge, Massachusetts. The name "Alcator" is a contraction of the Italian for "High Field Torus" and Alcator C-Mod is the the third tokamak in the Alcator series, after Alcator A and Alcator C. Discharges produced in C-Mod at toroidal field strengths of $\sim$5 T last typically for 2 seconds and are reproduced approximately every 15 minutes. For an 8-hour run day, there are typically about 30 plasma discharges ("shots"), each with a specific experimental purpose. Experiments are run 4 days per week during the run campaign, whose length is determined by annual funding levels, and has been approximately 15 weeks in recent years.

C-Mod is equipped with copper magnets which must be cooled between shots (a steady-state tokamak will require superconducting magnets to run continuously). They are capable of producing a toroidal magnetic field as high as 9 T (the highest field for tokamaks in operation today), though it is more typically run at 5.4 T [47]. C-Mod requires approximately 0.2-2 Volts per turn to drive up to 2 MA of plasma current, though a current of 1 MA is more typically used. During a pulse, the magnet system consumes approximately 100 MW of power (several hundred MJ over the course of a few seconds).

The densities routinely achieved on C-Mod ($2.5 \times 10^{19} - 3 \times 10^{20}$ particles/m^3 line-averaged) are also quite high when compared to other tokamaks. The combination of high magnetic fields and high densities on C-Mod make it ideal for performing exper-

iments that will guide the design and experimental program of the next generation of tokamaks, particularly ITER.

Plasmas in Alcator C-Mod are heated with a combination of ohmic and RF power in the ion cyclotron range of frequencies (ICRF). The plasma is heated using up to 5 MW of ICRF power that is coupled through three antennas. Electron temperatures of 5 keV (58 million degrees) are routinely reached in the center of the plasma. C-Mod has a 4.6 GHz lower hybrid system with up to 3 MW of source power coupled through a single LH grill, with plans to add a second grill. Power for the ICRF and LH systems is drawn directly off of the grid.

The tokamak itself is encased in a roughly 1 m thick layer of concrete called the "igloo," which shields diagnostics and electronics from neutrons. The entire assembly, including diagnostics located next to the reactor, is inside a room (the "cell") with 2 m thick concrete walls. The purpose of the concrete is to shield people from energetic neutrons that may be generated during operation, and all personnel need to evacuate the cell prior to the beginning of a run day. If an instrument in the cell needs to be adjusted manually between shots, the procedures for opening the cell door and allowing access in addition to the manned access time may delay the next experiment. Thus it is desirable to keep access to the cell to a minimum during the run day and to be able to control most aspects of a diagnostic remotely.

Though commercial fusion reactors are expected to run on a D-T fuel mix, C-Mod is run with a deuterium plasma. This avoids the safety and regulatory issues associated with having tritium on site. It also makes the dominant fusion reaction in the plasma a D-D reaction, which is far less likely to occur at the typical plasma temperatures achieved in the tokamak. This keeps the neutron flux down (relative to a D-T plasma), though it still occasionally reaches levels during operation that could activate components of the reactor, requiring personnel to take appropriate safety measures to minimize their radiation dosage.

An immense amount of data is taken during every shot: approximately 10 GB of data, depending on the experiment. This data is organized using MDSplus [48] software in a structure called the "tree." Each shot has several trees associated with

it, each associated with a related group of diagnostics. Trees contain addressable nodes, some of which are nested, and data is stored within the various nodes. The two trees used by the lower hybrid diagnostics and hard x-ray camera are "lh" (for lower hybrid and processed HXR data) and "hxrlocal" (for raw HXR data).

2.1 Lower Hybrid Current Drive on Alcator C-Mod

Alcator C-Mod is the only tokamak in the US and one of a handful of tokamaks in the world equipped with a lower hybrid current drive (LHCD) system. The present system has approximately 2.5 MW of source power, generated by a set of klystrons (microwave-producing vacuum tubes) at 4.6 GHz. The source power travels through WR-187 waveguides to the so-called "jungle gym," where the waveguide power is split before being fed into the launcher.

Two different lower hybrid launcher designs have been used on Alcator C-Mod. The first design, known as LH1, was used in the 2006, 2007, and 2008 campaigns [49]. In the LH1 design, the jungle gym split the waveguide power twice before feeding it into the rear wave guide (RWG) system, which split the waveguides again and then tapered the waveguides from the WR-187 size to the dimensions needed for the LH coupler (also known as the "grill"). The RWG fed into the forward wave guide (FWG) system, which fed into the coupler. The front of the coupler was only a few millimeters behind the last closed flux surface of the plasma. The coupler had 4 rows and 24 columns of waveguides. A schematic of the entire LH1 system is shown in Figure 2-1 and a picture of the LH1 launcher, as viewed from the inside of the tokamak is shown in Figure 2-2.

The second design, LH2, was used in the 2010 campaign [50]. The redesign of the launcher was meant to minimize costly losses of power from the klystrons to the grill. This was achieved through a simplification in the splitting scheme as well as by tapering the waveguides much closer to the plasma. In the LH2 design, the jungle

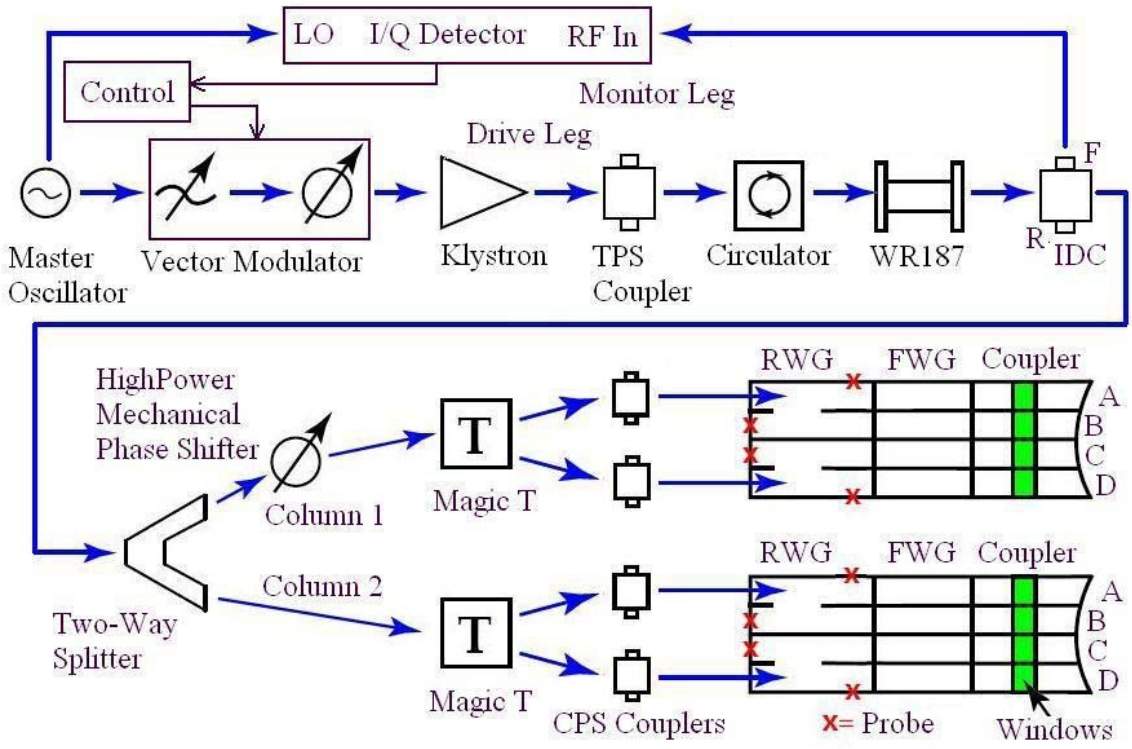

Figure 2-1: Schematic of LH1 System. [49]

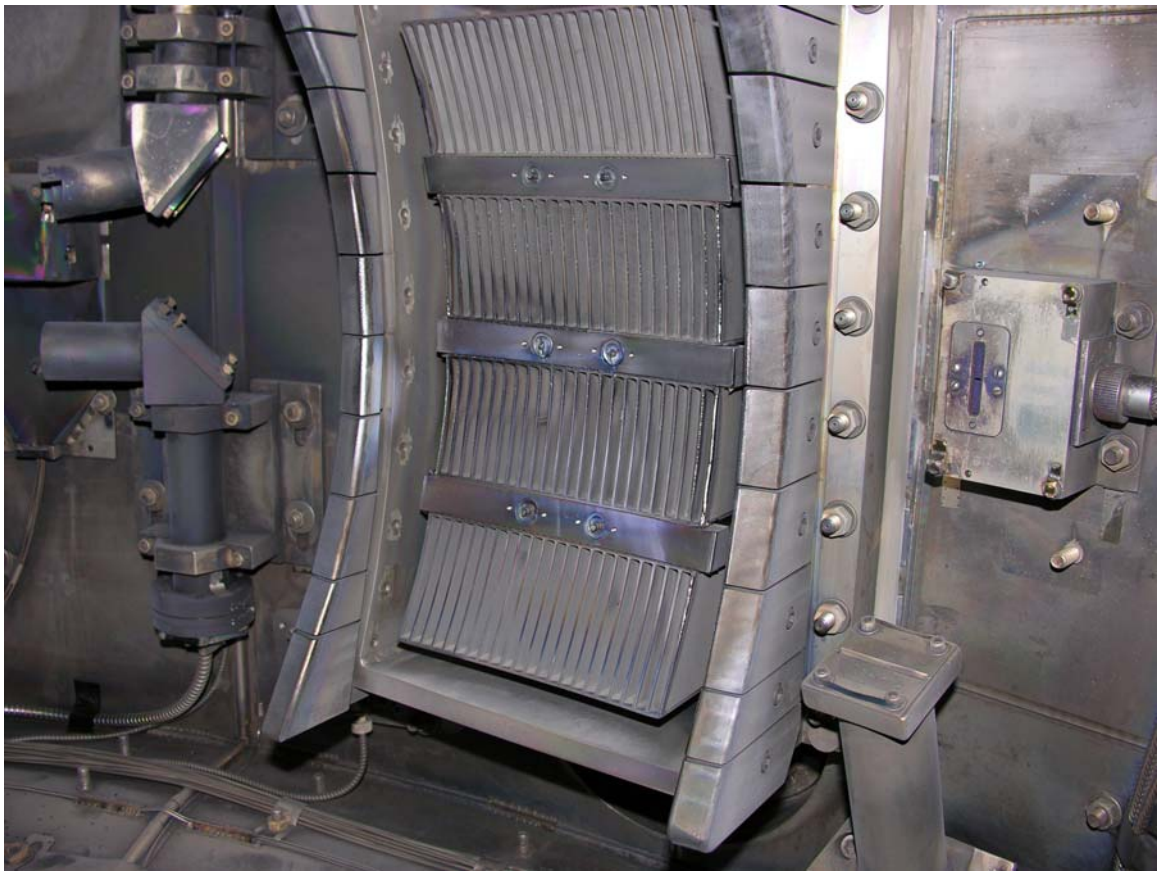

Figure 2-2: LH1 launcher installed in tokamak. [49]

gym only needs to split the waveguides once and then feeds them into the back of the launcher. The launcher tapers the waveguides and employs an innovative 4-way splitter design to feed power to each column of waveguides in the grill. The LH2 grill has 4 rows and 16 actively driven columns (it also has 2 passive columns). A drawing of the 4-way splitter design is shown in Figure 2-3 and the entire LH2 launcher is shown in Figure 2-4. Figure 2-5 showns a picture of the installed LH2 launcher.

The peak value of launched $n_{||}$ is determined by the phasing between columns of waveguides. The launched $n_{||}$ power spectra for LH1 and LH2 are shown in Figures 2-6 and 2-7, respectively.

Note that for each phasing, there are two main lobes of power. The dominant one, known as the forward lobe, is centered at a negative value of $n_{||}$ (anti-parallel to the magnetic field) and this is the part of the spectrum intended to create current drive. The smaller lobe, known as the reverse lobe, is centered at a positive value of $n_{||}$ of greater magnitude. In principle, the reverse lobe could drive current in the opposite direction, though in practice, counter-current is more difficult to drive. This is partly due to the direction of the toroidal electric field and partly due to the very low current drive efficiency for the high values of $n_{||}$ found in the reverse lobe (see Eq. 1.63). Because of its high value of $n_{||}$, the reverse lobe damps on the outer edge of the plasma where temperatures are lower. Clearly it is desirable to keep the reverse lobe as small as possible so as not to waste valuable power in the area of the spectrum that does not drive useful current.

Phasing between columns in LH1 and LH2 is controlled partly by the phasing of the low power drive signal input into individual klystrons. These drive signals can be controlled dynamically and remotely, allowing for multiple phasing segments during a single discharge as well as phase changes between discharges (without a cell access). However, in both LH1 and LH2, there are more columns of waveguides than there are klystrons. Since the power from some klystrons must be split between columns, the phase shift between those columns must be controlled mechanically.

In LH1, the phasing for each pair of columns was controlled by the klystron phasing. In the 2006 and 2007 campaigns and the early part of the 2008 campaign,

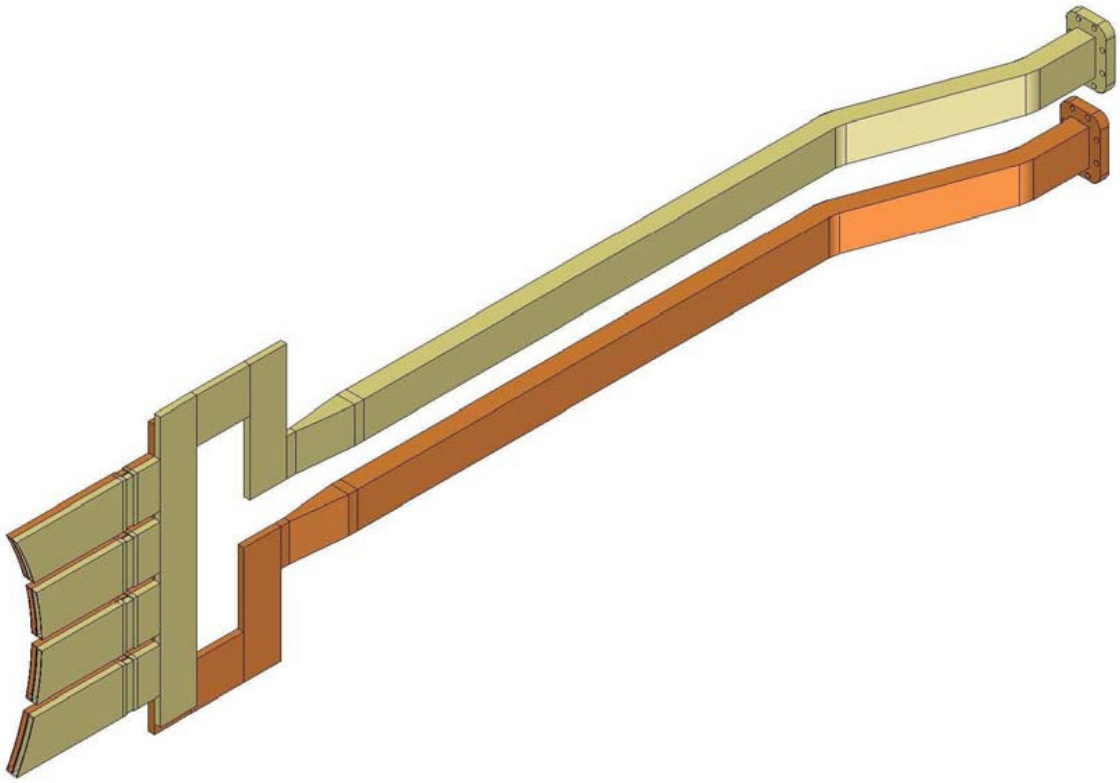

Figure 2-3: Schematic of 4-way splitter concept employed in LH2, showing two adjacent 4-way splitter plates. Each plate makes up a single column of 4 waveguides. Tapering of waveguide dimensions is done as close to the plasma as possible to minimize resistive losses. [50]

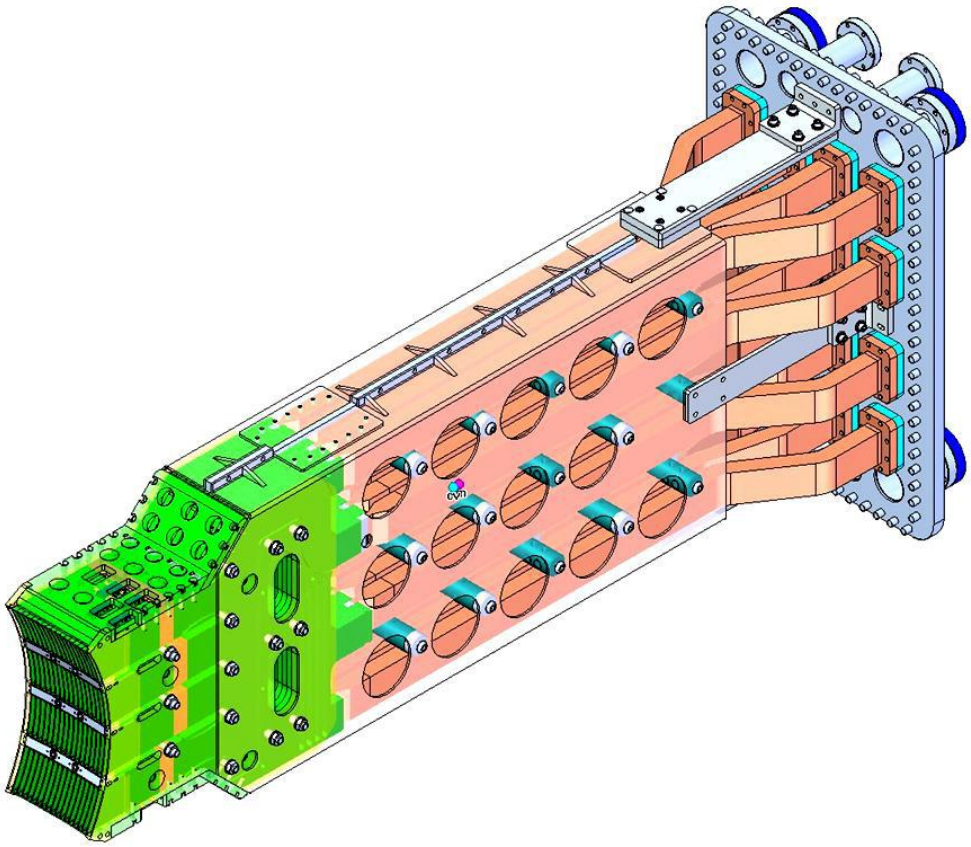

Figure 2-4: Full LH2 launcher. [50]

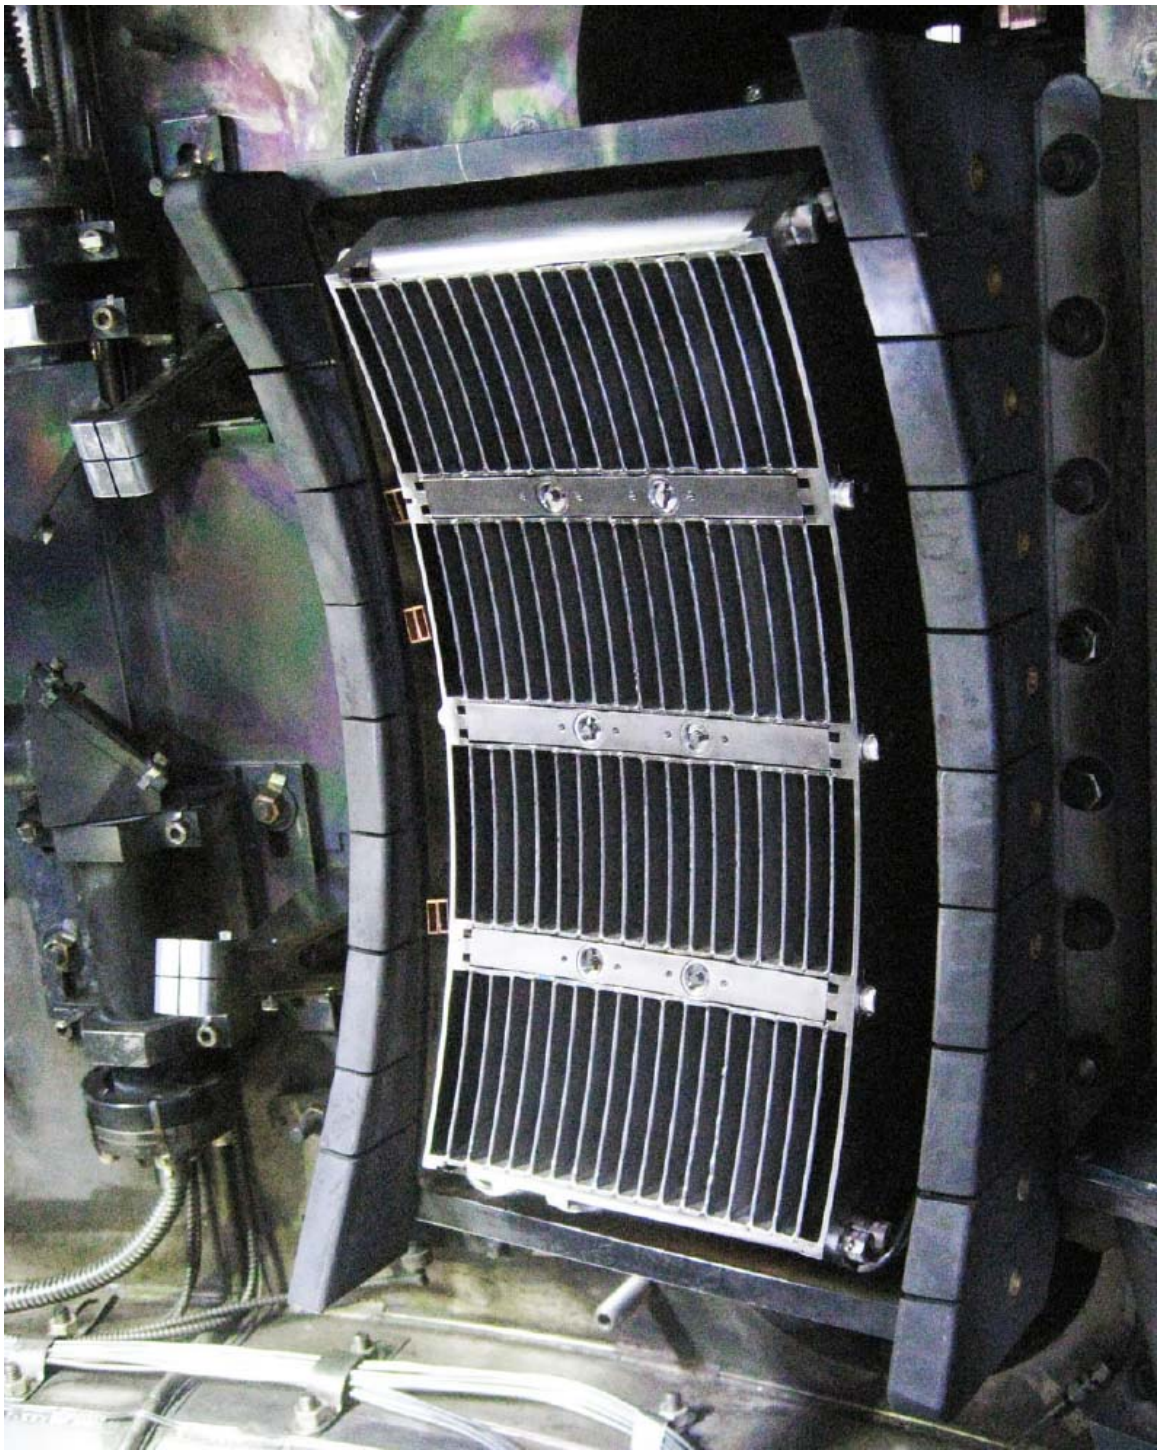

Figure 2-5: LH2 launcher installed in tokamak. [50]

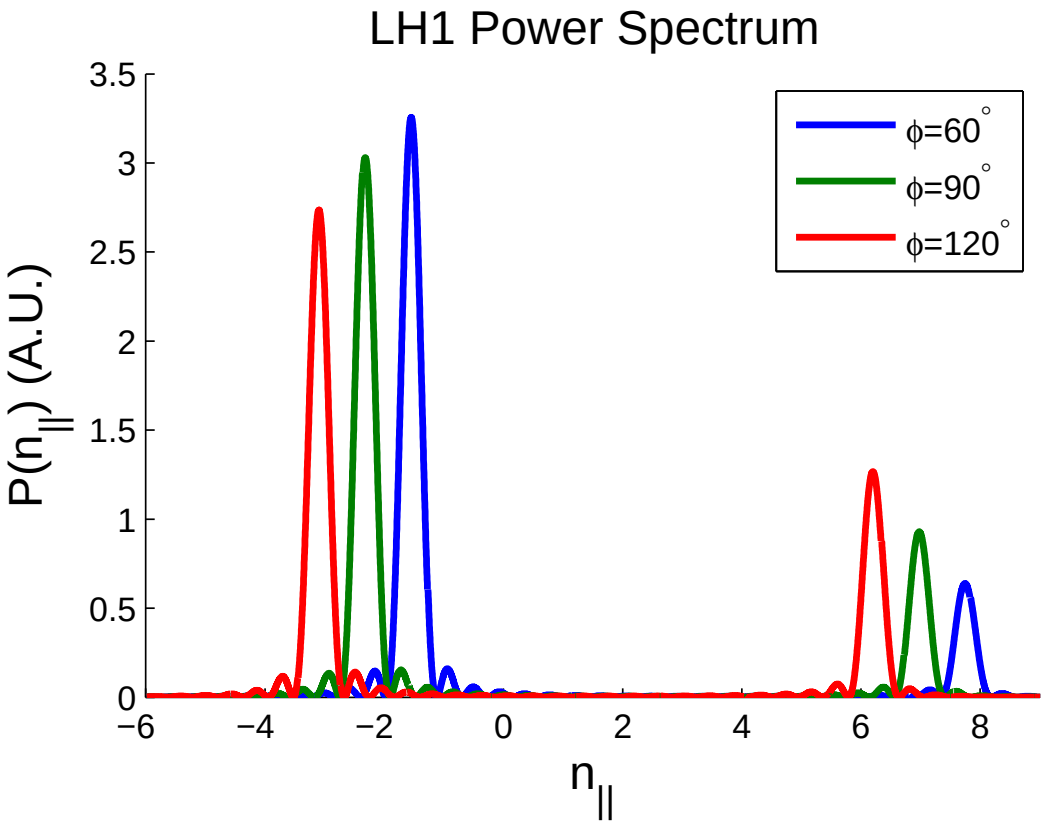

Figure 2-6: Launched $n_{||}$ spectra for LH1. Ideal phasing has been assumed. Launching into a vacuum (matched impedance) has been assumed. Spectrum calculation follows calculation method outlined in [51] (figure courtesy of Greg Wallace).

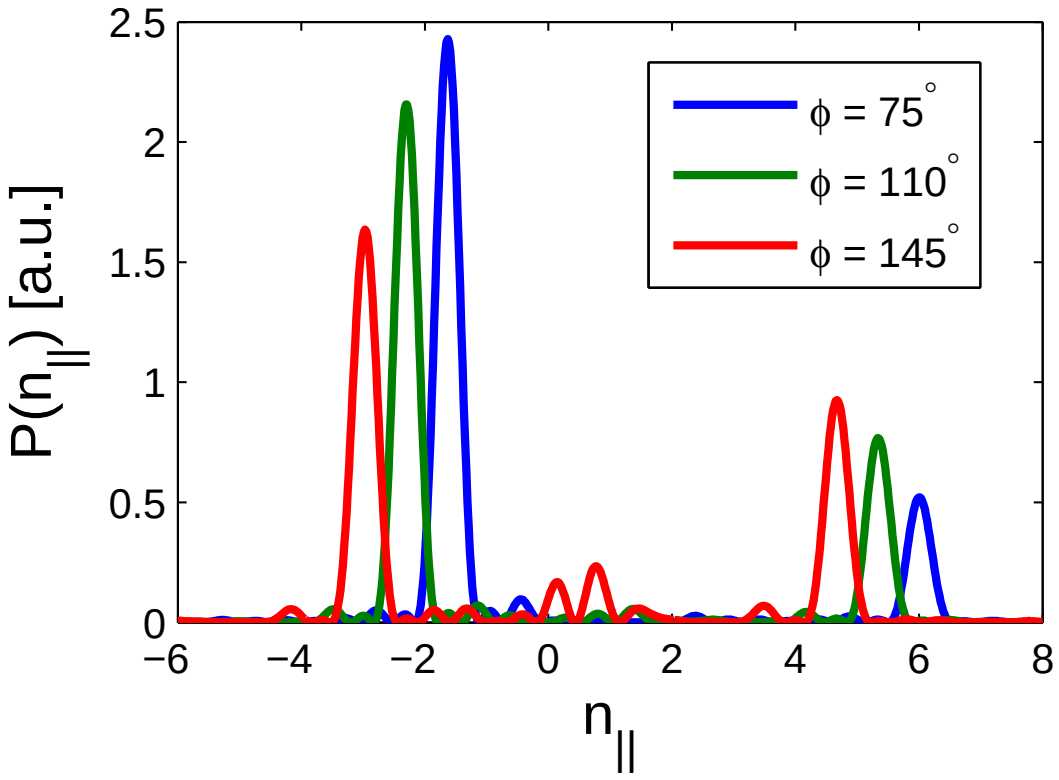

Figure 2-7: Launched $n_{||}$ spectra for LH2. Fixed phase shims at 90° has been assumed. Launching into a vacuum (matched impedance) has been assumed. Spectrum calculation follows calculation method outlined in [51] (figure courtesy of Greg Wallace).

the relative phasing between two columns in a pair were controlled by mechanical phase shifters, which were manually set to whatever phasing was desired (changing the mechanical phase shifters required a cell access). Early in the 2008 campaign, it was discovered that the mechanical phase shifters were creating arcing problems during high power operation. They were replaced by a fixed section of waveguide which set the phase shift between adjacent columns in a pair to 90° for the remainder of the campaign.

In LH2, each of the central 4 columns is driven by its own klystron and so their phasings are directly controlled by the klystron drive phasing. The remaining 12 columns are split into pairs, with each pair being driven by a single klystron. The relative phasing between two columns in a pair are controlled mechanically with "phase shims," which are sections of waveguide attached to the output of magic T splitters in the jungle gym, whose relative lengths determine the phasing between pairs of columns. In practice, the phasing between adjacent columns on LH2 has been fixed at 90°.

The LH wave couples to the plasma after tunneling through an evanescent region (of order 1 mm) where the plasma density is lower than the cut-off density of 2.6×10^{17} m^{-3} (the density required to make the electron plasma frequency equal to the source frequency, 4.6 GHz). In order to maintain a high vacuum inside the tokamak, ceramic windows must be brazed onto the individual coupler waveguides. Due to difficulties in brazing the ceramic windows at the edge columns in LH1, only 22 of the 24 columns were used in practice during operation. One important advantage of the LH2 design was that it called for brazing of individual windows before the launcher was assembled, which allowed successful brazing of every window and use of all 16 columns.

2.2 Hard X-Ray Camera on Alcator C-Mod

Alcator C-Mod has a pinhole hard x-ray (HXR) camera, designed to detect Bremsstrahlung emission from lower hybrid driven fast electrons [52]. The camera has 32 horizontally-viewing spatial chords that span the poloidal cross-section of the plasma and ter-

minate at cadmium zinc telluride (CZT) solid-state detectors. The chordal spatial resolution of the camera is 2 cm and the time resolution is 1 μs. The camera is designed to detect photons of energies 20-200 keV, with better than 1 keV energy resolution.

An incident photon that reaches a detector creates a small current, which is amplified and sent to a shaper circuit. The resulting voltage output is a pulse train of variable-height pulses, which is digitized and stored. The digitized signal can be post-processed for baseline subtraction and pulse height analysis. Each pulse is then binned, resulting in a histogram of photon counts as a function of time and energy. Extensive documentation of the electrical and mechanical design of the HXR camera, as well as the accompanying software, can be found in Reference [53].

2.2.1 Hard X-Ray Camera View

The hard x-ray camera is essentially a perpendicularly-viewing diagnostic, meaning that the viewing chords are mainly perpendicular to the tokamak magnetic fields. A schematic of the HXR camera and its viewing chords are shown in Figure 2-8.

As discussed in Section 1.4, bremsstrahlung emission is dominant in the direction of the velocity of the fast electron. While the lower hybrid current drive is designed to generate fast electrons moving primarily in the direction parallel to the magnetic field, the hxr viewing chords are viewing perpendicular emission from those current-carrying electrons. It therefore becomes important to consider the contribution to emission of pitch-angle scattered fast electrons for the current HXR camera view. This can be accomplished with bremsstrahlung cross-section data and a simulated electron distribution function.

Figure 2-9 shows relative levels of HXR emission as a function of electron pitch angle (angle between the electron velocity and the magnetic field) for a midplane, perpendicular view. These emissivities at each pitch angle were obtained by multiplying the effective bremsstrahlung cross sections (as seen from the camera, given the pitch angle under consideration) by the fast electron population at the given pitch angle from a simulated electron distribution function. It is clear from Figure 2-9 that

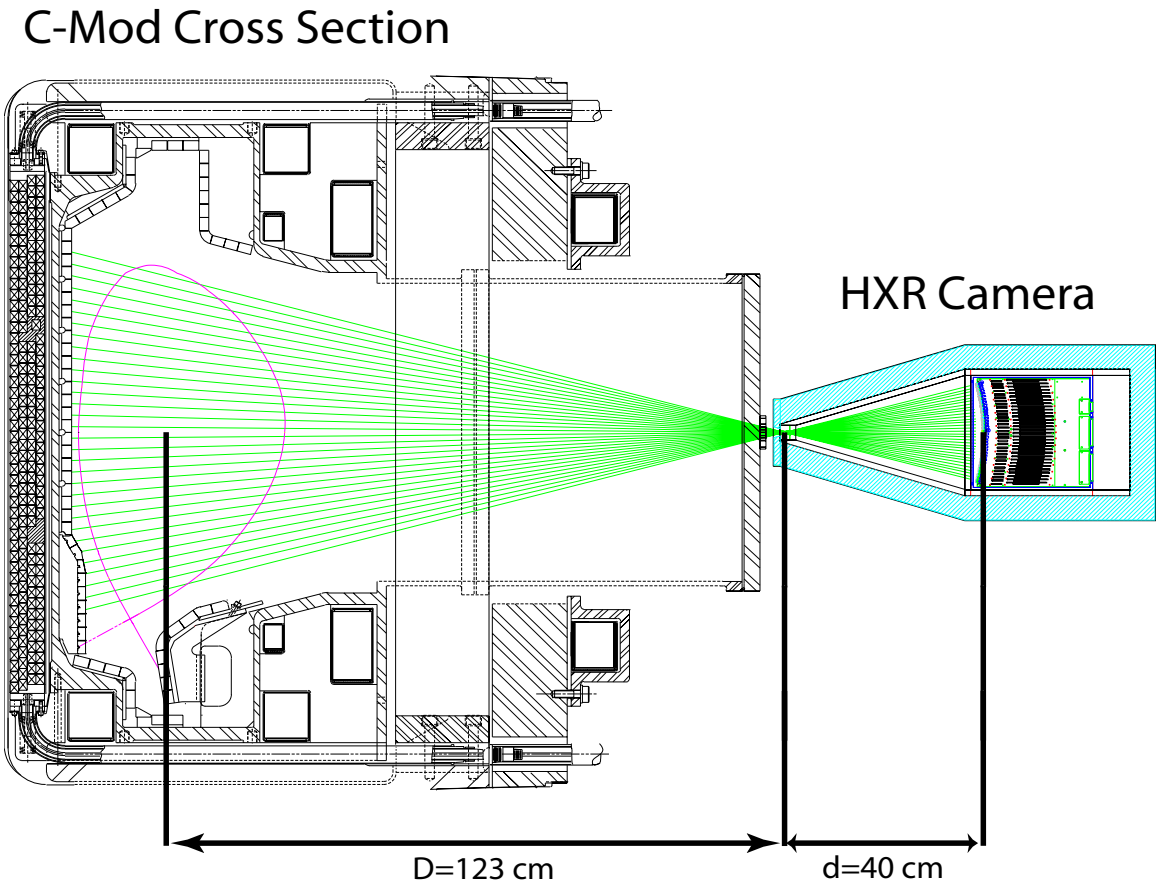

Figure 2-8: HXR camera schematic and viewing chords with respect to the poloidal cross-section of the tokamak. Figure courtesy John Liptac.

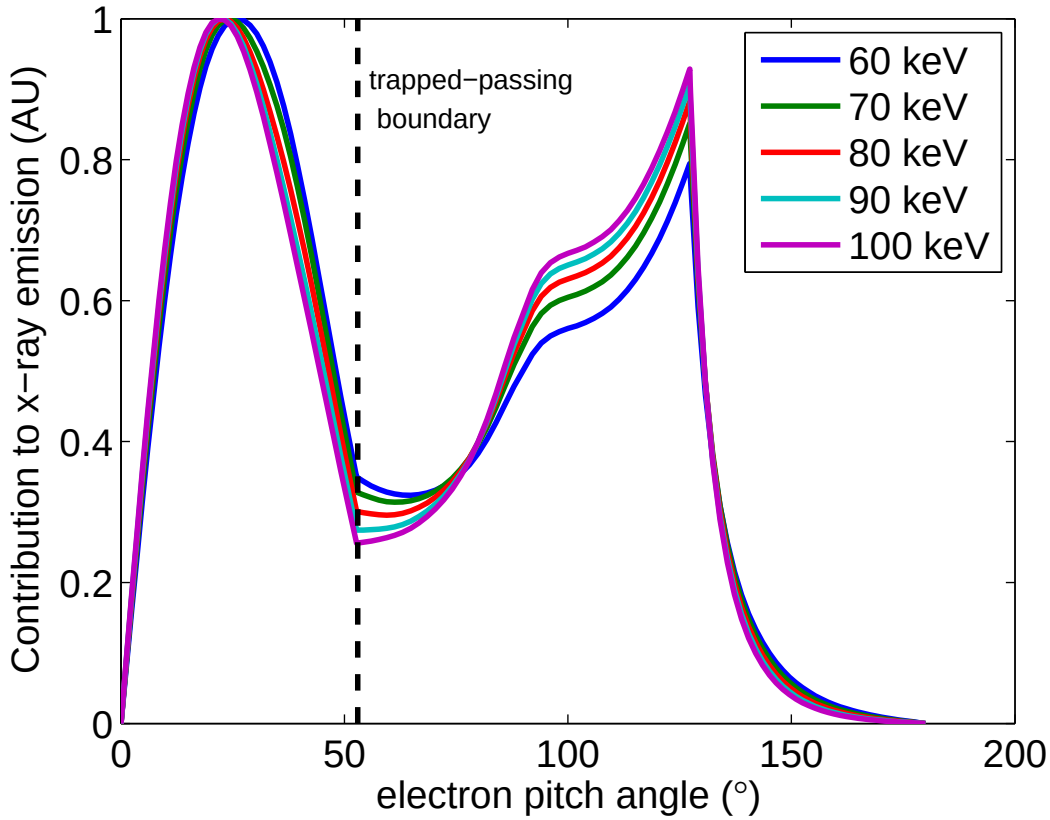

Figure 2-9: Pitch angle dependence of HXR emission, using bremsstrahlung cross sections and a simulated electron distribution function.

both forward-moving and pitch-angle-scattered fast electrons contribute substantially to the total HXR emission detected by the camera.

2.2.2 Electron Energy Dependence of HXR Emission

It is also an interesting exercise to use a simulated electron distribution function to calculate the relative contributions of various electron energies to a given detected photon energy. Figure 2-10 shows the relative contributions of various electron energies to HXR emission at specific photon energies. These curves were obtained by performing a convolution integral over all pitch angles of a simulated electron distribution function and effective bremsstrahlung cross-sections. They show that while emission from a specific photon energy is peaked at one electron energy, each photon

energy receives contributions from a broad range of electron energies.

2.2.3 Shielding

Many factors, most notably the plasma density, play a role in the overall HXR signal levels. There are certain experiments that can be anticipated to produce high enough signal levels that the detectors will become saturated. For such experiments, we have devised simple 316L stainless steel shields of varying thicknesses that can be placed between the aluminum vacuum window and the HXR Camera pinhole. The existing shields are .015", .030", and .058" thick and can be combined if desired and attached to the vacuum window with two set screws. A cell access is required to change out the shielding, so it is desirable to choose a shielding thickness that can be used for several consecutive discharges or the entire run day.

Stainless steel has the advantage that it attenuates HXR signal significantly in a reasonable thickness. To achieve the same effect with aluminum would require several centimeters of thickness, which is not feasible in the small space available. Lead shielding of any reasonable thickness would attenuate too much. Furthermore, lead shielding as a function of HXR energy would be discontinuous due to K shell edges.

The HXR attenuation by stainless steel is also a function of HXR energy, with lower energy photons being shielded the most. This bias turns out to work in our favor, since photons below 40 keV are typically ignored for all data analysis purposes (they are within the digitizer noise for some chords) and there is always more unattenuated signal at lower HXR energies than there is at higher energies. Transmission as a function of HXR energy for several stainless steel thicknesses is shown in Figure 2-11.

For every discharge, the shielding thickness in inches is recorded in the tree, under the node "\LH::TOP.HXR.PARAMS:SSTHICK". For any given photon energy, the transmission can be calculated and photon counts can be weighted by the inverse of transmission to reconstruct the unattenuated signal. This is particularly useful for comparing HXR from discharges with different thicknesses of shielding or for comparison with simulated HXR spectra.

