Observations of helical plasma dynamics using complementary x-ray diagnostics in the MST

by

Patrick D. VanMeter

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

(Physics)

At the

UNIVERSITY OF WISCONSIN-MADISON

2020

Date of final oral examination: September 23, 2020

The dissertation is approved by the following members of the Final Oral Committee:

Daniel Den Hartog, Advisor, Distinguished Scientist, Physics

John Sarff, Professor, Physics

Paul Terry, Professor, Physics

Carl Sovinec, Professor, Engineering Physics

Luis Felipe Delgado-Aparicio, Research Physicist, PPPL

Acknowledgments

First and foremost, I would like to thank Daniel Den Hartog for his invaluable guidance, assistance, wisdom, and patience through my years as a graduate student. I would not have made it this far without his support, and I am fortunate to have had him as my advisor. I would also like to offer my thanks to Lisa Reusch for her mentorship, for providing me with guidance and advice, and for teaching me how to be a good scientist and to Luis Felipe Delgado-Aparicio for welcoming me into the ME-SXR collaboration; providing me with support, advice, and opportunities; and introducing me to delicious Peruvian cuisine.

The MST team has changed significantly during my time as a graduate student, but it has always been full of helpful, thoughtful, and friendly people. I would like to thank the scientific staff, past and present, for their knowledge, assistance, advice, and support: John Sarff, Abdulgader Almagri, Karsten Mccollam, Brett Chapman, Jay Anderson, Mark Nornberg, and John Goetz. I would also like to thank Peter Weix for keeping the MST in working order and for always considering safety first; John Wallace, for his countless hours of assistance with the ME-SXR project; and the rest of the support team: Alex Squitieri, Paul Wilhite, and Steve Oliva. Last but not least, I would like to thank my colleagues who showed me the ropes in the control room, ran diagnostics, helped me collect data, listened to my ideas, and generally made graduate school more enjoyable: Steph Kubala, Anthony Xing, John Boguski, Takashi Nishizawa, Phil Bonofiglo, James Duff, Lucas Morton, Eli Parke, Bhavya Kenia, Marcel Granetzny, Zichao Li, Hui Lian, Nick Lauersdorff, and many others. And thank you to the many other MST'ers who, over the years, built this program into what it is today.

I would like to also thank the other people who have helped me, one way or another, throughout this project: Paolo Franz, for many hours of expertise, for reams of useful code, and for forcing me to learn IDL in spite of my better judgment; Paul Terry and Ian McKinney, for helping me make comparisons to their QSH model; and my friend Drake Jackson, for providing my thesis with some much-needed proofreading and writing advice.

Graduate school is so much more than a research project. It's been hard work, but also a lot of fun, and I have made some great friends. In particular, I would like to thank the members of my cohort who have kept me company throughout this wild ride: Adrian Fraser, Ben Lemberger, Bunheng Ty, Chad Bustard, Chris Yip, Josh Karpel, Neil Campbell, Sam Neyens, Steven Casper, and Zach Buckholtz. Our Friday dinners, weekly traditions, game nights, and general camaraderie have no doubt kept me sane during my graduate school years.

I could never fully express in words the level of gratitude I owe to my family for their many years of love and support. Thank you to my parents, who taught me to value education, stimulated my interest in science, encouraged my curiosity, and supported my accomplishments even as they've taken me so far away from home. Thank you to my sister Faith, who seems to have a near-endless capacity to offer assistance and give advice, and to my grandparents, who have always been supportive of my education. I would also like to thank the rest of my extended family, far too numerous to list here by name. You have all played a role in making me who I am today.

Finally, I have to thank my fiancée, Oindrila. Writing a thesis is hard work, and writing a thesis in the middle of a pandemic is even harder. But your constant love and support has made the process manageable.

Abstract

This is a dissertation for the completion of a Doctorate of Philosophy in Physics degree granted at the University of Wisconsin-Madison.

A new multi-energy soft x-ray (ME-SXR) diagnostic based on the PILATUS3 100K x-ray camera has been installed on the Madison Symmetric Torus (MST) reversed-field pinch. This photon-counting camera consists of a two-dimensional array of $\sim$100,000 pixels for which the lower photon absorption cutoff energy can be independently set. This allows it to be configured for a unique combination of simultaneous spatial and spectral resolution. The energy dependence of each pixel was calibrated by scanning the individual pixel thresholds while the detector was exposed to fluorescence emission of differing photon energies. The resulting data are then fit to a characteristic "S-curve." The statistical variation of this calibration from pixel-to-pixel was explored, and it was found that the discreteness of the threshold setting results in an effective threshold resolution of $\Delta E < 100$ eV for high-gain settings and $\Delta E < 200$ eV for medium gain. In order to properly interpret the data a full forward model has been developed which produces realistic chord-integrated ME-SXR synthetic measurement given the underlying T_e, n_e, neutral density, and impurity density profiles. A method for using this model within a Bayesian framework to extract equilibrium profiles is presented. This diagnostic forms part of a suite of complementary x-ray diagnostics on the MST, also including a two-color diode based tomography array, a Ross spectrometer, and a hard-xray detector. This suite was used to study the structure and evolution of the temperature profile during the saturated quasi single-helicity (QSH) state which forms spontaneously in non-reversed plasmas with high Lundquist number. During QSH, increased electron temperatures $T_e > 700$ eV were observed, with steep gradients $|\nabla T_e| > 3$ keV m^{-1}, suggesting the formation of a thermal transport barrier. A brief period of significantly enhanced confinement was observed during which secondary tearing activity was minimal, allowing a broad thermal structure to form. Runaway electrons were also found to be well-confined with a population energy $E_r > 18$ keV, suggesting the presence of restored flux surfaces. Connections are drawn to a theoretical model which relates the QSH state to shear in the magnetic or flow velocity profile. Observations of predator-prey-like oscillations between the dominant tearing mode amplitude, secondary mode amplitudes, and thermal emissions are consistent with this model.

Contents

List of Figures

Chapter 1

Introduction

The goal of the magnetic confinement research community, to which this work belongs, is to establish nuclear fusion as a viable source of abundant clean energy. For inspiration, one need only look up: every star in the night sky is an example of a functioning nuclear fusion reactor [1]. The problem is not whether it can work, but how to build it. As the common idiom says, the devil is in the details. And when the goal is to develop a system which can hold an ionized gas in steady-state at 10^8 K for long periods of time using only external magnets, there are many, many details.

This thesis, like all publications in this field, is about a small-but-important subset of those details. High temperature plasmas tend to emit substantial x-ray radiation. Although this radiation is a channel through which confined energy can be lost, it also provides a wealth of information about the plasma that emitted it. By examining a plasma's x-ray spectrum, we can learn about its temperature, elemental composition, and even its magnetic topology. This thesis focuses on the development of a new soft x-ray diagnostic for fusion-grade plasmas. This diagnostic forms part of a larger suite of complementary x-ray diagnostics on the Madison Symmetric Torus (MST) which, together, are able to thoroughly characterize the evolution of a plasma's properties. This suite is applied to study the physics of a self-organized helical equilibrium which spontaneously forms in the MST.

This chapter is intended not only to introduce the work discussed in the remainder of this thesis but also to situate it in the wider context of magnetic confinement fusion research. Section 1.1 provides a high-level motivation of the need for nuclear fusion energy and the development of a research community aimed at meeting this challenge. Section 1.2 narrows this scope somewhat to focus on the important topic of plasma diagnostics, the sometimes-unintuitive tools which allow us to determine a plasma's properties. Section 1.3 discusses one specific diagnostic, the multi-energy soft x-ray detector, which is the focus of much of this thesis. Section 1.4 provides an overview of the MST, the experimental facility where this detector was tested. Finally, Section 1.5 provides a general overview of the structure of this document and summarizes the key results.

1.1 The climate crisis and controlled thermonuclear fusion

In the year 2019, less than 20% of energy production in the United States came from sources which do not emit carbon dioxide into the atmosphere [2]. As shown in Figure 1.1, this consists of 11% from noncombustible renewables (hydroelectric, geothermal, wind, and solar) and another 8% from nuclear fission reactors. The remaining 81% of our energy production relies on the combustion of fossil fuels, releasing large quantities of carbon dioxide into the atmosphere. These greenhouse gases, released in large quantities by human society following the Industrial Revolution, give rise to a warming climate [3, 4]. This human-made climate change has already led to an increase in global temperatures, a reduction in the size of glaciers and a corresponding rise in sea levels, an increase in heat waves, and a reduction in the populations and ranges of numerous plants and animals [5]. Climate models show that the consequences could be far more severe if warming is not kept within an acceptable range, often estimated to be within 2°C [6]. This emphasizes the immediate need to invest in the development and deployment of carbon-free energy sources.

A 2019 report published by the Nuclear Energy Agency, a part of the Organisation for Economic Co-operation and Development, found that a decarbonization strategy which relies heavily on variable renewable energy[1] suffers a large cost penalty due to the requirement of excess capacity [7]. More economical approaches tend to rely heavily on nuclear energy to fill the void left behind by carbon-emitting fossil fuels. Although nuclear fission plants are currently far safer than fossil fuel-based energy production [8], there are still important concerns regarding overall safety, waste disposal, and proliferation [9]. This leaves an opening for energy production from thermonuclear fusion to contribute significantly to the global de-carbonization effort [10]. For a more thorough discussion of the safety concerns related to nuclear fusion energy, see chapter 14 of Reference [11].

Nuclear fusion is the process in which two or more atomic nuclei are combined to form an atom of a heavier element. When the mass of the product nuclei is less than that of the reactants, the binding energy given by $\Delta E = \Delta mc^2$ is released in the form of radiation [12]. This effect is most significant for very light reactant elements such as hydrogen. This process is the primary source of energy for all stars, including the Sun. In the 1950s this process was exploited by humanity for the first time in the form of thermonuclear weapons (i.e., the hydrogen bomb). The first test of this so-called *uncontrolled* nuclear fusion had an explosive power more than 500 times that of the bombs dropped on Nagasaki and Hiroshima, beginning a process of proliferation which spread a global fear of annihilation from nuclear war [13]. Since that time, there have been numerous attempts to tame the process for the peaceful goal of terrestrial energy production. However, this goal of achieving *controlled* thermonuclear fusion has turned out to be incredibly difficult and has inspired years of rigorous academic study.

[1]This category includes sources like wind and solar energy, whose production capacity depends on external factors like the weather and therefore cannot be deployed on-demand by energy companies.

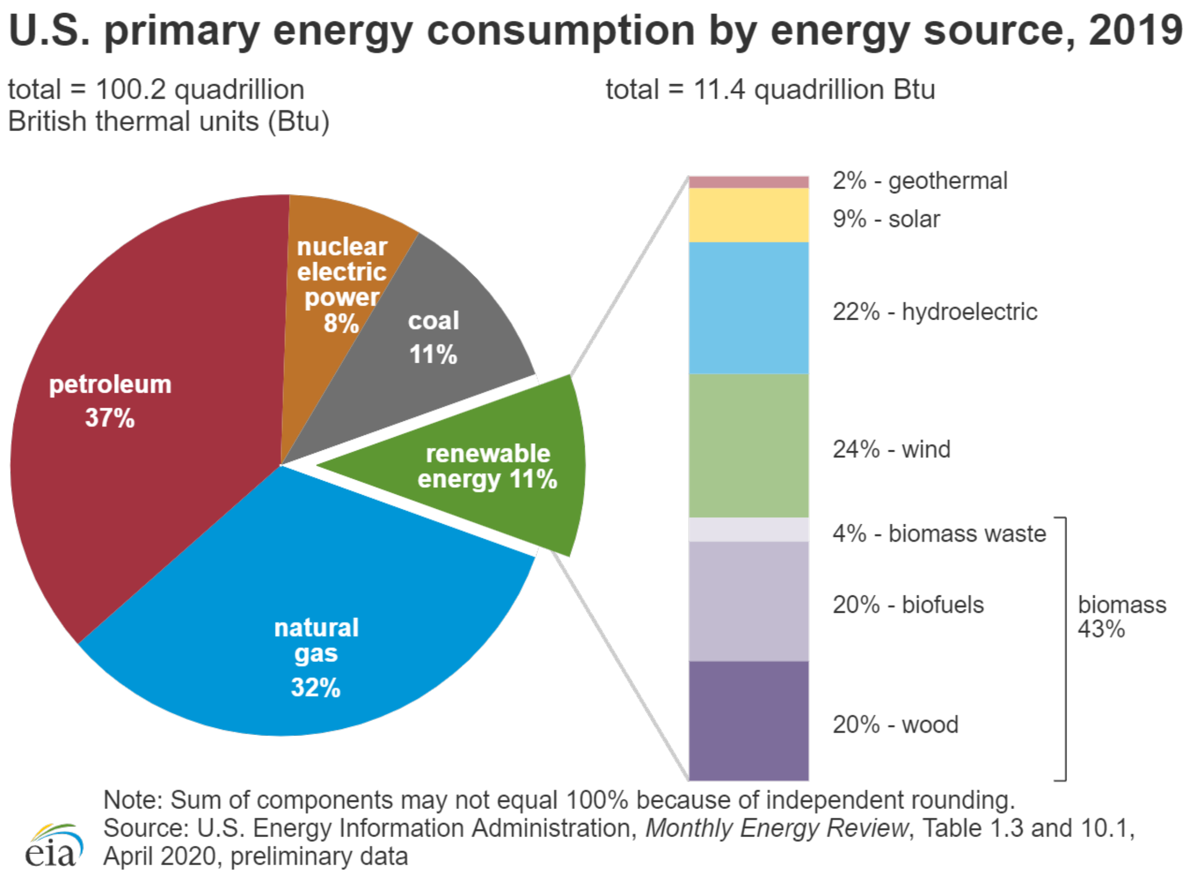

Figure 1.1: U.S. primary energy consumption by energy source for the calendar year 2019. Source: US EIA, eia.gov/energyexplained/us-energy-facts/.

Mastery of controlled thermonuclear fusion will necessarily require an understanding of plasma physics. Current designs for controlled thermonuclear reactors based on a deuterium-tritium mixture fuel source (widely regarded as the most efficient option) require an operating temperature of $T \sim 10$ keV, or approximately 10^8 Kelvin [14]. At such extreme temperatures the fuel source will necessarily be a plasma. This temperature far exceeds the melting point of tungsten (3,687 K), the highest of any naturally occurring element, meaning that confinement of a thermonuclear reaction cannot rely solely upon physical boundaries. There are currently two prominent approaches to solving this problem: inertial confinement fusion (ICF) and magnetic confinement fusion (MCF).

In the ICF approach, extremely high-energy beams (typically lasers) are used to implode a small D-T fuel pellet, thereby triggering a fusion reaction and releasing large amounts of energy [15]. The reaction occurs over very short time-scales, so external confinement is not necessary. In contrast, MCF (the focus of this thesis) uses electromagnetic forces to confine the plasma to a predetermined volume with steep gradients separating the hot core from the plasma wall [16]. The difficulty in the MCF approach arises from the fact that confined plasmas are subject to long-range interactions via the electromagnetic force, which leads to many modes of collective behavior that tend to resist stable confinement. This is the main hurdle which MCF research must overcome.