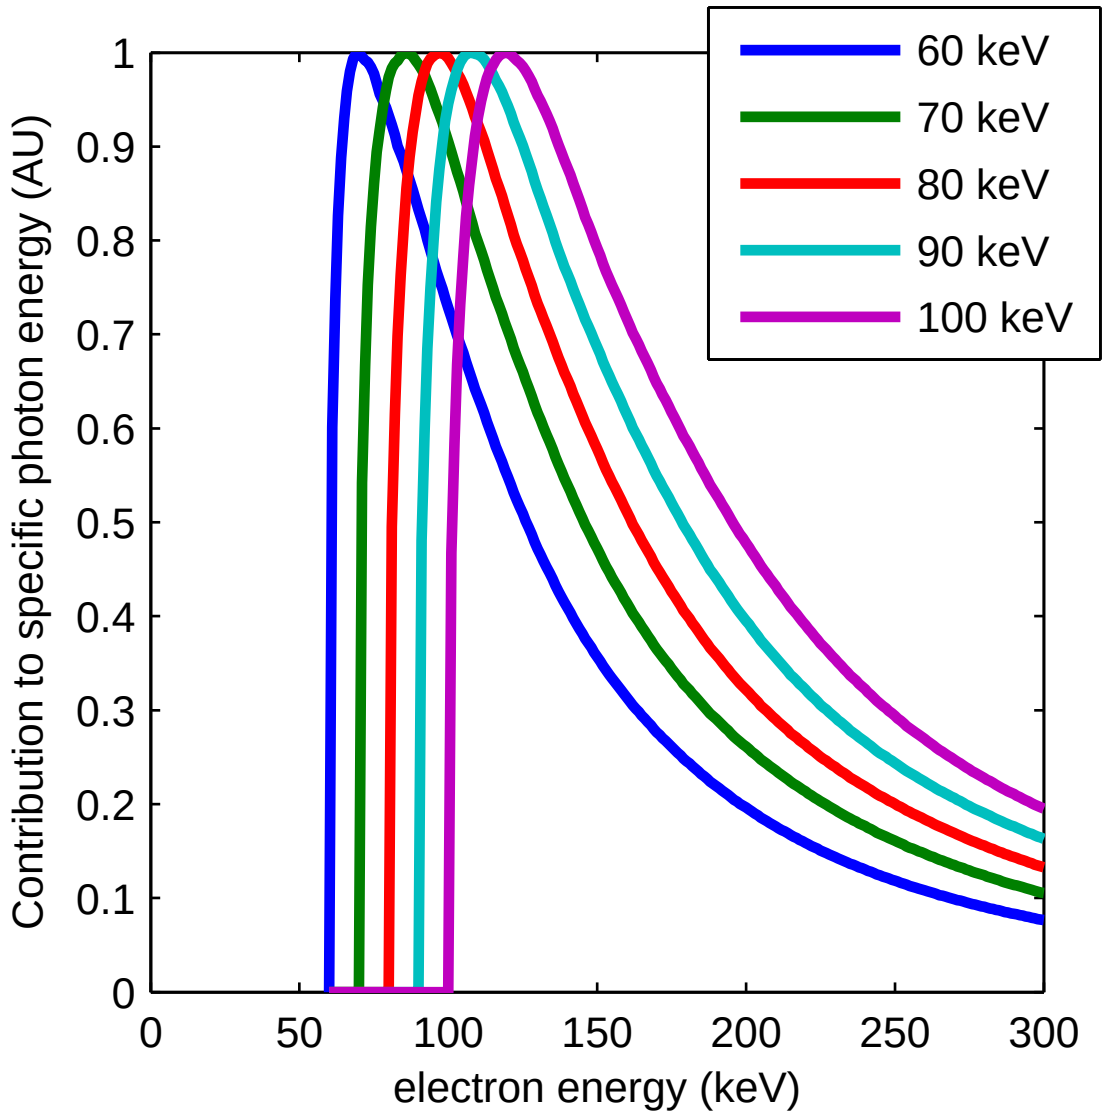

Figure 2-10: Electron energy dependence of HXR emission, using bremsstrahlung cross sections and a simulated electron distribution function. Each curve represents a different photon energy.

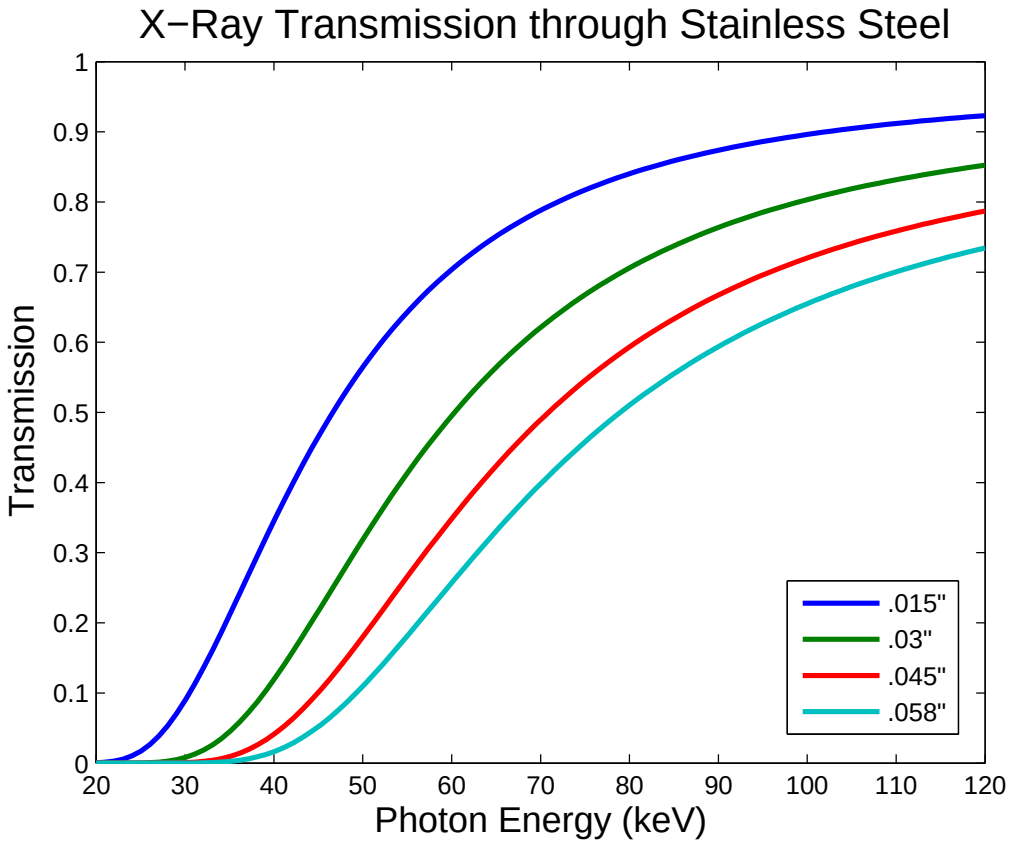

Figure 2-11: Transmission through stainless steel shielding of varying thicknesses, as a function of photon energy.

Another effect which can be accounted for when comparing simulated and experimental HXR data is the detector response function, which takes into account detector efficiency as well as the uncertainty in the photon energy measurement. A discussion of detector respose functions for 2mm CZT detectors can be found in References [53, 54].

2.3 Electron Cyclotron Emission Diagnostics on Alcator C-Mod

Alcator C-Mod has two types of outboard, midplane, horizontally-viewing electron cyclotron (ECE) diagnostics[1] that measure electron temperature during normal operation [55]. One is a Michelson Interferometer, which splits the incoming ECE spectrum into two beams that are later recombined. The path length of one of the beams is controlled by a moving mirror, which scans back and forth in time. The Fourier transform of the resulting interferogram reveals the frequency components of the incoming beam. The mirror scans position approximately 14 times per second, leading to a high frequency resolution, low time resolution measurement.

The other type of horizontally-viewing ECE diagnostic on C-Mod is a grating polychrometer (GPC). This diagnostic uses a diffraction grating to split the ECE spectrum into frequency components, with detectors physically spaced apart to measure the intensity of several different frequencies. Different diffraction gratings and diffraction grating angles can be used to accommodate various magnetic field operation points as well as to look for lower frequency emission during LH operation. There are actually two GPC systems on C-Mod, one with 9 detectors and one with 19 detectors. Thus the GPC diagnostics provide a high time resolution, low frequency resolution measurements of the ECE spectrum.

On Alcator C-Mod, the second ECE harmonic is optically thick and it is this harmonic which is used to measure electron temperature during normal operation.

[1] C-Mod has a third ECE diagnostic, FRCECE, which views at an oblique angle. However, none of the calculations of ECE spectra presented in this thesis correspond to the FRCECE view.

During LH operation, ECE is observed at frequencies below the second harmonic due to outboard relativistic electrons whose cyclotron frequency is downshifted.

Figure 2-12 shows the classical and relativistically downshifted ECE frequencies for the fundamental, 2^{nd} and 3^{rd} harmonics, assuming a central magnetic field of 5.4 T (typical C-Mod operation). The downshifted frequencies are based on $\gamma = 1.2$ (E=100 keV). Note how downshifted ECE originating from the inboard side of the tokamak can be re-absorbed by thermal electrons at a greater major radius. Thus in some plasma scenarios, the blackbody approximation is still valid for the classical 2^{nd} harmonic frequencies. However, sufficiently downshifted ECE from the outboard side of the tokamak travels to the detectors without the opportunity to be re-absorbed by the bulk electrons.

The result is that very high levels of ECE are observed at frequencies just below the 2^{nd} harmonic during LH operation (this has been referred to as "nonthermal ECE") [56, 37]. During LH operation, the thermal portion of the ECE spectrum may or may not be enhanced, depending on plasma conditions and the density and energy spectrum of the LH-driven fast electrons. For this reason, the Thomson Scattering diagnostic is often considered to be more reliable for electron temperature measurements during LH operation. It is best to use a synthetic ECE diagnostic to compare predicted and measured ECE for various plasma scenarios. Synthetic diagnostics are discussed in more detail in Chapter 4.

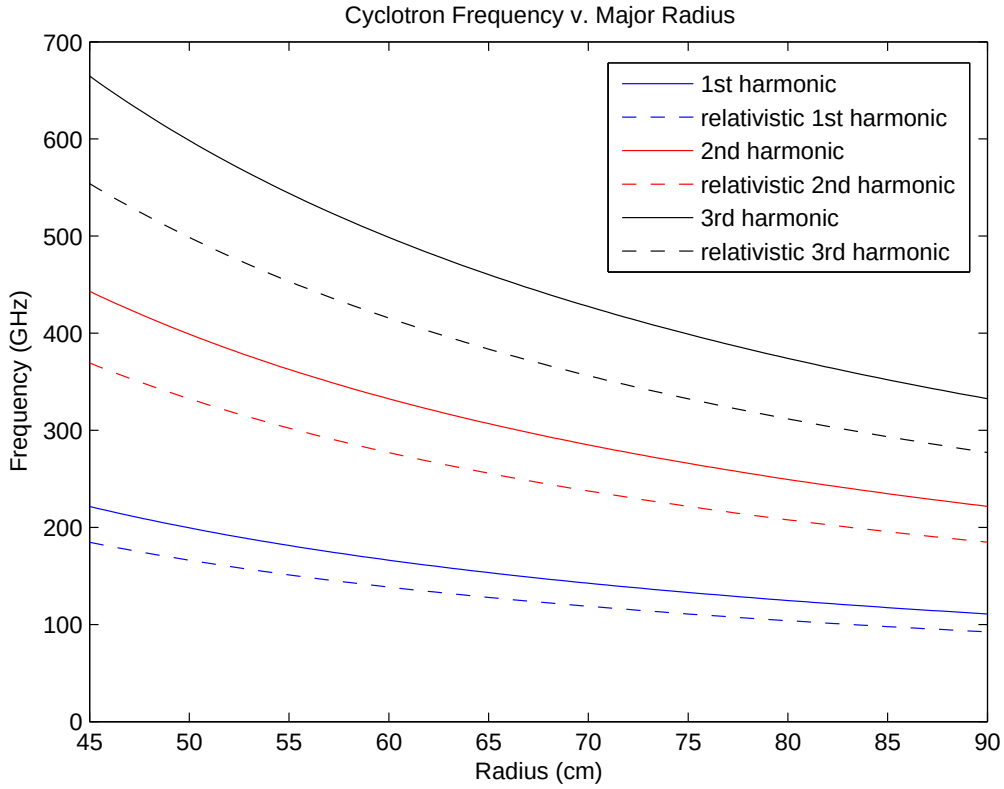

Figure 2-12: Classical and relativistically downshifted electron cyclotron frequencies for the fundamental and 2^{nd} and 3^{rd} harmonics, based on a central magnetic field of 5.4 T. For the downshifted frequencies, γ is taken to be 1.2.

Chapter 3

HXR Data Analysis Techniques

We have developed techniques that can be used to analyze HXR data for a variety of experiments. The two main techniques are energy fitting and spatial inversion, discussed in the following sections.

3.0.1 Energy Fitting

There are some experiments in which the binned HXR count rates are low, leading to poor counting statistics. For example, this can be the case when high time resolution (small time bins) is needed or when the experimentalist is interested in high energy x-ray photons (which are emitted less frequently than lower energy photons). It can also be the case when the total x-ray flux is low, as it might be during an experiment requiring low LH power or high electron density.

Poor counting statistics lead to noisy integrated profiles. These noisy profiles are difficult to spatially invert. Furthermore, it can be unclear whether changes in the HXR profiles as a function of time are due to changing plasma conditions or can simply be attributed to noise.

In these situations, some of the noise due to poor statistics can be removed through an energy fitting technique, in which a curve is fit to the x-ray data as a function of energy. The energy-fitted data can be much more accurately inverted. Figure 3-1 shows an example of such an energy fit.

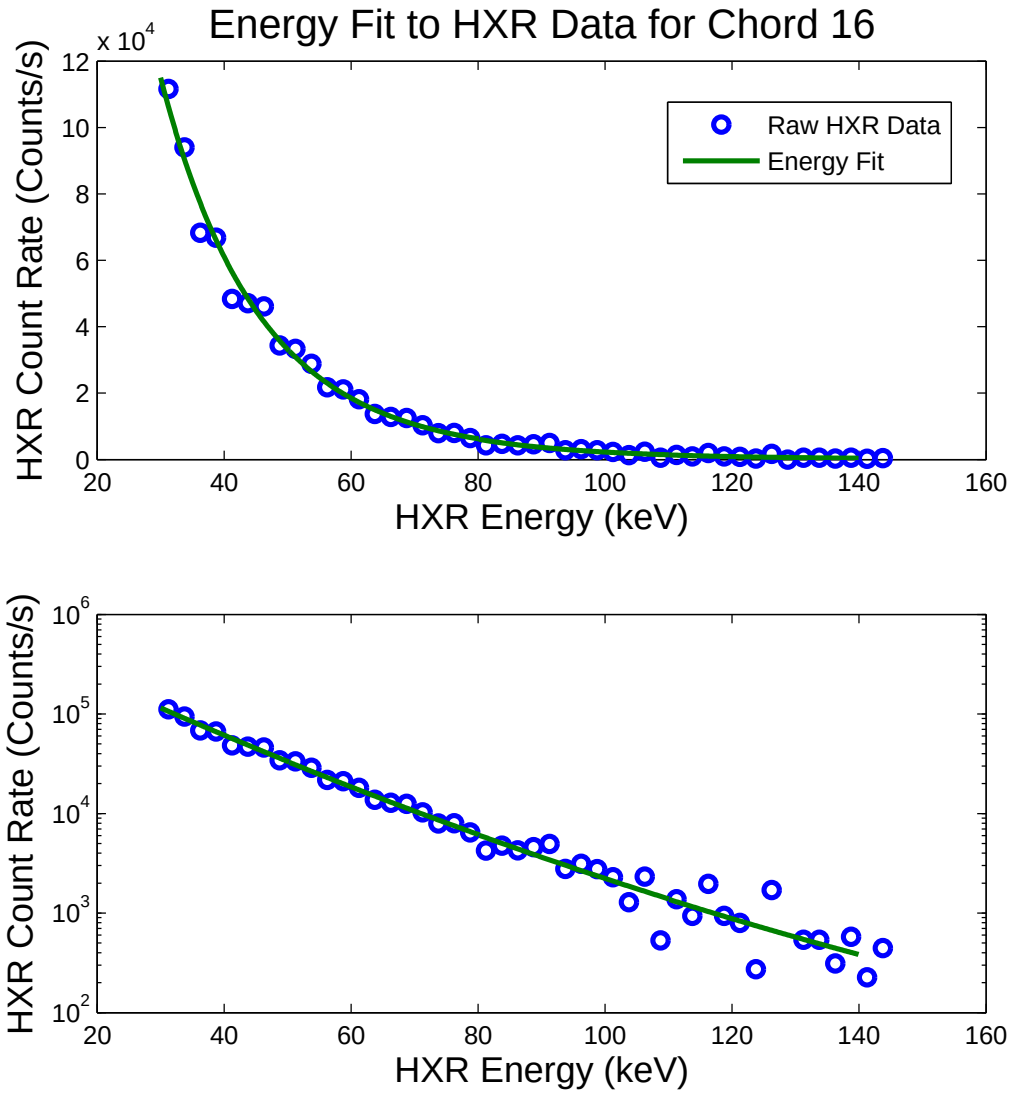

Figure 3-1: Sample energy fit to HXR data (top) on a linear scale and (bottom) on a log scale.

It is clear from the dependence of x-ray flux on photon energy that a curve must be fit to the log of the data, since the x-ray intensity decreases as a function of energy in an exponential-like fashion. However, it is not a pure exponential decrease and a linear fit through log data would not be sufficient to describe the data. In practice, a second order polynomial is usually fit to the log data, which allows for a slight curvature to the log data.

Sometimes a curve fit to the HXR data describes most of the data well, but does a poor job at the endpoints (the highest and lowest energies being fit). When energy fit data are used in place of raw data for further analysis, energies at the endpoints of the energy fit may be excluded from the analysis.

Figure 3-2 shows a series of HXR profiles before and after energy fitting (i.e. the raw count rates and the energy-fit count rates). Each profile corresponds to a different time slice in a lower hybrid power modulation experiment (discussed in more detail in Chapter 6). Notice how the energy fitted data points show much less scatter both from time bin to time bin and from chord to chord than the original data points, even though the data was fit in the energy dimension only. This is particularly noticeable for the time bins with the lowest count rates. Energy fitting can be a useful way to enable spatial profile inversions of low signal-to-noise data while preserving the small time or energy bin desired.

3.0.2 Spatial Inversion

It is often desirable to spatially invert HXR profiles to investigate the radial location of the fast electrons in the plasma. The spatial inversion routine divides the plasma into a given number of flux bins, which are spaced equally in toroidal flux space. In real space, they are closer together towards the edge and farther apart in the core. Few HXR chords go through the core due to its small volume, so it is not possible to make high spatial-resolution measurements of emissivity there. Furthermore, calculations of emissivity in the core are highly sensitive to noise because they depend on the (inferred) emissivity in the outer part of the plasma. A matrix describing the contribution of each flux surface to each chord is constructed, which is the basis for

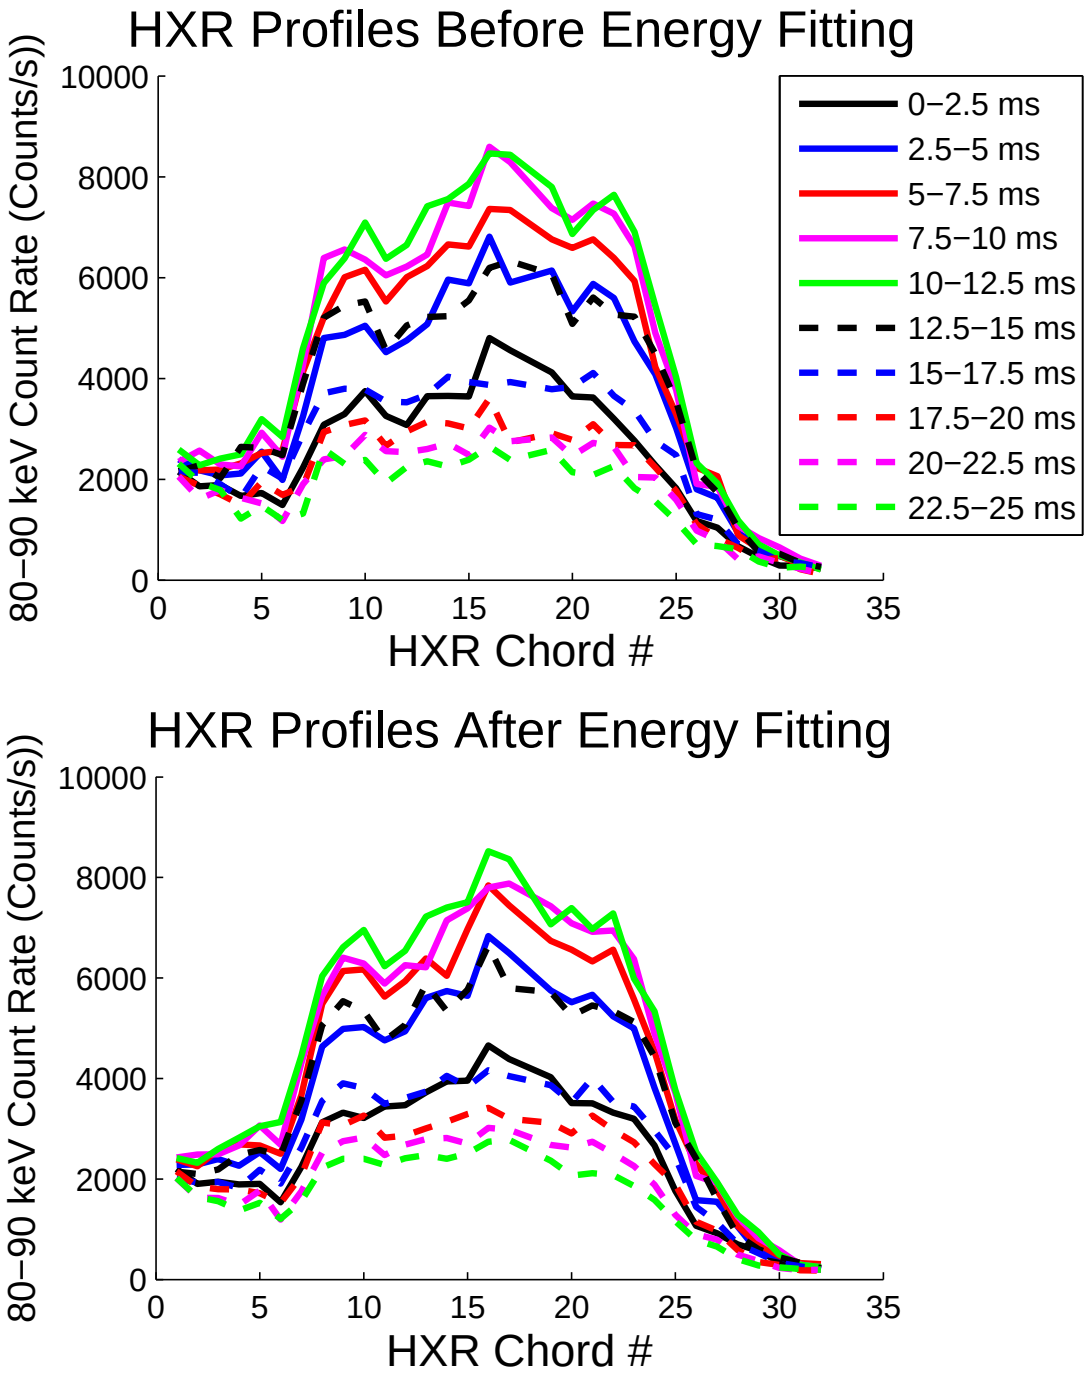

Figure 3-2: HXR profiles before (raw count rates) and after (energy-fit count rates) energy fitting. Each profile corresponds to a different time slice of the experiment. Note how the scatter in both time and space is reduced when data is fit in the energy dimension only. In this experiment, 90° waveguide phasing ($n_{||} = 2.3$) was used, Line-integrated density was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. Plasma current was $I_p = 800$ kA.

the inversion.

The outermost chords are sometimes excluded from the inversion, either because their signal is too low to be fit to an energy curve or because they are indicating high emission near the divertor region that is not thought to be originating from fast electrons inside the separatrix (this divertor region emission is discussed in more detail in Chapter 5). When outer chords are excluded from the inversion, the emissivity is constrained to vanish at the last closed flux surface.

Inverting the HXR profiles requires making some simplifying assumptions about the emission. The first assumption is that the electron distribution function is a flux function, which is a reasonable approximation, though drift orbits actually cause the distribution function to deviate along a flux surface. This assumption is not unique to HXR inversions and is usually needed to invert other integrated chordal measurements.

The second and more sweeping assumption involves the relationship between x-ray emission and viewing angle with respect to magnetic field. Recall that the HXR camera views the plasma perpendicularly. However, the poloidal field points away from the camera on the top half of the plasma and towards the camera on the bottom half of the plasma. Thus co-current LH-driven electrons are angled toward the camera on the top half of the plasma and away from the camera on the bottom half of the plasma. Since the Bremsstrahlung cross-sections are greatest in the direction of the velocity of the electron, the change in pitch angle from the top to the bottom of the plasma creates an up-down asymmetry in the HXR profiles.

Although the cross-sections as a function of angle are well known, they also depend on electron energy. Since we do not know the energy of the electrons contributing to a given HXR photon energy, it is not possible to predict the magnitude of the asymmetry a priori. However, it is experimentally clear that the up-down asymmetry increases with HXR photon energy, as expected. The HXR inversion makes the assumption that emission from a given point in the plasma is proportional to some generalized flux-function emissivity plus a perturbation which is proportional to the deviation in viewing angle from perpendicular to the magnetic field, as shown in

Eq. 3.1:

$$\frac{dn_k(\mathbf{x})}{dt\,dk\,d\Omega} = e(r)\left[1+\beta\left(\theta-\frac{\pi}{2}\right)\right], \tag{3.1}$$

where $\frac{dn_k(\mathbf{x})}{dt\,dk\,d\Omega}$ is the total photon density at position vector $\mathbf{x}$, $e(r)$ is the flux-surface emissivity that we wish to solve for, β is an adjustable parameter that controls the asymmetry of the reconstructed integrated HXR profiles, and θ is the angle between the local magnetic field and the HXR viewing chord.

The inversion is done using a least-squares fit that finds a radial emissivity profile that reproduces the experimental measurements of emissivity integrated along a chord with minimal error. In an overdetermined least-squares fit, there are more raw data points being fit to than there are fitting parameters. In this case, the fit parameters are radial locations where we wish to solve for emissivity. However, even if all 32 chords are included in the inversion, the spatial resolution of the emissivity profile would be very poor (it could be solved for at a maximum of 32 radial locations). Furthermore, noise in the integrated data is amplified in the inversion.

Figure 3-3 shows an example of such a simple inversion for two sample profiles for 3 different values of n_{flux}, the number of flux surfaces in the plasma where emissivity is solved for. These profiles were taken from the Fourier transformed data from the lower hybrid power modulation experiments, discussed in Chapter 6[1]. Note that the spatial resolution is poor and even so, the inversion is quite noisy. There is a trade-off between spatial resolution and noise reduction in such an inversion. For lower count rate signals, the noise becomes even worse.

One way to evaluate the quality of an inversion is to reverse the inversion and reconstruct the profile that was originally inverted from the inverted profile. Figure 3-4 shows such a reconstruction for the 3 inversions in Figure 3-3. Note that the higher the value of n_{flux}, the more closely the reconstructed profile follows the actual data. However, some features of the actual data that are being followed, are due to noise or

[1]The data was Fourier transformed in time in order to be used to solve for transport coefficients. However, the fact that the data has been Fourier transformed has no effect on the inversion routine. For the purposes of this discussion, the Fourier transformed profiles can simply be considered to be ordinary line-integrated HXR profiles.

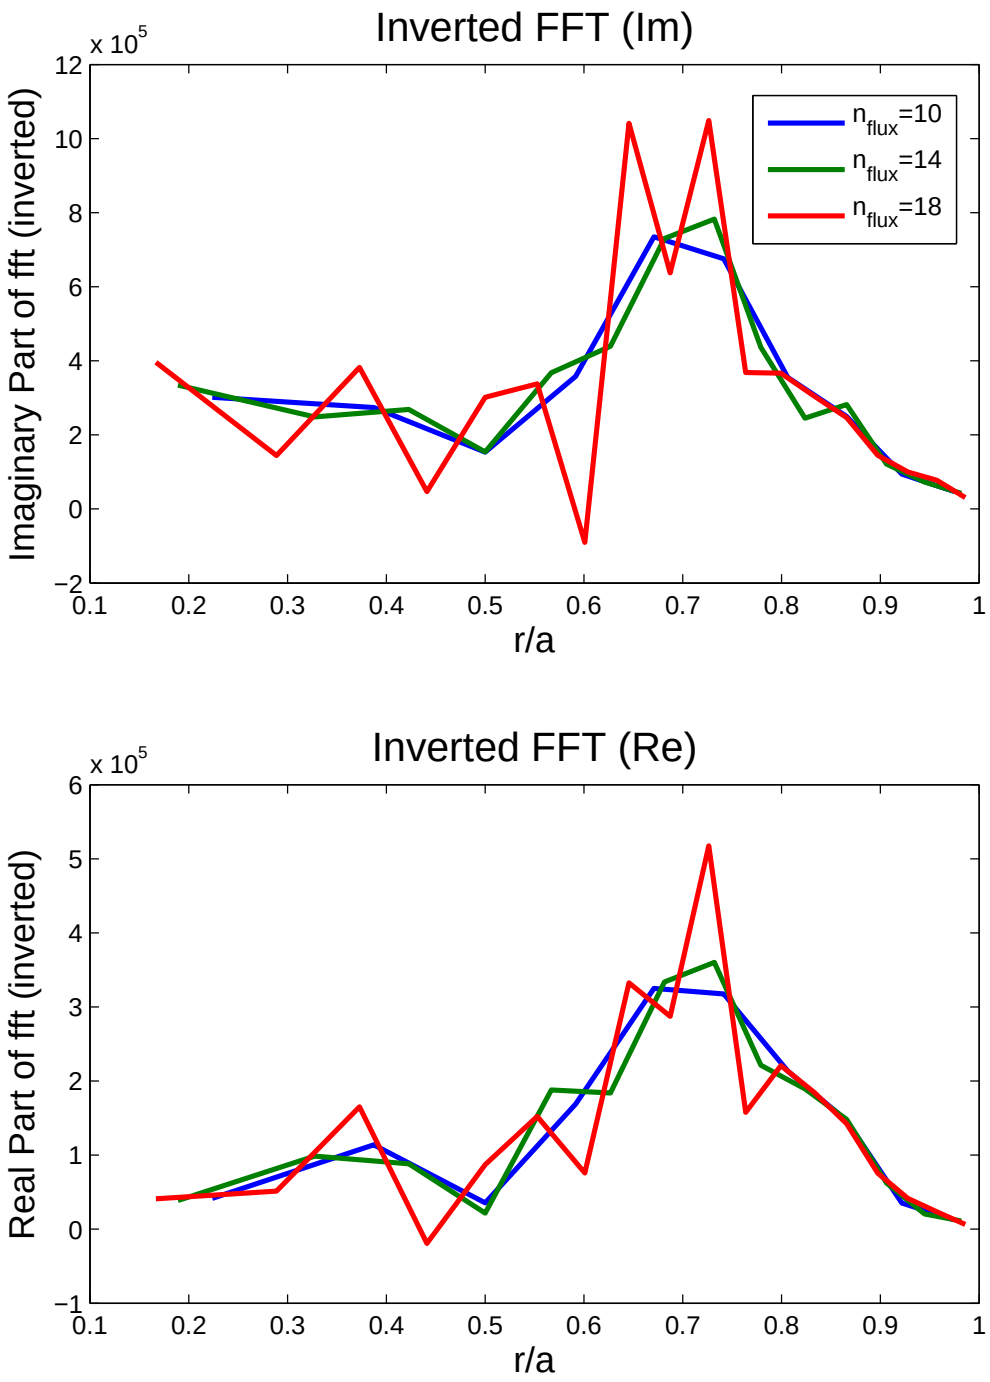

Figure 3-3: Example of a simple inversion of sample HXR profiles for 3 different values of n_{flux}, the number of flux surfaces in the plasma where emissivity is solved for.

chord-to-chord calibration errors and it is not necessarily a desirable result that the reconstructed profiles reproduce these features. Regardless of what type of inversion method is eventually chosen, there is always a trade-off between following the non-inverted data closely and finding an inversion which does not contain "features" that are probably measurement artifacts.

There are multiple ways to solve the problem of low spatial resolution and noise amplification. One is to force the least-squares solution to be a sum over a smooth function basis set, such as Bessel functions or polynomials, where the coefficients of the basis functions are the parameters being solved for in the least-squares fit. The problem with this approach is that it forces the solution to be expanded in the chosen basis set, which could be a very unrealistic model, depending on the shape of the emissivity profile. Furthermore, the number of terms included in the expansion is discrete and arbitrarily chosen. A 5^{th} order polynomial solution may look very different in certain locations from a 6^{th} order polynomial solution and it is rarely obvious which one is better.

Figure 3-5 shows a sample inversion using polynomial functions of a flux variable as a basis set. Note how more detail gets added in the inversion for higher order polynomials, but it is unclear what order of polynomial is the optimal one to use. In certain locations in the plasma, particularly the edge and the center, the profile is highly dependent on the order of polynomial used and does not appear to converge to any particular value. Other features, such as the peak location at r/a$\approx$0.7, also do not appear to converge. These same problems arise with sine/cosine and Bessel function basis functions as well.

One might hope that the reconstructed profiles would shed some light on which order of polynomial is best to use. The reconstructed profiles for the inversions in Figure 3-5 are shown in Figure 3-6. Unfortunately the reconstructed profiles give very little information about which order of polynomial is "best" and it is unclear whether any of them are particularly good. For example, note how the 7^{th} order polynomial and the 10^{th} order polynomial produce fairly different inverted profiles but very similar reconstructed profiles. It may be that expanding the solution in a

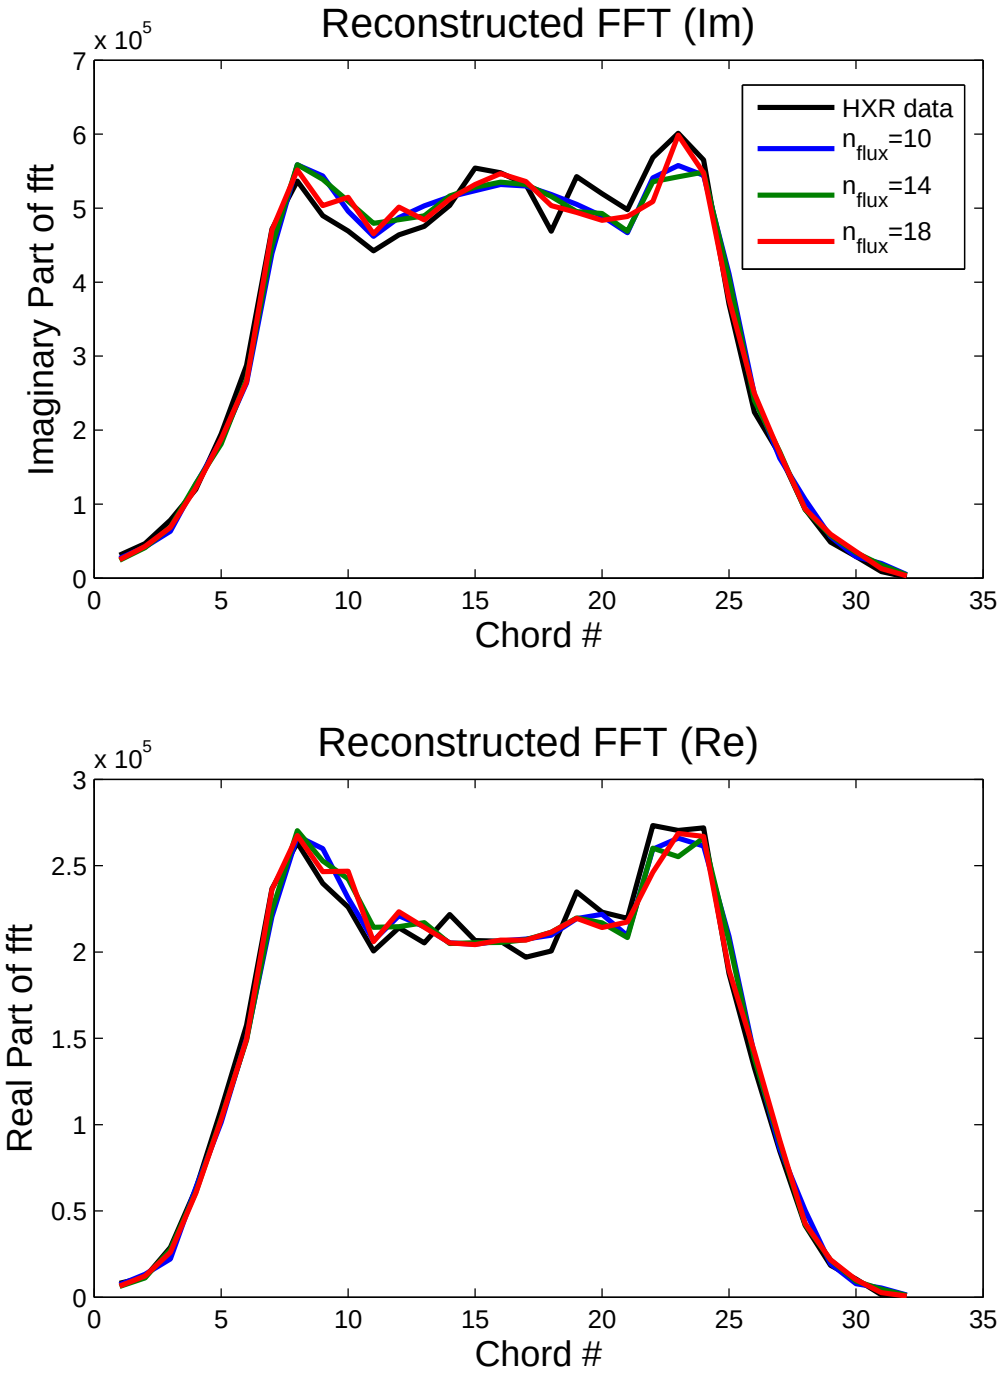

Figure 3-4: Reconstructed HXR profiles for the simple inversions shown in Figure 3-3.

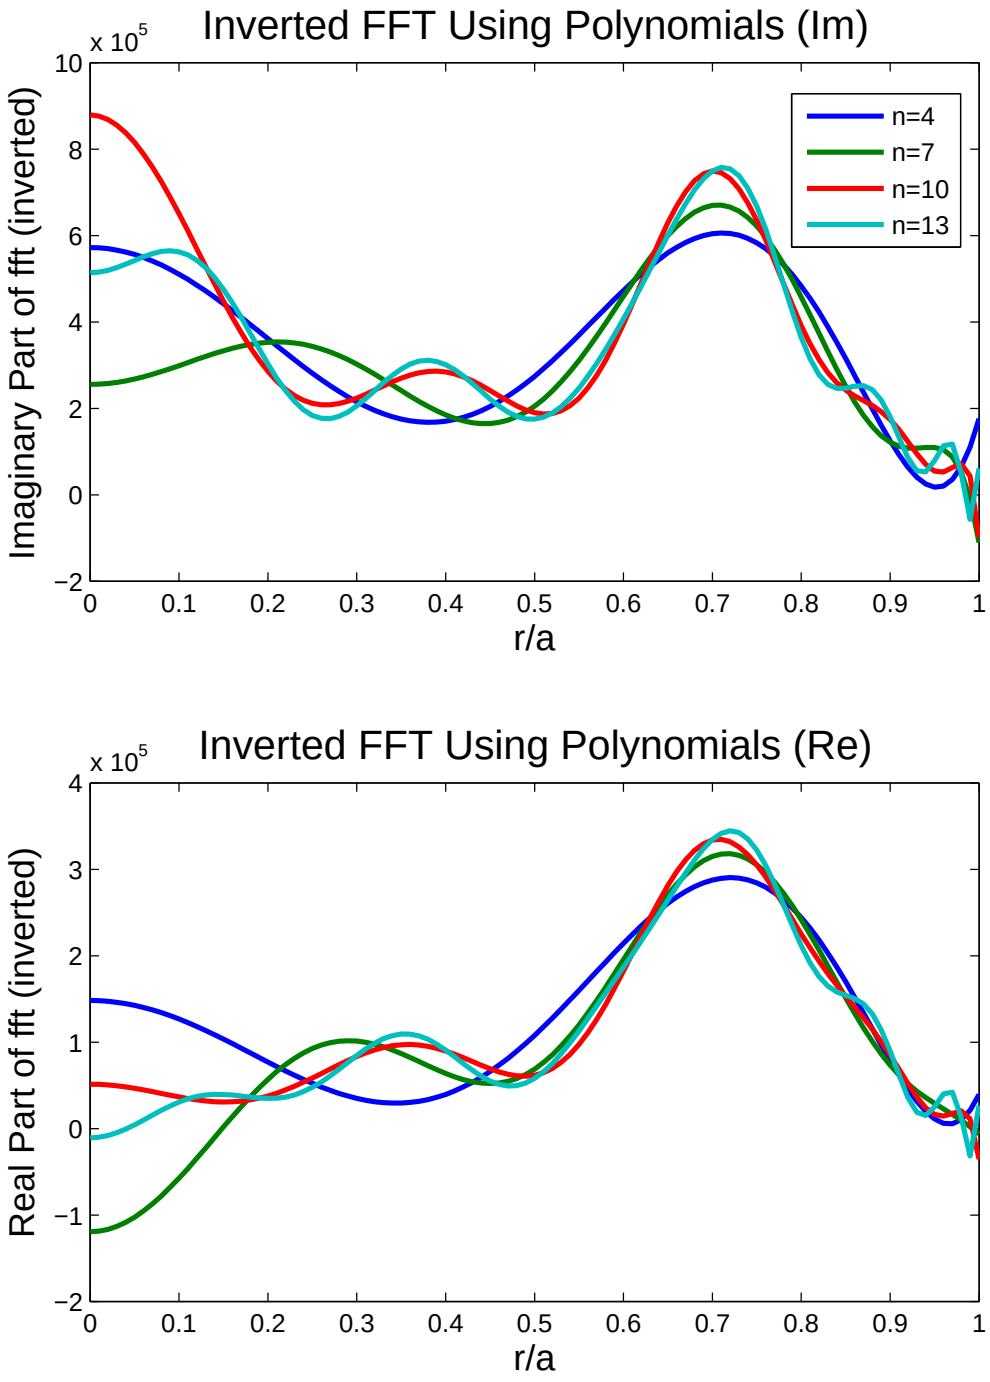

Figure 3-5: Example of an inversion using polynomials as a basis set. The legend indicates the order of the polynomial used.

set of basis functions is simply the wrong approach for this particular problem.

A better approach to this problem is to use regularization techniques, which in effect force the solution to be smooth, and allow the user to finely and continuously control the degree to which smoothness is imposed. The idea behind the most common form of regularization, Tikhonov regularization [57, 58], is to minimize the sum of the least-squares error and another linear operator on the solution, which is often a form of the second derivative of the solution. By requiring that the second derivative of the solution remain small, one forces the solution to be smooth to some degree. For a profile that is expected to be relatively continuous, such as x-ray emissivity, this adds another piece of information to the problem which results in much more noise-insensitive inversions and allows the user to use finer spatial resolution in the solution. In fact, the spatial resolution can be dialed up arbitrarily high, provided that the degree of regularization is adjusted accordingly. This does not mean that the inversion routine is necessarily resolving features that are as narrow as the resolution grid, but it does give the best guess for an arbitrarily fine grid based on neighboring points and an assumption of continuity.

The traditional least-squares problem begins with the assumption that the linear transformation of some vector of variables with unknown values, $\mathbf{x}$, is a close approximation to a vector of measured values, $\mathbf{y}$, within some error, ϵ:

$$A\mathbf{x} = \mathbf{y} + \epsilon. \tag{3.2}$$

For the HXR inversions, $\mathbf{y}$ is the measured profile, $\mathbf{x}$ is the inverted profile, and A is a weighting matrix that represents the contribution of each flux surface's emissivity to each integrated chord.

$$\overset{min}{\mathbf{x}} ||\epsilon||^2 = \overset{min}{\mathbf{x}} ||A\mathbf{x} - \mathbf{y}||^2, \tag{3.3}$$

$$\mathbf{x_{opt}} = (A^T A)^{-1} A^T \mathbf{y}. \tag{3.4}$$

If we wish to minimize ϵ as in Eq. 3.3, which is the same as finding the inverted

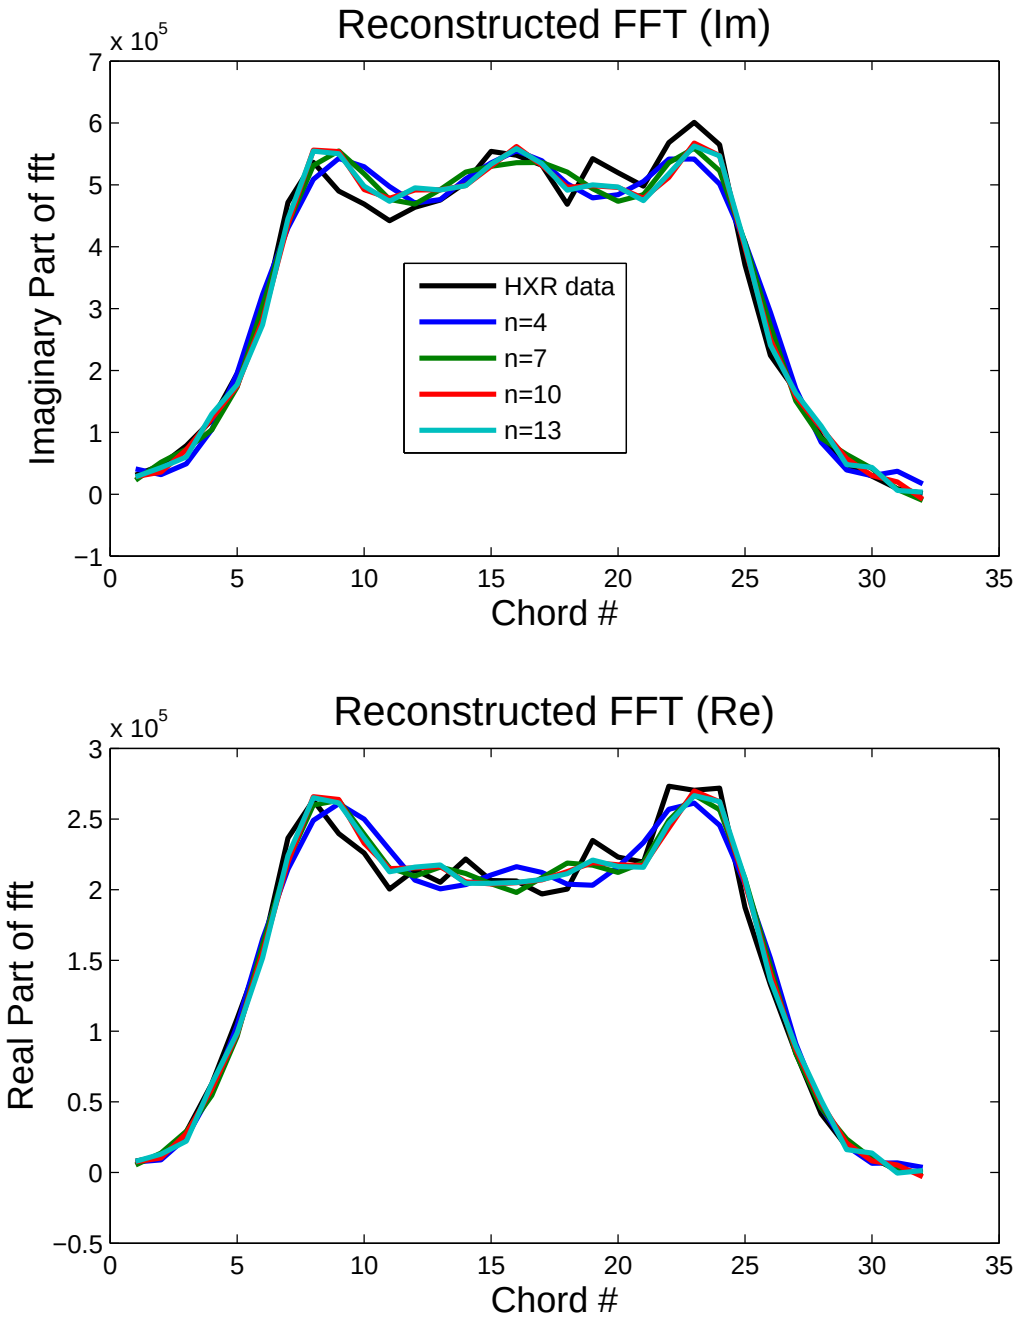

Figure 3-6: Reconstructed HXR profiles for the inversions in Figure 3-5.

profile that most closely reconstructs the integrated chord measurements, we use the simple formula in Eq. 3.4 to find the optimal value of $\mathbf{x}$. The double vertical line notation "$\|$" indicates the norm of the vector enclosed on either side by it.

However, we already know that in the case of an inversion, the solution which reconstructs most closely the integrated chord measurements is not necessarily the optimal solution. So instead of trying to minimize ϵ alone, we can try to minimize the sum of ϵ and another linear operator on $\mathbf{x}$:

$$\overset{min}{\mathbf{x}} \{||\epsilon||^2 + \alpha||L\mathbf{x}||^2\} = \overset{min}{\mathbf{x}} \{||A\mathbf{x} - \mathbf{y}||^2 + \alpha||L\mathbf{x}||^2\}. \tag{3.5}$$

In Eq. 3.5, L is a linear operator acting on $\mathbf{x}$, in this case a numerical second derivative of $\mathbf{x}$. The parameter α is called the regularization parameter and the magnitude of α controls the relative emphasis put on minimizing ϵ versus minimizing the second derivative of the solution. The solution to this minimization problem is:

$$\mathbf{x_{opt}} = (A^T A + \alpha L^T L)^{-1} A^T \mathbf{y}. \tag{3.6}$$

It is clear from Eq. 3.6 that Tikhonov regularization is not difficult or computationally intensive to perform. The most challenging part is choosing an appropriate value for α.

Figure 3-7 shows an example of an inversion performed with regularization, using three different values of α. High values of α produce inversions that are nearly flat and featureless (an example is shown in the red curve). Low values of α produce inversions that contain features which may be due to the noise in the integrated data (an example is shown in the blue curve). Some of the bumps and wiggles in the blue trace span a normalized radius of 0.05, approximately 1 cm for C-Mod, which is the spatial resolution of the HXR camera. Clearly these features in the inverted data should not be resolvable based on the spatial resolution of the instrument, and therefore are an artifact of the inversion. The green trace is an example of an inversion with an intermediate value of regularization parameter chosen to obtain a smooth profile while still resolving the main peak.

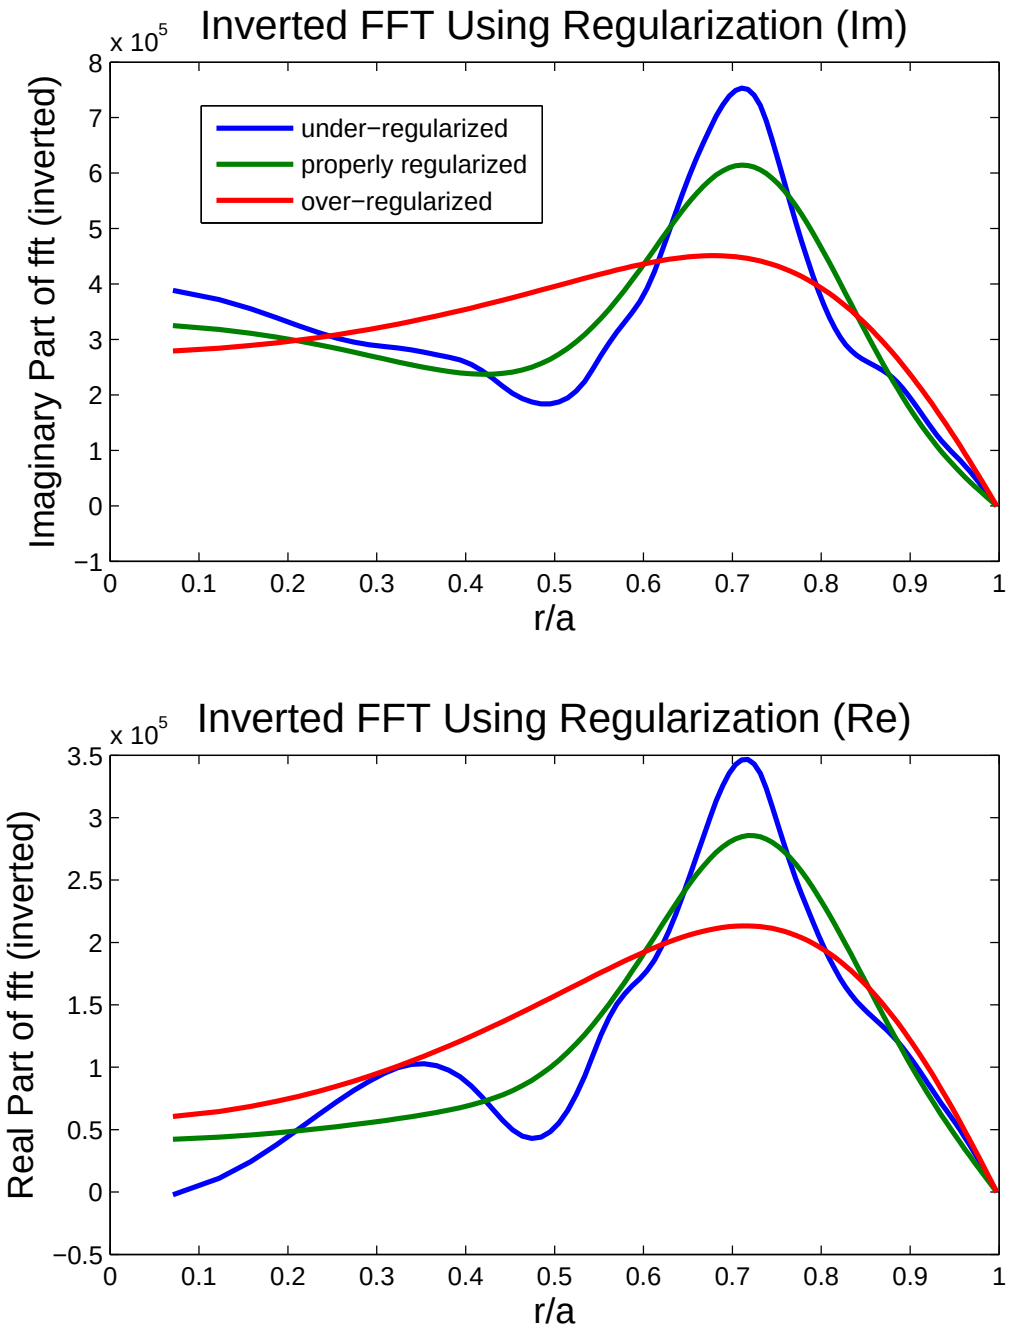

Figure 3-7: Example of an inversion using regularization techniques, Three different regularization parameters have been chosen to demonstrate over- and under-regularization as well as proper regularization.

Figure 3-8 shows reconstructed HXR profiles for the regularized inversion in Figure 3-7. Note how the degree of regularization needed can also be partially judged from this plot. The red trace is obviously over-regularized, since its reconstructed HXR profile does not follow the actual data well at all. The blue trace is not very obviously under-regularized, but the blue inversion is definitely reconstructing features of the HXR profile that could be attributed to noise. The green trace represents a good compromise between the two, following the main features of the HXR profile, but not conforming to the little details that may be attributed to noise.

Choosing the correct regularization parameter, which is the number that controls how much emphasis is put on smoothing vs. minimizing least-squares error, is a very complex topic. The HXR inversions find an optimal regularization parameter using a method called "generalized cross-validation" (GCV) [59]. The GCV technique evaluates how well a particular regularization parameter can reproduce an individual data point when that point is not included in the inversion. An inaccurate outlying data point is not typically well reproduced in a reconstruction (regardless of the regularization parameter) if it is not included in the inversion. However, an accurate data point should be well reproduced in a reconstruction even if it is excluded from the inversion because in principal it is consistent with the rest of the data.

Generalized cross-validation picks the regularization parameter which best reproduces removed data points from the inversion based on remaining data points. Using this method, under-regularization that follows anomalous features that are likely due to measurement error is avoided because an anomalous feature is not typically reproduced if the data point exhibiting it is removed. Over-regularization is also avoided because an inversion that is over-regularized does not reconstruct the data well.

Unfortunately, the GCV objective function to be optimized does not always have a local minimum with respect to regularization parameter. This is more often a problem when signal is low and noisy, causing several of the points to be outliers. In these cases, a regularization parameter can be chosen subjectively, by looking at the inverted and reconstructed profiles with different degrees of regularization. An alternate approach is to find regularization parameters for neighboring (in time or

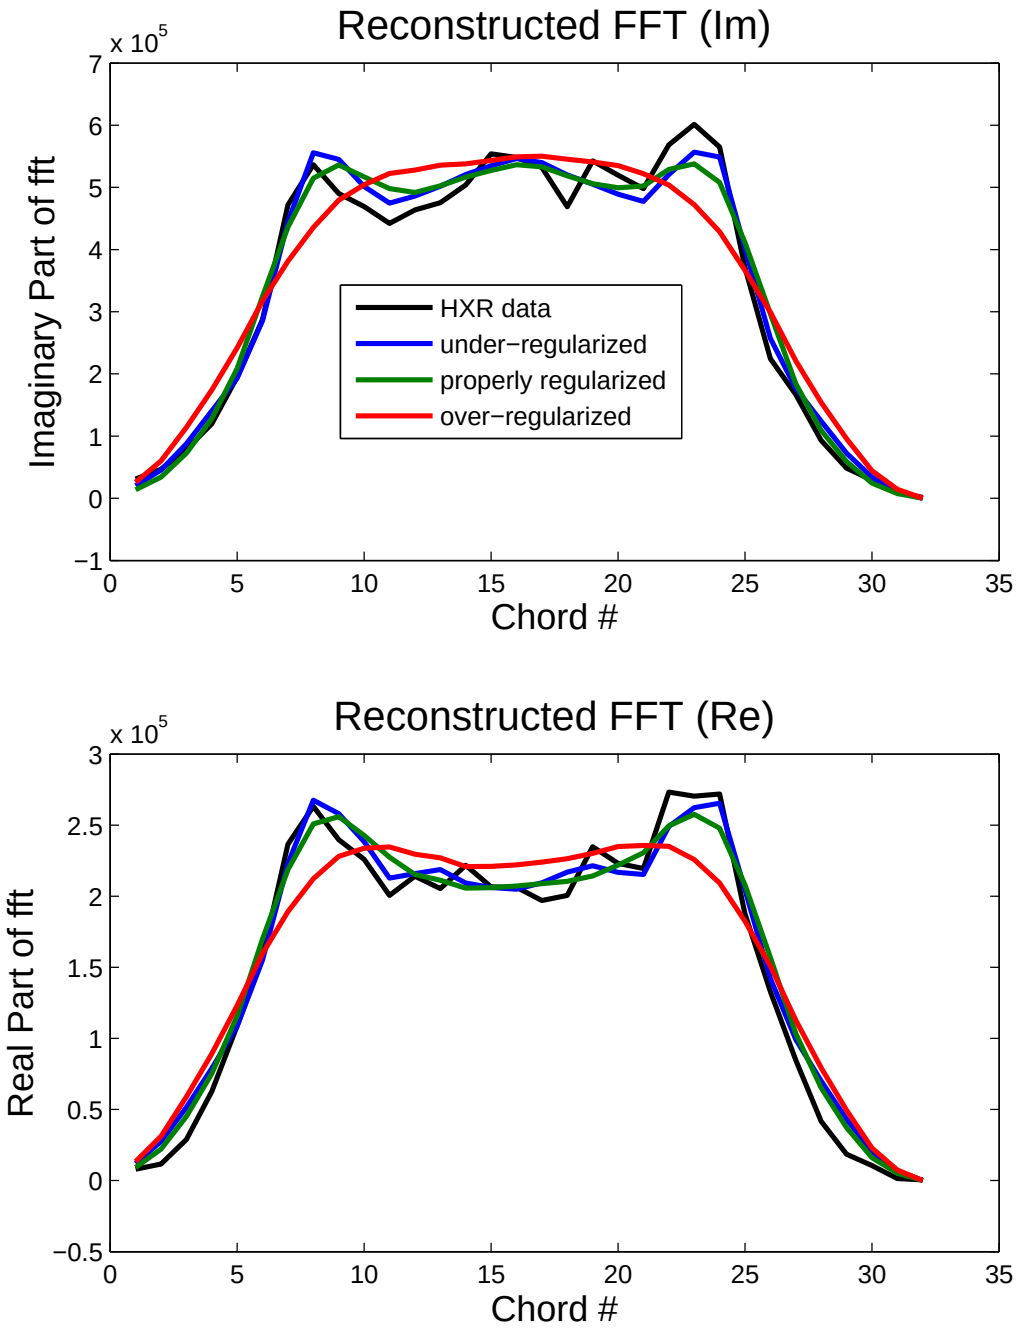

Figure 3-8: Reconstructed HXR profiles for the inversions in Figure 3-7.

energy) profiles using GCV and to use an intermediate parameter for the difficult-to-invert profile. Some extremely noisy profiles do not contain enough information to be inverted reliably regardless of the method used.

Regularized HXR profile inversions shown in this thesis were performed using an extensive package of regularization tools developed in Matlab [59].

Chapter 4

Modeling Tools for LHCD Scenarios on Alcator C-Mod

It is often desirable to model the plasma response to LHCD to compare with experimental measurements. The modeling presented in this thesis had many purposes, chief among them:

1. To validate the model being used.

2. To simulate plasma scenarios that we could not create experimentally at the time of the modeling.

3. To guide the direction of experiments that we might want to perform in the future.

4. To look in depth at (predicted) plasma properties that could not be directly obtained experimentally, such as the electron distribution function or location of power damping.

The modeling was performed on the PSFC computing clusters Marshall and Loki. Marshall had 24 nodes with 2 processors each, for a total of 48 processors, while Loki has 75 nodes with 8 processors each, for a total of 600 processors.

4.1 Ray Tracing and Full Wave Solvers

There are two major components to modeling current drive in a plasma. The first component is a calculation of the electric field due to the LH wave inside the plasma, which is used in the quasilinear diffusion coefficient to quantify momentum transfer from the wave to the particles. This calculation is typically performed either by ray tracing or full wave solution techniques, as discussed below.

The second component is the calculation of the electron distribution function that balances the momentum input from the wave and the ohmic electric field with the effects of particle collisions. This calculation is typically done using a code which employs a Fokker-Planck collision operator [60]. Since the electron distribution function affects the wave damping, the two calculations must be iterated several times, beginning with a Maxwellian plasma, to converge on a steady-state distribution function.

It can be computationally intensive and very complicated to calculate the electric field due to the wave everywhere in the plasma. The traditional method used in the past and still in use today is a geometrical optics calculation [61, 62] to trace rays through a plasma via a dispersion relation. In 3D ray tracing in a toroidal geometry [22], rays are launched at a variety of locations and initial vector wave numbers in order to simulate the RF power coupled into the plasma. The ray paths and vector wave numbers are allowed to evolve in space along with the fraction of power left in the ray as it traverses its path in the plasma. Rays are stopped when they damp to a sufficiently small fraction of their initial power or when they reach a user-specified maximum path length. Figures 4-1 and 4-2 show poloidal and toroidal cross sections of a set of LH rays traced in the C-Mod plasma using the Genray [63, 64] code.

Another more accurate method of calculating the electric field in the plasma is the so-called full wave approach, which involves solving Maxwell's equations everywhere inside the plasma, subject to boundary conditions [65, 66]. This is a far more computationally intensive calculation than ray tracing. Until recently, full wave codes had been used only to model RF sources of much longer wavelength than lower hybrid

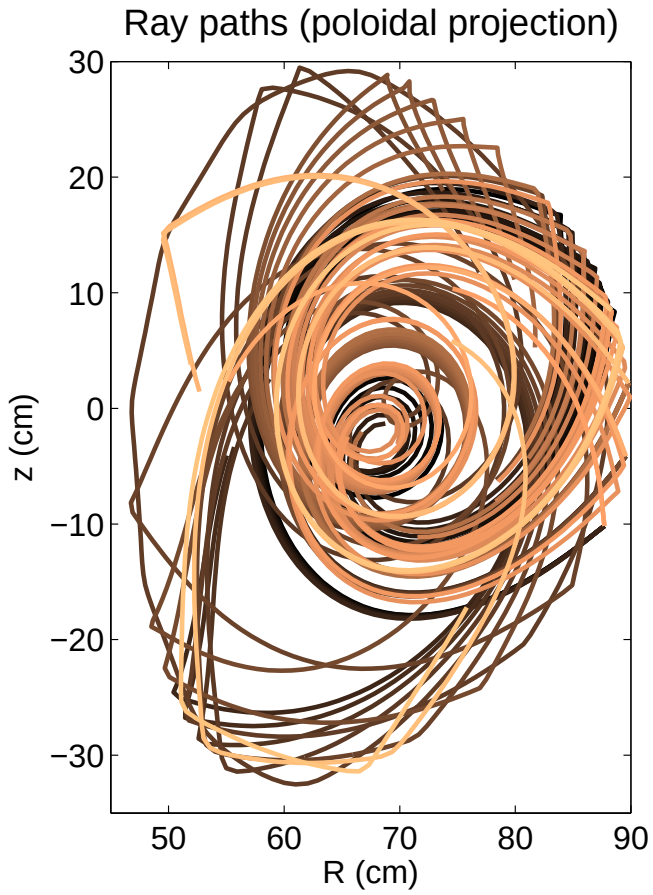

Figure 4-1: Example of rays traced with Genray [63, 64] code in simulated C-Mod nonthermal plasma (poloidal projection). These rays are taken from a simulation presented in Chapter 6. Only 1 out of 5 rays used in the simulation is shown here (different colors represent different rays). In this simulation, plasma parameters were B=5.4 T, T_{e0}=2.3 keV, $\bar{n}_e = 9 \times 10^{19}$ m^{-3}, and I_p=800 kA. A peak launched $n_{||}$=3.1 was used for the forward spectrum and $n_{||}$=6.2 for the reverse spectrum.

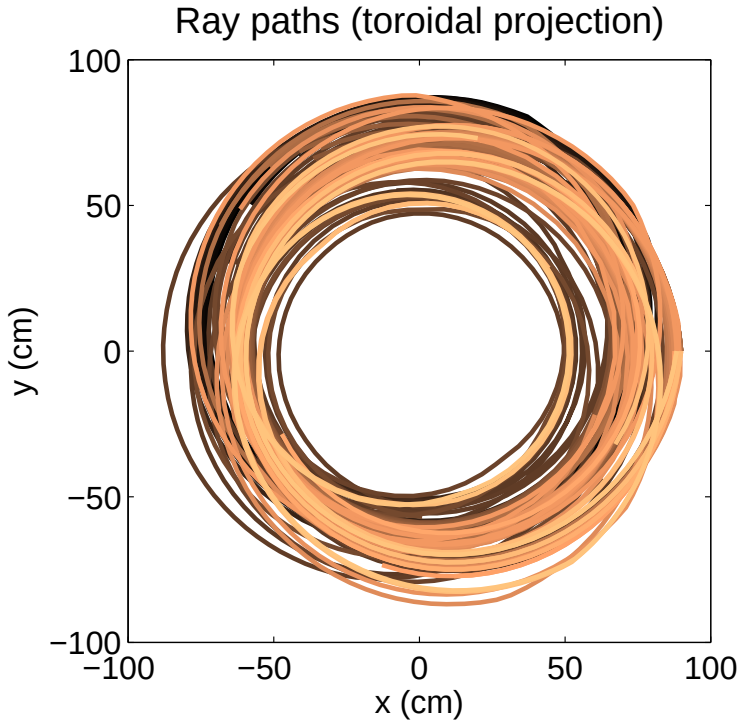

Figure 4-2: Example of rays traced with Genray [63, 64] code in simulated C-Mod nonthermal plasma (toroidal projection). These rays are taken from a simulation presented in Chapter 6. Only 1 out of 5 rays used in the simulation is shown here (different colors represent different rays). In this simulation, plasma parameters were B=5.4 T, T_{e0}=2.3 keV, $\bar{n}_e = 9 \times 10^{19}$ m^{-3}, and I_p=800 kA. A peak launched $n_{||}$=3.1 was used for the forward spectrum and $n_{||}$=6.2 for the reverse spectrum.

waves because the spatial resolution needed to resolve such waves was much coarser and required less computing power.

In recent years, increasing computing power has enabled a large effort to use full-wave codes to resolve the electric fields of the much smaller wavelength lower hybrid waves. Two major modeling projects at the PSFC are aimed at integrating full-wave solvers with Fokker-Planck codes for the purpose of modeling lower hybrid waves. One effort, described in Ref. [65], was to modify a full-wave solver called TORIC [67] to simulate LH waves and to couple it with a bounce-averaged Fokker-Planck code called CQL3D [68]. TORIC uses spectral decomposition in the poloidal dimension and finite elements in the radial dimension to represent EM fields and was originally designed to model ICRF waves. A sample electric field calculation performed using TORIC for LH waves is shown in Figure 4-3.

The second effort, described in Ref. [66], was to write a finite-element full-wave solver called LHEAF. LHEAF employs the COMSOL Multiphysics software and can simultaneously model the LH wave fields from the coupler in the scrape off layer out to the vacuum vessel, and into the plasma core. An example electric field calculation performed in LHEAF is shown in Figure 4-4.

The major advantage of ray tracing over the full wave approach is the significant reduction in computing requirements. Thus ray tracing is preferable in plasma scenarios where ray tracing and full wave are expected to agree. It is also the only practical way to model time-dependent scenarios which require many iterations between the quasi-linear diffusion solver and the Fokker-Planck solver.

The weakness of ray tracing lies in the accuracy of calculating the electric fields in the plasma. The first reason why ray tracing may be inaccurate is that it does not take into account focusing and diffraction of the wave. It is not currently known whether diffraction plays a large role in wave damping location, but diffraction effects are not included in ray tracing physics. It is certainly true that focusing can affect the intensity of the wave fields used to evaluate the quasilinear diffusion coefficient.

A second important reason that ray tracing can be inaccurate is related to how wave propagation is handled during ray reflections in weak single pass absorption

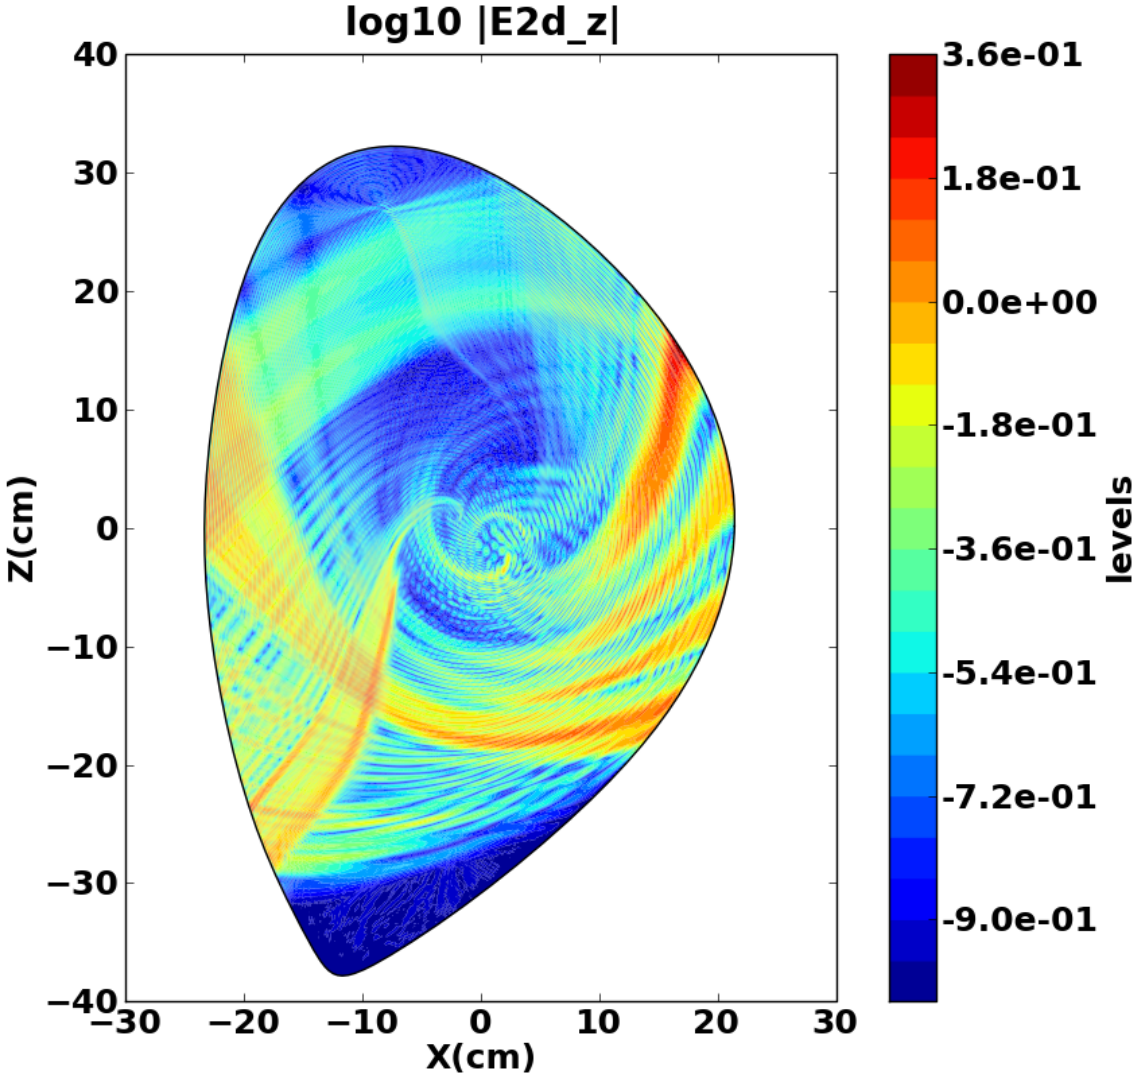

Figure 4-3: Example of electric field calculation in Alcator C-Mod geometry from ToricLH [65]. The log of the parallel electric field from a launched wave ($n_{||} = 1.55$) is shown here in arbitrary units. In this simulation, 1023 poloidal modes and 980 radial modes were used. This was a simulation of shot 1060728011, a nearly fully non-inductive discharge. Plasma parameters in the simulation are $T_{e0} = 2.3$ keV, $n_{e0} = 7 \times 10^{19}$ m^{-3}, $B = 5.4$ T, and $I_p = 540$ kA. Four iterations were performed with CQL3D, yielding a self-consistent distribution function. Figure courtesy of John Wright and Paul Bonoli.

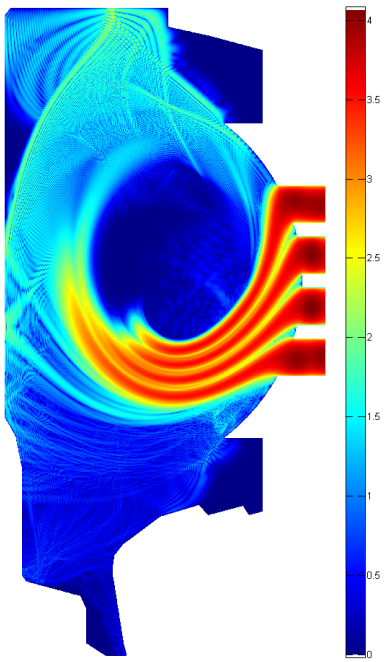

Figure 4-4: Example electric field calculation from LHEAF. The log of the electric field from a launched wave ($n_{||} = 1.9$) is shown here in arbitrary units. Because LHEAF uses finite elements, it can solve for the electric field inside the launcher and through the entire vacuum vessel. The C-Mod geometry is modeled here for shot 1080320017. The plasma parameters used were $n_{e0} = 8 \times 10^{19}$ m^{-3}, $B = 5.4$ T, and $I_p = 800$ kA. The simulated plasma had $T_{e0} = 4$ keV with a small high-energy seed population of electrons. No iteration was performed with a Fokker-Planck code in this case. Figure courtesy of Orso Meneghini and Syun'ichi Shiraiwa.

regimes. Lower hybrid waves do not propagate at densities that are below the lower hybrid cutoff density. When the rays reach the location at the edge of the plasma where the density is such that $\omega_{pe} = 4.6$ GHz, they reflect back into the plasma. In reality some of the wave power can tunnel beyond the cut-off into the evanescent region and this effect is only described accurately by a full-wave analysis. In the cutoff region, the WKB approximation made in ray tracing breaks down and in this respect full-wave simulations may be able to eventually provide also a more accurate description of the wave physics at cut-off.

Furthermore there can be numerical challenges with integrating rays through these cutoff reflections because the perpendicular wavenumber is changing rapidly as the cut-off is approached. As a result, the predictions from a ray tracing model can be highly sensitive to plasma conditions near the cutoff. Because the full wave approach is fundamentally a boundary value problem, it is not expected that its predictions would depend strongly on edge conditions, but a thorough side by side comparison of the two approaches has not been performed to date.

Numerical integration near cutoffs is not a problem for ray tracing in all scenarios. We have found while modeling scenarios of Alcator C-Mod that the sensitivity to edge conditions is weak for scenarios with strong absorption. Quasilinear theory predicts that LH waves will damp when their phase velocity is 2.5 times the electron thermal velocity [69, 28], where $v_{th} = \sqrt{2T_e/m_e}$. This requires that LH waves penetrate into the plasma until they reach a sufficiently high temperature to damp, about $T_e(kev) = 40/n_{||}^2$. Thus in a high single-pass absorption regime, theory predicts lower $n_{||}$ waves to damp farther in towards the core, where the plasma temperature is higher.

C-Mod plasmas are typically in the multi-pass absorption regime, meaning that rays do not encounter a temperature of $T_e(kev) = 40/n_{||}^2$ on their first pass through the plasma. Instead, LH waves will reflect from the vacuum vessel walls or a cut-off layer at least once before they damp. As the LH wave propagates in the plasma, the $n_{||}$ of the wave can vary due to the variation in toroidal and poloidal magnetic fields and the wave does not damp until $n_{||}$ increases enough relative to the local electron temperature.

Ray-tracing simulations that require a single reflection per ray for the power to damp (either through a high enough temperature or a high enough $n_{||}$) do not show much sensitivity to edge conditions. On Alcator C-Mod, we find that for $n_{||}$=3.1 discharges, the damping is strong enough for ray tracing models to work well.

In modeling high temperature tokamaks like ITER, rays are expected damp on the first pass through the plasma and therefore ray tracing should be valid[1]. Indeed the use of full-wave solvers to simulate LH wave propagation in ITER will be extremely challenging because its large linear dimensions relative to an LH wavelength would make the electric field calculation very large (computationally).

Even in cases where there is weak single pass damping, it is possible that ray tracing could be a valid approach if done properly. Specific numerical integration methods could potentially alleviate the integration problem. The inclusion of a scrape-off layer in the model[2] might allow for gently sloping profiles that would render ray tracing far less sensitive to changes in the plasma edge conditions. As the full wave codes being developed today become standard tools, it will become increasingly clear how their predictions differ (if at all) from the ray tracing approach.

4.2 Introduction to Genray-CQL3D

All simulations shown in Chapters 5 and 6 were performed using the modeling package Genray-CQL3D. Genray [63, 64] is a generalized 3D optical ray-tracing code that can be used to trace LH rays as they propagate through a given plasma.

CQL3D [68] is a bounce-averaged Fokker-Planck code. It takes as an input the flux-surface-averaged quasi-linear diffusion coefficient from Genray[3] and calculates the

[1]It is expected in ITER that the launcher will be far from the last closed flux surface of the plasma and thus even though ray tracing should be valid, it may be challenging to properly model the coupling of LH waves to the plasma.

[2]Some ray tracing and full wave codes only model the plasma inside the last closed flux surface, though it is becoming increasingly more apparent that scrape-off layer models are needed to describe the physics more accurately. As we shall see in Chapter 5, a scrape-off layer model is needed to explain LH behavior at high densities.

[3]As mentioned in the previous section, CQL3D can also be passed the quasi-linear diffusion coefficient from the full-wave solver Toric, though the modeling presented in this thesis has been done exclusively with Genray.

electron distribution function resulting from the balance between RF sources, toroidal electric field, collisional slowing-down, and pitch-angle scattering. The distribution function is solved for in three dimensions: radius, velocity, and pitch-angle.

CQL3D advances the electron distribution function in time. At each time step, the ray paths calculated by Genray are used to re-calculate the quasilinear diffusion coefficient based on the most recent CQL3D-calculated electron distribution function. The ray paths and vector wave numbers do not need to be recalculated between time steps, as they do not depend strongly on the details of the distribution function. This is because there is only a small (<1%) number of particles in the non-thermal electron tail and thus the real part of the plasma response that determines wave propagation is unchanged. A sample output distribution function from CQL3D is shown in Figures 4-5 and 4-6. Figure 4-5 shows contours of the electron distribution function in $v_{\perp}$-$v_{||}$ space and Figure 4-6 shows cuts of the distribution function at different pitch angles.

In addition to the electron distribution function, CQL3D calculates many useful quantities including the radial power damping and current profiles. Genray-CQL3D also contains synthetic diagnostics that can be useful for comparing simulations with experimental measurements.

4.2.1 Synthetic Diagnostics in Genray-CQL3D

Synthetic diagnostics are software modules that calculate the expected measurement that a real diagnostic would see based on the simulated plasma response (in this case, the simulated electron distribution function). These synthetic diagnostics can be a powerful tool for comparing simulations with experiment and for interpreting experimental measurements [70].

CQL3D contains an x-ray synthetic diagnostic which calculates photon flux due to Bremsstrahlung emission along user-specified chords [71]. The calculation integrates the emission from the perturbed distribution function against a background of stationary scatterers along the line of sight of the detector. The pitch angle dependence of the distribution function is approximated as a finite sum of Legendre polynomials.

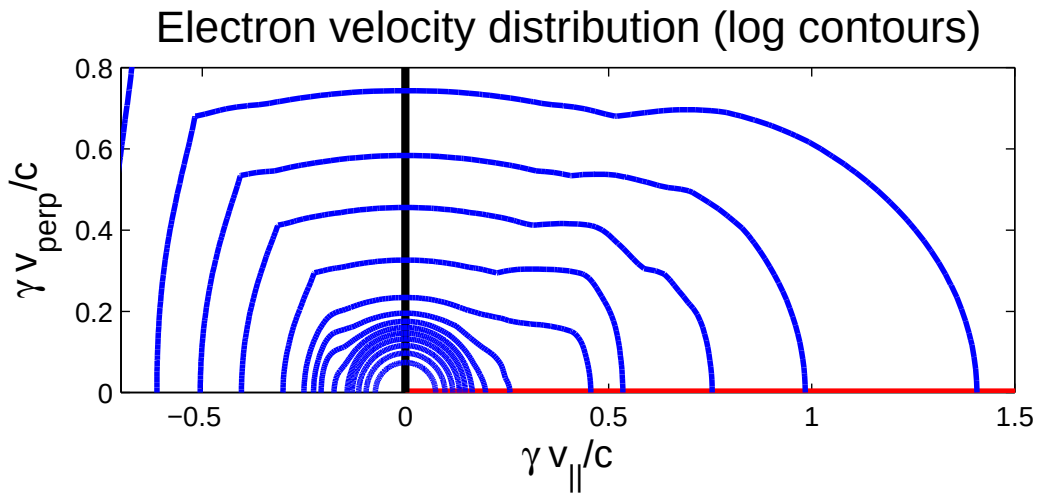

Figure 4-5: Example of a contour plot of distribution function calculated by CQL3D. The central semi-circular (Maxwellian) contours, which are spaced very closely together, have been removed from the plot. The main non-thermal tail can be seen in the positive $v_{||}$ direction. In this simulation, plasma parameters were B=5.4 T, T_{e0}=2.3 keV, $\bar{n}_e = 9 \times 10^{19}$ m^{-3}, and I_p=800 kA. A launched $n_{||}$=3.1 was used

The electron-ion and electron-electron Bremsstrahlung photon intensities are taken from the Bethe-Heitler formula [42] with Elwert correction [43] and the Haug formula [44] with Elwert correction respectively. The resulting flux can be combined with detector geometry for comparison with experimentally observed x-ray emission.

In addition to tracing LH rays, Genray can be run in another mode in order to calculate ECE spectra resulting from an electron distribution function [72]. This ECE synthetic diagnostic calculation is run as a post-processing step after a distribution function has been calculated by CQL3D[4].

The x-ray synthetic diagnostic in CQL3D and the ECE synthetic diagnostic in Genray provide a powerful way to compare code with experimental measurements. Traditionally, current drive codes have been benchmarked by comparing the code's predictions of more macroscopic quantities, like total current driven, with experiment (as is done, for example, in Ref. [28]). The recent use of synthetic diagnostics has

[4]Genray can also predict ECE emission from a Maxwellian plasma without a CQL3D-calculated distribution function.