For over 50 years, the tokamak has been the dominant approach for MCF reactors. Developed in 1968 in the Soviet Union, tokamaks are toroidal devices which make use of a strong central magnetic field to confine and stabilize the plasma [16]. Progress has been made over generations of increasingly large and energetic devices [17], culminating in the next-generation ITER project currently under construction in the south of France [18]. This international mega-project aims to build the first tokamak that produces more energy than it consumes, thus demonstrating the viability of the concept as a power plant. An illustration of the planned vacuum vessel is shown in Figure 1.2 demonstrating the vast size of the device. Currently, first plasmas are planned for 2025, with a slow ramp-up to full D-T operations by 2035 [19]. Following the success of ITER, plans are currently being developed to build a follow-up demonstration power plant (or several), tentatively referred to as DEMO [20]. However, this is not the only path to fusion energy that is currently being explored. Scientists from MIT have proposed the ARC reactor ("Affordable, Robust, Compact") which proposes to deliver fusion energy on an expedited timeline by trading out ITER's large volume for extremely high-strength magnetic fields, based on recent innovations in high-temperature superconducting magnets [21]. Likewise there is significant interest in the stellerator, an alternative magnetic configuration based on using complex helical magnetic geometries to confine the plasma and permit current-free operation. A prominent example is Wendelstein 7-X in Germany [22]. It has been argued that the stellerator might provide a superior path to fusion energy due to its inherent stability and scalability [23].

The MCF plasma physics community consists of much more than these large international collaborations. For decades, much of the work of understanding the physical behavior of magnetically-confined plasmas has occurred in small- and medium-scale devices, often associated with universities. These devices provide a more flexible environment to explore questions of fundamental physics, test theoretical predictions, benchmark computational models, and develop technology. The Madison Symmetric Torus (MST), a reversed-field pinch on which the experiments described in this thesis were conducted, is one such device.

Much of this thesis is concerned with the development of a new multi-energy soft x-ray (ME-SXR) diagnostic concept for magnetically confined plasmas. As will be discussed in Section 1.3, this diagnostic will be well-suited for addressing the difficulties arising from the extreme conditions and limited access that will exist on devices like ITER. However, it is important that the technology first be developed and tested on less restrictive plasma sources like the MST. Before discussing the ME-SXR project further, we must first discuss both the role that plasma diagnostics play in the field of experimental

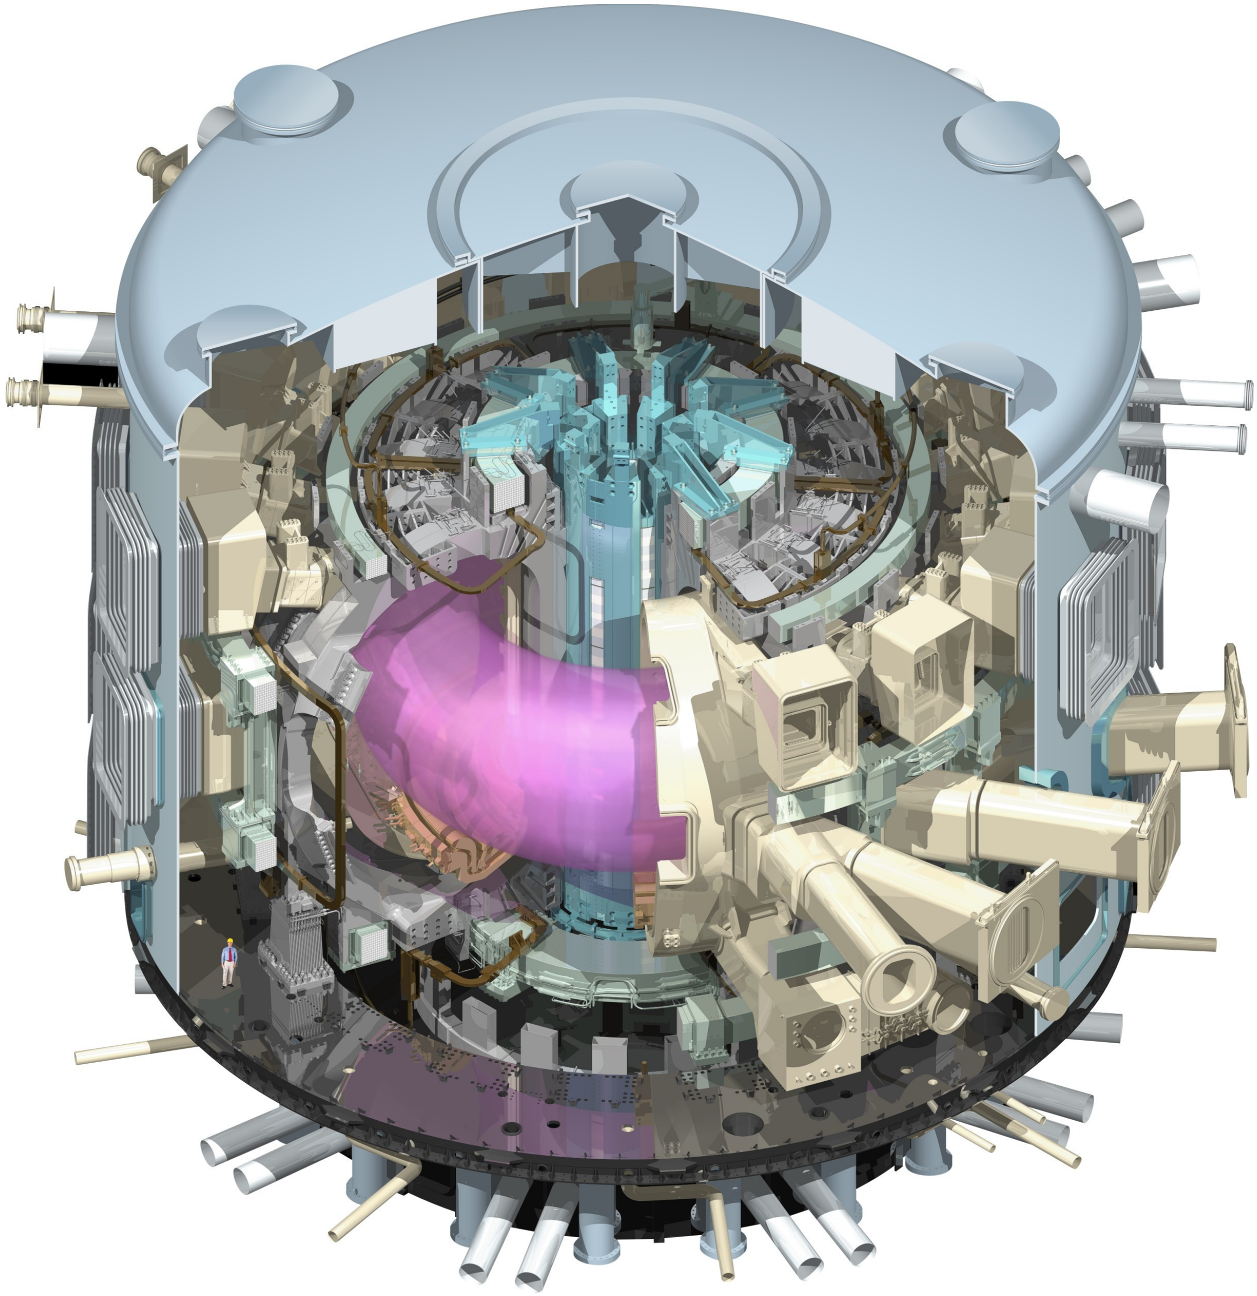

Figure 1.2: Conceptual rendering of the ITER tokamak and surrounding magnetic field and cryogenic subsystems. ITER aims to be the first controlled fusion device. A person is drawn in the lower left side for scale. Source: ITER Organization, http://www.iter.org/.

plasma physics and also the MST itself.

1.2 Diagnosing the properties of a confined plasma

In order to understand the physical principles driving the evolution of a confined plasma, it is first necessary to measure the plasma's properties. This task is made difficult by the very nature of plasma confinement. Fusion plasmas are typically very hot and can evolve on very short time scales, eliminating the possibility of using traditional physical temperature sensors. In fact, any internal probes can disturb the magnetic equilibrium and therefore degrade confinement. Likewise, confined plasmas typically produce strong and rapidly varying magnetic fields, leaving any unshielded components vulnerable to unexpected electromagnetic interference. In a sense, the extreme environment of a fusion plasma begets extreme requirements for diagnostic instrumentation. As such, the study of plasma diagnostics has become an active area of research in its own right alongside plasma physics.