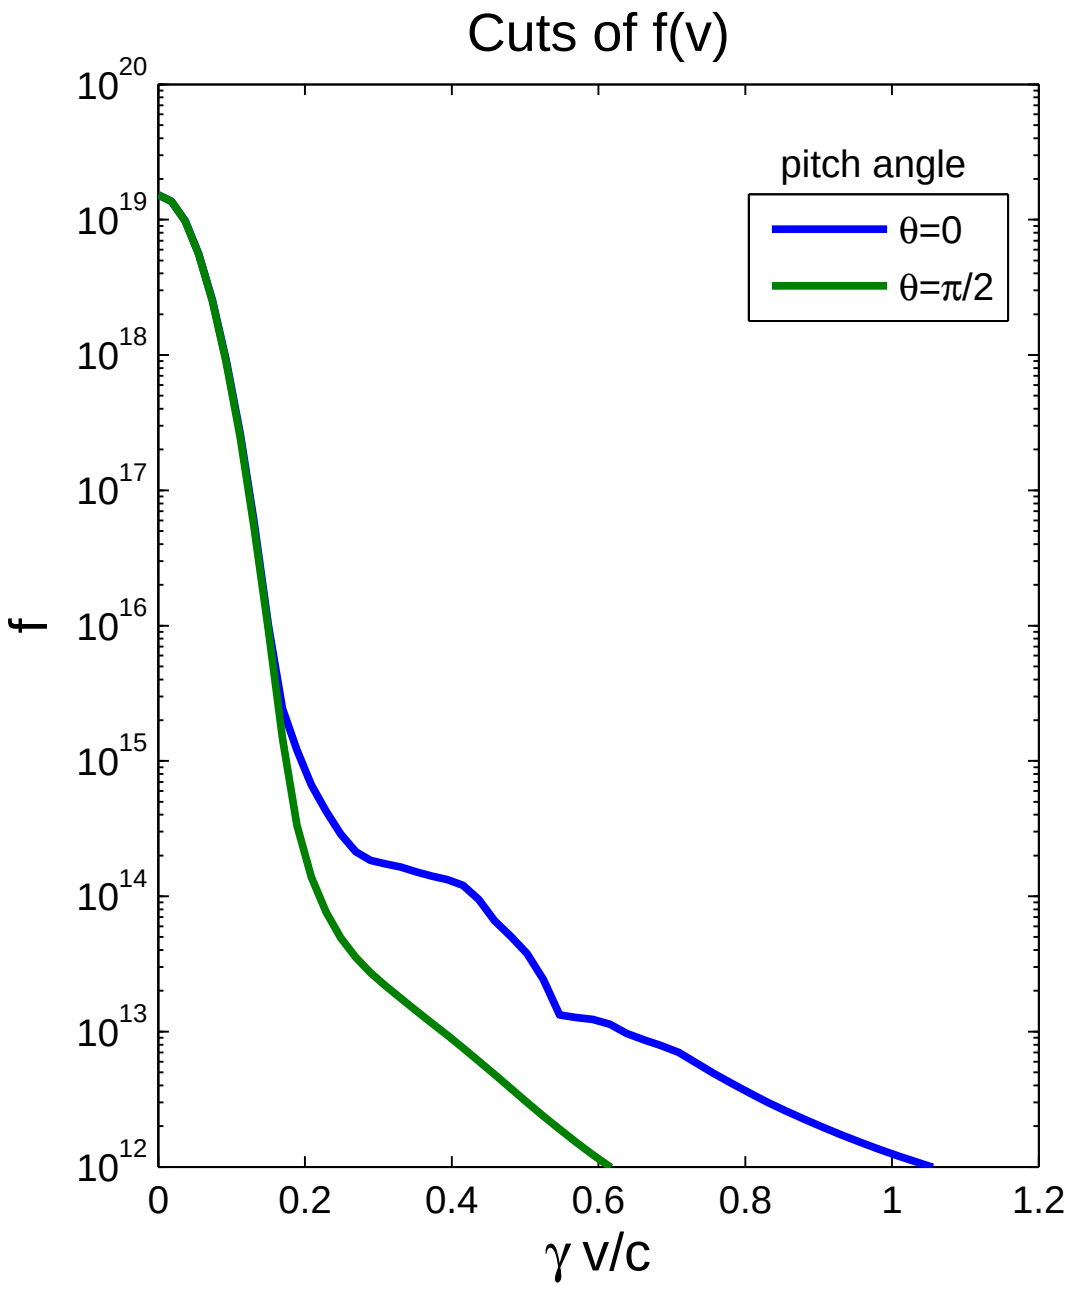

Figure 4-6: Cuts of CQL3D distribution function at constant pitch angle. The thermal part of the distribution function can be seen to the far left, where $\beta\gamma$ is less than 0.2. The non-thermal tail can be seen at greater values of $\beta\gamma$ for the $\theta = 0$ cut. Pitch angle scattered fast electrons can be seen on the $\theta = \frac{\pi}{2}$ cut for $\beta\gamma > 0.2$. In this simulation, plasma parameters were B=5.4 T, T_{e0}=2.3 keV, $\bar{n}_e = 9 \times 10^{19}$ m^{-3}, and I_p=800 kA. A launched $n_{||}$=3.1 was used

allowed for a much more detailed comparison that significantly raises the bar for "agreement" between code and experiment.

For example, agreement between predicted and measure x-ray profiles indicates that the distribution function calculated by the code is likely a realistic approximation to the experimental distribution function. The electron distribution function cannot be measured directly [37], so a reliable simulated distribution function can be used to better understand the details of the nonthermal electron population.

4.2.2 Inputs to Genray-CQL3D

Both Genray and CQL3D employ the plasma density profile, the electron and ion temperature profiles, Z_{eff} profiles, and a magnetic equilibrium file in the eqdsk format[5]. Additionally, the LH launched power and $n_{||}$ spectrum is input into Genray and the toroidal DC electric field is input into CQL3D.

In the modeling shown in this thesis, the density and electron temperature profiles were obtained using the IDL fits routine [73, 74], which runs a regression and then a spline fit on various density and temperature diagnostics. The spline fit was interpolated to the $\sqrt{\Psi_t}$ (square root of normalized toroidal flux) grid used by Genray and CQL3D. The ECE diagnostics were excluded from the temperature fitting since they are considered inaccurate measurements of temperature during LH operation.

The ion temperature profiles were taken to be 2/3 of the electron temperature profiles. The ion temperature is not well known on C-Mod in the lower density discharges where LHCD is performed and this is a reasonable approximation for our purposes. The ions do not play a large role in LH wave propagation and damping in this regime because the ion thermal speed is much lower than the LH wave phase velocity. The bremsstahlung emission is not significantly affected by the ion temperature because the ion thermal speed is small compared to the speed of LH-driven fast electrons. CQL3D has an optional input for a 3^{rd} species, such as energetic alphas or an impurity species, but this optional input was not used in the simulations presented

[5]Extensive documentation on eqdsk file formats can be found at https://fusion.gat.com/theory/Efitoutputs

in this thesis. The main effect of impurities in the C-Mod LH experiments was to raise Z_{eff}, which is another input to CQL3D (discussed later in this section).

For the modeling presented in this thesis, the density profiles used in Genray do not include the scrape-off layer. Thus the density profiles were artificially lowered at the edge of the plasma to be below the LH cut-off density inside the last closed flux surface. This is done because the implementation of Genray without the scrape-off layer does not trace beyond the last closed flux surface and thus if a ray reaches the last closed flux surface and has not yet reached a cut-off, the code performs a specular reflection.

Cut-off reflections and specular reflections are different physical processes. In a cut-off reflection, the radial index of refraction goes to 0. A specular reflection is a mirror-like reflection, in which n_r simply reverses direction at the reflection surface. Thus it is undesirable to have rays specularly reflecting in the code when we believe that in the experiment, they generally reflect at a cut-off before reaching the vessel wall (where a specular reflection is then possible)[6].

C-Mod employs a visible bremsstrahlung diagnostic to measure Z_{eff} in thermal plasmas [55]. The average Z_{eff} calculation employed routinely after each discharge assumes a flat density profile, which is a resonable approximation for H-mode discharges, but a poor approximation for the L-mode discharges modeled in Chapters 5 and 6. It also uses GPC temperature measurements (which are artificially elevated during LH operation) to calculate Z_{eff}, though it could in principle be directed to use Thomson scattering temperature measurements instead. Finally, the nonthermal electrons will have a small effect on the visible Bremsstrahlung levels. All three effects result in a higher calculated average Z_{eff} than the true value. For this reason, the average Z_{eff} measurement is not always used in simulations.

In all of the modeling shown in this thesis, the Z_{eff} profiles were taken to be constant in radius. In the modeling in Section 5.1, Z_{eff} was taken from the visible bremsstahlung diagnostic measurements. In Sections 5.2 and 6.5, Z_{eff} was taken

[6]In some C-Mod discharges, the density in the scrape-off layer does not fall below the LH cut-off in some poloidal locations, such as parts of the inner wall. In this case, some LH power may be specularly reflected.

from a routine which finds a value of Z_{eff} consistent with the neoclassical plasma resistivity, based on the plasma current and loop voltage (taken from a time slice right before LH power is turned on) [75].

In the simulations presented in this thesis, the toroidal DC electric fields input into CQL3D correspond to a constant loop voltage. Namely, the loop voltage is assumed to be the same at every major radius and the radial electric field scales with 1/R. This is of course not true while the current profile is evolving, but we do not have a direct measurement of toroidal electric field. Therefore we make this approximation and try to model steady-state plasmas. If the loop voltage is not well known (for example, it is changing in time), we try to find a reasonable range of measured loop voltage values and then use the one which predicts the plasma current correctly, as discussed in the next section.

4.2.3 Using the Plasma Current as a Constraint

While some plasma parameters are not well known in the experiment, the total plasma current is generally used as a constraint in the modeling presented in this thesis. Input parameters that are not well known are adjusted until the total current predicted by the code matches the experimentally known plasma current. In principle, with a direct measurement of plasma current profiles, we could go a step further and try to match the entire profile. This idea is explored further in Section 5.2.

Unfortunately, there are three parameters that are not well known in the modeling, all of which can have a significant effect on the total amount of current predicted by the model. The first is Z_{eff}, which is not easily measured during LH operation. The Z_{eff} figures prominently into both the ohmic current drive and the LH current drive efficiency.

The second parameter that is not well known is the loop voltage, often denoted V_{loop}, which characterizes the strength of the toroidal electric field inside the plasma. In theory V_{loop} is well known in a steady-state plasma, based on the ramp rate of the central solenoid and direct measurement of V_{loop} on the outboard side of the plasma. In practice, however, we do not know V_{loop} inside the plasma while the current profiles are

evolving and the plasma current is relaxing. This can be problematic because keeping the LH power and other plasma conditions stable through an entire or several current relaxation times is not easy to do.

Finally, the fast electron spatial diffusion coefficient is a free parameter in the code which is also unknown (experimentally) and which can affect the total amount of predicted plasma current. Generally speaking, even a small diffusion coefficient in the code will significantly reduce the total amount of current driven and a large diffusion coefficient will reduce the current driven to nearly ohmic values[7]. In Chapter 6, a method for experimentally determining a fast electron diffusion coefficient will be presented.

Because the total current is affected by all three of these parameters, we typically use the most accurate available measurement for two out of three and use the plasma current as a constraint on the third. Typically, Z_{eff} is constrained to the most reliable available Z_{eff} measurement. If V_{loop} is steady for the discharge, its measured value is used. If V_{loop} is drifting significantly, the total plasma current is used as a constraint to find a reasonable value for V_{loop}. For simulations that include radial diffusion, the diffusion coefficient is typically adjusted to match some sort of smoothness or shape of a profile (either the current profile or HXR profile). The diffusion coefficient is never adjusted to match the total current, even if Z_{eff} and V_{loop} appear to be well-known, because of the great uncertainty in measuring Z_{eff} and V_{loop}.

Our experience with Genray-CQL3D is that when reasonable physical values for Z_{eff} and V_{loop} are used in the code, it has always predicted total plasma current that is similar to the experimentally measured value.

4.2.4 Representing the LH Grill and $n_{||}$ Spectrum in Genray

Every LH ray in Genray has an initial location and $n_{||}$ value associated with it. The launched value of $n_{||}$ for each ray has a big impact on the location of its power absorption and how far it propagates before being absorbed. The poloidal location of a launched ray also has an impact on its trajectory and greatly affects the evolution

[7]Section 4.2.5 will address the magnitudes of diffusion coefficients used in the code.

of the $n_{||}$ of the ray. The launched $n_{||}$ spectrum is a continuum and the launcher spans a wide poloidal location, so ray-tracing by nature is a discretization of a continuous problem.

Due to memory limitations[8], there is a trade-off between launching rays with spatial (poloidal launch location) resolution and spectral resolution. In the simulations presented in this thesis, the LH grill is modeled with four poloidal launch locations, corresponding to the poloidal center of each LH launcher row. In principle, one could model it more accurately with more poloidal launch locations, with the power in each ray scaled like the square of a sine function over the poloidal extent of each row, but this would increase the memory needed and the running time of the code in proportion to the number of rays used.

Spectral resolution is treated by launching several rays of varying $n_{||}$. Two spectral bundles of rays are launched from each poloidal location, one to represent the main lobe of the forward spectrum and one to represent the main lobe of the reverse spectrum. The power assigned to each ray is a sinc function of $n_{||}$. Side lobes are not included in the simulation although in principle they could be simulated. The number of rays representing the two lobes varies from simulation to simulation, but a typical model includes 35 rays launched forward and 10 rays launched in the reverse direction (at each of the 4 poloidal locations).

Since in practice, some of the power into the plasma is in the side lobes of the $n_{||}$ spectra, the power input into the simulation is scaled down so that only the power going into the main forward and reverse lobes is included in the model. Thus the model includes less total power than the plasma absorbs experimentally, but it is not fictitiously modeling the power as if it were all in the two largest spectral lobes.

[8]Genray is written in Fortran 77, which does not have dynamic array sizing. This can lead to arrays which are sized larger than necessary (and correspondingly use up memory inefficiently) because making the arrays smaller requires recompiling the code. Additionally, the version currently run has not been parallelized, so it is typically run on a single processor.

4.2.5 Inclusion of Radial Diffusion in Modeling

It is possible to turn on spatial particle diffusion in CQL3D, in effect coupling the distribution function calculations from neighboring flux surfaces which would otherwise be independent. When diffusion is turned on, very small time steps (fractions of a millisecond) must be used in the time advance. The most efficient way to run a simulation that includes spatial diffusion in CQL3D is to first compute the steady-state distribution function without diffusion and then to turn diffusion on and decrease the time step size.

The diffusion coefficient velocity scaling used in the modeling in this thesis is shown in Eq. 4.1:

$$D(v_{||}) = D_0 \frac{v_{||}/v_{te}}{\gamma^3}, \tag{4.1}$$

where γ is the relativistic factor and D_0 is a user specified diffusion coefficient. This velocity scaling is based on a magnetic turbulence model [76]. Typical values of D_0 used in the modeling presented here are .01 to .04 m^2/s. Larger diffusion coefficients typically predict a loss of fast electrons (and hence current driven) that is clearly inconsistent with experimental observation. Excluding diffusion effects entirely often leads to highly localized "spiky" power deposition and current profiles that are not physically realistic.

The greater number of time steps needed to include diffusion in the model increases the computing time of a single simulation from approximately 1 hour to several hours.

4.3 ECE Predictions

This section presents some results from purely predictive modeling. The other modeling presented in this thesis was performed in conjunction with experiments as a means of both benchmarking the code and interpreting the data. Simulations done in support of experimental data will be presented in Chapters 5 and 6 along with the experimental data that they were simulating.

4.3.1 Modeling the Effect of LHCD on ECE Spectra

Before the Lower Hybrid system was operational on C-Mod, we simulated the effect that the fast electrons would have on ECE spectra. As outlined in Section 2.3, we expected that the fast electrons would have relativistically downshifted ECE frequencies, but it was unclear how they would affect the overall spectrum. A high density H-Mode discharge, shot 1030612013 at t=1 s, was used as the target plasma for this set of simulations. The peak electron temperature in this plasma was 5.5 keV, with a peak electron density of $2.2 \times 10^{20} m^{-3}$ and 3.5 MW of ICRF heating.

In this set of simulations, the peak $n_{||}$ was set at 2.5, the toroidal electric field was turned off, and the LH power was scanned from 200 kW to 2 MW. The resulting ECE spectra are shown in Figure 4-7. The main result is that the second harmonic ECE spectrum appears virtually unaffected by the fast electrons, but there is a frequency region below the second harmonic, where some emission appears only in the presence of the fast electrons. We refer to this emission as the "non-thermal" emission and it was predicted to occur at appreciable levels in this plasma scenario when the LH power reaches 600 kW.

The result that the fast electrons do not affect the "thermal" spectrum (the frequencies that correspond to an electron cyclotron second harmonic somewhere in the plasma) makes physical sense because downshifted emission traveling towards the ECE diagnostic should in principle be re-absorbed when it encounters the EC resonance for non-relativistic electrons. The result that "non-thermal" emission escapes the plasma is also easily explained. Presumably fast electrons at the outboard edge of the plasma emit relativistically downshifted radiation which has no resonance layer to absorb it, and it reaches the ECE detectors.

This result turned out to be very important because it guided our use of the ECE diagnostic during LH operation, even during early operation of the LH system. Based on the prediction that there would be non-thermal emission below the second harmonic, some of the GPC channels were set to look at those frequencies. And indeed, strong non-thermal emission was observed at these frequencies during LH

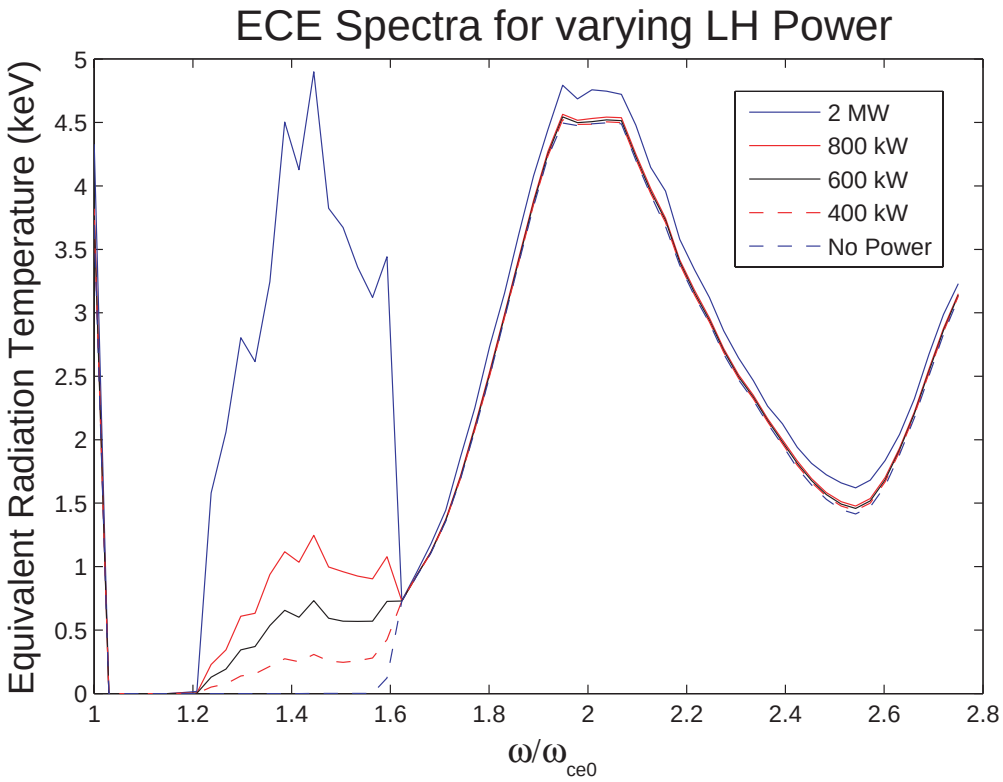

Figure 4-7: Predicted ECE spectra for varying LH power in H-mode discharge 1030612013 at t=1 s. Plasma parameters for this discharge were B=5.4 T, T_{e0}=5.5 keV, $\bar{n}_e = 1.5 \times 10^{20}$ m^{-3}, and I_p=600 kA..

operation. Experimentally, we have found that the thermal part of the spectrum is often affected (sometimes dominated) by the non-thermal emission, particularly at low densities. Modeling of ECE spectra for a low density discharge is presented in Section 5.1. In that case, it was found that the Genray-CQL3D model significantly over-predicts the thermal emission level. Since the Thomson Scattering diagnostic shows no change or much more modest changes in electron temperature than the ECE diagnostics, we typically consider the ECE diagnostic data during LH operation to be an unreliable measurement of temperature.

Chapter 5

Measurements from Lower Hybrid Current Drive Experiments and Comparison with Simulations

The results described in this chapter are from lower hybrid experiments run on Alcator C-Mod from the 2006 campaign, when LHCD experiments first produced useful measurements, to the present. Section 5.1 presents the first comparisons of the HXR and ECE synthetic diagnostics in Genray-CQL3D with experimental measurements. In the two discharges modeled, which were nearly fully non-inductive, the simulations are predicting much higher levels of HXR and ECE than observed experimentally. However, historically it was the first such benchmarking of Genray-CQL3D against measurements from an LH-driven plasma.

Section 5.2 presents efforts to use information from the motional stark effect (MSE) diagnostic to compare experimentally measured current profiles with modeled current profiles. In one case, an inferred current profile from MSE is compared directly with a predicted current profile. It is seen in this case that a small amount of fast electron diffusion is needed in the simulation to produce a current profile that is as smoothly varying as that inferred from the MSE diagnostic.

An antenna phasing scan is then presented, in which predicted current profiles are compared with MSE-constrained EFIT current profiles. In this case it is found

that predicted current profiles qualitatively agree with the EFIT generated current profiles only for high antenna phasing [77].

Finally, Section 5.3 contains an overview of the observed drastic reduction in HXR count rates at high density [78]. While the published literature on the LH density limit typically presents the HXR data agregated over many chords, the data is presented here as spatial profiles. The HXR profiles for single null discharges show increased counts in the area of the active divertor, supporting the notion postulated in [78] that a significant fraction of LH power is being absorbed in the SOL in high density diverted discharges.

5.1 Nearly Fully Non-Inductive Discharges

There is a feedback system in C-Mod that instructs the central solenoid to ramp its current at such a rate as to ohmically drive the plasma current to a pre-programmed value. Meanwhile, we measure the surface loop voltage, denoted V_{loop}, encircling the central solenoid as a function of time for every discharge. In a steady-state plasma, the surface loop voltage reflects the loop voltage everywhere in the plasma. Depending on plasma conditions, a discharge with no LH power generally needs 1-2 volts to drive 500 kA - 1 MA of current ohmically. When sufficient LH power is applied to the plasma, the loop voltage drops[1], indicating that some fraction of the ohmic current has been replaced with non-inductively driven LH current.

At the end of the 2006 campaign, two discharges were obtained with 800 kW of coupled LH power, nearly sufficient LH power to drive the plasma current non-inductively. These two discharges were modeled extensively for a comparison between the observed ECE and HXR measurements and the output of the ECE and HXR synthetic diagnostics in Genray/CQL3D.

Figure 5-1 shows time-traces for the two discharges. Shots 1060728011 and 1060728014, hereafter referred to as "shot 11" and "shot 14" respectively, both exhibit low plasma

[1]It should be noted that the sign convention used in the Alcator C-Mod loop voltage is such that loop voltage is typically negative; a loop voltage "drop" refers to the loop voltage becoming less negative.

density. In shot 11, the requested plasma current was 500 kA and in shot 14, it was 1 MA. Both discharges had approximately 800 kW of coupled LH power, with a peak $n_{||}$ of 1.55 (60° phasing on the C-Mod LH1 launcher). During the LH pulse in both shots, the loop voltage drops significantly, and in shot 14, the loop voltage is transiently driven to zero and then positive.

It is difficult to tell from the HXR count rate time trace, but the raw HXR data for shots 11 and 14 indicate major problems with baseline drift, pile-up, and detector saturation. A simplified peakfinding routine which does not account for baseline drift and does not count piled-up pulses can make the HXR count rates appear artificially low[2]. In the following comparison between measured and modeled HXR, a peakfinding routine is used in which baseline drift is corrected and piled-up pulses are fit to multiple Gaussians [79].

These two shots were modeled extensively in CQL3D for the time-slice $t = 1.1s$ for shot 11 and $t = 1.0s$ for shot 14. Measured density and temperature profiles and Z_{eff} were used. An eqdsk file generated by relaxing q_0 constraints in EFIT [80] was used as the magnetic equilibrium input (note that this equilibrium was not self-consistent with CQL3D current profile calculations). A loop voltage of 0.2 V was used to model shot 11, and this value came directly from the V_{loop} measurement, which remained fairly steady throughout the LH pulse. A loop voltage of 0.1 V was used to model shot 14, which was the value (found through iteration) that predicted the correct amount of total plasma current in the model.

Figures 5-2 and 5-3 show a comparison between the experimentally observed and modeled HXR count rates for shots 11 and 14. Clearly, the predicted x-ray count rates are much higher (by approximately a factor of 5) than those observed experimentally, for both discharges. Furthermore, the spatial distribution of x-rays in the experiment is much broader than the model predicts. Though the discrepancies between experimental and modeled HXR is rather large in these cases, this modeling

[2]Pile-up and the associated baseline drift, when treated simply, cause the measured count rate to appear lower than it actually is. Baseline drift makes HXR photons appear lower energy than they actually are, causing many of them to miss the threshold for being counted. "Piled-up" photons (double-peaks) are thrown away and not counted.

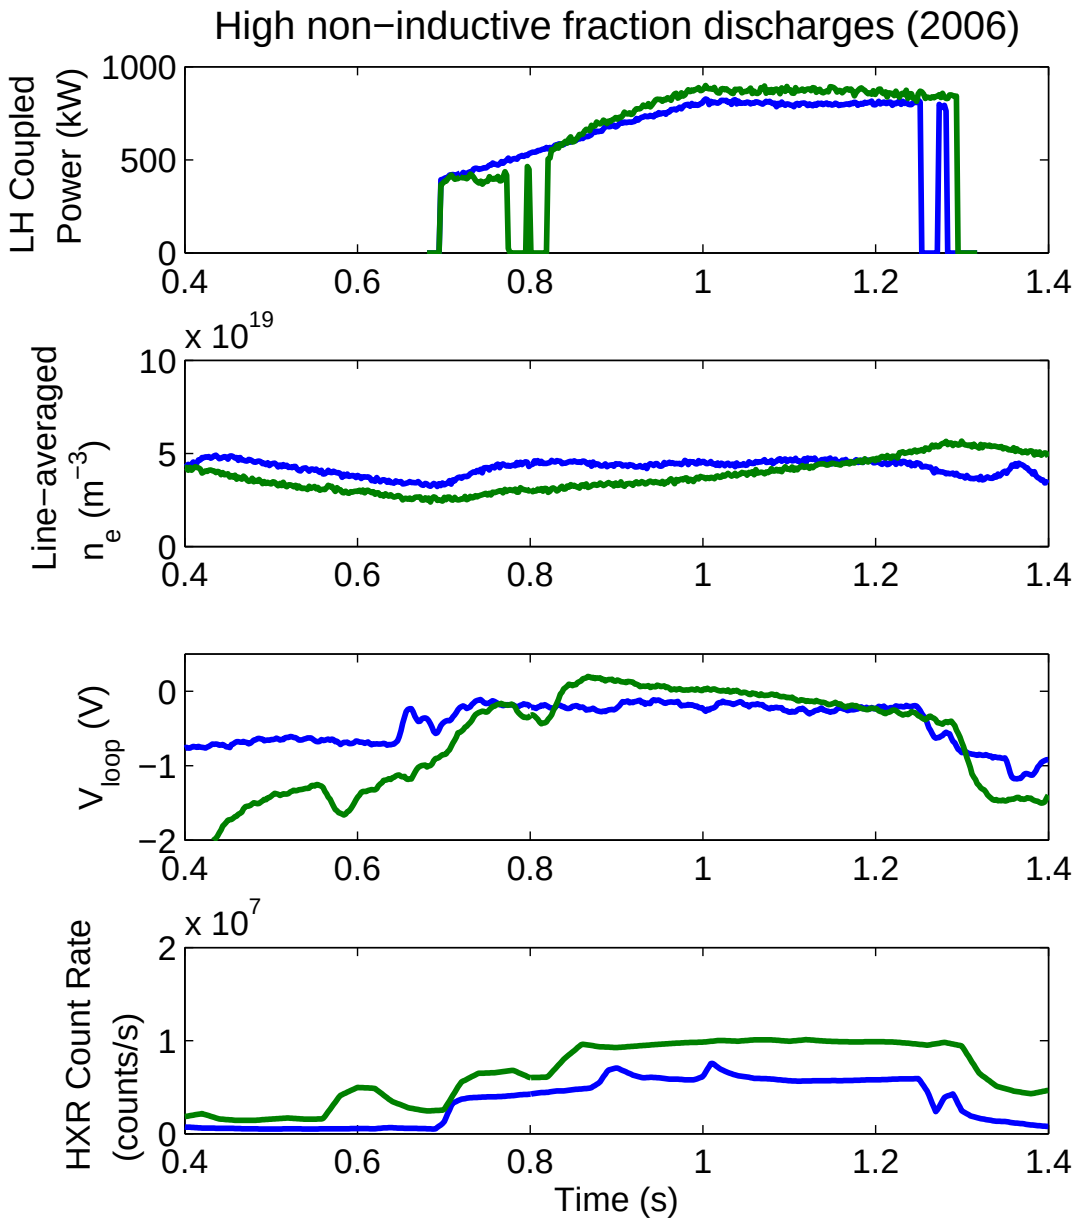

Figure 5-1: Time traces for nearly full non-inductive discharges, 1060728011 (blue traces) and 1060728014 (green traces). Note that the loop voltage for shot 11 remains consistently low, around -0.2 V, while the loop voltage for shot 14 initially reverses sign when the LH power is turned on, and then eventually drifts to a low (negative) value, indicating that it is not fully non-inductive. HXR count rates shown are summed over all spatial chords and summed over all photon energies from 20-200 keV.

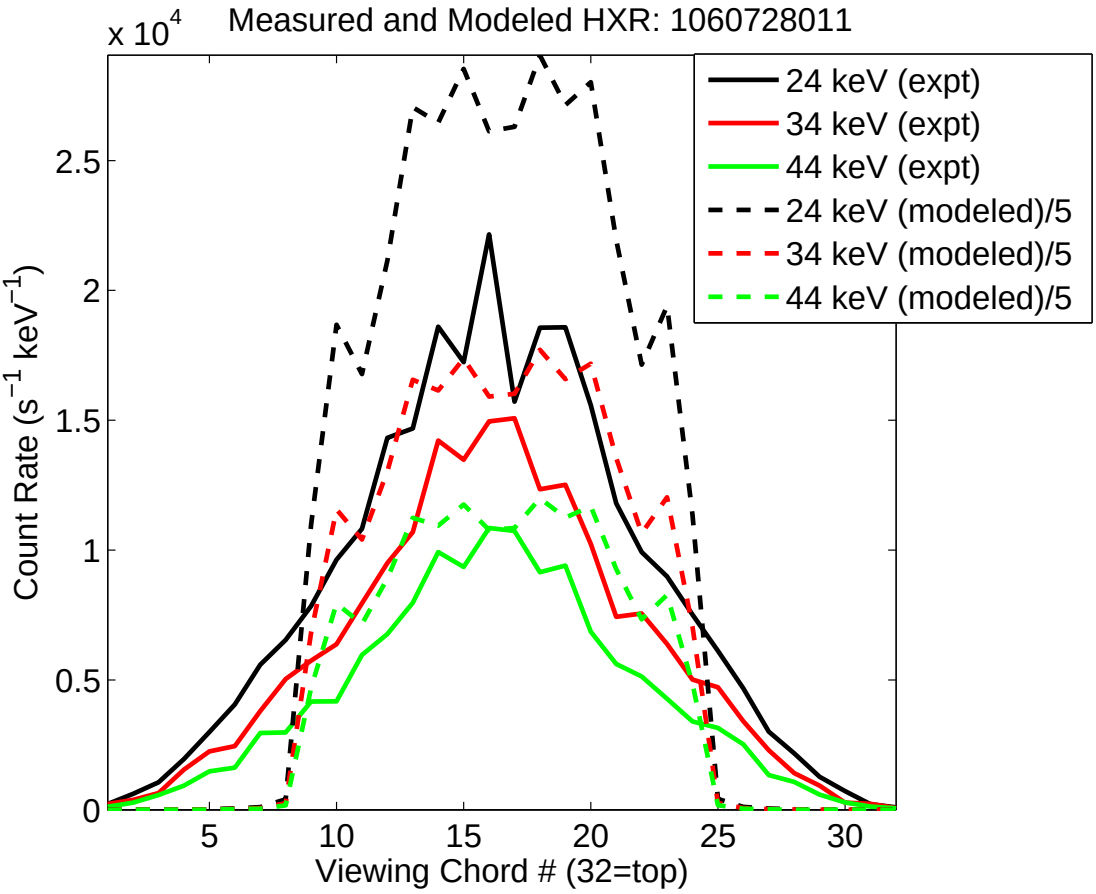

Figure 5-2: Comparison of measured and modeled HXR profiles for shot 1060728011. Modeled profiles have been scaled by a factor of 1/5, and are much narrower in spatial extent than experimental profiles.

was the first of its kind to compare the synthetic HXR diagnostic in CQL3D with experimental profiles.

The cause of the disagreement in the HXR profile width has been the topic of much speculation in the past few years. At one point, it was thought that perhaps radial diffusion of fast electrons (not included in the model producing Figures 5-2 and 5-3) could be playing a role in this spatial broadening. Fast electron diffusion was later turned on in the model, and it very clearly could not account for the difference in the width of the experimental vs. modeled HXR spectra [81]. Furthermore, as we shall see in Chapter 6, measurements of off-axis fast electron diffusion confirm that it could not account for these differences.

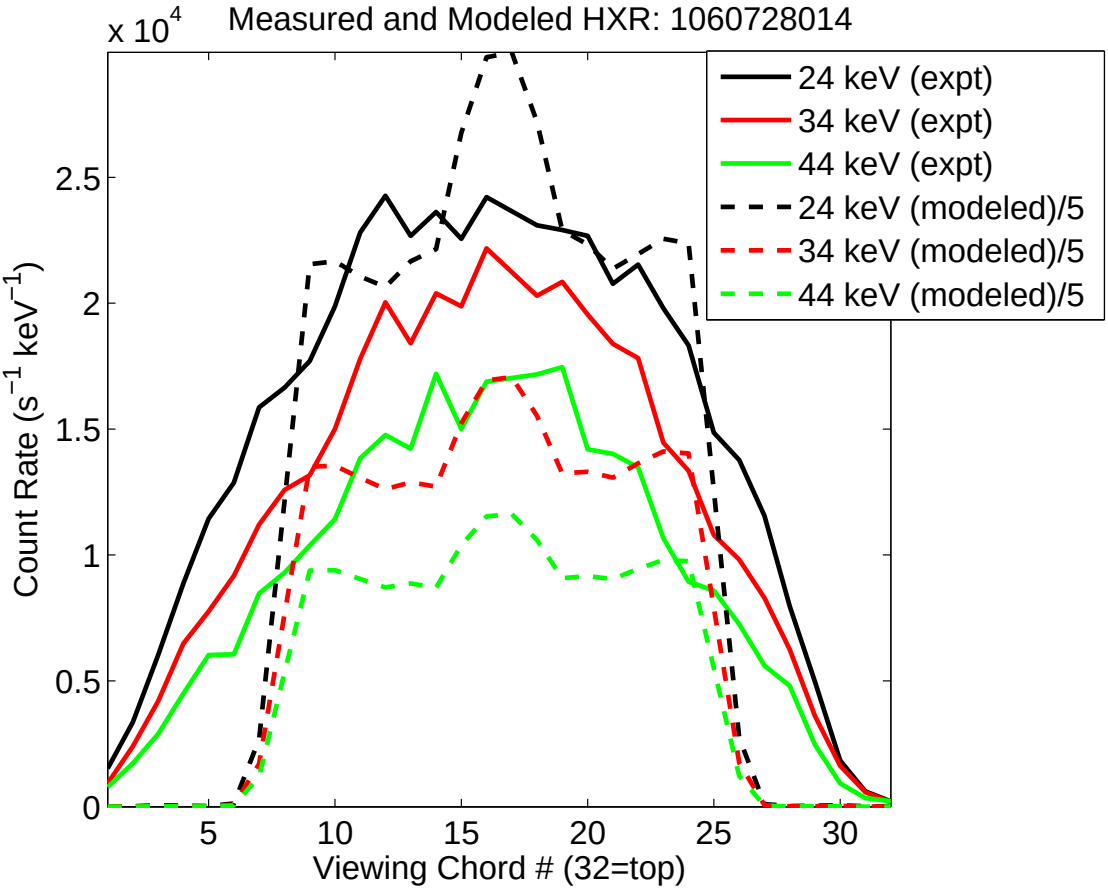

Figure 5-3: Comparison of measured and modeled HXR profiles for shot 1060728014. Modeled profiles have been scaled by a factor of 1/5, and are much narrower in spatial extent than experimental profiles.

Another possible explanation for spatial broadening was diffraction of LH waves that would be seen in a full-wave code, but not in a ray-tracing code. This discrepancy and other motivations launched the two projects to integrate full-wave solvers with Fokker-Planck codes described in Section 4.1. Currently, neither full-wave code is showing significantly better agreement with experiment than the ray tracing approach, though they do predict different results than ray-tracing does. So it is likely that diffraction is also not responsible for the broadening of the HXR profiles seen in experiment.

As we shall see in Section 6.5, it was later discovered that Genray-CQL3D makes poor predictions of both the magnitude and shape of experimental x-ray spectra in cases with a low launched $n_{||}$. Our interpretation of this observation is related to the ray path lengths. Rays with low $n_{||}$ tend to propagate longer distances than rays with high $n_{||}$ because they need greater $n_{||}$ upshifts to damp in a plasma of a given temperature. Rays in a $n_{||}$=1.6 simulation typically undergo many reflections before they damp. In the experiment, when rays reflect in the scrape-off-layer, they may lose some of their power there. Indeed, such a loss of power in the SOL is needed to explain the lower hybrid density limit [78] (see Section 5.3). The SOL is not included in the simulations shown here, which may partially account for the magnitude of discrepancy in HXR response between the experiment and the simulations.

Shot 11 was also used to compare the experimental ECE spectrum with the predicted ECE spectrum from Genray's synthetic ECE diagnostic. Figure 5-4 shows the measured and modeled ECE spectrum, with the modeled spectrum reduced by a factor of 5. The shape of the two spectra agree quite well, but the magnitude clearly differs in the two cases.

The overprediction of both ECE and x-rays indicates that the code is not correctly predicting the electron distribution function in this case. Though the discrepancies between measurements and both synthetic diagnostics is rather large, these were the first comparisons of CQL3D predictions in an LH plasma with experimental measurements.

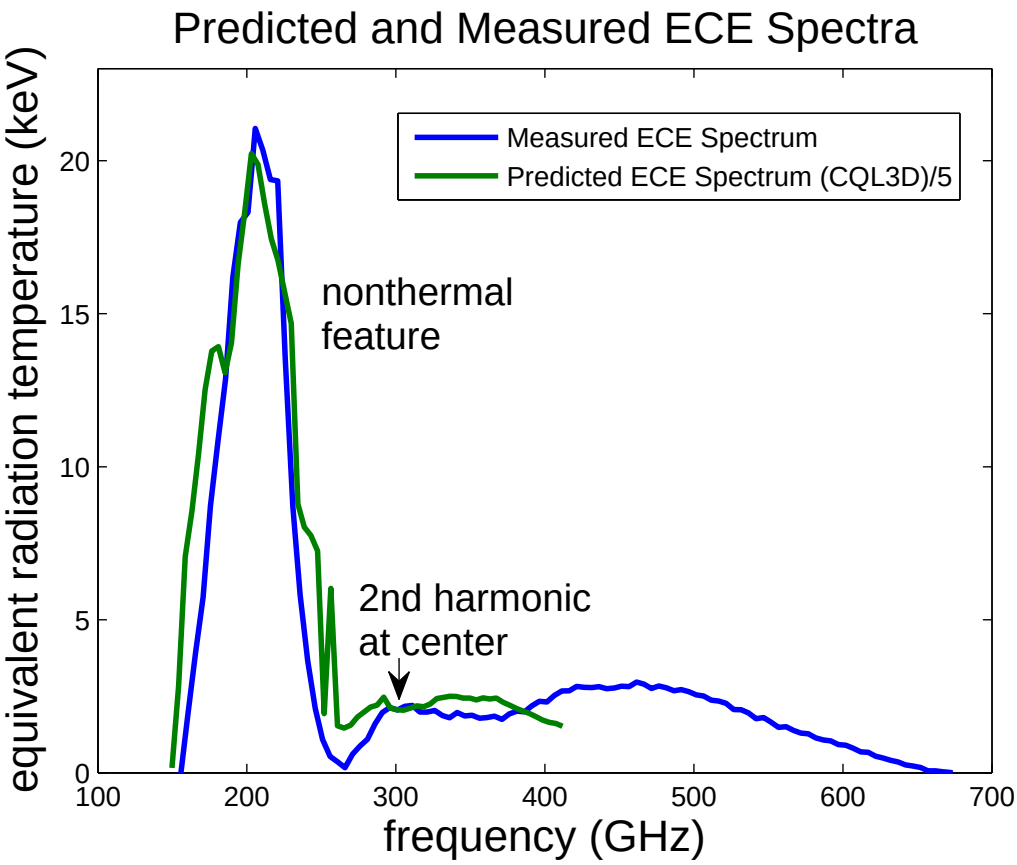

Figure 5-4: Comparison of measured and modeled ECE spectra for shot 1060728011. Modeled spectrum has been scaled by 1/5 to overlay the experimental measurement.

5.2 Motional Start Effect Diagnostic Comparison

A Motional Stark Effect (MSE) diagnostic was implemented on Alcator C-Mod in order to measure the plasma current density profile [82]. This diagnostic is particularly germaine to LH physics studies because the main goal of LHCD is to modify and control the current profile. C-Mod's MSE diagnostic was a very challenging measurement with many difficult technical problems, the most significant of which was a signal drift due to anisotropic heating of in-vessel optical components [82].

At the conclusion of the 2008 campaign, the MSE diagnostic was able to infer a current profile during LH operation by using the equilibrium solver EFIT to calculate the current profile during the ohmic phase of a discharge and observing changes in the MSE signal during the LH pulse [77]. In order to measure appreciable changes in the MSE signal, we coupled very high levels of LH power into a low density plasma. We simulated one such discharge, shot 1080320017, hereafter called "shot 17." Time traces of relevant plasma parameters for this discharge are shown in Figure 5-5. The LH power in shot 17 had a peak launched $n_{||} = 1.96$ (75° phasing).

Shot 17 was modeled both with and without radial diffusion of fast electrons. Figure 5-6 shows the modeled current profiles in both cases, as well as the current profile inferred from MSE measurements. The diffusion coefficient used in this simulation (with a velocity dependence based on a magnetic turbulence model [76]) is:

$$D(v_{||}) = .04\frac{v_{||}/v_{te}}{\gamma^3}, \tag{5.1}$$

where γ is the relativistic factor. The HXR signal for this set of experiments was characterized by many piled-up pulses as well as baseline drift due to the extremely low density, making it difficult to extrapolate reliable HXR spectra. Thus a comparison of modeled and measured HXR is not presented here.