The philosophy of plasma diagnostics is laid out well on the first page of I. H. Hutchinson's seminal work *Principles of Plasma Diagnostics*:

> The overall objective of plasma diagnostics is to deduce information about the state of the plasma from practical observations of physical processes and their effects. This usually requires a rather elaborate chain of deduction based on an understanding of the physical processes involved. In more mundane situations the same is true of other diagnostic measurements; for example, a mercury/glass thermometer relies on the physical process of thermal expansion of mercury, which determines the height of the mercury column observed. However, since plasmas have properties that are often rather different from the more familiar states of matter met in everyday life, the train of reasoning is sometimes more specialized and may seem more obscure, especially since

plasma diagnostics are rarely routine [24].

Plasma diagnostics are designed to exploit the unintuitive characteristics of a plasma in order to measure its properties. These measurements are then used to further our understanding of the field of plasma physics. This enhanced understanding can in turn present new characteristics which may be further exploited for diagnostic innovation.

Although an experiment's diagnostic suite must necessarily be customized to the peculiarities of each individual device, some standard diagnostic approaches have been developed. Hutchinson's book breaks these into eight categories based on the underlying physical mechanism that is being exploited: magnetic measurements, particle flux measurements, measurements based on refractive index, electromagnetic emission from free electrons, electromagnetic emission from bound electrons, scattering electromagnetic waves, neutral atom diagnostics, and fast ion measurements [24]. A brief overview of the standard diagnostics on MST which are relevant to this thesis is given in Chapter 2, and the development of a new multi-energy soft x-ray diagnostic is a central focus of this thesis.

The categories described by Hutchinson include all but the most specialized instrumentation in use today, but they are not always independent. For instance, the ME-SXR diagnostic which is the focus of much of this thesis measures electromagnetic emission from both free and bound electrons and is also sensitive to the geometry of internal magnetic fields and particle densities. In the view where one diagnostic is intended to provide an independent measurement of one plasma property, these complex interdependencies are often seen as a nuisance. However, if properly understood and handled, these interdependencies form the basis of a powerful approach to data analysis. By exploiting the differing ways in which multiple plasma diagnostic measurements are correlated to one another via the underlying plasma parameters, it is possible to produce measurements more accurate than is possible with any individual diagnostic [25]. This methodology, termed *integrated data analysis* (IDA) [26], is also capable of measuring properties which cannot be measured by any individual diagnostic [27]. This approach requires extensive modeling and substantial computational resources, but the end results often justify the effort. A detailed methodology for IDA based on Bayesian inference is presented and applied in Section 5.5.

The need for this kind of sophisticated methodology is becoming more apparent as we prepare to enter the era of burning plasmas. Diagnostics on ITER will face far more extreme conditions (temperature, exhaust power, neutron fluence) than any magnetic confinement device that has been built before. These conditions in turn will pale in comparison to those encountered in a future DEMO-scale device [28]. All the while, access to the plasma will be much more limited than ever before, with unprecedented amounts of space required for large auxiliary systems and tritium breeding blankets. Routine maintenance will be all but impossible due to irradiation concerns. In such an environment, it is critical to extract the maximum amount of information from the minimum number of diagnostics. In this aim, IDA will be indispensable. Diagnostics will also need to be non-perturbative, robust, and versatile. As a result, measurements of electromagnetic radiation emitted from the plasma, and x-rays in particular, will be critical to the future of MCF diagnostics. These requirements provide a strong motivation for the multi-energy soft x-ray project currently under development.

1.3 The multi-energy soft x-ray project

The multi-energy soft x-ray (ME-SXR) diagnostic concept is based on the use of hybrid photon counting detector (HPCD) technology to combine the broadband energy sensitivity of multi-foil diode arrays with the versatility of pulse height analysis detectors. The PILATUS series of HCPDs allows the user to individually adjust the lower threshold energy for photon detection over a wide range on a per-pixel basis. This results in an unprecedented combination of spatial, spectral, and temporal sensitivity from a single diagnostic [29]. This sensitivity allows for simultaneous measurements, using a single compact diagnostic, of the core electron temperature and impurity densities (and therefore Z_{eff}) as well as the non-Maxwellian population (i.e., runaway electrons). Technical details about the ME-SXR diagnostic implementation on MST, including installation, calibration, and analysis, can be found in Chapter 3.

An ME-SXR detector based on the DECTRIS PILATUS2, a 450 μm silicon HPCD sensitive to photons between 1.5-30 keV, has been previously installed and tested on the Alcator C-mod tokamak [30]. The detector consists of an array of approximately 100,000 individual pixels configured with thirteen different thresholds. It was demonstrated that the detector was able to provide simultaneous core T_e and tungsten concentration measurements over the duration of a plasma discharge (Figure 1.3). These measurements increased as expected during periods of heating via lower hybrid current drive and tungsten impurity injection. The diagnostic presented in this thesis is a direct continuation of that work. The new ME-SXR diagnostic on the MST is based on the PILATUS3 [31], an updated version of the PILATUS2 which features reduced dead time. A multi-energy hard x-ray (ME-HXR) diagnostic based on a CdTe PILATUS detector [32] is also being developed and tested on the WEST tokamak [33].

The ME-SXR project is part of a larger push to develop real-time diagnostic and control systems to combat the accumulation of high-Z impurities in the core of next-generation tokamak systems [34]. The system is envisioned as complementing rather than replacing existing x-ray diagnostics, including an existing x-ray crystal imaging spectrometry [34] and a traditional diode-based multi-color system [35]. The first ME-SXR diagnostic was installed and tested on Alcator C-Mod [30]. The next generation diagnostic was intended for use on the NSTX-U tokamak, but after technical difficulties forced that device offline the collaboration with UW-Madison was established. The MST produces a fusion-grade plasma with sufficient soft x-ray emission for temperature diagnosis, and the presence of aluminum ions with line emissions ~ 2 keV makes the MST

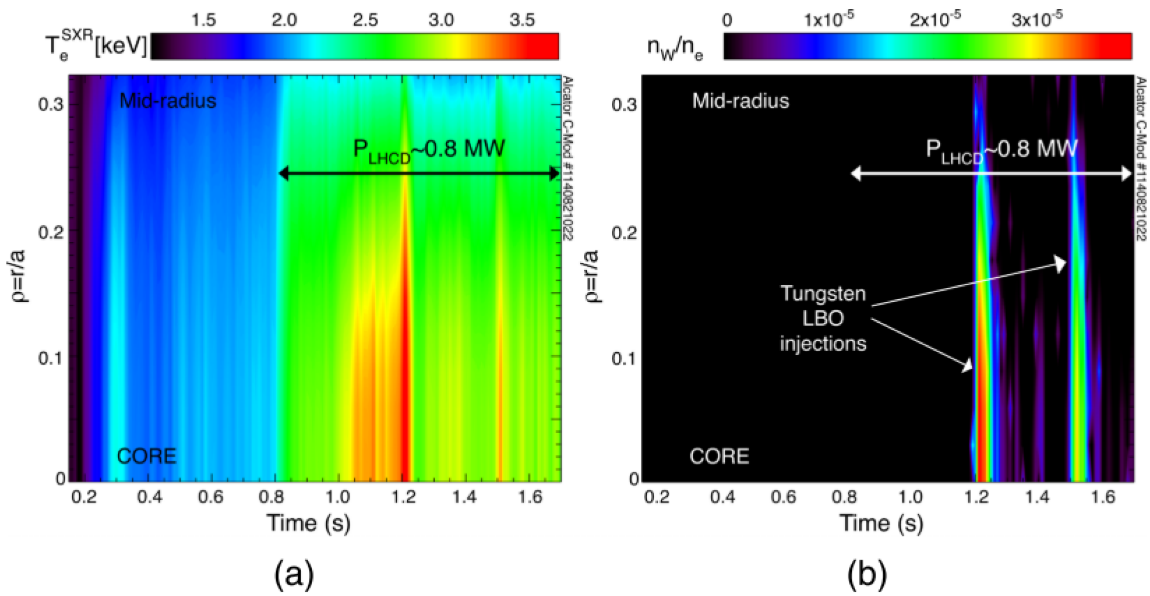

Figure 1.3: Results on Alcator C-Mod from a single ME-SXR measurement: (a) Reconstructed T_e during Ohmic and LHCD heating phases; (b) Core impurity concentrations from tungsten laser-blow-off injections. Figure reproduced from L.F. Delgado-Aparicio, *et al.* [29]

an interesting test case for the ME-SXR technique. Future iterations of the ME-SXR diagnostic are planned for installation on WEST, NSTX-U, and JT60-SA. The diagnostic is being developed with an eye toward ITER, although additional developments regarding the radiation hardness of Si-based detectors will be necessary [36].

The explicit goals of the ME-SXR collaboration between UW-Madison and PPPL is outlined in Figure 1.4. This thesis demonstrates significant progress on all of the diagnostic goals, and all three core physics goals are addressed. The diagnostic goals are addressed in Chapter 5, which considers both thermal and impurity ion profiles. The evaluation of Z_{eff} using an IDA methodology is one of the main results of Section 5.5. The physics of quasi-single helicity (QSH) plasmas is primary focus of Chapter 6. Impurity transport studies are not the focus of this thesis, but the topic is addressed at various points. The impact of hard x-rays on ME-SXR measurements, and using this effect to

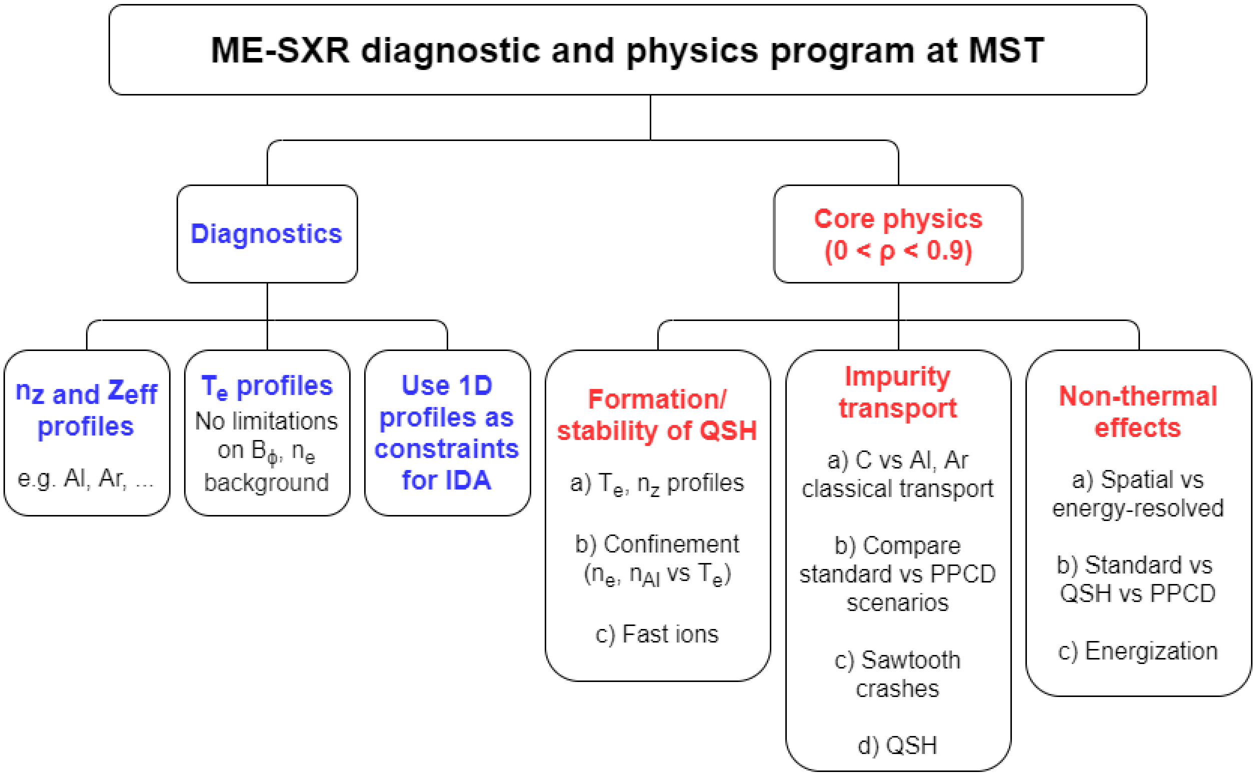

Figure 1.4: Overview of the physics goals of the ME-SXR project on MST. Courtesy of L. F. Delgado-Aparicio.

quantify the runaway population during QSH, is discussed in Section 6.3. Significant progress has also been made in diagnostic development, including characterizing the pixel-to-pixel variation of the energy calibration (Chapter 3) and developing a physics-based quantitative forward model of the diagnostic (Chapter 4).

Before continuing on to these results, however, it is worthwhile to discuss the details of the Madison Symmetric Torus itself. Although similar in many ways to a tokamak, the distinctions become important when addressing the specific physics goals discussed above.

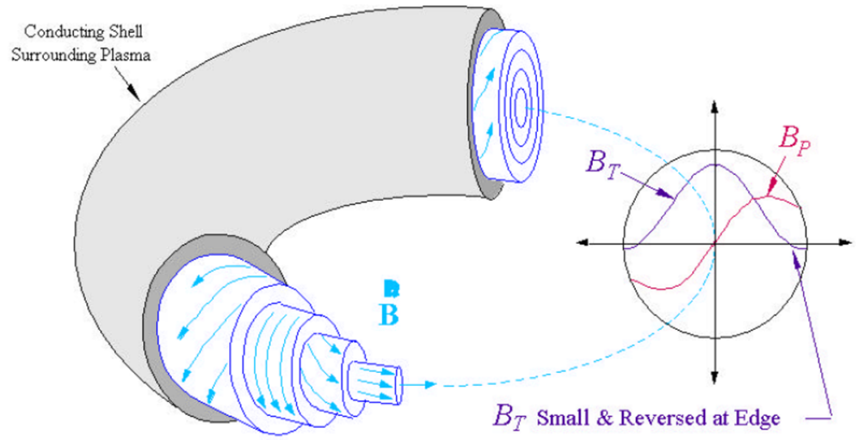

Figure 1.5: The magnetic field in an RFP is predominantly toroidal in the core but reverses near the edge.

1.4 The Madison Symmetric Torus Reversed-Field Pinch

The history of the RFP concept stretches back to the 1960s [37]. Observations on the ZETA ("Zero Energy Thermonuclear Assembly") toroidal pinch device showed that the direction of the toroidal field would occasionally reverse near the wall, leading to a brief period of improved stability which was referred to as the "quiescent" period [38]. Attempts to explain this phenomenon led to the development of relaxation theory by J. Taylor [39]. According to this theory, the conservation of the global magnetic helicity,

$$K = \int \boldsymbol{A} \cdot \boldsymbol{B} \, d^3\boldsymbol{x}, \tag{1.1}$$

in the presence of a boundary shell with finite conductivity causes the plasma to spontaneously relax into a minimum energy state described by the equilibrium

$$\boldsymbol{\nabla} \times \boldsymbol{B} = \lambda \boldsymbol{B}, \tag{1.2}$$

where $\lambda = \mu_0 J_{\|}/B$ is a global constant. This equilibrium is commonly referred to as a *Taylor state*.