In order to compare the predictions of CQL3D with the measurements of the MSE diagnostic, we also modeled two other LH phasings, 90° and 105°, which corresponded to similar discharges in the same run day, shots 23 and 25 respectively. To isolate the effects of the phasing and to make the modeling as simple as possible, we used the

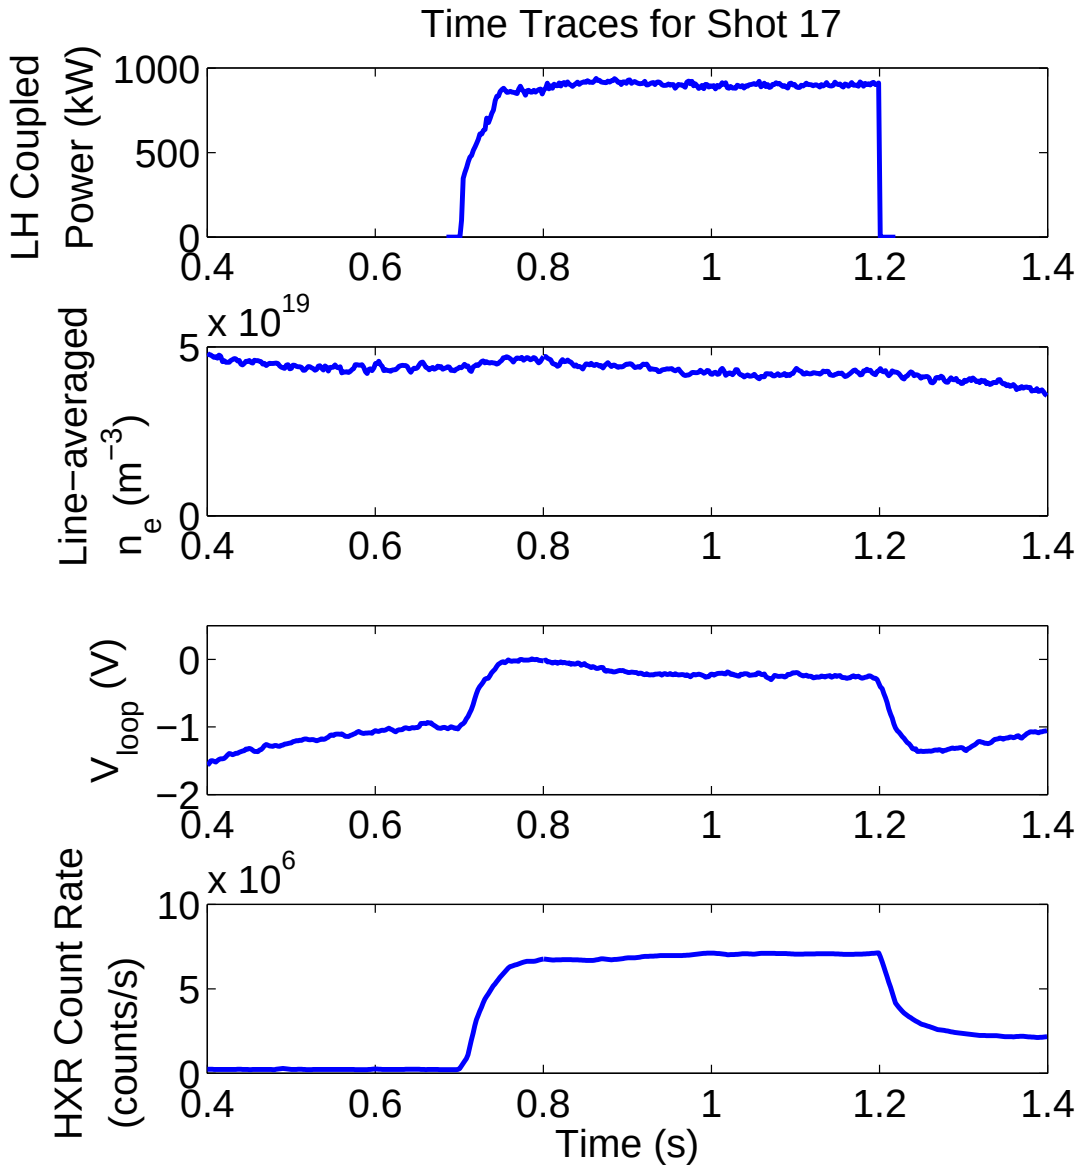

Figure 5-5: Time traces for Shot 1080320017: Coupled LH power, line-averaged density, loop voltage, and HXR count rates. In this discharge, $B = 6$ T and $I_p = 800$ kA.

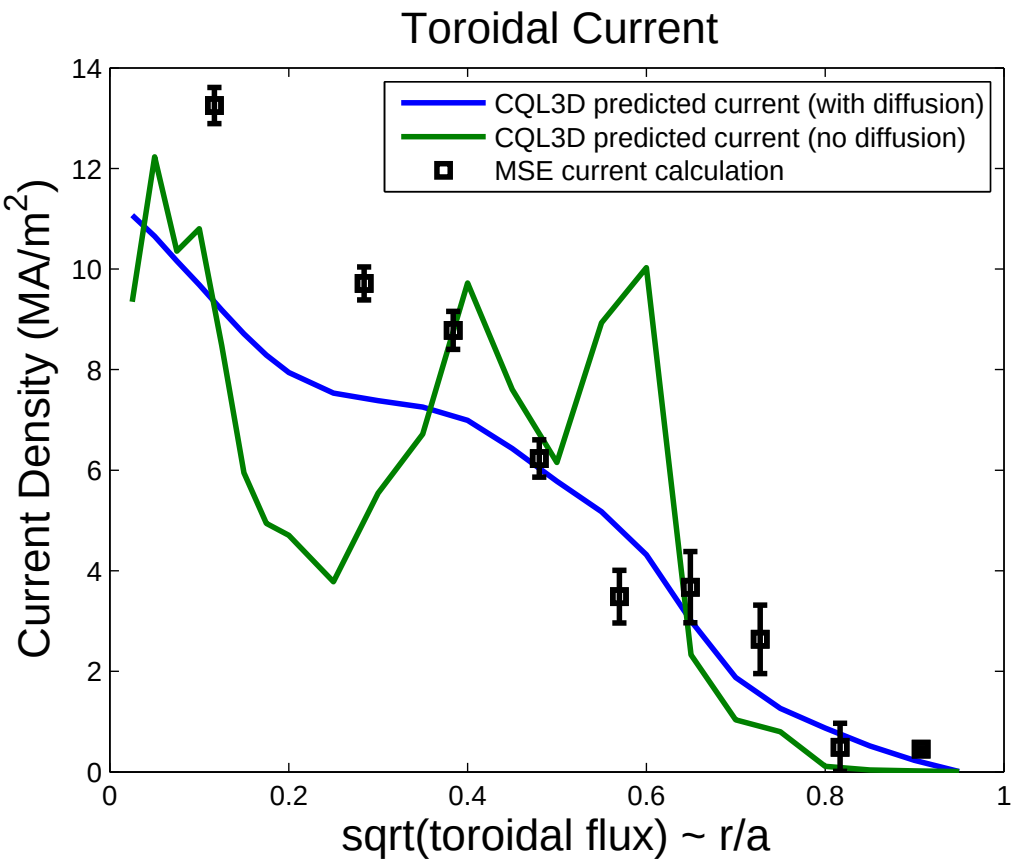

Figure 5-6: Comparison of measured and modeled current profiles for shot 1080320017. Modeled current profiles are shown both with and without radial diffusion of fast electrons. Measurements are taken from the MSE diagnostic where EFIT is used to calculate the current profile during the ohmic part of the discharge and the current profile during LH operation is inferred from the change in MSE signal.

same plasma parameters as used in the modeling of shot 17. Only the LH phasing and the loop voltage (as necessary to keep the total plasma current in the model consistent with experiment) were changed in the simulation.

The power deposition and current profiles from the LH phasing scan are shown in Figure 5-7. Current profiles generated by MSE-constrainted EFIT are shown in Figure 5-8. The agreement between experimental measurements and CQL3D modeling is relatively poor. The MSE measurements show LH current drive location that is independent of phasing, near r/a=0.8. As LH power is applied, the EFIT current profiles are showing an increase in the non-inductively driven edge current and a decrease in the ohmically driven core current.

By contrast, the modeling predicts LH-driven current whose location is highly dependent on phasing. However, both the modeling and the experiment indicate that higher LH phasing corresponds to a lower fraction of inductively driven current. This can be seen in the EFIT current profiles simply by looking at the magnitude of the increase in edge driven current as a function of phasing. It is not evident from Figure 5-7 that the inductively driven fraction of current trends upward at lower phasing, but can instead be seen from the input parameters to the simulations themselves: a greater loop voltage is needed for higher phasings to achieve the measured plasma current in the simulations.

It is quite possible that the qualitative discrepancy between modeling and MSE measurements is due to the same problem that existed in the simulations in Section 5.1. That is, low $n_{||}$ cases cause the rays to reflect multiple times, making the model highly sensitive to changes in the edge plasma conditions. Indeed, the 105° modeled case is the one that agrees most closely with the MSE measurements of LH driven current location. It is also possible that EFIT is not calculating the correct current profile during the ohmic phase of the discharge and that the current profile calculation for the LH phase of the discharge is correspondingly incorrect.

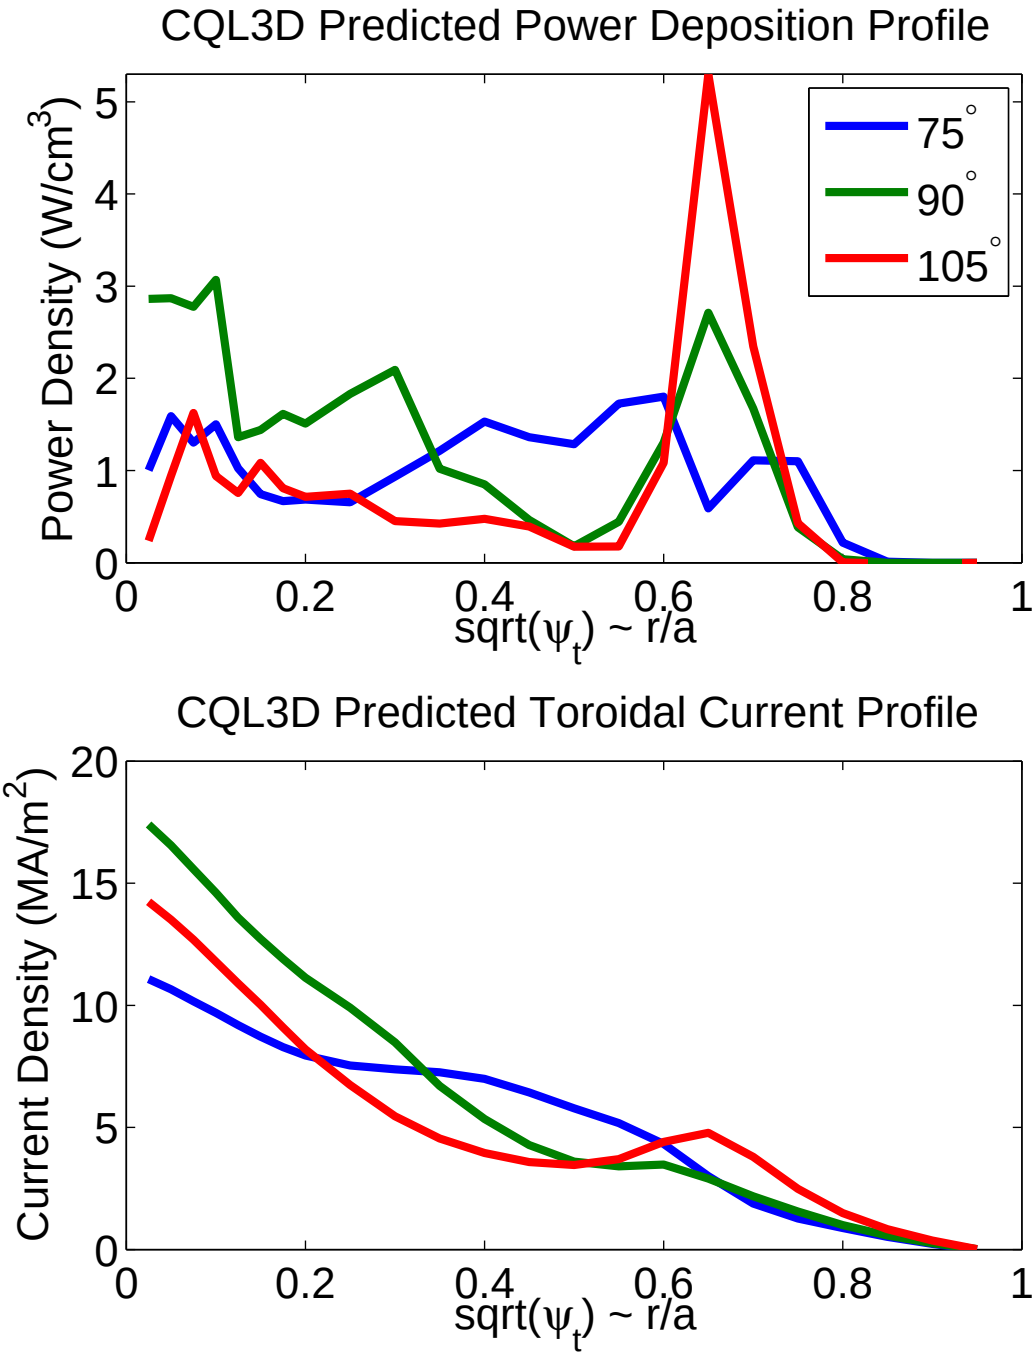

Figure 5-7: Modeled power deposition and current profiles for 3 LH phasings to compare with MSE measurements under similar plasma conditions. In these simulations, $\overline{n}_e = 5 \times 10^{19}$ m^{-3}, $P_{LH} = 900$ kW, $B = 6$ T and $I_p = 800$ kA. Simulations are shown for 75° ($n_{||} = 1.9$), 90° ($n_{||} = 2.3$), and 105° ($n_{||} = 2.7$) phasings.

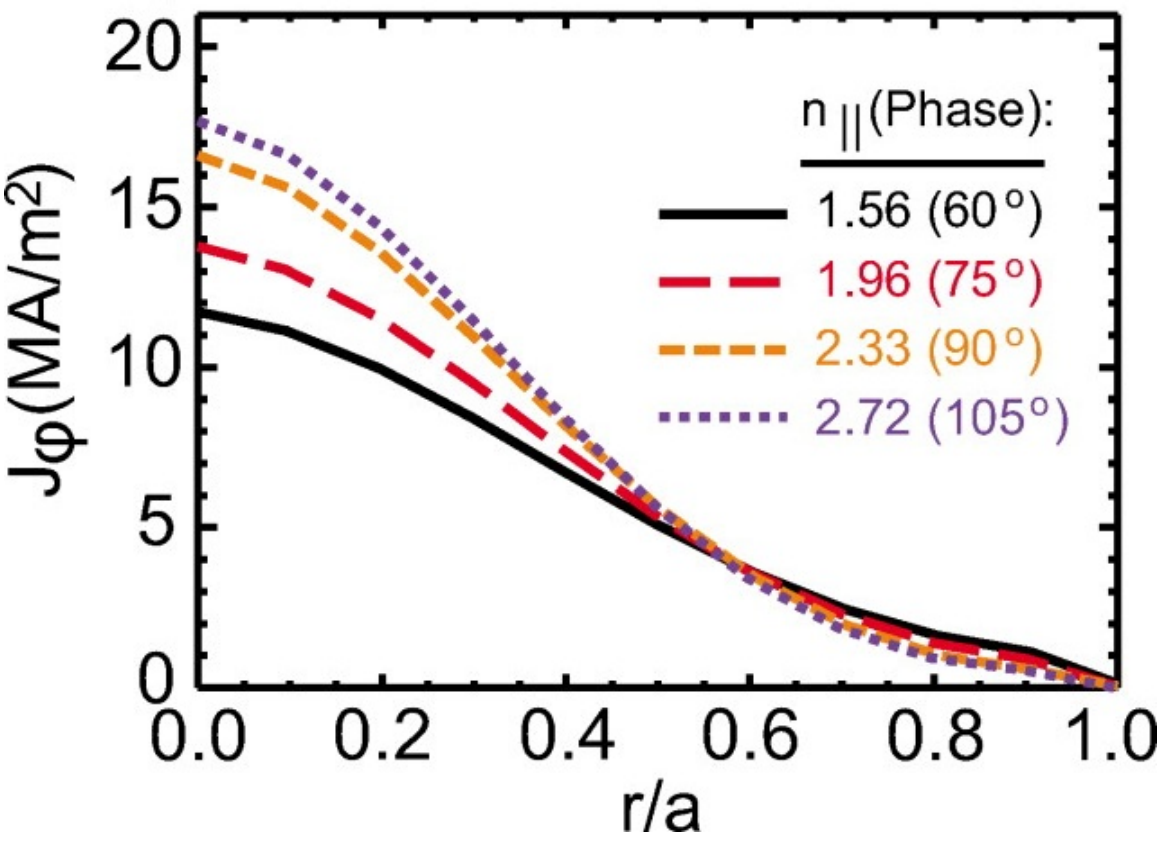

Figure 5-8: Current profiles produced by MSE-constrained EFIT for 4 LH antenna phasings. Here current profiles are represented by low order polynomials. In these discharges, $\overline{n}_e = 5 \times 10^{19}$ m^{-3}, $P_{LH} = 900$ kW, $B = 6$ T and $I_p = 800$ kA. Current profiles are shown for 60° ($n_{||} = 1.6$), 75° ($n_{||} = 1.9$), 90° ($n_{||} = 2.3$), and 105° ($n_{||} = 2.7$) phasings. Current profile data is the same as Figure 6 of Ref. [77].

5.3 Observed Lower Hybrid Density Limit

Late in the 2008 campaign, it was observed that the total HXR count rate was highly sensitive to plasma density in diverted discharges [78]. More specifically, an exponential drop in count rates as a function of density is observed. Near $\overline{n_e} = 1 \times 10^{20} m^{-3}$, the HXR signal drops at an even faster rate. Experiments in 2010 found much higher HXR count rates at high densities in inner wall limited discharges [83]. Figure 5-9 shows the HXR count rates as a function of line-averaged density for ramped density discharges. Diverted discharges with a large and small inner gap are shown, as well as limited discharges.

This density limit cannot be explained by the lower hybrid accessibility condition, because the values of launched $n_{||}$ were greater than the accessible $n_{||}$ everywhere in the plasma. Furthermore, the scaling of x-ray flux with magnetic field indicates that accessibility is not the main cause of the observed density limit [78].

Other tokamak experiments, such as Alcator C and FTU, cite parametric decay instabilities as the cause of their observed LH density limits [84, 85], albeit at density limits typically twice as high as in the C-Mod diverted discharges. This is not thought to be the cause of the C-Mod LH density limit because neither strong downshifted side bands nor broadening of the pump wave are observed [78].

Based on observed changes in SOL parameters during LH operation and simulations using a version of Genray that includes a SOL model, it is postulated that at high densities the majority of lower hybrid power in diverted discharges is deposited in the SOL. In high density simulations, many rays do not propagate to the hotter core plasma region where they can be absorbed via collisionless Landau damping and instead propagate for significant distances in the plasma periphery and SOL [78]. Some rays intersect with the divertor plates, where they are terminated (in the model) and never penetrate the core plasma.

Figure 5-10 shows the count rates predicted by Genray/CQL3D in diverted discharges using three different treatments of the SOL: no SOL model (all ray tracing within the last closed flux surface), a SOL with no collisional damping, and a SOL

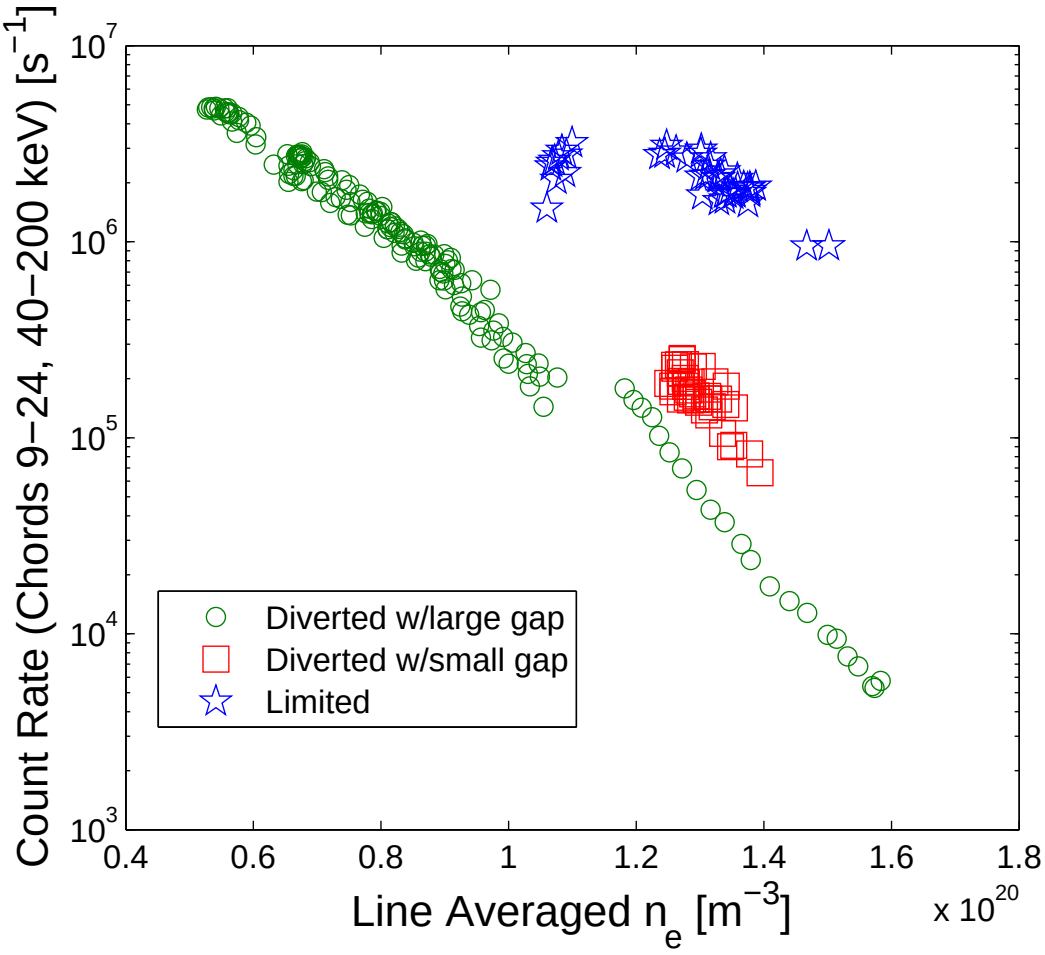

Figure 5-9: HXR count rates as function of density for USN, LSN, and limited ramped density discharges. In these discharges, $n_{||} = 1.9$, $I_p = 800$ kA and $B = 5.4$ T. Both USN and LSN diverted discharges (green circles) show an exponential decrease in count rates as a function of density. Limited discharges (blue stars) show significantly higher count rates at high density. Some x-ray flux is recovered in diverted discharges with a small (3-5 mm) inner gap (red squares). Figure courtesy of Greg Wallace.

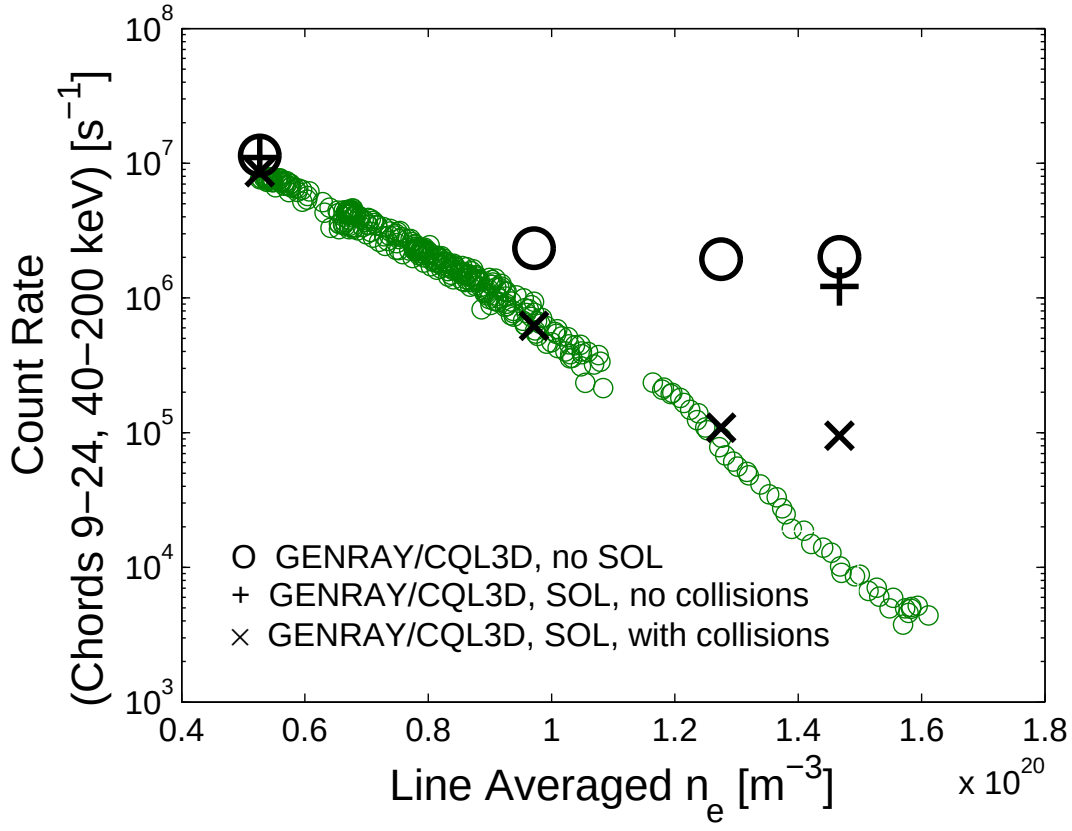

Figure 5-10: Predicted HXR count rates using Genray with no SOL (black circles), Genray with a SOL model (black + symbols) and Genray with collisional absorption in the SOL (black × symbols). Green circles are experimental count rates. All simulations and experimental data correspond to $n_{||} = 1.9$, $I_p = 800$ kA and $B = 5.4$ T. Figure courtesy of Greg Wallace.

model that includes collisional damping. It is clear from this figure that the model is in much better agreement with experimentally observed count rates when a damping mechanism is included in the SOL model. Suggested mechanisms for power deposition in the SOL include scattering from the wall, scattering from density fluctuations in the SOL, collisional damping, or interactions with neutral atoms [78].

The HXR spatial profiles for several densities are shown in Figures 5-11 (USN), 5-12 (LSN), and 5-13 (limited). In each figure, the upper plot shows the x-ray spatial profiles integrated over all energies for several densities. The lower plot shows the same profiles normalized to total countrate.

These data clearly show the decrease in x-ray signal with increasing density for all conditions shown. The overall shape of the spatial profiles also changes with density for diverted discharges. In upper single null (Figure 5-11), a "shoulder" feature is evident on the chords passing through the upper part of the plasma, while in lower single null (Figure 5-12), the shoulder feature appears on the lowermost chords. At higher densities, the prominence of the shoulder feature increases (relative to the rest of the profile). In limited discharges (Figure 5-13), there is no shoulder on either side of the plasma and the HXR profile shapes do not change much with changing density.

The cause of the shoulder feature is unclear, though the fact that it is on the same end of the plasma as the active x-point for diverted discharges and absent in limited discharges suggests that it may come from activity in the divertor.

The cause of the sharp decrease in x-rays at high density on C-Mod is still not fully understood. It is clear from the overall magnitude of x-ray flux and the shape of the HXR profiles that limited and diverted discharges respond very differently to LH power at high densities.

Proposed methods of avoiding or partially avoiding the density limit are optimization of plasma topology and use of a higher electron temperature (which is guaranteed in ITER) to guarantee single pass absorption in LH experiments. Further experiments and modeling must be carried out to optimize LH performance at high densities on C-Mod and future experiments.

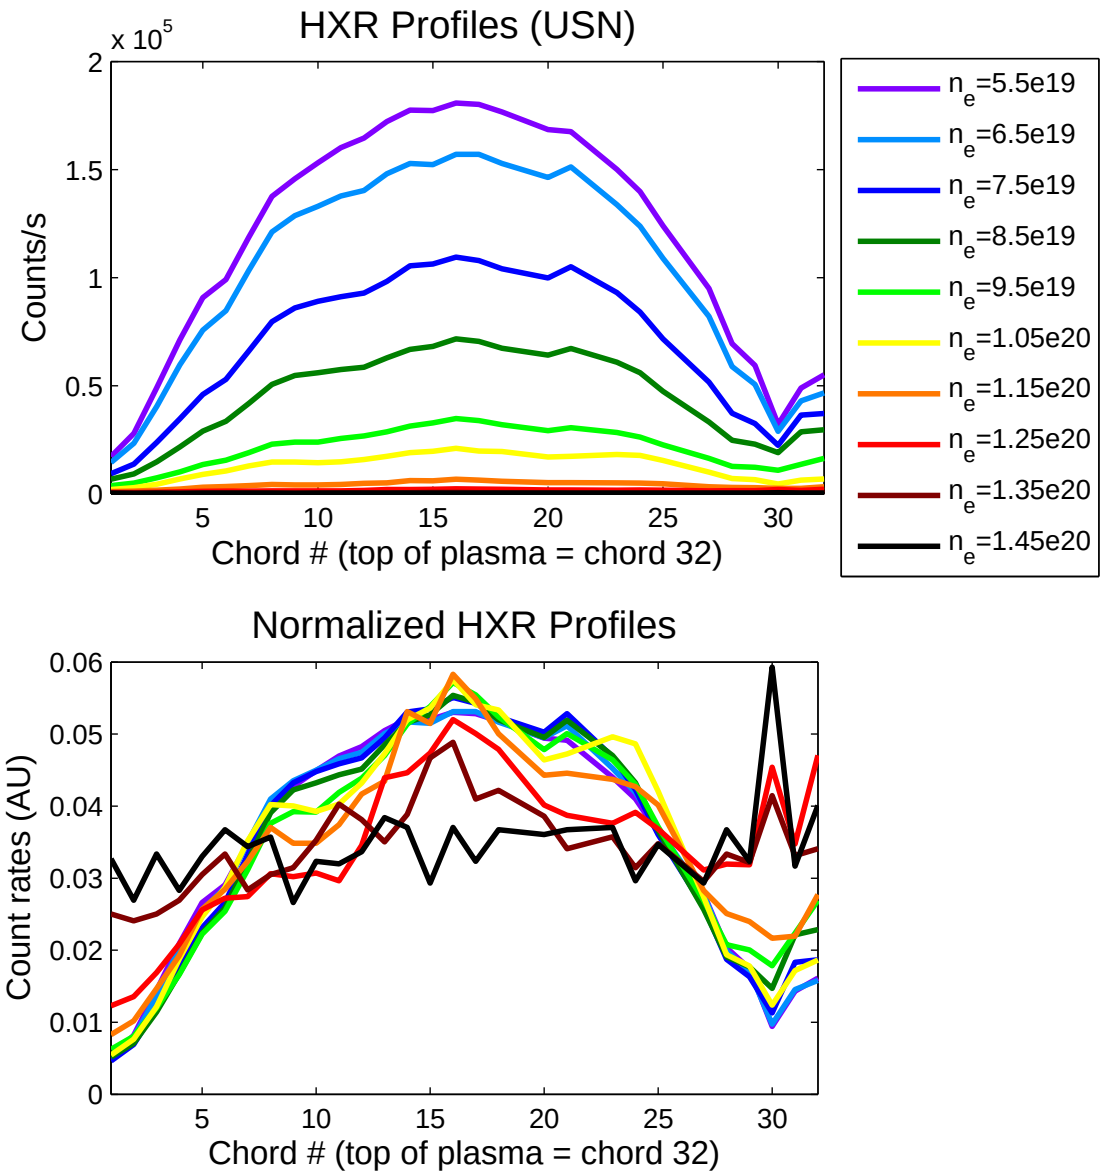

Figure 5-11: HXR profiles for varying density (USN), taken from shots 1080429003/4/7/8/9/14/15/16/17/18, and 1080513006/7/15/16/22/23/24. Densities in the legend are line-averaged. In these discharges, launched $n_{||}$=1.9, B =5.4 T, and I_p=800 kA. The upper plot shows count rates and the lower plot shows count rates normalized to the total count rate (on all chords). In USN, a "shoulder" feature appears in the HXR profiles at the top of the plasma. This shoulder becomes more prominent in the normalized HXR profiles at higher densities.

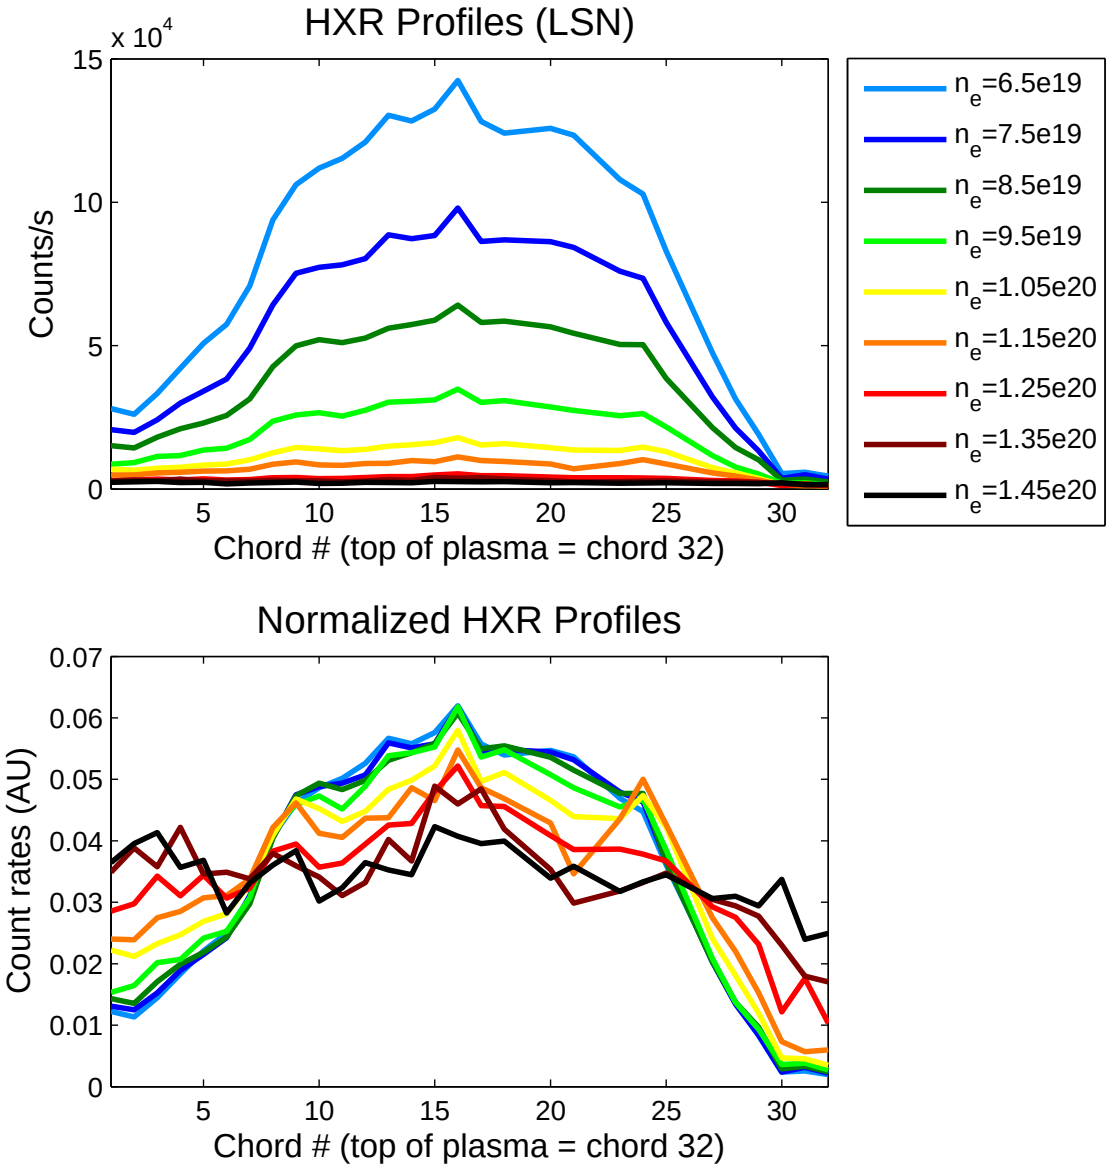

Figure 5-12: HXR profiles for varying density (LSN), taken from shots 1100806007/19/21/23. Densities in the legend are line-averaged. In these discharges, launched $n_{||}$=1.9, B =5.4 T, and I_p=800 kA. The upper plot shows count rates and the lower plot shows count rates normalized to the total count rate (on all chords). In LSN, a "shoulder" feature appears in the HXR profiles at the bottom of the plasma (in this case, it is not as prominent as the shoulder at the top of the plasma in the USN case). This shoulder becomes more prominent in the normalized HXR profiles at higher densities.

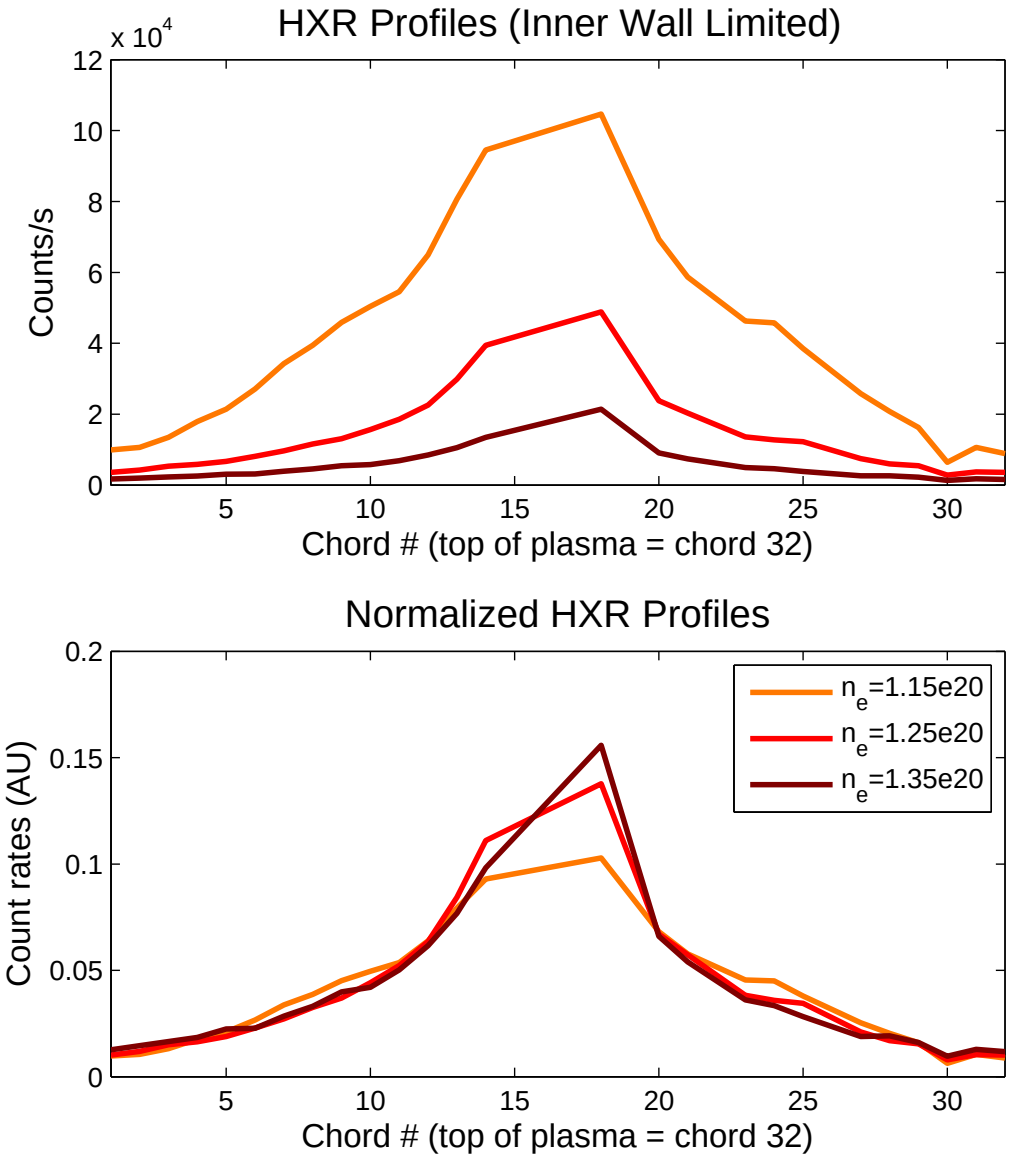

Figure 5-13: HXR profiles for varying densities (inner wall limited), all above the "density limit" exhibited by diverted discharges. These data are taken from shots 1100806027, 1100806028, and 1100806030. Densities in the legend are line-averaged. In these discharges, launched $n_{||}$=1.9, B =5.4 T, and I_p=800 kA. The upper plot shows count rates and the lower plot shows count rates normalized to the total count rate (on all chords). In the limited plasma, the shape of the HXR profile remains fairly consistent at different densities. No "shoulder" feature appears at the top or bottom of the plasma. In this experiment, chords 15-17 on the HXR camera had electronics problems, which we may have been due to a high x-ray flux from thick target bremsstrahlung on the inner wall. HXR data for chords 15-17 have been removed in these plots. There are not HXR data for the entire density range for this case.

5.4 Discussion of Ray Tracing/Fokker-Planck Model

The majority of simulations presented in this chapter corresponded to a launched $n_{||}$ in the range of 1.6 to 1.9, and have shown moderate agreement with experiment. In Section 6.5, simulations of a phasing scan are presented, showing that much better agreement between the model and experiment is obtained at higher $n_{||}$. We posit that this better agreement at higher $n_{||}$ is related to the shorter ray paths (or, equivalently, the fewer passes into and out of the plasma) necessary for those rays to damp. An analysis of the validity of the ray tracing/Fokker Planck model based on the simulations in Chapters 5 and 6 is presented in Chapter 7.

Chapter 6

Lower Hybrid Power Modulation Experiments

In order to use LHCD as an effective tool for controlling the current profile, we must be able to predict and control the radial location of driven current. Furthermore, we must be reasonably confident that the fast electrons do not significantly diffuse or convect away from the radial location where they are born before they thermalize.

To this end, a series of lower hybrid power modulation experiments were performed in order (1) to quantify fast electron transport (see Sections 6.1-6.3), (2) to examine trends in LH power deposition location (see Section 6.4), and (3) to validate RF modeling tools (see Section 6.5). A schematic of the modulated LH power trace is shown in Figure 6-1.

Power modulation is necessary for an experiment that will quantify fast electron transport, since it is not possible to separate out the effects of power deposition and transport in a steady-state experiment. A data analysis method has been developed for determining diffusivities and convection velocities of fast electrons for hollow inverted profiles. The methodology is described in Sections 6.1-6.2 and the measured transport coefficients are presented in Section 6.3.

Steady-state experiments are difficult to perform because plasma parameters must be kept constant for a long period of time. Even if plasma parameters to which the HXR is extremely sensitive (such as density) are successfully kept constant, the

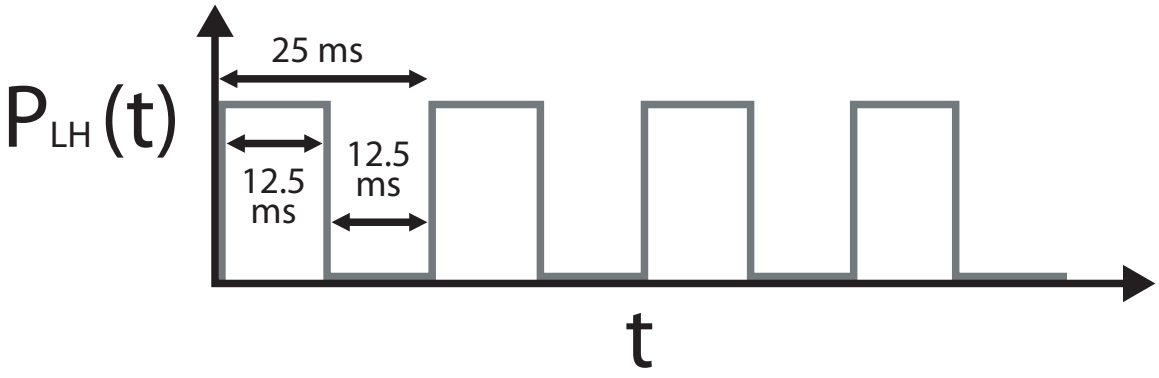

Figure 6-1: Schematic of lower hybrid power modulation.

current profile will evolve when the LHCD pulse exceeds the current diffusion time. It can be difficult to interpret the HXR spectra when plasma parameters are changing in time.

One advantage of a dynamic experiment over a steady-state experiment is that the HXR data are averaged over many RF pulses, in effect averaging out some of the variation due to fluctuations in plasma parameters. Additionally, relatively low power can be used in these experiments because the HXR count statistics can be built up through averaging over many pulses. At these low power levels, very little current is driven non-inductively, leaving the plasma conditions and magnetic geometry relatively unchanged[1].

The LH power modulation experiments are an effective way to look at the response of the x-rays to parameter scans for the reasons outlined above, namely, (1) because the LH power does not significantly perturb the plasma parameters and equilibrium, and (2) because small variations in plasma parameters are averaged over many pulses. A phasing scan and current scan are presented in Section 6.4.

The LH modulation experiments also lend themselves particularly well to modeling for the same reasons. Realistic plasma parameters and magnetic geometries can be input into the codes precisely because the experiment does not significantly alter the

[1]The data presented in this chapter come from experiments in which the peak LH power was $\approx$ 400 kW and thus average LH power was $\approx$ 200 kW. Though the loop voltage does respond to the modulation, the change in loop voltage during modulation is $\approx$ 10%, so most of the current is still being driven inductively.

plasma. Simulations of the phasing scan are presented in Section 6.5.

In the experiments described in this chapter, the LH power was square-wave modulated with a period of 25 ms: 12.5 ms of LH power at peak power and 12.5 ms at minimum power. This time-scale was purposefully chosen to be longer than the predicted slowing-down time of a few ms and shorter than the current relaxation time of a few hundred ms, so that the plasma is in a quasi-equilibrium state at the end of each LH pulse.

Figure 6-1 indicates that the power goes to 0 during the minimum. During the campaigns prior to 2010, the power was set to drop to a low, non-zero value during the "power off" phase of the modulation. In the 2010 campaign, the power was actually dropped to 0 during the "power off" phase.[2]

Lower Hybrid modulation experiments have been carried out on C-Mod for multiple launched peak $n_{||}$ values (1.55, 2.33, and 3.1), as well as at different plasma currents. Though these experiments have been carried out at multiple plasma densities, only the lowest density experiments provided enough x-ray signal for transport coefficient analysis. Data were taken for several discharges at each condition, with 20 to 40 modulations per discharge. The HXR signal was then averaged over several LH pulses to create a highly time-resolved, high signal-to-noise x-ray time trace of the rise, flat-top, and slowing-down of the fast electrons. Time traces of relevant plasma parameters for one group of similar discharges are shown in Figure 6-2.

6.1 Transport Model and Assumptions

The HXR data from the LH modulation experiments can be fit to a simplified transport model in order to solve for relevant transport parameters. Let $n(r,t)$ be a fast electron number density. The spatially inverted HXR signal should be proportional

[2]In both cases, the decision to drop the power to a low value or completely to 0 was made in order to avoid nuisance trips. In the LH1 launcher protection system, trips on the turn-off were avoided by keeping the power at a small but finite value. During LH2 launcher operation, the system was intentionally tripped after each pulse turn-off in order to avoid trips during the power turn-on because reflection coefficients are masked during restrikes. Masking reflection coefficients during changes in power level is a functionality that the LH2 coupler protection system was not explicitly designed to handle.

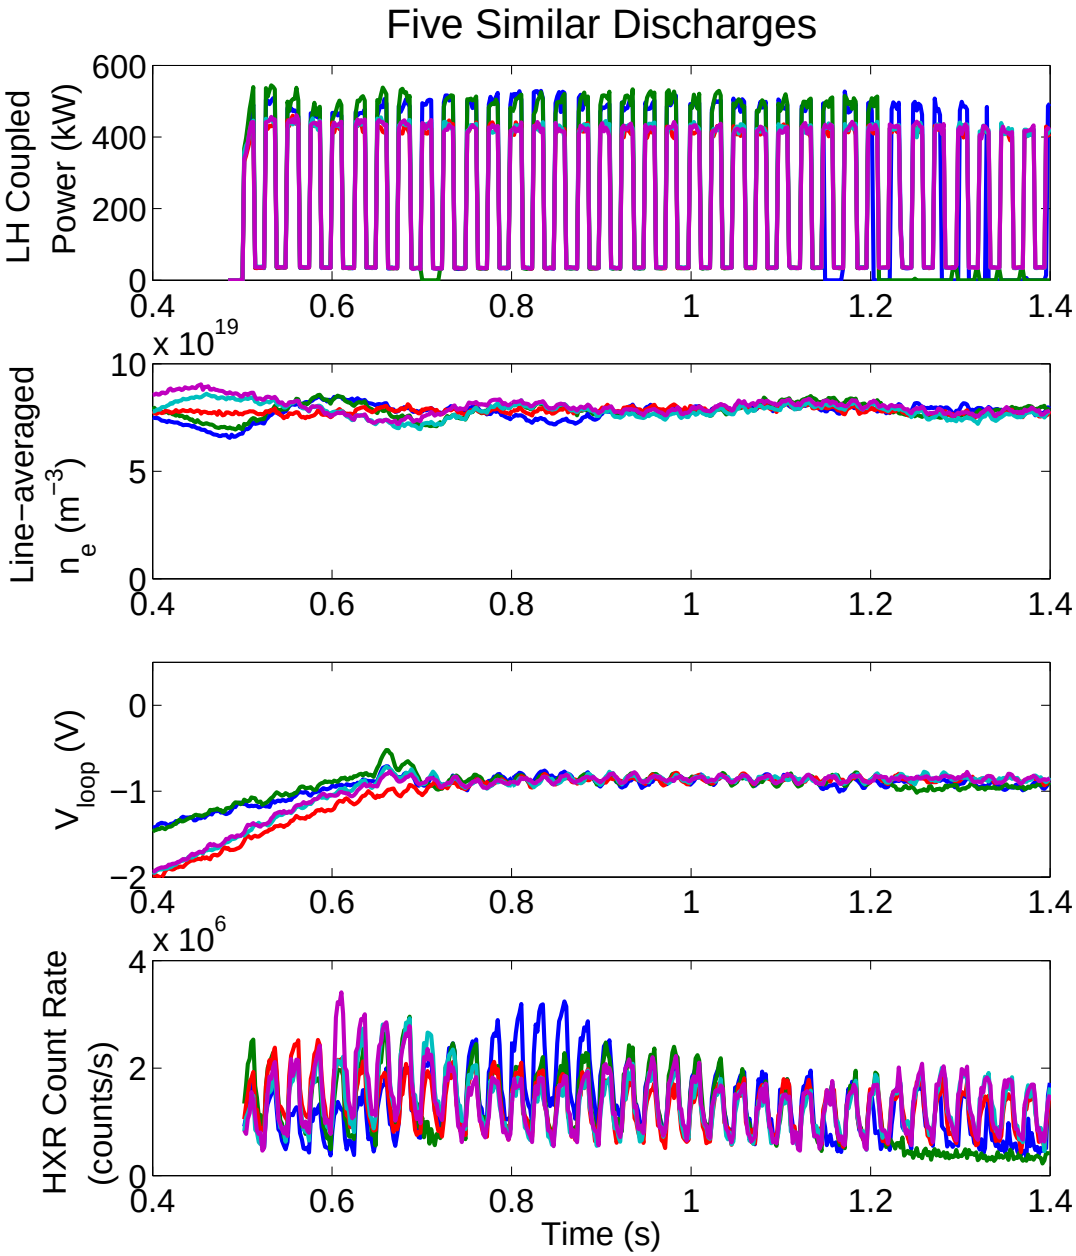

Figure 6-2: Time-traces of relevant plasma quantities during LH modulation experiments for several similar shots. Note that the line-averaged density and the loop voltage remain relatively constant during the experiment. In these discharges, 60° phasing was used, which corresponds to a peak $n_{||} = 1.55$, line-integrated density was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. $I_p = 800$ kA, B_T=5.4 T.

to $n(r,t)n_e(r)$, where $n_e(r)$ is the bulk electron density profile, so that $n(r,t)$ can be obtained from the inverted HXR signals simply by dividing by the bulk density profile.

We propose that $n(r,t)$ obeys the model transport equation:

$$\frac{\partial n}{\partial t} = -\frac{n}{\tau} - \nabla \cdot (nv) + \nabla \cdot (D \nabla n) + S(r,t), \tag{6.1}$$

where τ is a slowing-down time, D is a spatial diffusion coefficient, v is a drift velocity, and S is a source of fast electrons.

Because the LH modulation experiments are periodic, we should be able to decompose $n(r,t)$ into its Fourier components, as shown in Eq. 6.2, where ω is the experimental modulation frequency:

$$\begin{aligned} n(r,t) = n_0(r,t) + n_{1i}(r)\sin(\omega t) + n_{3i}(r)\sin(3\omega t) + n_{5i}(r)\sin(5\omega t) + \ldots \\ + n_{1r}(r)\cos(\omega t) + n_{3r}(r)\cos(3\omega t) + n_{5r}(r)\cos(5\omega t) + \ldots . \end{aligned} \tag{6.2}$$

Note that we could have included even harmonics of ω in Eq. 6.2, but there is no drive at those frequencies and therefore we expect and indeed find that there is very little HXR signal at those frequencies as well. In practice, we include only the fundamental frequency in the analysis because there is not enough signal-to-noise at higher harmonics.

We can also assume a form for the fast electron source term, $S(r,t)$, where we stipulate that the fast electron source term is proportional to the LH power trace (an odd function in time modulo a constant term):

$$S(r,t) = S(r)[const. + \sin(\omega(t)) + \frac{1}{3}\sin(3\omega(t)) + \frac{1}{5}\sin(5\omega(t)) + \ldots]. \tag{6.3}$$

We can substitute Eqs. 6.2 and 6.3 into Eq. 6.1 and collect sine and cosine terms, resulting in:

$$l\omega n_{lr}(r) = \frac{n_{li}(r)}{\tau(r)} + \nabla \cdot [n_{li}(r)\mathbf{v}] - \nabla \cdot (D\nabla n_{li}(r)) - \frac{S(r)}{l}, \tag{6.4}$$

$$l\omega n_{li}(r) = -\frac{n_{lr}(r)}{\tau(r)} - \nabla \cdot [n_{lr}(r)\mathbf{v}] + \nabla \cdot (D\nabla n_{lr}(r)), \tag{6.5}$$

where l refers to the harmonic number and $n_{lr}(r)$ and $n_{li}(r)$ are the real and imaginary components of the inverted, Fourier transformed HXR signal at the lth harmonic.

Unfortunately the power deposition profile, $S(r)$, is not well known. Nor is the power deposition highly localized, as it is in the case of ECCD modulation experiments [86]. To get around this problem, we use only Eq. 6.5, which relates the real and imaginary profiles to each other, and allows us to solve for the free parameters τ, v, and D without any knowledge of $S(r)$.

As noted already, only the fundamental frequency is used in practice, because the higher odd harmonics do not have high enough signal-to-noise for proper inversions. The imaginary (sine) component of the inverted signal can be thought of as the direct response to the drive (LH power time trace) since the source term, $S(r,t)$, is a sum of sine terms. The real (cosine) component of the inverted signal can be thought of as a phase lagged response to the drive. The ratio of the two amplitudes, therefore, is a measure of the time-response of the system, and can be used to calculate the slowing down time, τ.

In order to calculate v and D, we need reasonable estimates of the first and second spatial derivatives of the inverted x-ray profiles. As we shall see in section 6.2.5, this will require that both the real and imaginary inverted profiles are hollow.

6.2 Fitting Process

It takes a significant amount of processing to extract the parameters that we are solving for from the raw data. The raw data are a list of x-ray counts for several discharges, each count associated with both a time and an x-ray energy. The HXR data are binned in time bins that are synchronized with the LH pulses and binned

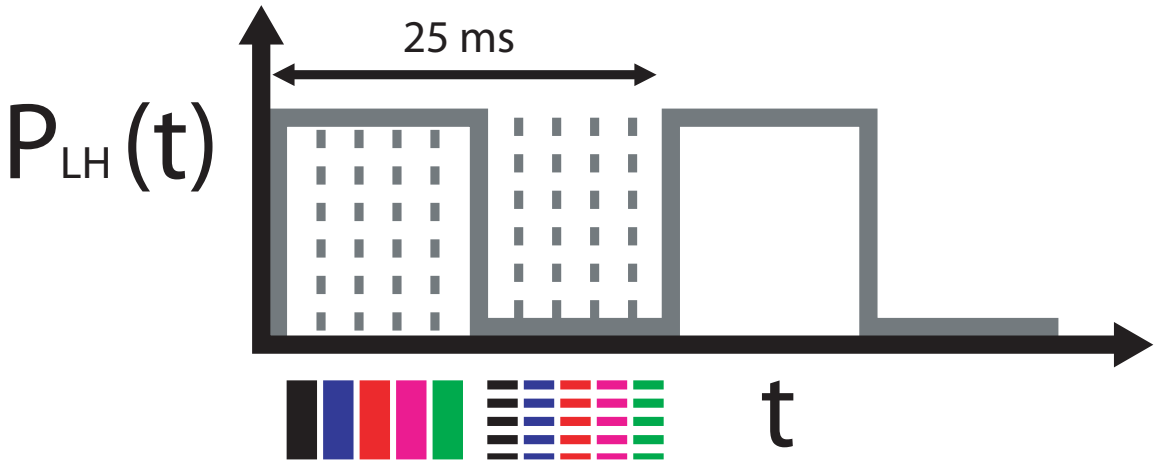

Figure 6-3: Schematic of boxcar binning with 10 time bins. The first 5 time bins correspond to times when the LH power was on and the last 5 time bins correspond to times when the LH power was off. Each LH pulse from each discharge with similar plasma conditions is binned in this way and the binned x-rays are averaged over every LH pulse, leading to (in this case) a high signal-to-noise x-ray signal with several energy bins and 10 total time bins. In the data analysis, more time bins (typically 50) are used for accurate Fourier transforming.

in several energy bins. HXR data associated with LH pulses that tripped early or never fired are removed from the data set. The remaining data are averaged over all pulses to obtain an averaged HXR time trace corresponding to a single pulse. The data are then Fourier transformed in time and the real and imaginary components of the resulting signal are fit to a curve in the energy dimension, which removes much of the noise due to count statistics. The fit profiles are then spatially inverted. The inverted Fourier components are sent through an algebraic solver that computes τ, v and D. This process is performed for each energy bin. The details of the data analysis described above are discussed in this section.

6.2.1 Boxcar Binning

After HXR data associated with misfired LH pulses is removed, it is "boxcar binned," or averaged over each LH pulse with some number of time bins per pulse (typically 50) that determines the time-resolution of the data. A schematic of boxcar binning with only 10 time bins is shown in Figure 6-3. The first 5 time bins correspond to

times when the LH power is on, during the fast electron tail build-up. The last 5 times bins correspond to times when the LH power is off, during the decay of the fast electron tail. Each LH pulse from each discharge with similar plasma conditions is binned in this way and the binned x-rays are averaged over every LH pulse, leading to (in this case) a high signal-to-noise x-ray signal with several energy bins and 10 total time bins.

The HXR chordal measurements for these 10 time bins for energies 40-60 keV are shown in Figure 6-4. In this experiment, 120° phasing was used, which corresponds to a peak $n_{||} = 3.12$, line-integrated density was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. The plasma current was $I_p = 800$ kA. It is clear from this figure that 25 ms is long enough to resolve the build-up of fast electrons to a quasi-steady-state and the complete decay of fast electrons to a baseline level. Though 10 time bins is a good number to illustrate the build-up and decay of fast electrons, for the actual data analysis we typically use 50 time bins to allow for accurate Fourier transforming.

6.2.2 Fourier Transform

After the boxcar binning is performed, the binned data are Fourier transformed in time. This is done using the FFT routine in Matlab, which assumes that data start at $t = 0$. The Fourier transformed data are then phase shifted so that the time base actually corresponds to the middle of each time bin, rather than the beginning. The real and imaginary parts of the Fourier transformed data are shown in the top halves of Figures 6-5 and 6-6. Since the drive (LH power time trace) is an odd function of time, we expect the imaginary (sine) component of the Fourier transform to be larger than the real (cosine) component, at least at the fundamental frequency. Indeed, this is what we see; the imaginary profiles are about 50% greater in peak magnitude than the real profiles.

In some modulation experiments, the raw data are detrended [86], a process that removes the baseline signal level before processing begins. Because of the very poor signal-to-noise in small time bins, it turns out to be rather difficult to detrend the

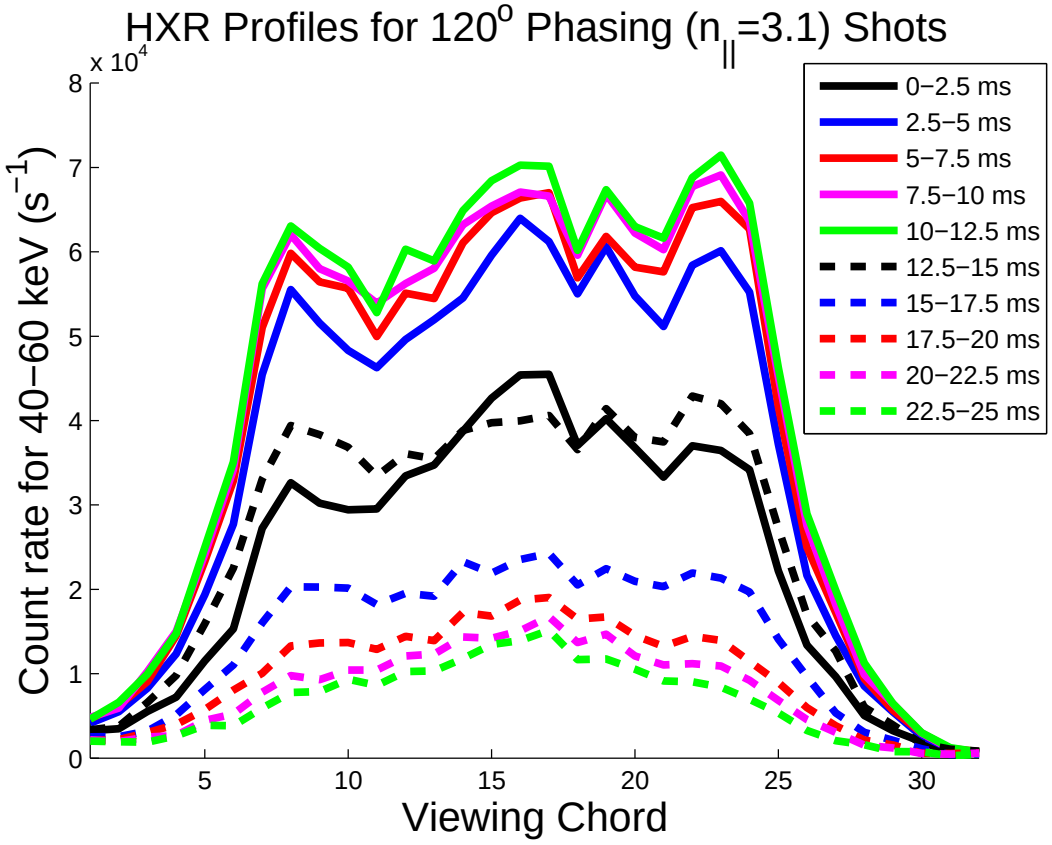

Figure 6-4: Boxcar binned chordal HXR measurements for 10 time bins in photon energies ranging from 40-60 keV. In this experiment, 120° phasing was used, which corresponds to a peak $n_{||} = 3.12$, line-integrated density was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. The plasma current was $I_p = 800$ kA. Solid lines correspond to times when the LH power was on and dashed lines to times when LH power was off.

HXR data in a reasonable way. Boxcar binning and Fourier transforming the data is mathematically equivalent to detrending the signal by simply removing the time average of each signal corresponding to a single spatial chord and single energy bin. This is because we ignore the 0th harmonic (constant term) in the Fourier transform, which is simply the time-averaged signal.

6.2.3 Energy Fitting

After the Fourier transform is performed, the real and imaginary part of each spatial chord is fit to an energy curve in order to remove statistical scatter from the measurement, as discussed in Section 3.0.1. This is particularly important for the actual analysis, where up to 80 energy bins might be used, dramatically reducing the signal-to-noise for any given bin. The energy fit, rather than the actual data, is then used in the rest of the data analysis. Figures 6-5 and 6-6 show the real and imaginary part of the Fourier transformed profiles before and after the energy fitting is performed.

6.2.4 Spatial Inversion

The Fourier transformed data are then spatially inverted, as described in Section 3.0.2. Spatial inversions for the real and imaginary components of the Fourier transformed data for the $n_{||} = 3.12$, $I_p = 800$ kA case is shown in Fig. 6-7

It could be debated whether the Fourier transforming or the spatial inversion should be performed first. If Fourier transforming is done first, then the real and imaginary components must be inverted separately, resulting in 2 inversions per energy bin. If the spatial inversion is done first, then every single time bin must be inverted separately, which results in 50 inversions per energy bin. Not only would it be time-consuming to perform that many inversions, but the low signal-to-noise in each individual time bin has a deleterious effect on the spatial inversion, resulting in anomalous spatial features.

Since both the inversion and the Fourier transform are linear operations and they operate along different dimensions of each energy bin (one along the spatial dimension

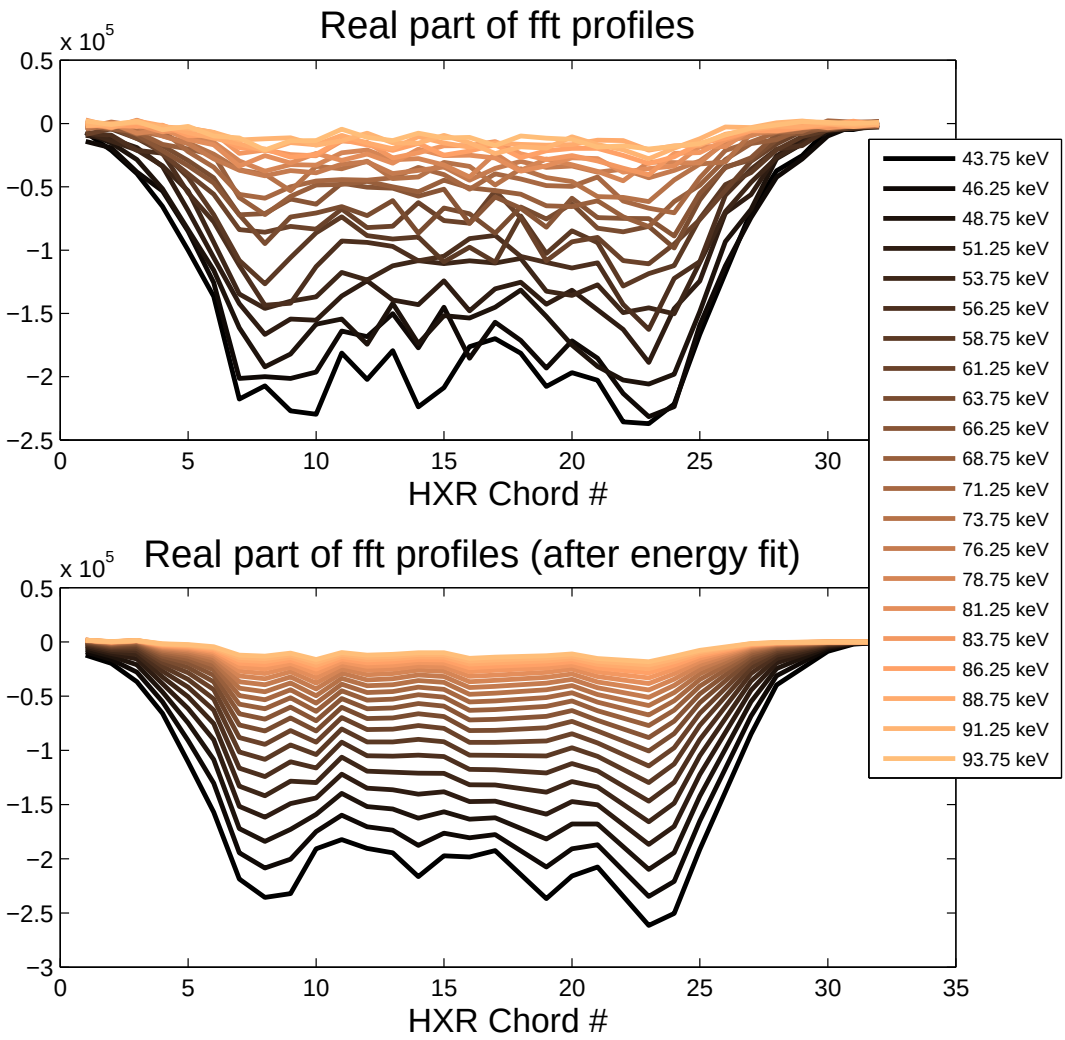

Figure 6-5: Real part of FFT profiles before and after energy fitting. In this experiment, 120° phasing was used, which corresponds to a peak $n_{||} = 3.12$, line-integrated density was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. $I_p = 800$ kA, B_T=5.4 T.

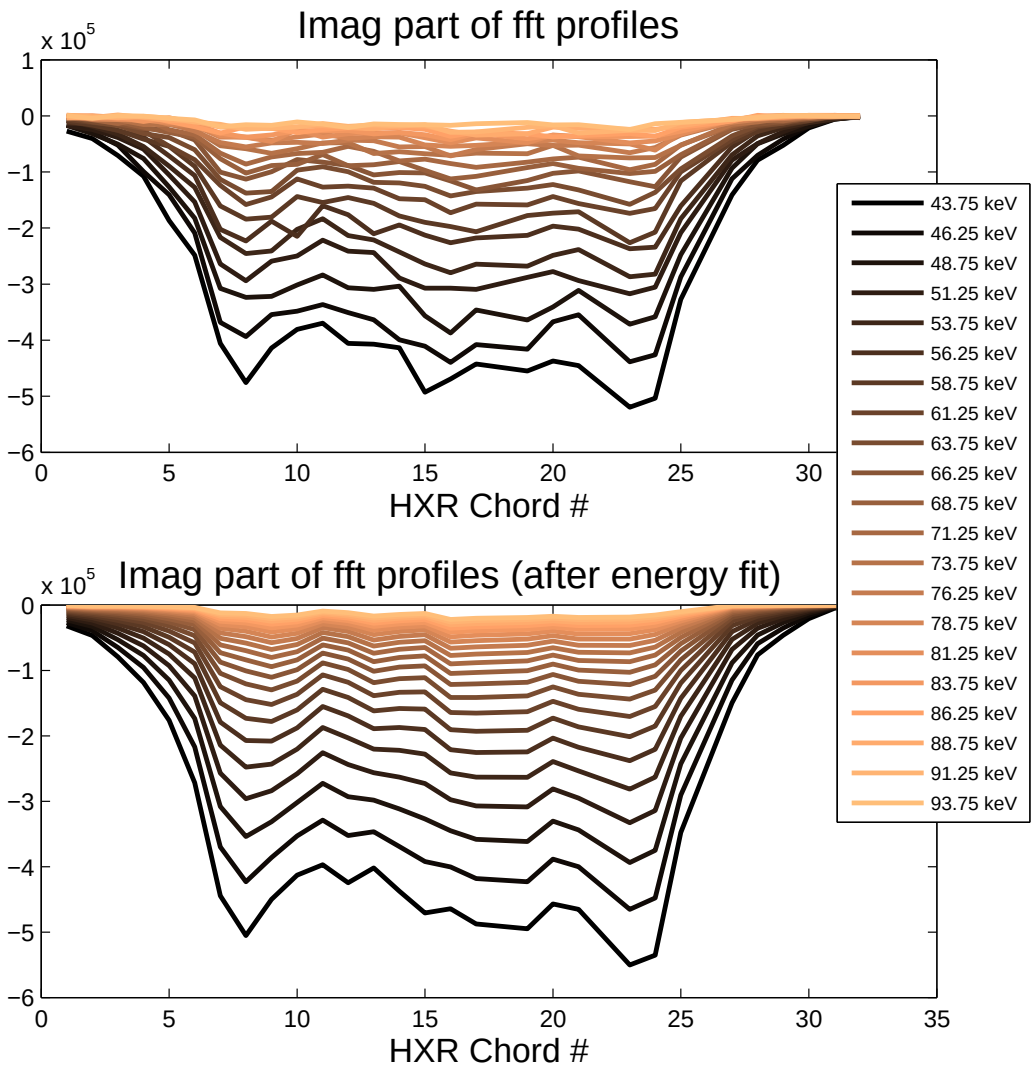

Figure 6-6: Imaginary part of FFT profiles before and after energy fitting. In this experiment, 120° phasing was used, which corresponds to a peak $n_{||} = 3.12$, the line-integrated density was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. $I_p = 800$ kA, B_T=5.4 T.

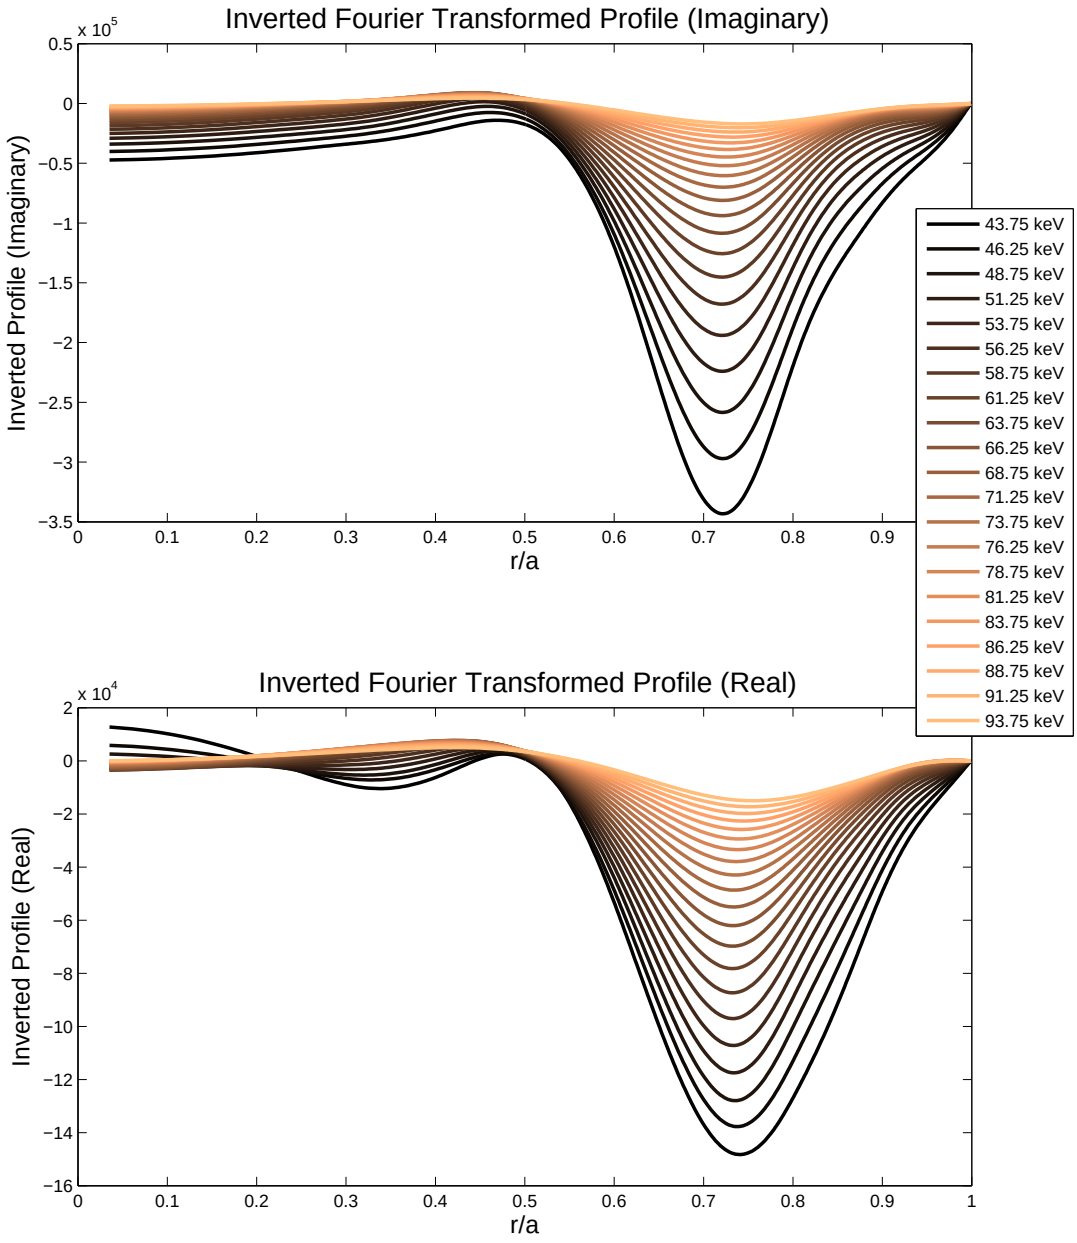

Figure 6-7: Inverted real and imaginary components of FFT profiles. In this experiment, 120° phasing was used, which corresponds to a peak $n_{||} = 3.12$, Line-integrated density was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. Plasma current was $I_p = 800$ kA.

and one along the time dimension), it seems reasonable that either order of operations should be a valid approach and so we choose the order which is better for our particular data set, with the Fourier transform first. It may naïvely seem that by the associative property of matrix multiplication, the order of operations should not matter. It turns out that the degree of regularization involved in the inversion is highly dependent on the signal-to-noise ratio and thus the final state of the data does depend on the order used.

6.2.5 Algebraic Fitting

Finally, after the data has been Fourier transformed and spatially inverted, it is ready to be used to calculate fast electron transport. This is performed in Matlab with an algebraic algorithm that relates the real and imaginary inverted profiles to each other.

In the derivation that follows, we relate the real and imaginary components of the FFT profiles, taking into account the effects of slowing down, convection, and diffusion. Taking $l = 1$ (the fundamental frequency) in Eq. 6.5 yields:

$$\omega n_i(r) = -\frac{n_r(r)}{\tau(r)} - \nabla \cdot [n_r(r)\mathbf{v}] + \nabla \cdot (D\nabla n_r(r)), \tag{6.6}$$

where we have dropped the subscript 1 in n_{1r} and n_{1i}.

We can now use this relation to solve for τ, v and D. In order to do this, we must assume a functional form for $\tau(r)$. Since the electron slowing down time is inversely proportional to the bulk plasma density, we assume the form:

$$\tau(r) = \frac{\tau_0 n_0}{n_e(r)}, \tag{6.7}$$

where n_0 is the central bulk plasma density and τ_0 is the expected slowing down time for a plasma of density n_0. The dependence of τ on Z_{eff} is not accounted for, under the assumption that Z_{eff} is relatively constant over the region of interest.

Combining Eq. 6.6 with Eq. 6.7 results in:

$$\omega n_i(r) = -\frac{n_r(r)n_e(r)}{\tau_0 n_0} - \nabla \cdot [n_r(r)\mathbf{v}] + \nabla \cdot (D\nabla n_r(r)). \tag{6.8}$$

In principle, v and D are also profiles as a function of r. However, we do not know their functional forms and do not have enough information to solve for them. Instead, we can solve for each at a single radial location, based on the local derivatives of the off-axis peak of hollow real and imaginary inverted profiles. In this local region, we assume $\mathbf{v} = v\hat{\mathbf{r}}$ and D are slowly varying in space and therefore can be pulled out of the spatial derivatives, as done in Eq 6.9.

$$\omega n_i(r) = -\frac{n_r(r)n_e(r)}{\tau_0 n_0} - v\nabla \cdot [n_r(r)\hat{\mathbf{r}}] + D\nabla \cdot (\nabla n_r(r)). \tag{6.9}$$

Now n_r and n_i and their derivatives can be obtained from the inverted HXR profiles, and $n_e(r)/n_0$ can be obtained from a spatially-resolved density diagnostic, such as the Thompson Scattering diagnostic. Then the radial function on the left sides of Eq. 6.9 must be linear combinations of the radial functions on the right side of Eq. 6.9 with the proper values of the parameters τ_0, v, and D.

Recall that the imaginary component of the FFT is the direct response of the HXR to the LH power trace, while the real component of the FFT is the time-lagged response. We are looking for differences between the real and imaginary components of hollow inverted profiles to see how the peak changes in time. Changes in the magnitude of the peak are indicative of slowing down time. Changes in the location of the peak are indicative of convection. Changes in the width of the peak are indicative of diffusion.

This simplified intuitive explanation does not tell the whole story of course. For example, changes in peak location can also be due to longer slowing down times at the edge of the plasma due to lower bulk density. Changes in the magnitude of the peak can also be caused by particles diffusing away from the peak, not just slowing down. All three processes (slowing down, convection, and diffusion) are interlinked and must be properly accounted for when solving for the various transport coefficients of interest.

To solve for τ_0, v, and D, we first determine the radial location of the peak of the imaginary inverted profile, which we shall call r_p. We shall refer to the slowing down time, convection velocity, and diffusivity at r_p as τ_p, v_p, and D_p respectively. In the $n_{||} = 3.1$ phasing case from the previous sections, r_p/a is at $r/a \sim 0.7$. Then the real and imaginary inverted HXR profiles as well as the bulk density profile can be expanded in a Taylor series centered at r_p:

$$n_i(r) = P_{i0} + P_{i2}(r - r_p)^2 + P_{i3}(r - r_p)^3 + ..., \tag{6.10}$$

$$n_r(r) = P_{r0} + P_{r1}(r - r_p) + P_{r2}(r - r_p)^2 + P_{r3}(r - r_p)^3 + ..., \tag{6.11}$$

$$n_{e,bulk}(r) = P_{e0} + P_{e1}(r - r_p) + P_{e2}(r - r_p)^2 + P_{e3}(r - r_p)^3 + ... \tag{6.12}$$

Substituting Eqs. 6.10, 6.11, and 6.12 into Eq. 6.9, and keeping all quantities to second order in $(r - r_p)$, yields on the left side:

$$\omega[P_{i0} + P_{i2}(r - r_p)^2]. \tag{6.13}$$

On the right side, there are 3 terms. The first, associated with slowing down, is:

$$-\frac{(P_{r0} + P_{r1}(r - r_p) + P_{r2}(r - r_p)^2 + ...)(P_{e0} + P_{e1}(r - r_p) + P_{e2}(r - r_p)^2) + ...}{\tau_p P_{e0}} \tag{6.14}$$

$$= -\frac{P_{r0}P_{e0} + (P_{r0}P_{e1} + P_{r1}P_{e0})(r - r_p) + (P_{r0}P_{e2} + P_{r1}P_{e1} + P_{r2}P_{e0})(r - r_p)^2 + ...}{\tau_p P_{e0}}. \tag{6.15}$$

The second term on the right side of Eq. 6.9, associated with convection, is:

$$-v_p \nabla \cdot [n_{lr}(r)\hat{\mathbf{r}}] = -v[\frac{dn_r}{dr} + \frac{1}{r}n_r] \tag{6.16}$$

$$\begin{aligned} = \quad & -v_p[P_{r1} + 2P_{r2}(r-r_p) + 3P_{r3}(r-r_p)^2... \\ & +(\tfrac{1}{r_p} - \tfrac{r-r_p}{r_p^2} + \tfrac{(r-r_p)^2}{r_p^3} + ...)(P_{r0} + P_{r1}(r-r_p) + P_{r2}(r-r_p)^2)...] \end{aligned} \tag{6.17}$$

$$= -v_p[(P_{r1}+\frac{P_{r0}}{r_p})+(2P_{r2}+\frac{P_{r1}}{r_p}-\frac{P_{r0}}{r_p^2})(r-r_p)+(3P_{r3}+\frac{P_{r2}}{r_p}-\frac{P_{r1}}{r_p^2}+\frac{P_{r0}}{r_p^3})(r-r_p)^2]+.... \tag{6.18}$$

The third term on the right side of Eq. 6.9, associated with diffusion, is:

$$D_p \nabla \cdot (\nabla n_{lr}(r)) = D[\frac{d^2 n_r}{dr^2} + \frac{1}{r}\frac{dn_r}{dr}] \tag{6.19}$$

$$\begin{aligned} = \quad & D_p[2P_{r2} + 6P_{r3}(r-r_p) + 12P_{r4}(r-r_p)^2... \\ & +(\tfrac{1}{r_p} - \tfrac{r-r_p}{r_p^2} + \tfrac{(r-r_p)^2}{r_p^3} + ...)(P_{r1} + 2P_{r2}(r-r_p) + 3P_{r3}(r-r_p)^2)...] \end{aligned} \tag{6.20}$$

$$= D_p[(2P_{r2}+\frac{P_{r1}}{r_p})+(6P_{r3}+\frac{2P_{r2}}{r_p}-\frac{P_{r1}}{r_p^2})(r-r_p)+(12P_{r4}+\frac{3P_{r3}}{r_p}-\frac{2P_{r2}}{r_p^2}+\frac{P_{r1}}{r_p^3})(r-r_p)^2]+.... \tag{6.21}$$

By grouping terms in powers of $(r - r_p)$, we can conveniently equate the left and right sides of Eq. 6.9 in matrix form:

$$M \begin{bmatrix} -\frac{1}{\tau_p} \\ -v_p \\ D_p \end{bmatrix} = \omega \underbrace{\begin{bmatrix} P_{i2} \\ 0 \\ P_{i0} \end{bmatrix}}_{\mathbf{P}_i}, \tag{6.22}$$

where

$$M = \begin{bmatrix} \frac{P_{r0}P_{e2}+P_{r1}P_{e1}+P_{r2}P_{e0}}{P_{e0}} & 3P_{r3} + \frac{P_{r2}}{r_p} - \frac{P_{r1}}{r_p^2} + \frac{P_{r0}}{r_p^3} & 12P_{r4} + \frac{3P_{r3}}{r_p} - \frac{2P_{r2}}{r_p^2} + \frac{P_{r1}}{r_p^3} \\ \frac{P_{r0}P_{e1}+P_{r1}P_{e0}}{P_{e0}} & 2P_{r2} + \frac{P_{r1}}{r_p} - \frac{P_{r0}}{r_p^2} & 6P_{r3} + \frac{2P_{r2}}{r_p} - \frac{P_{r1}}{r_p^2} \\ P_{r0} & P_{r1} + \frac{P_{r0}}{r_p} & 2P_{r2} + \frac{P_{r1}}{r_p} \end{bmatrix}. \tag{6.23}$$

We can then solve for τ_p, v_p, and D_p by the relation:

$$\begin{bmatrix} -\frac{1}{\tau_p} \\ -v_p \\ D_p \end{bmatrix} = \omega M^{-1} \begin{bmatrix} P_{i2} \\ 0 \\ P_{i0} \end{bmatrix}. \tag{6.24}$$

In order to construct M and $\mathbf{P}_i$, we must find the coefficients $P_{i0}, ...P_{i2}$, $P_{r0}, ...P_{r4}$, and $P_{e0}, ...P_{e2}$. In practice, we do this by first obtaining r_p, the radial location of the peak in the imaginary profile. We then fit 4^{th} order polynomials to the real and imaginary components of the Fourier transformed HXR profiles. Using polynomial fits to the data not only provides the model with coefficients, but also conveniently avoids the errors associated with taking numerical derivatives, since all the derivatives in this derivation were taken analytically.