The concept of the RFP attracted attention as an alternative to the burgeoning tokamak configuration, in which a large toroidal field $B_\phi >> B_\theta$ is used to stabilize the plasma. In contrast, the RFP has $B_\phi \sim B_\theta$, and B_ϕ reverses direction at a point near the vacuum vessel wall (see Figure 1.5). The extent of the reversal is quantified by the *reversal parameter*,

$$F = \frac{B_\phi(a)}{\langle B_\phi \rangle}. \tag{1.3}$$

The chosen equilibrium greatly affects the MHD instabilities, called *tearing modes*, present in each type of configuration. This can be illustrated by considering a single helical perturbation $\tilde{\boldsymbol{b}}$ to the equilibrium magnetic field $\boldsymbol{B}$ with a wave vector $\boldsymbol{k} = \frac{m}{r}\hat{\boldsymbol{e}}_\theta + \frac{n}{R}\hat{\boldsymbol{e}}_\phi$. The condition for resonance is given by

$$\boldsymbol{k} \cdot \boldsymbol{B} = 0, \tag{1.4}$$

which simplifies to

$$\frac{m}{r}B_\theta + \frac{n}{R}B_\phi = 0 \rightarrow q = -\frac{m}{n}, \tag{1.5}$$

where r is the minor radius, R is the major radius, and we have defined the *safety factor* q to be

$$q = \frac{rB_\phi}{RB_\theta}. \tag{1.6}$$

The criteria $q = -m/n$ specify that tearing modes are resonant on *rational surfaces*, which are toroidal surfaces within the plasma volume characterized by closed magnetic field lines. For all other flux surfaces within the plasma, q is irrational, meaning that field lines cover the surface ergodically. The closed nature of the field lines on rational surfaces, combined with shearing in the magnetic field ($dq/dr \neq 0$), makes these surfaces vulnerable to significant topological changes as a result of small perturbations in the radial magnetic field. This process is illustrated in Figure 1.8 and is discussed further below.

The application of a large B_ϕ in the tokamak generates an equilibrium in which $q > 1$ for all radii, meaning that no $m = 1$ tearing modes are resonant. In contrast, the RFP has $q < 1$ for all radii, meaning that $m = 1$ instabilities are resonant in the core and that $m = 0$ instabilities are resonant at the reversal surface (where $B_\phi = 0$). A typical q profile for MST is shown in Figure 1.6. This results in substantial differences in stability and transport behavior between the two configurations. The remainder of this section will focus on the properties of a specific reversed-field pinch device. For more information about the tokamak and other modern magnetic confinement configurations, I refer the reader to standard textbooks [16] and review articles [40].

The experimental work presented in this document was performed on the Madison Symmetric Torus reversed-field pinch device. The MST began plasma operations in 1988 with the stated goals of studying the effect of large plasma size on confinement and exploring self-organized phenomena such as turbulence, transport, and the MHD dynamo in the RFP [41]. The device itself consists of a toroidal aluminum conducting shell (with a minor radius of 0.52 m and a major radius of 1.5 m). The plasma current is driven by a single iron-core transformer, with typical values of $I_p \sim 200 - 500$ kA. Plasma heating is entirely Ohmic, with no requirement for auxiliary heating schemes like neutral beam injection or RF. This property has made the RFP appealing as a fusion concept. Due to the properties of self-organization resulting from the plasma's interaction with the conducting shell, fewer external magnetic field coils are required to maintain the equilibrium, another significant advantage over the tokamak and the stellerator. However, the prevalence of core-resonant $m = 1$ tearing modes leads to substantial stochastic transport of

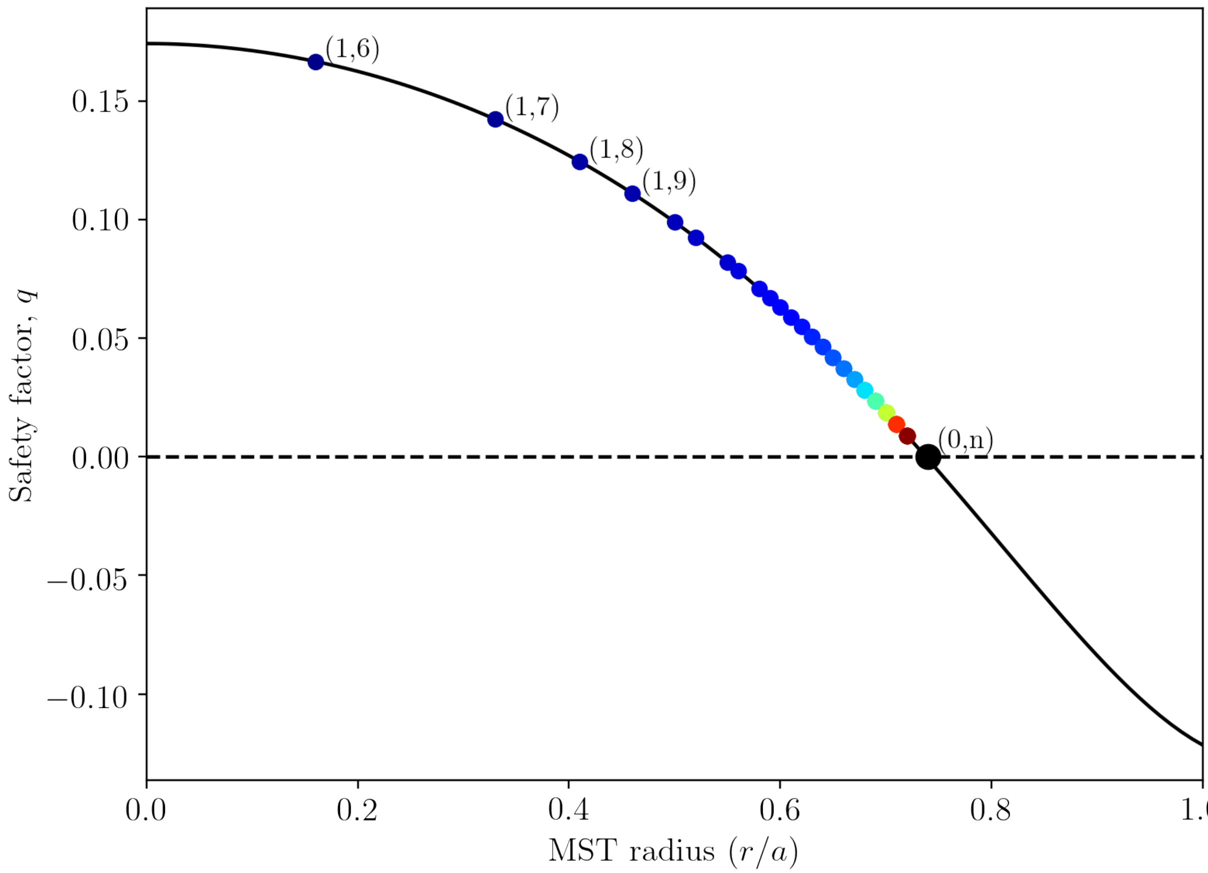

Figure 1.6: Reconstruction of the q profile for MST shot #1191204074 (400 kA PPCD). The $m = 1$, $n = 6$ mode is resonant in the core, with several higher-n modes (7, 8, 9) resonant along the mid-radius. Approaching the reversal surface, infinitely many higher-n resonant surfaces are densely packed. On the reversal surface, designated by $(0, n)$, all $m = 0$ modes are simultaneously resonant. q becomes negative in the outer radius, indicating field reversal.

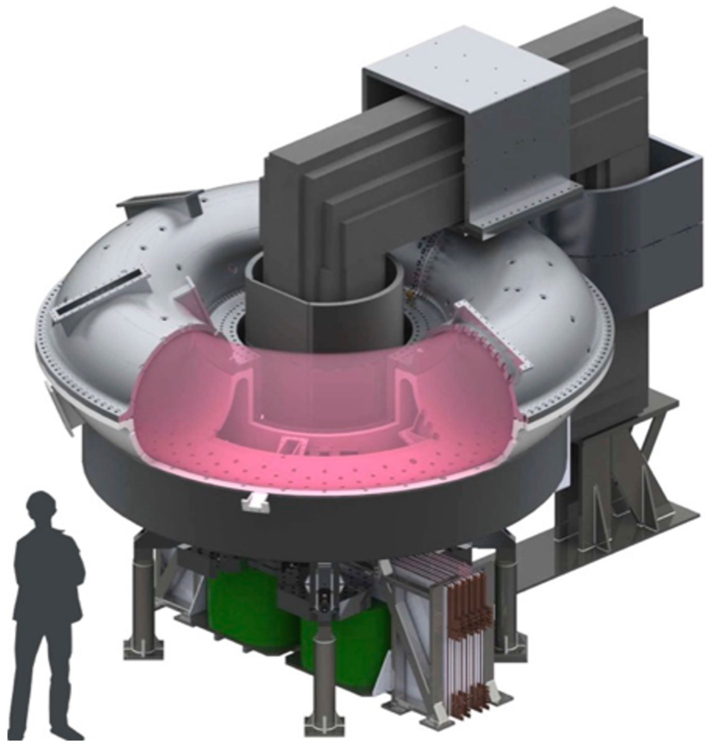

Figure 1.7: Illustration of the Madison Symmetric Torus vacuum vessel, TF, and PF systems. A person is also drawn for scale.

confined particles. This presents a significant challenge for the RFP configuration and will be further discussed later in this section. A summary of MST's basic properties is provided in Table 1.1, and the geometry of this system is illustrated in Figure 1.7.