Each polynomial fit is examined to ensure that a 4^{th} order polynomial is sufficient to describe the peak of the curve. Sample polynomial fits for both the real and imaginary profiles are shown in Fig. 6-8. In this example, the polynomial fit lies right on top of the original data. Indeed, even the quadratic terms of the polynomial fit do a good job of describing the experimental profiles without the higher order terms. It is a good idea to observe the closeness of the fit of the quadratic terms only since they are the only terms of $n_i(r)$included in the model.

Once M and $\mathbf{P}_i$ are constructed using the coefficients from the polynomial fit, τ_p, v_p, and D_p can be solved for using Eq. 6.24. In two simpler calculations, submatrices of M can be used to solve for τ_p only or τ_p and v_p only. In the former case, the magnitude of the two peaks are used to calculate τ_p (utilizing only the 0^{th} order term in Eq. 6.15. In the later case, both the magnitude and the peak location are used to

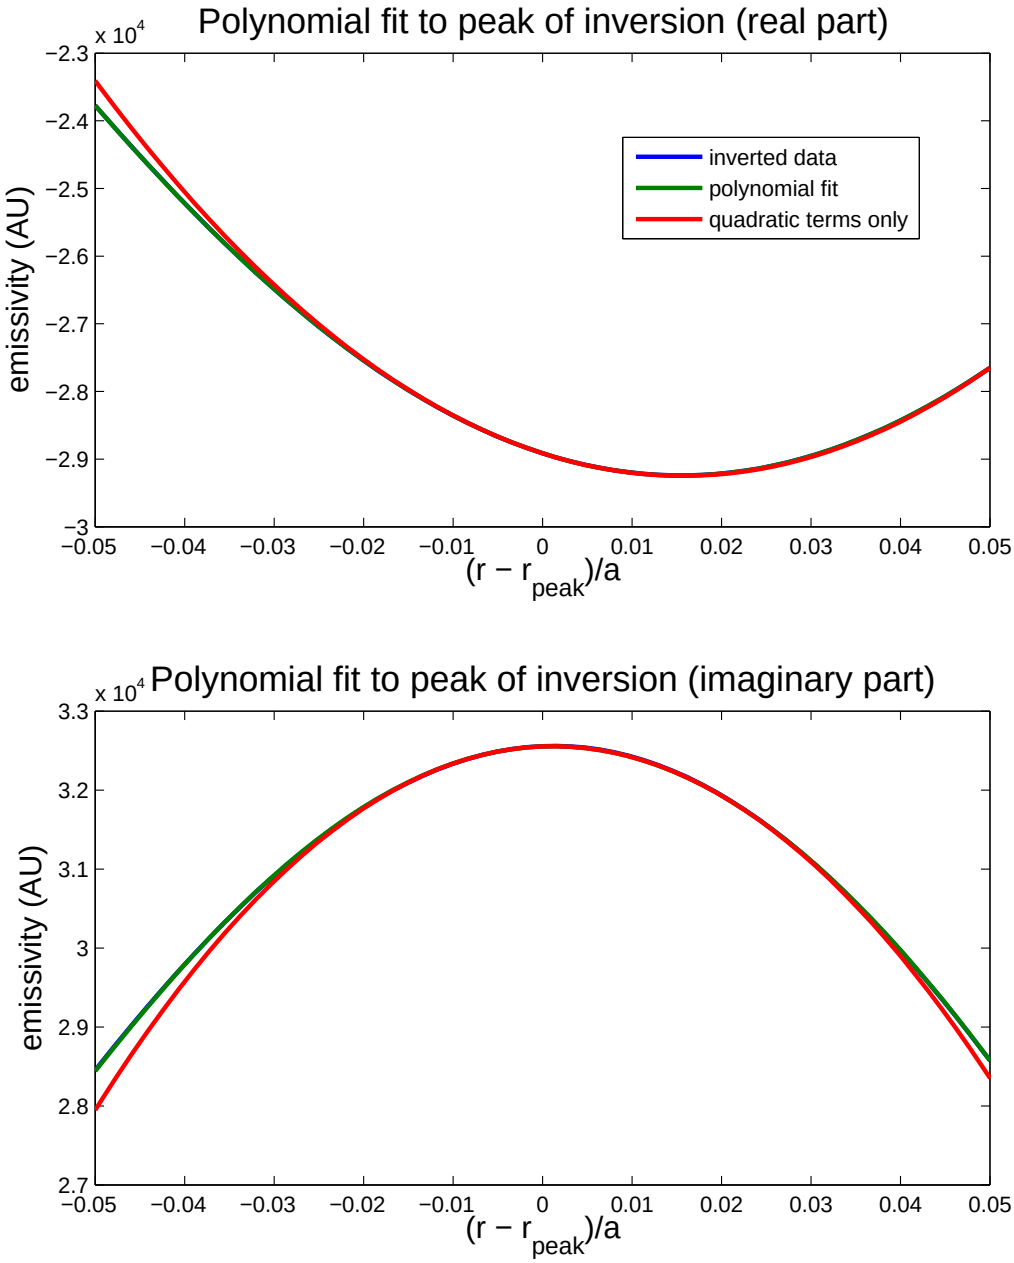

Figure 6-8: Sample polynomial fit to HXR profile peaks. Profiles are taken from the 94 keV energy bin of the $n_{||} = 3.12$, $I_p = 800$ kA case.

calculated τ_p and v_p (using the 0^{th} and 1^{st} order terms in Eqs. 6.15 and 6.18.

The calculated values of τ_p, v_p, and D_p can then be used to reconstruct the imaginary part of the profile. Each reconstruction is examined for accuracy of reproducing the original profile. A sample reconstruction is shown in Fig. 6-9. The 3 different reconstructions demonstrate the increasing sophistication of the model as we successively include the effects of convection and then diffusion. When only the slowing down effect is accounted for, the two profiles peak in different locations. When the effect of convection is added, the imaginary profile and the reconstruction are constrained to peak in the same location, but they have mismatching widths. Finally, when diffusion is added, the widths of the imaginary profile and reconstruction are most closely matched.

This analysis is performed for each energy bin of HXR data, resulting in several transport parameters as a function of HXR energy. Strictly speaking, Eq. 6.1 is not decoupled for each HXR energy bin, since the loss of fast electrons of one energy contributes to the gain of fast electrons of another energy. However, for simplicity they are treated here as decoupled.

6.2.6 Validation of Algebraic Solver

The algebraic solver was verified to ensure there were no errors in the analytic derivations of the equations to be solved or errors in the coding of those equations. In order to verify the solver, we use a Matlab PDE solver to evolve a "fast electron density" profile in time, given a source term, a slowing-down time, a convection velocity, and a diffusivity. The initial condition given to the PDE solver is zero density everywhere. The resulting time-dependent fast electron density profile is then Fourier transformed and run through the algebraic solver.

It is found that the algebraic algorithm recovers the transport coefficients entered into the PDE solver to within 2-5%. The PDE solver has the option of taking as input a Gaussian, Lorentzian, or parabolic source term as a function of space. The algebraic solver accurately recovers the input transport coefficients independent of the shape of source term used. A sample time-evolved fast electron density profile

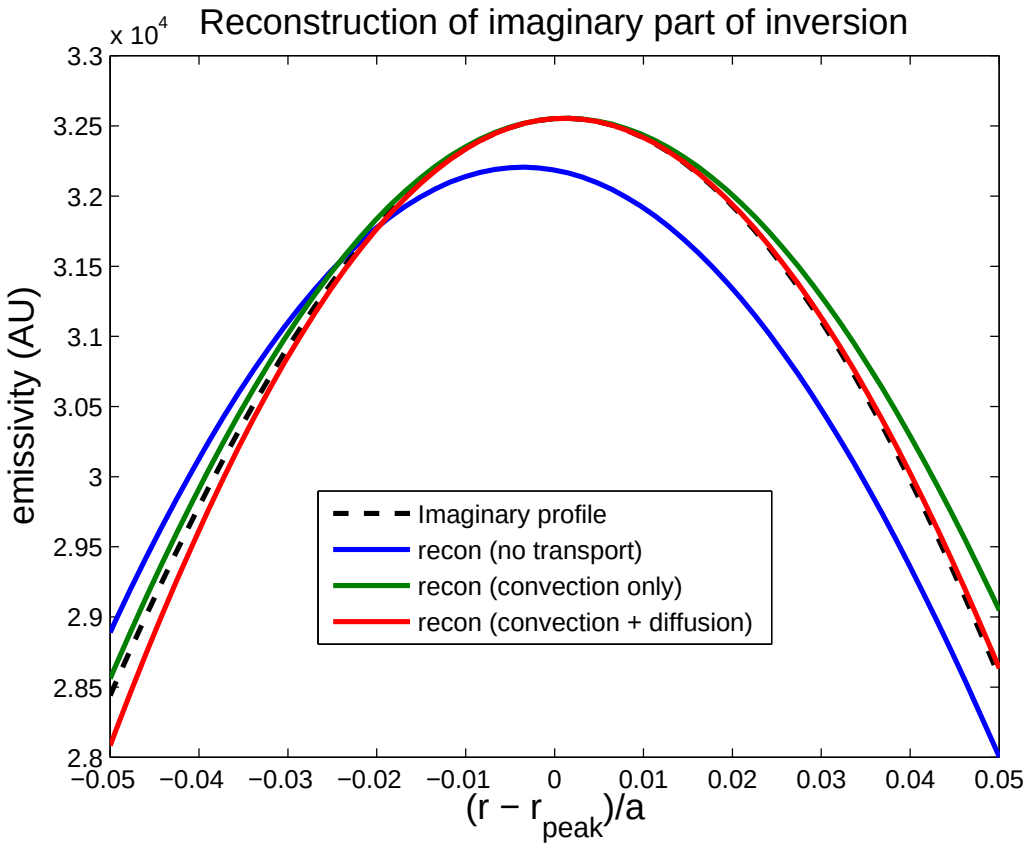

Figure 6-9: Sample reconstruction of imaginary profile based on calculated values of τ_p, v_p, and D_p. Profile and transport coefficients are taken from the 94 keV energy bin of the $n_{||} = 3.12$, $I_p = 800$ kA case. When both convection and diffusion are included in the model, $\tau_p = 3.7$ ms, $v_p = -0.1$ m/s, and $D_p = 0.0076$ m^2/s.

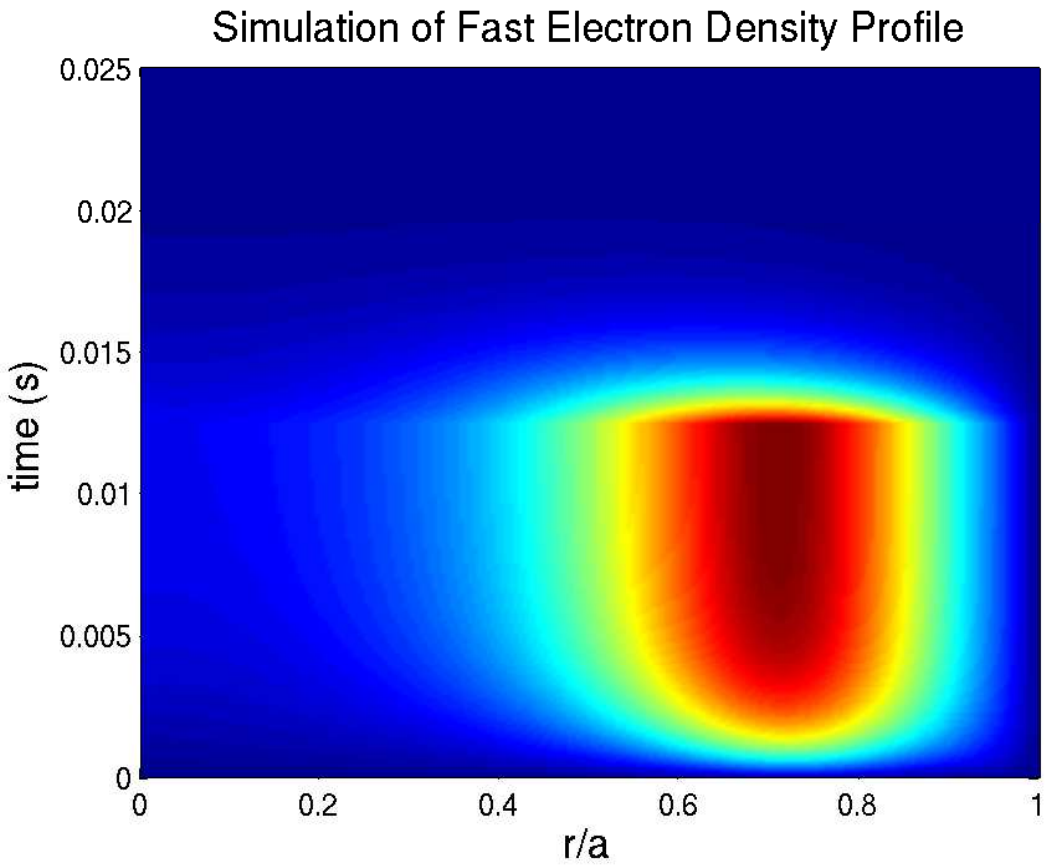

Figure 6-10: Sample time-evolved "fast electron density" from Matlab PDE solver. In this case, the PDE solver was given a diffusion coefficient of $D = 0.2$ m^2/s, a convection velocity of $v = -1$ m/s (the negative sign indicates an inward pinch), and a slowing-down time of $\tau = 2$ ms, with a Lorentzian-shaped source term. The algebraic solver recovered $D = 0.19$ m^2/s, $v = -1.00$ m/s, and $\tau = 1.94$ ms.

from the PDE solver is shown in Figure 6-10.

6.3 Measured Transport Coefficients

Earlier calculations of LH-driven fast electron diffusion coefficients have been deduced from experiments and modeling on Alcator C [36], PBX-M [46], Tore Supra [56, 87], ASDEX [88, 87], and JET [87]. These analyses yielded diffusivities ranging from 0.1 to 10 m^2/s. The wide variety of analysis methods used as well as the wide range of target plasmas contribute to the differences in measured diffusivities on various tokamaks. In

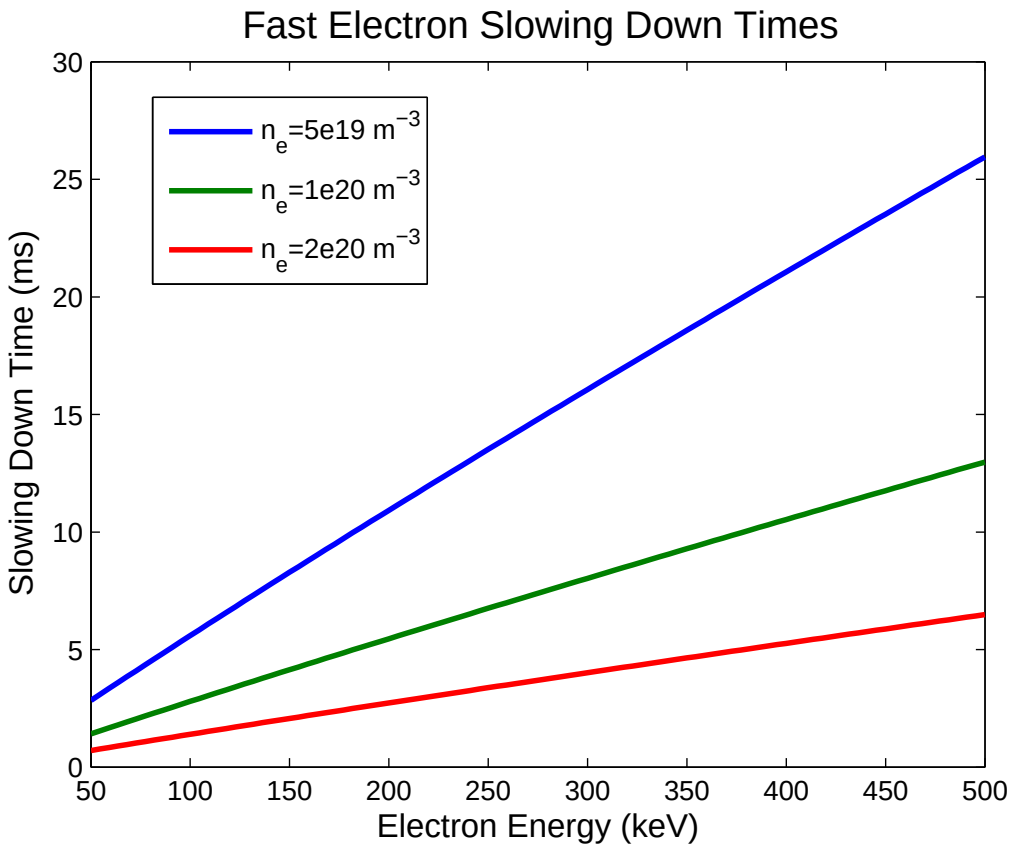

Figure 6-11: Fast electron slowing down times as a function of electron energy for various densities, with relativistic effects included. Slowing down times are computed based on Eq. 25 of Ref [89].

each of these previous transport experiments, LH power deposition was primarily in the core, due to the low electron temperatures and/or low launched $n_{||}$. By contrast, the transport coefficients presented in this paper were deduced from hollow, off-axis inverted x-ray profiles, peaked near r/a of 0.7. Thus these measurements of fast electron transport are relevant to LHCD operation in burning plasmas, in which LH power absorption and driven current are expected to be strongly peaked in the outer half of the plasma.

The data analysis described in Section 6.2 was applied to the hollow x-ray profiles produced by the $n_{||}$=3.1 cases, at plasma currents of 600 kA and 800 kA. The results described in this section are measurements of three fast electron transport properties: slowing down time, convection velocity, and diffusivity. Figure 6-11 shows the theo-

retical slowing down time for fast electrons as a function of energy for various bulk densities. Because the slowing-down time is relatively short in Alcator C-Mod, the fast electrons slow down before they have time to convect or diffuse very far.

Figure 6-12 shows the measured values of τ_p, v_p, and D_p for the 800 kA case, using all 3 models described in the previous section. The line-averaged density in this experiment was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. The toroidal magnetic field was $B_t = 5.4$ T. The peak coupled LH power was approximately $P_{LH} = 400$ kW (average $P_{LH} = 200$ kW).

Recall that the HXR energy, while correlated with the electron energy, does not correspond 1-to-1 with electron energy. For any given HXR energy bin, electrons of multiple energies above the HXR energy have contributed to the radiation.

As expected, the slowing down times shown in Fig. 6-12 are all in the range of a few milliseconds, and increasing with energy. When convection and/or diffusion are added to the model, the calculated value of the slowing down time does not change significantly.

The calculated convection velocities are shown in the middle plot of Fig. 6-12 for two models: one without diffusion (green) and one with diffusion (red). In the sign convention used here, a positive velocity indicates outward convection and a negative velocity indicates a pinch. The two models yield different values of convection velocity, though for both models, they are 0 within error bars for most energies. This indicates that convection, if it is occurring, is very small.

The calculated diffusion coefficients shown in Fig. 6-12 are also small ($D \leq .01$ m^2/s), and consistent in magnitude with the diffusion coefficients typically used in the Fokker-Planck modeling (see Section 5.2). They increase as a function of energy, in qualitative agreement with the magnetic turbulence model [76] and the velocity dependence of the spatial diffusion coefficients used in the modeling.

Typical measured slowing down times for this case are 3 ms. With an upper bound of $|v| \leq 0.5$ m/s on convection and $D \leq 0.01$ m^2/s on diffusion, a fast electron could convect or diffuse a maximum of 2 mm or 5 mm respectively on a slowing down time scale. These are relatively small displacements for C-Mod, whose minor radius is 23

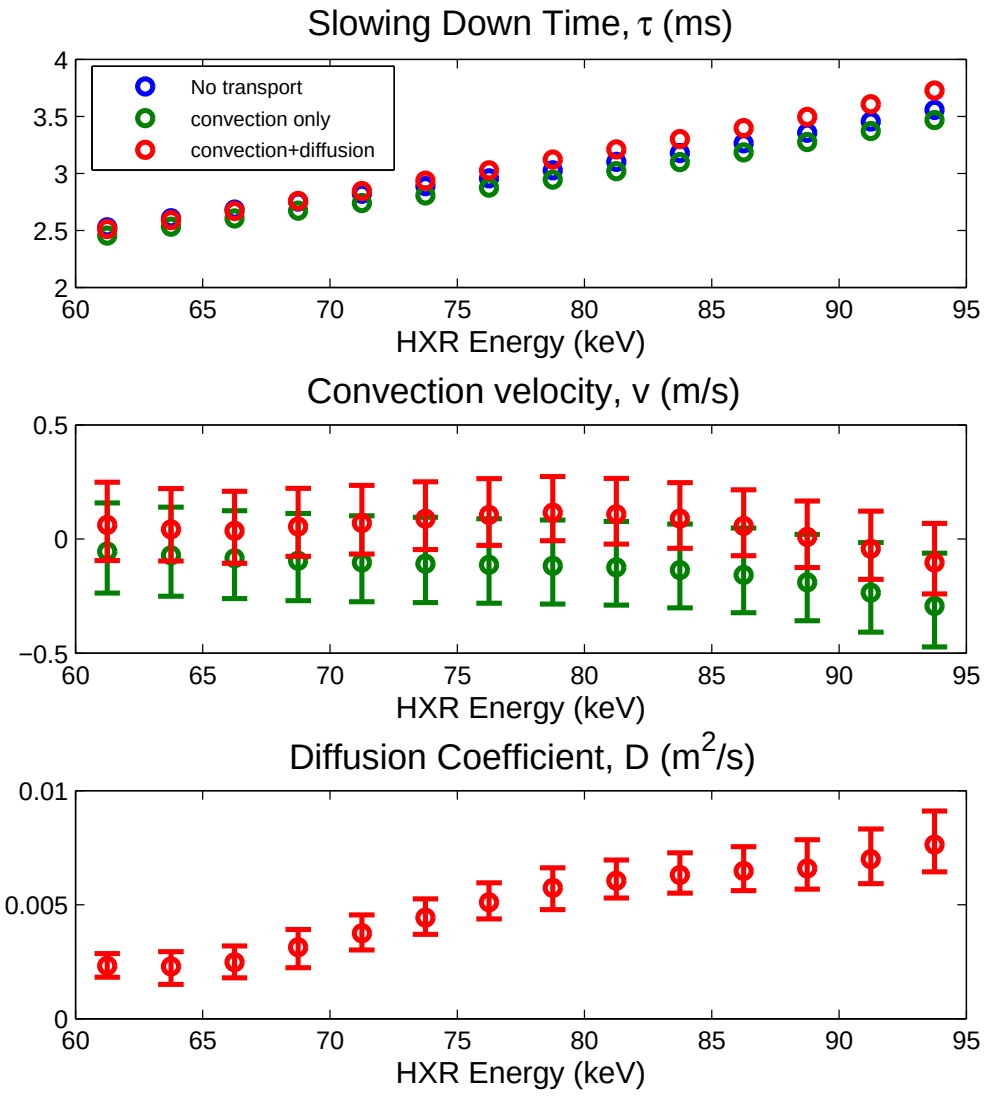

Figure 6-12: Measured transport parameters solved for in data analysis. The line-averaged density in this experiment was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. The toroidal magnetic field was $B_t = 5.4$ T and the plasma current was $I_p = 800$ kA. The peak coupled LH power was approximately $P_{LH} = 400$ kW.

cm.

Figure 6-13 shows the measured convection velocities and diffusivities for fast electrons emitting at various photon energies for a similar plasma with only 600 kA of plasma current. The line-averaged density in this experiment was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. The toroidal magnetic field was $B_t = 5.4$ T. In this case, the model was unable to distinguish between the effects of convection and diffusion. As a result, the two quantities were not solved for simultaneously. Instead, one model excludes diffusion and calculates convection, while the other excludes convection and calculates diffusion.

As in the previous case, the measured convection and diffusion are small. Typical measured slowing down times for this case are 2 ms. With an upper bound of $|v| \leq 1$ m/s on convection and $D \leq 0.02$ m^2/s on diffusion, a fast electron could convect or diffuse a maximum of 2 mm or 6 mm respectively on a slowing down time scale.

6.4 Power Deposition Trends

6.4.1 Variation of X-ray Profiles with Launched Spectrum

For fixed target plasma conditions, the main parameter for controlling the driven current location by LHCD is the launched value of $n_{||}$. Quasilinear theory predicts that LH waves will damp when their phase velocity is 2.5 times the electron thermal velocity [69, 28], where $v_{th} = \sqrt{2T_e/m_e}$. This requires that LH waves penetrate into the plasma until they reach a sufficiently high temperature to damp, about $T_e(kev) = 40/n_{||}^2$. Thus in a single-pass absorption regime, theory predicts lower $n_{||}$ waves will damp farther in towards the core, where the plasma temperature is higher.

However, due to the low temperatures (relative to $40/n_{||}^2$) in C-Mod, LH waves will not be absorbed on their first pass into the plasma center. Instead, they propagate back out to the plasma scrape off layer where they reflect from either the vacuum vessel walls or a cut-off layer at least once before damping on a later pass into the plasma core. As the LH wave propagates in the plasma, the $n_{||}$ of the wave can

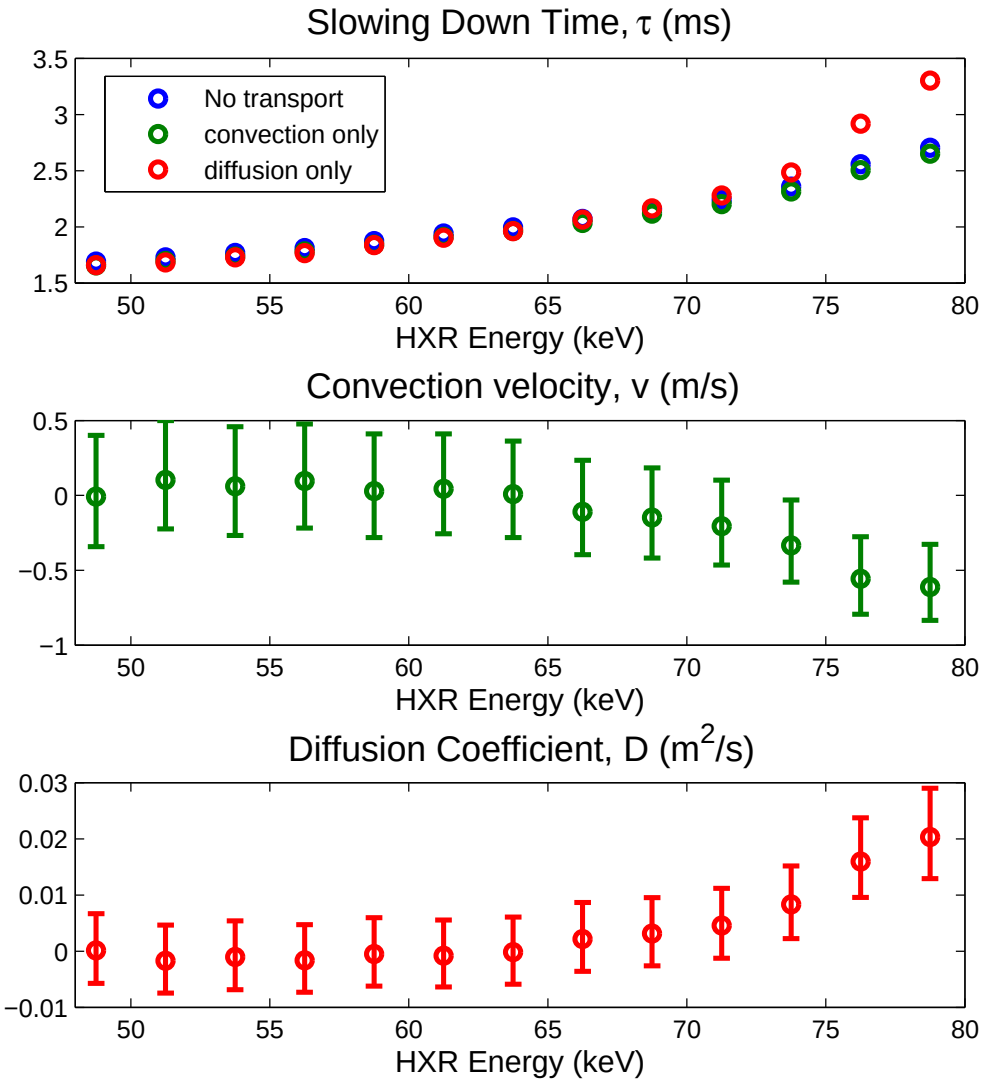

Figure 6-13: Measured transport parameters solved for in data analysis. The line-averaged density in this experiment was $\overline{n}_e = 9 \times 10^{19}$ m^{-3}, and peak electron temperature was $T_{e0} = 2.3$ keV. The toroidal magnetic field was $B_t = 5.4$ T and the plasma current was $I_p = 600$ kA. The peak coupled LH power was approximately $P_{LH} = 400$ kW. In this case, the model has difficulty distinguishing between the effects of convection and diffusion. The coefficients shown here were obtained by excluding (one at a time) diffusion and convection in the model and solving for the remaining transport quantity.

vary due to toroidal geometry effects [22] and the wave does not damp until $n_{||}$ increases enough to damp at the local electron temperature. In these multi-pass absorption regimes, the power damping location depends on the evolution of $n_{||}$ as the wave propagates, not necessarily the launched value. The ray tracing/Fokker-Planck simulations shown in Section 6.5 show that higher $n_{||}$ waves damp farther off-axis than lower $n_{||}$ waves because they undergo an $n_{||}$ upshift sufficient to damp at lower temperatures.

The hard x-ray data shown in this section were taken for 3 different values of $n_{||}$: 1.5, 2.3, and 3.1. Figure 6-14 shows the inverted HXR profiles for these 3 experiments, taken from the last boxcar binned time step during which the LH power was on. In these experiments, the peak electron temperature was 2.3 keV, the line-averaged electron density was 9×10^{19} m^{-3}, the plasma current was 800 kA, and the toroidal B field was 5.4 T. The bremsstrahlung emission location indicates that the wave power damps most strongly on axis for the lowest $n_{||}$ case and most strongly off axis for the highest $n_{||}$ case, in agreement with theory.

The fact that power deposition location is strongly dependent on $n_{||}$ indicates that current profile control is attainable with LHCD, especially in single pass damping regimes, using the launched value of $n_{||}$ as a control parameter. This is the first demonstration of power deposition location control using the launched spectrum on an experiment with the capability of launching a wide range of $n_{||}$. The radial dependence on $n_{||}$ of LH power deposition has also been measured on Tore Supra (whose LH system launches a narrower range of $n_{||}$), using real time feedback control of antenna phasing [90]. Modeling of the phasing scan on C-Mod is presented in Section 6.5.

6.4.2 Variation of X-ray Profiles with Plasma Current

Lower hybrid power modulation experiments were performed with 3 different plasma currents to examine the dependence of driven fast electron location on plasma current. Figure 6-15 shows the inverted HXR profiles for these 3 experiments, also taken from the last boxcar binned time step during which the LH power was on. The central electron temperature for the 600 kA, 800 kA, and 1 MA discharges was 1.9, 2.3,

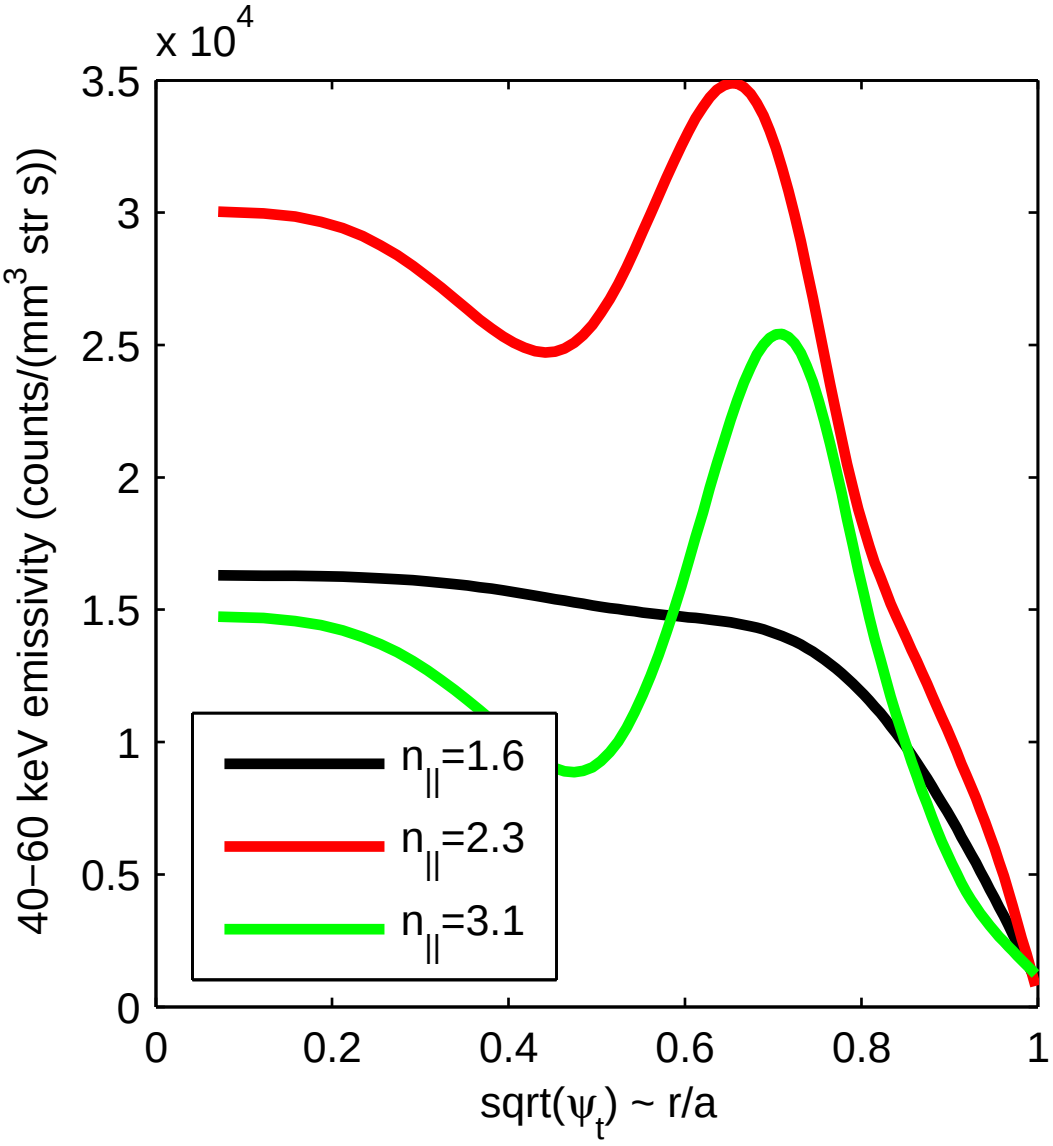

Figure 6-14: Inverted HXR spectra from discharges with 3 different launched $n_{||}$: 1.6, 2.3, and 3.1. Plasma parameters for this experiment were B=5.4 T, T_{e0}=2.3 keV, $\bar{n}_e = 9 \times 10^{19}$ m^{-3}, and I_p=800 kA. Peak coupled power in these experiments was approximately 455 kW ($n_{||}$=1.6), 415 kW ($n_{||}$=2.3), and 370 kW ($n_{||}$=3.1).

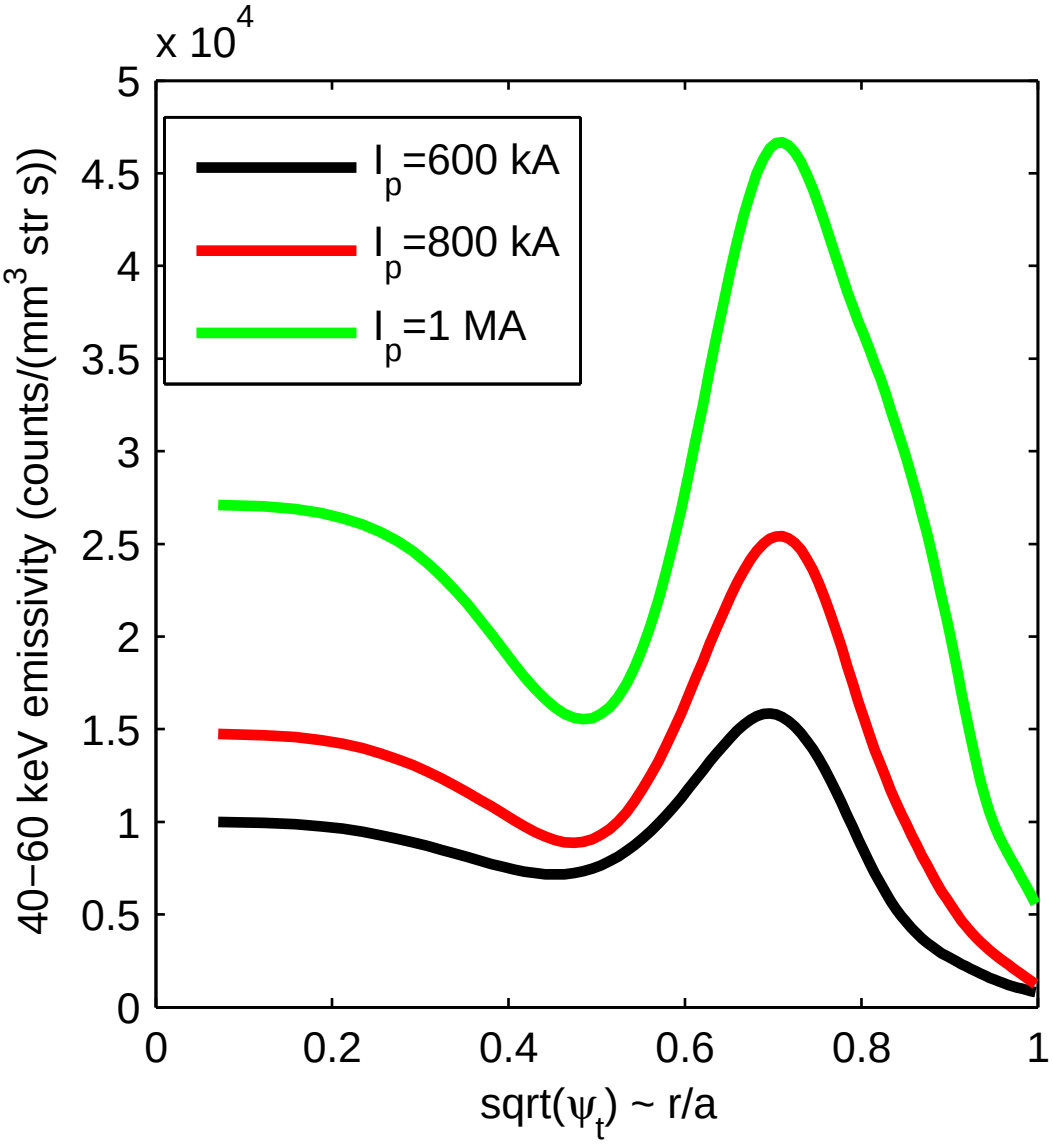

Figure 6-15: Inverted HXR spectra from discharges with 3 different plasma currents: 600 kA, 800 kA, and 1 MA. Plasma parameters for this experiment were B=5.4 T, $\bar{n}_e = 9 \times 10^{19}$ m^{-3}. T_{e0} for the 600 kA, 800 kA, and 1 MA discharges was 1.9, 2.3, and 2.1 keV respectively. As the current increased from 600 kA to 1 MA, the loop voltage increased from 0.8 to 1.1 V. The launched $n_{||}$=3.1 for these discharges. Peak coupled power in these experiments was approximately 360 kW (600 kA), 370 kW (800 kA), and 350 kW (1 MA).

and 2.1 keV respectively. As the current increased from 600 kA to 1 MA, the loop voltage increased from 0.8 to 1.1 V. For all 3 cases, the peak launched $n_{||}$ was 3.1, the line-averaged electron density was 9×10^{19} m^{-3} and the toroidal B field was 5.4 T.

The HXR emission location indicates that the plasma current does not strongly affect the wave power damping location. The x-ray intensity clearly varies for these 3 different scenarios. These variations in flux are likely due to the variation in the DC electric field, which is highest for the 1 MA discharges. However, they could be partially due to differences in the poloidal magnetic field strength as well. It is possible that higher poloidal field results in larger upshifts in the parallel wavenumber and consequently stronger wave absorption. Increases in x-ray count rates at higher plasma currents have also been observed on Tore Supra [38].

6.5 Modeling of Lower Hybrid Modulation Experiments

The LH power modulation experiments discussed in this chapter lend themselves particularly well to comparison with modeling. One reason is that the average power in these experiments is rather low and thus most of the plasma current is ohmic. Magnetic equilibrium files, created by EFIT [80], include calculations of the poloidal magnetic field using the assumption that the plasma current is ohmically driven[3]. Thus EFIT-calculated equilibrium files will likely be more accurate for a low power experiment than they would be for a predominantly non-inductive experiment, leading to more self-consistent modeling.

Another reason that these experiments are interesting to model is that they inherently involve averaging the HXR signals in time. Parameters to which the HXR signal is extremely sensitive, such as electron density and loop voltage, remain fairly constant in these experiments. Even so, the x-rays are averaged over many time bins so that effects due to small fluctuations in plasma parameters should get averaged

[3]Current profile measurements from the MSE diagnostic can be used as a constraint in EFIT [77]. The EFIT used in the modeling presented here was not constrained by MSE measurements.

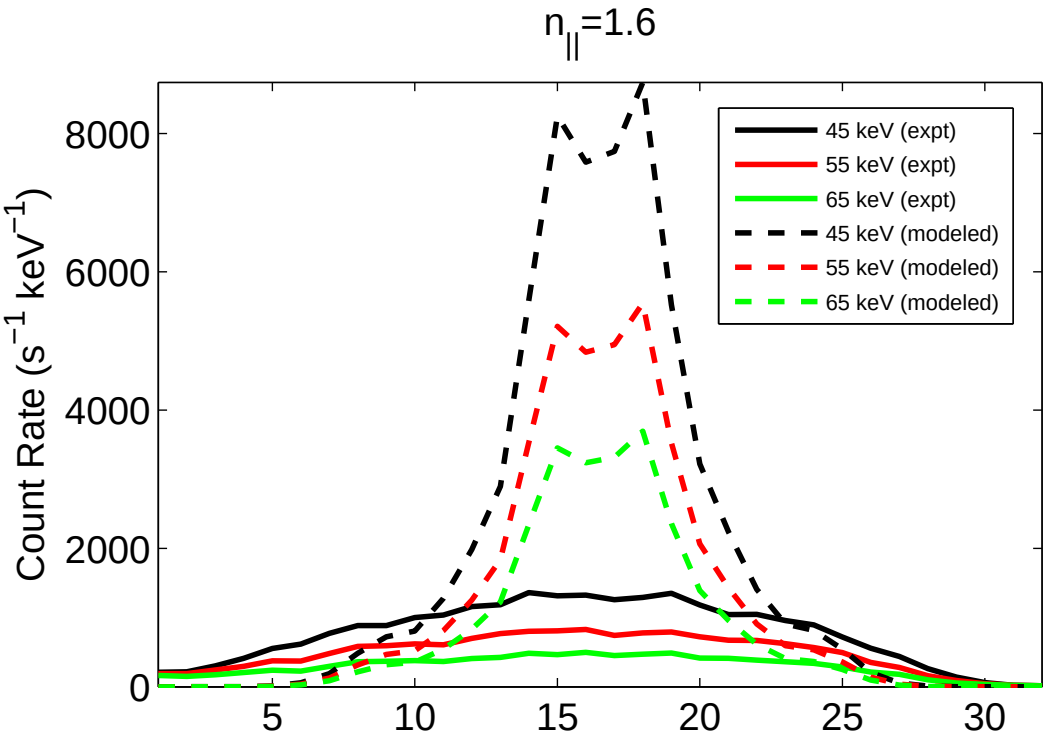

Figure 6-16: Modeled and measured HXR profiles for $n_{||}$=1.6 (60° phasing) LH modulation. A fast electron diffusivity of $0.01 v_{||}/(v_{th}\gamma^3)$ m^2/s was used.

out.

In the set of simulations described in this section, the experimental HXR profiles were taken from the last time bin during which the LH was on (the 5^{th} out of 10 time bins). The simulations were run with 50 time steps of 0.25 ms each, in order to simulate the evolution of the experiment during the 12.5 ms that the LH power was on.

Recall that the LH modulation experiments were performed at 3 different peak launched $n_{||}$ values (1.55, 2.33, and 3.1). We used Genray/CQL3D to model this phasing scan and compared the measured x-ray counts in the experiment with the synthetic x-ray diagnostic calculation. Figures 6-16, 6-17, and 6-18 show comparisons of HXR profiles generated by CQL3D with the experimental boxcar binned HXR profiles for $n_{||}$=1.55, 2.33, and 3.1 respectively. In each simulation, a fast electron diffusivity of $0.01 v_{||}/(v_{th}\gamma^3)$ m^2/s was used, where γ is the relativistic factor.

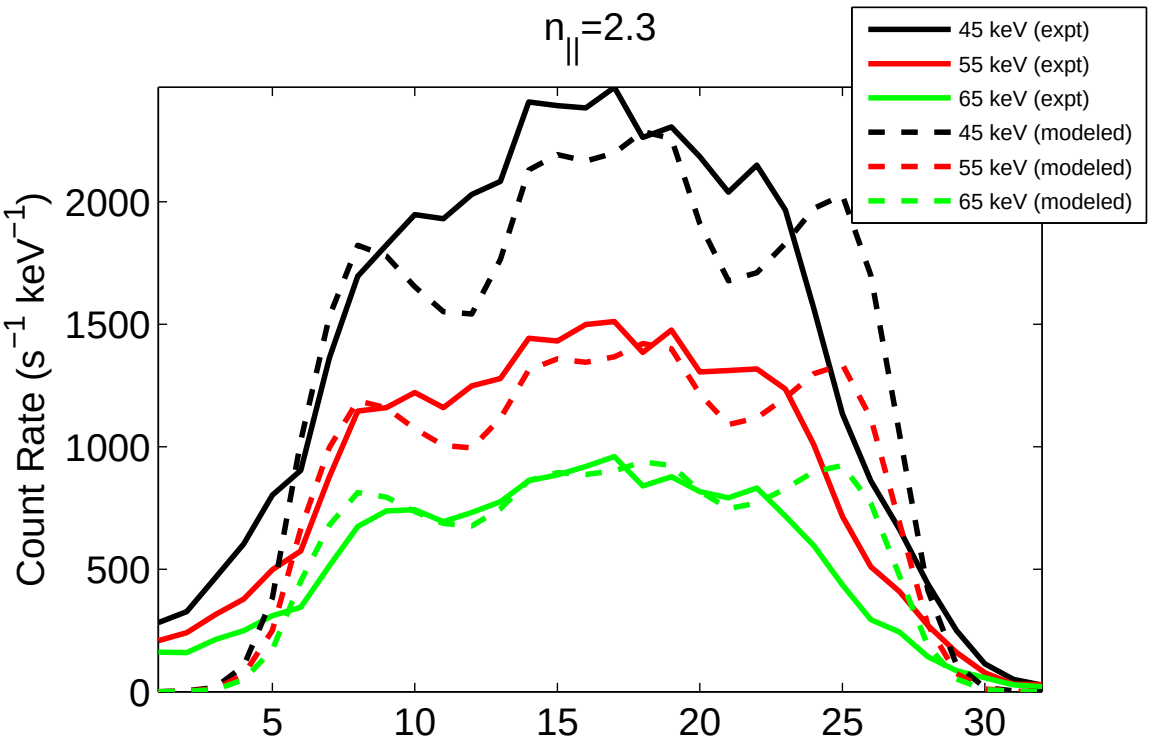

Figure 6-17: Modeled and measured HXR profiles for $n_{||}$=2.3 (90° phasing) LH modulation. A fast electron diffusivity of $0.01 v_{||}/(v_{th}\gamma^3)$ m^2/s was used.

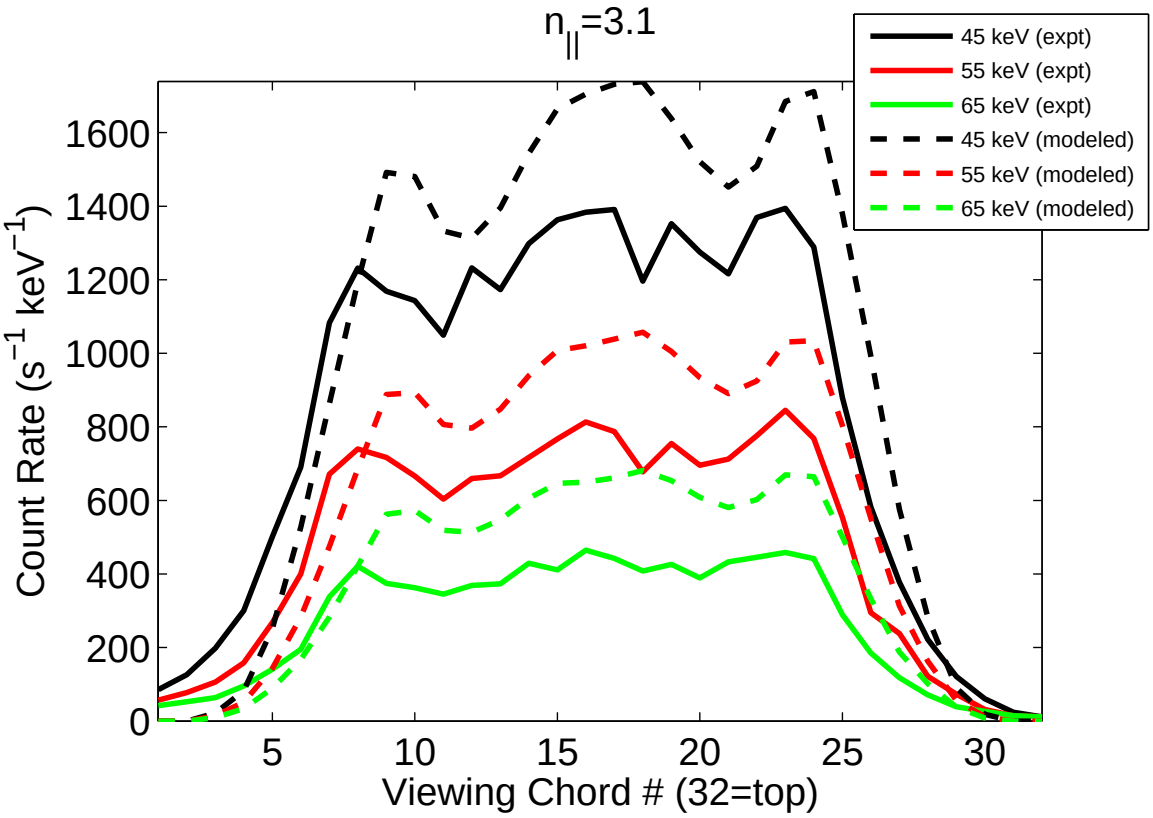

Figure 6-18: Modeled and measured HXR profiles for $n_{||}$=3.1 (120° phasing) LH modulation. The simulated HXR profiles mimic the fairly flat central region of the profiles as well as the side peaks. The simulated profiles are very similar in width to the experimental profiles, though they appear to be shifted slightly towards the top of the plasma (closer to chord 32). A fast electron diffusivity of $0.01 v_{||}/(v_{th}\gamma^3)$ m^2/s was used.

It is clear that the poorest agreement is obtained in the $n_{||} = 1.6$ case (Figure 6-16), where the simulation predicts that the power damps predominantly on axis but the experimental x-rays are much broader. We attribute the poor agreement for the low $n_{||}$ case to the multiple passes that the rays undergo before they damp because the initial parallel phase velocity of the wave is very large compared to the electron thermal velocity. This idea can be further explored by looking at the ray power damping in each case.

Figure 6-19 shows the fraction of power remaining in all rays representing the forward wave spectrum as a function of poloidal distance traveled in the plasma for the 3 different antenna phasings modeled. The fractional ray power is shown for rays damping on a Maxwellian electron distribution (before the first time step in CQL3D) and on a non-thermal electron tail (during the last time step in CQL3D).

It is clear from Figure 6-19 that the rays damp most strongly in the $n_{||} = 3.1$ case and most weakly in the $n_{||} = 1.6$ case, regardless of whether they are damping on a Maxwellian plasma or a non-thermal plasma. In particular, the rays for the $n_{||}$=1.6 case damp very weakly in the first 100-150 cm of poloidal propagation distance, where the rays undergo their first few reflections. Because ray tracing is an initial value problem, we would expect the code's predictions to be least accurate in cases where many passes of the rays are necessary for the rays to damp. In particular, the ray paths and $n_{||}$ can be modified significantly at or near reflections, in a region of the plasma (the SOL) that is not being modeled. Significant modification of the LH wave front can also occur at cut-offs due to full-wave effects such as diffraction and scattering that are not included in the ray tracing.

It is important for accurate predictions that the waves damp significantly in a minimal number of passes on both the Maxwellian plasma and the non-thermal plasma. Ray damping on a Maxwellian plasma calculated by Genray seeds the fast electron population in a particular radial location. Ray damping during subsequent time steps in CQL3D will tend to be strongest in radial locations where a fast electron tail has already been populated from previous time steps.

While the approximation of moving the cut-off layer inside the separatrix ap-

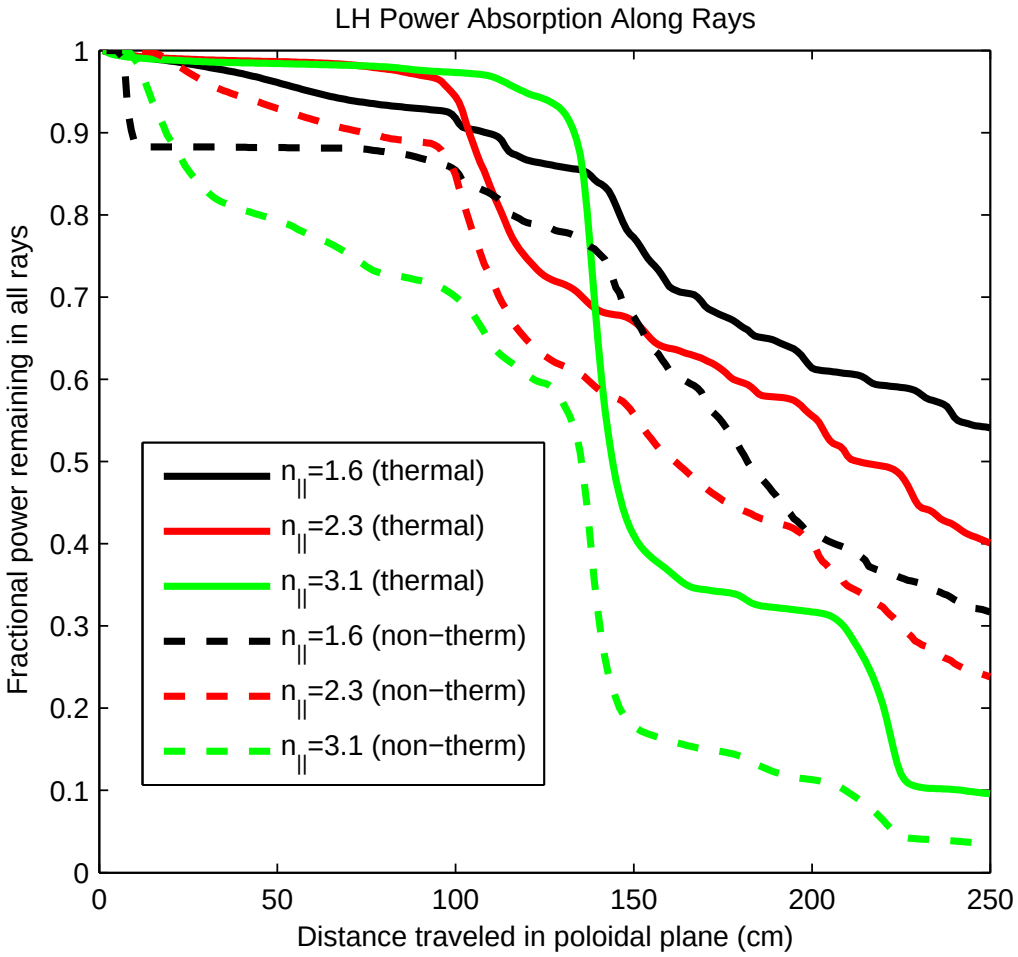

Figure 6-19: Fractional power remaining in all rays representing the forward wave spectrum as a function of poloidal distance traveled in the plasma for 3 different antenna phasings. Solid lines show power remaining when rays are damped on a thermal (Maxwellian) plasma. Dashed lines show power remaining when rays are damped on a non-thermal distribution function, during the last time step of CQL3D.

pears to be valid for the two higher $n_{||}$ cases, where not as many ray passes are needed for damping, it may not be sufficient for modeling the lowest $n_{||}$ case. In this case, ray tracing with a SOL model or full-wave model might more accurately predict the wave damping. Since it is expected that LH waves will exhibit single pass absorption at reactor relevant temperatures, modeling presented here suggests that ray-tracing/Fokker-Planck calculations may be valid in temperature regimes relevant to ITER, at least at the plasma densities used here.

Chapter 7

Summary and Conclusions

Lower hybrid current drive is a promising approach to non-inductive current drive for steady-state tokamak operation. The results presented in this thesis represent an advance in our understanding of lower hybrid physics and RF simulations.

The simulated ECE spectra shown in Section 4.3.1 guided our use of the electron cyclotron nonthermal emission as a signature of LH-driven fast electrons. The comparison of experimental HXR and ECE spectra with simulated spectra in Section 5.1 was the first benchmarking test of Genray/CQL3D synthetic diagnostics against an LHCD experiment and one of the first detailed comparisons of experiment with the predictions of any 3D Fokker-Planck code. In this case, the discrepancies between experiment and simulations led to an investigation of the effects of fast electron radial diffusion in the modeling and the experimental measurement of fast electron diffusion. It also brought about a strong interest in studying full wave effects in the modeling, an ongoing effort in our group.

MSE current profile measurements during nearly full non-inductive discharges were compared with a simulated current profile in Section 5.2, with reasonable agreement. Observations of HXR profiles as a function of density were presented in Section 5.3 for three magnetic topologies: USN, LSN, and inner-wall limited. These profiles show increased HXR activity in the vicinity of the active x-point for diverted discharges and self-similar HXR profiles for limited discharges.

The lower hybrid power modulation experiments on Alcator C-Mod led to several

interesting physics results. Boxcar binning of the resultant HXR signal produced high signal-to-noise x-ray profiles that are well resolved in time. An $n_{||}$ scan showed clear variation in power deposition location with launched spectrum, in agreement with theoretical expectations. A current scan showed little variation in power deposition location with changes in plasma current. Modeling of the line-integrated HXR count rates with a combined ray-tracing/Fokker-Planck code showed that agreement with experiment for high $n_{||}$ launched rays was excellent. Agreement for the lowest $n_{||}$ case, in which rays undergo many radial reflections at the plasma edge, was significantly poorer. A simple transport model has been used to determine upper bounds on fast electron transport of $D \leq 0.02m^2/s, |v| \leq 0.75m/s$, indicating that fast electrons do not stray far from the flux surface where they are born before slowing down.

7.1 Assessment of Ray Tracing/Fokker-Planck Simulations on Alcator C-Mod

Ray-tracing in a tokamak can be a highly stochastic process [22]. Thus we should expect that longer ray paths involving more than a single reflection could result in a ray being in a very different part of phase space with just a small change in conditions at the cut-off location. The simulated phasing scan and comparison with experiment in Section 6.5 provides further evidence that the simulations follow the experimental data more closely when the ray paths are shorter, as they are in higher $n_{||}$ cases. It is possible that discrepancies in magnitude of the simulated and measured HXR spectra in the $n_{||} = 1.6$ case could be partially resolved by including parasitic edge absorption mechanisms, such as collisions, which are known to be important for high density regimes [78].

It is becoming increasingly clear that reliable SOL models are needed to model multi-pass absorption regimes, such as those that occur at high density or when the launched $n_{||}$ is low. Further work is needed to clarify whether ray tracing with a SOL model is sufficient for modeling such multi-pass absorption regimes. The density

profile in the C-Mod SOL is not well known and is poloidally asymmetric. Thus even if we calculate ray trajectories in the SOL, we may not be able to employ accurate enough SOL profiles in the model to correctly predict the propagation and absorption along long ray paths. Furthermore, CQL3D can only calculate ray damping and electron distribution functions on closed flux surfaces, so it will be challenging to model resonant power dissipation if it does occur in the SOL.

Benchmarking of full-wave codes against experiment and against ray tracing is an ongoing effort that may provide further insight into this complex problem. Full-wave codes include diffraction and focussing effects which are absent in the ray-tracing treatment and which can significantly impact the evolution of the parallel wavenumber through spectral broadening. More investigation is needed to determine if full-waves codes are the solution to reliably modeling multi-pass regimes.

The agreement between experimental measurements and the synthetic x-ray diagnostic for moderate-to-strong damping cases is encouraging evidence that ray-tracing/Fokker-Plank modeling has strong predictive power for the next generation of tokamaks, where strong single-pass damping of the waves is expected due to higher electron temperatures. In particular, it may be very useful for simulating LHCD scenarios on ITER.

7.2 Implications of LH Modulation Experiments for Current Profile Control

The results from the modulation experiments are encouraging for the use of LHCD on the next generation of tokamaks, which are expected to operate at higher electron temperatures. The variation in bremsstrahlung emission location based on the launched $n_{||}$ spectrum implies that LH power deposition location can be controlled through antenna phasing. The small magnitude of diffusion and convection in the outer half of the plasma suggests that the LH-driven fast electrons remain localized. The combination of these two results indicates that LHCD may be used effectively for

current profile control in the next generation of reactors. Furthermore, this analysis was the first of its kind to be performed at reactor-relevant densities and magnetic fields with LH power deposition dominantly in the outer half of the plasma.

Bibliography

[1] International Energy Outlook 2010. Technical report, Energy Information Administration, U.S. Department of Energy, Washington, DC, Jul 2010.

[2] K.S. Krane. *Introductory Nuclear Physics*. John Wiley & Sons, Inc., 1988.

[3] S. H. Glenzer, B. J. MacGowan, P. Michel, N. B. Meezan, L. J. Suter, S. N. Dixit, J. L. Kline, G. A. Kyrala, D. K. Bradley, D. A. Callahan, E. L. Dewald, L. Divol, E. Dzenitis, M. J. Edwards, A. V. Hamza, C. A. Haynam, D. E. Hinkel, D. H. Kalantar, J. D. Kilkenny, O. L. Landen, J. D. Lindl, S. LePape, J. D. Moody, A. Nikroo, T. Parham, M. B. Schneider, R. P. J. Town, P. Wegner, K. Widmann, P. Whitman, B. K. F. Young, B. Van Wonterghem, L. J. Atherton, and E. I. Moses. Symmetric inertial confinement fusion implosions at ultra-high laser energies. *Science*, 327(5970):1228–1231, 2010.

[4] Allen H. Boozer. Physics of magnetically confined plasmas. *Rev. Mod. Phys.*, 76(4):1071–1141, Jan 2005.

[5] M. Shimada, D.J. Campbell, V. Mukhovatov, M. Fujiwara, N. Kirneva, K. Lackner, M. Nagami, V.D. Pustovitov, N. Uckan, J. Wesley, N. Asakura, A.E. Costley, A.J.H. Donn, E.J. Doyle, A. Fasoli, C. Gormezano, Y. Gribov, O. Gruber, T.C. Hender, W. Houlberg, S. Ide, Y. Kamada, A. Leonard, B. Lipschultz, A. Loarte, K. Miyamoto, V. Mukhovatov, T.H. Osborne, A. Polevoi, and A.C.C. Sips. Chapter 1: Overview and summary. *Nuclear Fusion*, 47(6):S1, 2007.

[6] John Wesson. *Tokamaks*. Oxford University Press, 2004.

[7] H. Zushi, S. Itoh, K. Hanada, K. Nakamura, M. Sakamoto, E. Jotaki, M. Hasegawa, Y. D. Pan, S. V. Kulkarni, A. Iyomasa, S. Kawasaki, H. Nakashima, N. Yoshida, K. Tokunaga, T. Fujiwara, M. Miyamoto, H. Nakano, M. Yuno, A. Murakami, S. Nakamura, N. Sakamoto, K. Shinoda, S. Yamazoe, H. Akanishi, K. Kuramoto, Y. Matsuo, A. Iwamae, T. Fuijimoto, A. Komori, T. Morisaki, H. Suzuki, S. Masuzaki, Y. Hirooka, Y. Nakashima, and O. Mitarai. Overview of steady state tokamak plasma experiments in TRIAM-1M. *Nuclear Fusion*, 43:1600–1609, December 2003.

[8] B J D Tubbing and JET Team. On the operation cycle of tokamak fusion reactors. *Plasma Physics and Controlled Fusion*, 35(SB):B55, 1993.

[9] J.P. Freidberg. *Plasma Physics and Fusion Energy*. Cambridge University Press, 2007.

[10] N. J. Fisch. Theory of current drive in plasmas. *Reviews of Modern Physics*, 59:175–234, January 1987.

[11] T. Ohkawa. New methods of driving plasma current in fusion devices. *Nuclear Fusion*, 10(2):185–188, 1970.

[12] Wang Jinfang, Wu Bin, and Hu Chundong. Simulation on heating and current drive using neutral beam injection with variable injection angle on east. *Plasma Science and Technology*, 12(3):289, 2010.

[13] M. Murakami, J.M. Park, C.C. Petty, T.C. Luce, W.W. Heidbrink, T.H. Osborne, R. Prater, M.R. Wade, P.M. Anderson, M.E. Austin, N.H. Brooks, R.V. Budny, C.D. Challis, J.C. DeBoo, J.S. deGrassie, J.R. Ferron, P. Gohil, J. Hobirk, C.T. Holcomb, E.M. Hollmann, R.M. Hong, A.W. Hyatt, J. Lohr, M.J. Lanctot, M.A. Makowski, D.C. McCune, P.A. Politzer, J.T. Scoville, H.E. St John, T. Suzuki, T.S. Taylor, W.P. West, E.A. Unterberg, M.A. Van Zeeland, and J.H. Yu. Off-axis neutral beam current drive for advanced scenario development in diii-d. *Nuclear Fusion*, 49(6):065031, 2009.

[14] J. M. Park, M. Murakami, C. C. Petty, W. W. Heidbrink, T. H. Osborne, C. T. Holcomb, M. A. Van Zeeland, R. Prater, T. C. Luce, M. R. Wade, M. E. Austin, N. H. Brooks, R. V. Budny, C. D. Challis, J. C. DeBoo, J. S. deGrassie, J. R. Ferron, P. Gohil, J. Hobirk, E. M. Hollmann, R. M. Hong, A. W. Hyatt, J. Lohr, M. J. Lanctot, M. A. Makowski, D. C. McCune, P. A. Politzer, H. E. St John, T. Suzuki, W. P. West, E. A. Unterberg, and J. H. Yu. Validation of on- and off-axis neutral beam current drive against experiment in diii-d. *Physics of Plasmas*, 16(9):092508, 2009.

[15] N. J. Fisch and A. H. Boozer. Creating an Asymmetric Plasma Resistivity with Waves. *Physical Review Letters*, 45(9):720–722, Sep 1980.

[16] V Erckmann and U Gasparino. Electron cyclotron resonance heating and current drive in toroidal fusion plasmas. *Plasma Physics and Controlled Fusion*, 36(12):1869, 1994.

[17] Nathaniel J. Fisch. Confining a tokamak plasma with rf-driven currents. *Phys. Rev. Lett.*, 41(13):873–876, Sep 1978.

[18] S. C. Jardin, C. E. Kessel, C. G. Bathke, D. A. Ehst, T. K. Mau, F. Najmabadi, and T. W. Petrie. Physics basis for a reversed shear tokamak power plant. *Fusion Engineering and Design*, 38(1-2):27 – 57, 1997.

[19] The JET Team (presented by F X Sldner). Shear optimization experiments with current profile control on jet. *Plasma Physics and Controlled Fusion*, 39(12B):B353, 1997.

[20] C.F.F. Karney and N.J. Fisch. Numerical studies of current generation by radio-frequency traveling waves. *Physics of Fluids*, 22(9):1817–1824, 1979.

[21] Thomas H. Stix. *Waves in Plasmas*. Springer, 1992.

[22] Paul T. Bonoli and Edward Ott. Toroidal and scattering effects on lower-hybrid wave propagation. *Physics of Fluids*, 25(2):359–375, 1982.

[23] V. E. Golant. Plasma Penetration near the Lower Hybrid Frequency. *Soviet Physics Technical Physics*, 16:1980, June 1972.

[24] F. Troyon and F.W. Perkins. Lower hybrid heating in a large tokamak. *Proceedings of the 2^{nd} Topical Conference on RF Plasma Heating*, Paper B4, June 1974.

[25] F. Wagner, G. Fussmann, T. Grave, M. Keilhacker, M. Kornherr, K. Lackner, K. McCormick, E. R. Müller, A. Stäbler, G. Becker, K. Bernhardi, U. Ditte, A. Eberhagen, O. Gehre, J. Gernhardt, G. v. Gierke, E. Glock, O. Gruber, G. Haas, M. Hesse, G. Janeschitz, F. Karger, S. Kissel, O. Klüber, G. Lisitano, H. M. Mayer, D. Meisel, V. Mertens, H. Murmann, W. Poschenrieder, H. Rapp, H. Röhr, F. Ryter, F. Schneider, G. Siller, P. Smeulders, F. Söldner, E. Speth, K. H. Steuer, Z. Szymanski, and O. Vollmer. Development of an edge transport barrier at the h-mode transition of asdex. *Phys. Rev. Lett.*, 53(15):1453–1456, Oct 1984.

[26] A.A. Vedenov. Theory of a Weakly Turbulent Plasma. *Reviews of Plasma Physics*, 3:229, 1967.

[27] M. Porkolab, J. J. Schuss, B. Lloyd, Y. Takase, S. Texter, P. Bonoli, C. Fiore, R. Gandy, D. Gwinn, B. Lipschultz, E. Marmar, D. Pappas, R. Parker, and P. Pribyl. Observation of Lower-Hybrid Current Drive at High Densities in the Alcator *C* Tokamak. *Physical Review Letters*, 53(5):450–453, Jul 1984.

[28] Paul T. Bonoli and Ronald C. Englade. Simulation model for lower hybrid current drive. *Physics of Fluids*, 29(9):2937–2950, 1986.

[29] C. F. Kennel and F. Engelmann. Velocity Space Diffusion from Weak Plasma Turbulence in a Magnetic Field. *Physics of Fluids*, 9(12):2377–2388, 1966.

[30] Nicholas A. Krall and Alvin W. Trivelpiece. *Principles of Plasma Physics*. San Francisco Press, 1986.

[31] J. D. Jackson. *Classical Electrodynamics*. Wiley, 1998.

[32] I. H. Hutchinson. *Principles of Plasma Diagnostics*. Cambridge University Press, 2002.

[33] Schweickhard E. von Goeler, James E. Stevens, Stefano Bernabei, Manfred L. Bitter, Tsu-Kai Chu, Phillip C. Efthimion, Nathaniel J. Fisch, William M. Hooke, Kenneth W. Hill, Joel C. Hosea, Forrest C. Jobes, Charles F. F. Karney, J. A. Mervine, E. B. Meservey, Robert W. Motley, Phyllis G. Roney, Steven S. Sesnic, K. Silber, and Gary Taylor. Angular distribution of the bremsstrahlung emission during lower hybrid current drive on PLT. *Nucl. Fusion*, 25(11):1515–1528, November 1985.

[34] S. Texter, S. Knowlton, M. Porkolab, and Y. Takase. High energy x-ray measurements during lower hybrid current drive on the alcator c tokamak. *Nuclear Fusion*, 26(10):1279–1291, October 1986.

[35] S. Texter, M. Porkolab, P. T. Bonoli, S. Knowlton, and Y. Takase. Energy confinement of the high energy tail electrons during lower hybrid current drive experiments on the alcator c tokamak. *Physics Letters A*, 175(6):428 – 432, 1993.

[36] S. Knowlton, M. Porkolab, and Y. Takase. Experimental studies of the radial diffusion of energetic electrons in lower-hybrid wave heated plasmas in the alcator-c tokamak. *Physics of Plasmas*, 1(4):891–900, 1994.

[37] M. Brusati, D.V. Bartlett, A. Ekedahl, P. Froissard, A. Airoldi, G. Ramponi, R.P. da Silva, and Y. Peysson. Determination of the suprathermal electron distribution function during lower hybrid current drive in jet. *Nuclear Fusion*, 34(1):23, 1994.

[38] Y Peysson, the TORE SUPRA Team, and Y Peysson. Progress towards high-power lower hybrid current drive in tore supra. *Plasma Physics and Controlled Fusion*, 42(12B):B87, 2000.

[39] Y. Peysson and Tore Supra Team. High power lower hybrid current drive experiments in the Tore Supra tokamak. *Nuclear Fusion*, 41:1703–1713, November 2001.

[40] J. Decker and Y. Peysson. DKE: A fast numerical solver for the 3D drift kinetic equation. Technical report, Report EUR-CEA-FC-1736, Euratom-CEA, 2004.

[41] H. W. Koch and J. W. Motz. Bremsstrahlung cross-section formulas and related data. *Reviews of Modern Physics*, 31(4):920–955, Oct 1959.

[42] H. Bethe and W. Heitler. On the Stopping of Fast Particles and on the Creation of Positive Electrons. *Proceedings of the Royal Society of London. Series A*, 146(856):83–112, 1934.

[43] Gerhard Elwert and Eberhard Haug. Calculation of bremsstrahlung cross sections with sommerfeld-maue eigenfunctions. *Phys. Rev.*, 183(1):90, Jul 1969.

[44] E. Haug. Bremsstrahlung and pair production in the field of free electrons. *Zeit. Naturforsch.*, 30a:1099, 1975.

[45] M. Bornatici, R. Cano, O. De Barbieri, and F. Engelmann. Electron cyclotron emission and absorption in fusion plasmas. *Nuclear Fusion*, 23(9):1153, 1983.

[46] S. E. Jones, J. Kesner, S. Luckhardt, F. Paoletti, S. von Goeler, S. Bernabei, R. Kaita, PBX-M Team, and G. Giruzzi. Fast electron transport and lower hybrid absorbed power profiles from hard x-ray imaging in the Princeton Beta Experiment-Modified. *Physics of Plasmas*, 2:1548–1560, May 1995.

[47] I. H. Hutchinson, R. Boivin, F. Bombarda, P. Bonoli, S. Fairfax, C. Fiore, J. Goetz, S. Golovato, R. Granetz, M. Greenwald, S. Horne, A. Hubbard, J. Irby, B. LaBombard, B. Lipschultz, E. Marmar, G. McCracken, M. Porkolab, J. Rice, J. Snipes, Y. Takase, J. Terry, S. Wolfe, C. Christensen, D. Garnier, M. Graf, T. Hsu, T. Luke, M. May, A. Niemczewski, G. Tinios, J. Schachter, and J. Urbahn. First results from Alcator C-MOD. *Physics of Plasmas*, 1(5):1511–1518, 1994.

[48] www.mdsplus.org.

[49] S. Bernabei, J. Hosea, C. Kung, G. Loesser, J. Rushinski, J. Wilson, R. Parker, and M. Porkolab. Design of a Compact Lower Hybrid Coupler for Alcator C-Mod. *Fusion Science and Technology*, 43:145–152, March 2003.

[50] P. Koert, P. MacGibbon, R. Vieira, D. Terry, R. Leccacorvi, J. Doody, and W. Beck. Waveguide splitter for lower hybrid current drive. *Fusion Science and Technology*, 56(1):109–113, July 2009.

[51] G. M. Wallace. *Behavior of Lower Hybrid Waves in the Scrape Off Layer of a Diverted Tokamak*. PhD thesis, Massachusetts Institute of Technology, 2010.

[52] J. Liptac, R. Parker, V. Tang, Y. Peysson, and J. Decker. Hard x-ray diagnostic for lower hybrid experiments on Alcator C-Mod. *Review of Scientific Instruments*, 77(10):103504, 2006.

[53] J. Liptac. *Lower Hybrid Modeling and Experiments on Alcator C-Mod*. PhD dissertation, Massachusetts Institute of Technology, Department of Nuclear Science and Engineering, 2006.

[54] Yves Peysson and Frédéric Imbeaux. Tomography of the fast electron bremsstrahlung emission during lower hybrid current drive on tore supra. *Review of Scientific Instruments*, 70(10):3987–4007, 1999.

[55] NP Basse, A. Dominguez, EM Edlund, CL Fiore, RS Granetz, AE Hubbard, JW Hughes, IH Hutchinson, JH Irby, B. LaBombard, et al. Diagnostic systems on Alcator C-Mod. *Fusion Science and Technology*, 51(3):476–507, 2007.

[56] D. Moreau TORE SUPRA Team. Lower-hybrid current drive experiments in tore supra. *Physics of Fluids B: Plasma Physics*, 4(7):2165–2175, 1992.

[57] David L. Phillips. A technique for the numerical solution of certain integral equations of the first kind. *J. ACM*, 9(1):84–97, 1962.

[58] A. N. Tikhonov and V. Y. Arsenin. *Solutions of Ill-Posed Problems*. Winston and Sons, 1977.

[59] Per Hansen. Regularization tools version 4.0 for matlab 7.3. *Numerical Algorithms*, 46:189–194, 2007. 10.1007/s11075-007-9136-9.

[60] G. D. Kerbel and M. G. McCoy. Kinetic theory and simulation of multispecies plasmas in tokamaks excited with electromagnetic waves in the ion-cyclotron range of frequencies. *Physics of Fluids*, 28(12):3629–3653, 1985.

[61] S. Weinberg. Eikonal Method in Magnetohydrodynamics. *Physical Review*, 126:1899–1909, June 1962.

[62] L.D. Landau and L.M. Lifshitz. *The Classical Theory of Fields*, volume 2. Pergamon Press, New York, 1974.

[63] A. P. Smirnov and R.W. Harvey. Calculations of the Current Drive in DIII-D with the GENRAY Ray Tracing Code. *Bulletin of the American Physical Society*, 40(11), 1995.

[64] A.P. Smirnov and R.W. Harvey. Genray manual. Technical Report Report CompX-2000-01, CompX, P.O. Box 2672, Del Mar, CA 92014, available through www.compxco.com, March 2003. This is a full TECHREPORT entry.

[65] J. C. Wright, P. T. Bonoli, A. E. Schmidt, C. K. Phillips, E. J. Valeo, R. W. Harvey, and M. A. Brambilla. An assessment of full wave effects on the propagation and absorption of lower hybrid waves. *Physics of Plasmas*, 16(7):072502, 2009.

[66] S. Shiraiwa, O. Meneghini, R. Parker, P. Bonoli, M. Garrett, M. C. Kaufman, J. C. Wright, and S. Wukitch. Plasma wave simulation based on a versatile finite element method solver. *Physics of Plasmas*, 17(5):056119, 2010.

[67] M Brambilla. Numerical simulation of ion cyclotron waves in tokamak plasmas. *Plasma Physics and Controlled Fusion*, 41(1):1, 1999.

[68] R. W. Harvey and M. G. McCoy. The CQL3D Code. *Proceedings of the IAEA Technical Committee Meeting on Simulation and Modeling of Thermonuclear Plasmas*, 1992.

[69] M. Brambilla. . *Proc. Physics of Plasmas Close to Thermonuclear Conditions*, page 291, 1980.

[70] Martin Greenwald. Verification and validation for magnetic fusion. *Physics of Plasmas*, 17(5):058101, 2010.

[71] R.W. Harvey. X-Ray Energy Spectra Calculated from CQL3D Fokker-Planck Electron Distributions. Technical report, CompX Co., February 2001.

[72] R. W. Harvey, M. R. O'Brien, V. V. Rozhdestvensky, T. C. Luce, M. G. McCoy, and G. D. Kerbel. Electron cyclotron emission from nonthermal tokamak plasmas. *Physics of Fluids B: Plasma Physics*, 5(2):446–456, 1993.

[73] D.R. Ernst, P.T. Bonoli, P.J. Catto, W. Dorland, C.L. Fiore, R.S. Granetz, M. Greenwald, A.E. Hubbard, M. Porkolab, M.H. Redi, J.E. Rice, K. Zhurovich, and the Alcator C-Mod Group. Role of trapped electron mode turbulence in internal transport barrier control in the Alcator C-Mod tokamak. 11:2637, 2004. APS invited talk (2003).

[74] D. R. Ernst, N. Basse, W. Dorland, C. L. Fiore, L. Lin, A. Long, M. Porkolab, K. Zeller, and K. Zhurovich. Identification of TEM turbulence through direct comparison of nonlinear gyrokinetic simulations with phase contrast imaging density fluctuation measurements. In *Proc. 21st Int'l. Atomic Energy Agency Fusion Energy Conference, Chengdu, China, 2006*. IAEA-CN-149/TH/1-3.

[75] IDL code provided by M. Greenwald.

[76] H. E. Mynick and J. D. Strachan. Transport of runaway and thermal electrons due to magnetic microturbulence. *Physics of Fluids*, 24(4):695–702, 1981.

[77] Jinseok Ko, Steve Scott, Syun'ichi Shiraiwa, Martin Greenwald, Ronald Parker, and Gregory Wallace. Intrashot motional stark effect calibration technique for lower hybrid current drive experiments. *Review of Scientific Instruments*, 81(3):033505, 2010.

[78] G. M. Wallace, R. R. Parker, P. T. Bonoli, A. E. Hubbard, J. W. Hughes, B. L. LaBombard, O. Meneghini, A. E. Schmidt, S. Shiraiwa, D. G. Whyte, J. C. Wright, S. J. Wukitch, R. W. Harvey, A. P. Smirnov, and J. R. Wilson. Absorption of lower hybrid waves in the scrape off layer of a diverted tokamak. *Physics of Plasmas*, 17(8):082508, 2010.

[79] Peakfinding code provided by O. Meneghini.

[80] L.L. Lao, J.R. Ferron, R.J. Groebner, W. Howl, H. St. John, E.J. Strait, and T.S. Taylor. Equilibrium analysis of current profiles in tokamaks. *Nuclear Fusion*, 30(6):1035, 1990.

[81] P. T. Bonoli, J. Ko, R. Parker, A. E. Schmidt, G. Wallace, J. C. Wright, C. L. Fiore, A. E. Hubbard, J. Irby, E. Marmar, M. Porkolab, D. Terry, S. M. Wolfe, S. J. Wukitch, the Alcator C-Mod Team, J. R. Wilson, S. Scott, E. Valeo, C. K. Phillips, and R. W. Harvey. Lower hybrid current drive experiments on Alcator C-Mod: Comparison with theory and simulation. *Physics of Plasmas*, 15(5):056117, 2008.

[82] Jinseok Ko, Steve Scott, Manfred Bitter, and Scott Lerner. Design of a new optical system for alcator c-mod motional stark effect diagnostic. *Review of Scientific Instruments*, 79(10):10F520, 2008.

[83] G. M. Wallace et. al. *Nuclear Fusion*, 2011.

[84] Y. Takase, M. Porkolab, J. J. Schuss, R. L. Watterson, C. L. Fiore, R. E. Slusher, and C. M. Surko. Observation of parametric instabilities in the lower-hybrid range of frequencies in the high-density tokamak. *Physics of Fluids*, 28(3):983–994, 1985.

[85] R. Cesario, A. Cardinali, and C. Castaldo. Lower Hybrid Current Drive at ITER-relevant high plasma densities. *RADIO FREQUENCY POWER IN PLASMAS: 18th Topical Conference on Radio Frequency Power in Plasmas*, 2009.

[86] C. C. Petty, W. A. Cox, C. B. Forest, R. J. Jayakumar, J. Lohr, T. C. Luce, M. A. Makowski, and R. Prater. Modulated current drive measurements. *AIP Conference Proceedings*, 787(1):265–268, 2005.

[87] Y Peysson. Transport of fast electrons during lhcd in ts, jet, and asdex. *Plasma Physics and Controlled Fusion*, 35(SB):B253, 1993.

[88] R. Bartiromo, E. Barbato, L. Gabellieri, A. Spaziani, A.A. Tuccillo, F. Leuterer, F.X. Soldner, O. Gehre, and H.D. Murmann. Fast electron confinement during lower hybrid experiments in asdex. *Nuclear Fusion*, 33(10):1483, 1993.

[89] K. Hizanidis and A. Bers. Steady-state rf-current drive theory based upon the relativistic Fokker-Planck equation with quasilinear diffusion. *Physics of Fluids*, 27:2669–2675, November 1984. Equation 25 has relativistic slowing down collision frequencies.

[90] O Barana, D Mazon, L Laborde, and F Turco. Feedback control of the lower hybrid power deposition profile on tore supra. *Plasma Physics and Controlled Fusion*, 49(7):947, 2007.