As previously mentioned, both the tokamak and the RFP are subject to a class of instabilities referred to as tearing modes [14]. These global resistive kink instabilities, characterized by poloidal and toroidal mode numbers m and n respectively, are driven by a gradient in $J_{||}$ and are resonant on the corresponding $q = m/n$ rational surfaces within the plasma volume. In the RFP these instabilities can cause magnetic reconnection events near the corresponding rational surfaces whereby large amounts of energy are rapidly released as the current profile is suddenly relaxed [42]. This occurs quasi-periodically in standard RFP operation and is the mechanism by which the plasma relaxes toward

MST typical parameters		
Parameter	**Symbol**	**Value**
Major radius	R_0	1.5 m
Minor radius	a	0.52 m
Plasma current	I_p	200-500 kA
Pulse length	τ	< 100 ms
Electron density	n_e	$0.5 - 1.0 \times 10^{19}$ m^{-3}
Electron temperature	T_e	200-2000 eV
Core magnetic field	B_0	≤ 0.5 T
Plasma beta	β	10%

Table 1.1: Typical parameters for the MST discharges considered in this thesis. Note the temperature ranges $T_e > 500eV$ are found only in improved-confinement scenarios 1.4.1.

a Taylor state. In the RFP literature individual relaxation events are often referred to as "sawtooth" events due to the characteristic shape featured on many waveforms, and the cyclic process is the "sawtooth cycle."

Magnetic reconnection occurs when, in the presence of finite dissipation, magnetic field line topology suddenly changes in a way which is forbidden in ideal MHD. This results in the formation of enclosed topological structures called magnetic islands on the rational surfaces, as illustrated in Figure 1.8. These structures occur periodically with the same helicity as the associated tearing mode. The presence of magnetic islands profoundly affects the transport properties of the RFP. A reduction in electron thermal transport, resulting in a localized increase in T_e, has been found inside magnetic island structures [43]. In many cases, however, the presence of magnetic islands leads to a substantial reduction in plasma confinement.

As the tearing modes grow in amplitude, magnetic islands from adjacent rational surfaces can begin to overlap. When islands overlap, the underlying good flux surfaces are destroyed, meaning that the trajectory of a magnetic field line becomes stochastic throughout the region. Particle transport thus becomes stochastic, resulting in significantly reduced confinement [45]. In some cases overlap with the core $n = 6$ island is only partial, allowing a small region of good flux surfaces to survive. It has also been

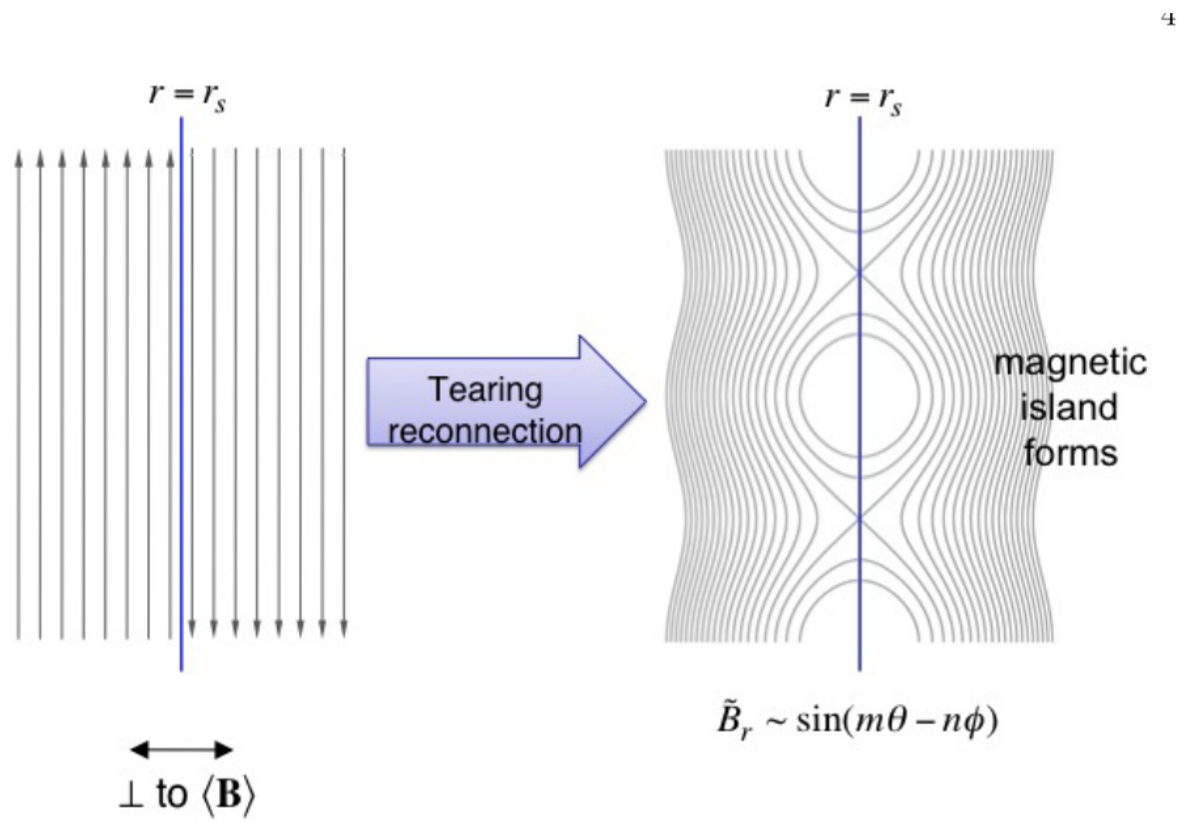

Figure 1.8: Illustration showing the generation of a magnetic island on a rational surface r_S as a result of a radial perturbation in the magnetic field. Figure reproduced from Reference [44].

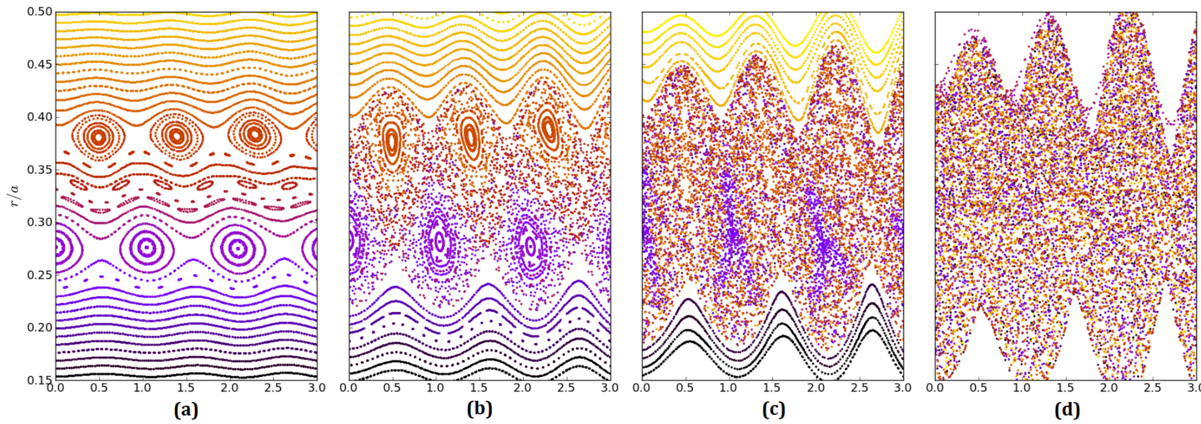

Figure 1.9: Poincare plots illustrate the development of magnetic stochasticity as the width of two magnetic islands with different helicities increase and begin to overlap. In a), tearing mode amplitudes are low, and the region contains good flux surfaces. In b) and c), as the amplitudes increase, the good flux surfaces are destroyed, but some residual island structure remains. In d) the mode amplitudes are sufficiently high that the magnetic structure has become totally chaotic. Figure reproduced from Reference [47].

observed that even in cases where overlap is significant some residual island structure may survive, resulting in a regime where transport is only partially stochastic [46]. The effect of overlapping islands inducing magnetic stochasticity is illustrated in Figure 1.9.

1.4.1 Enhanced confinement scenarios

The effect of core-resonant tearing modes in degrading confinement in the RFP is one of the configuration's biggest drawbacks as a fusion reactor concept. As a result, multiple techniques have been devised to attempt to mitigate this effect and improve confinement times to a reactor-relevant regime. Two such techniques are detailed in this section. The first involves using active current control to flatten the current profile, thereby suppressing the gradient which drives tearing modes. The other approach involves guiding the plasma into a self-organized helical state in which energy transfer between tearing modes is naturally suppressed.

The most successful strategy that has been used to reduce transport in the MST is Pulsed Parallel Current Drive (PPCD) [48]. This technique involves increasing $E_{\|}$ in the plasma edge in order to manually flatten the $J_{\|}$ profile, thereby greatly reducing the energy available to the core-resonant tearing modes. This technique leads to a substantial reduction in stochasticity and corresponding improvement to confinement. The core can reach temperatures up to $T_e = 2$ keV, substantially higher than the $T_e \sim 500$ eV achievable in standard RFP operations. If PPCD is engaged immediately after an energetic reconnection event (a process called "crash heating"), ion temperatures $T_i \sim 1$ keV can be achieved [49]. With a global energy confinement time of $\tau_E \sim 12$ ms, this mode of operation is comparable to the tokamak H-mode confinement time scaling for a device of comparable size and operating parameters (see Figure 1.10). This improved confinement also leads to the generation of a runaway electron population localized to the plasma core [50].

The major drawback of PPCD is that it is inherently transient, as it makes use of a constant change in flux to induce the parallel current drive [48]. As a result, the improved confinement period only lasts until the PPCD capacitor banks have fully discharged, which is currently a duration of ~ 10 ms. An alternative scheme for achieving improved confinement using oscillating current drive have been tested [51, 52], but no such technique is yet available for routine use on the MST.

An alternative method for improving plasma confinement properties is based upon the self-organized quasi-single helical (QSH) state [53]. Under conditions of high current and low density, the plasma tends to spontaneously transition into a regime where the core-most resonant $m = 1$ tearing mode grows to very large amplitude while other tearing modes are strongly suppressed. Given sufficient amplitude, the core magnetic island can grow to encompass the magnetic access, producing a three-dimensional helical equi-

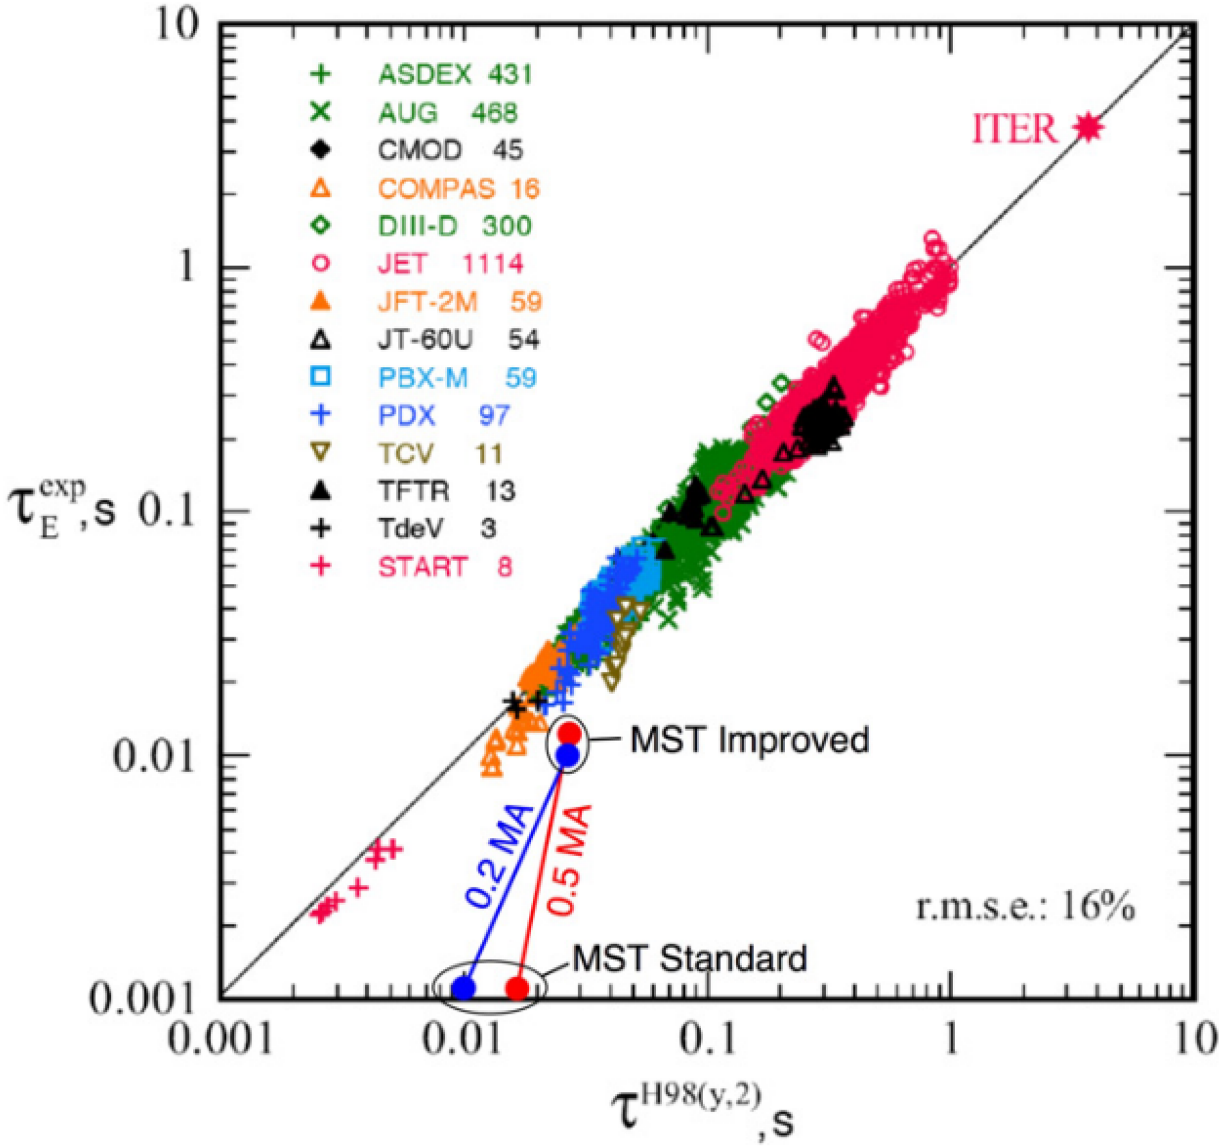

Figure 1.10: Comparison of standard RFP and crash-heated PPCD confinement times with the empirical ELMy H-mode scaling specified by IPB98(y,2). Figure reproduced from B. Chapman, *et al.* [49].

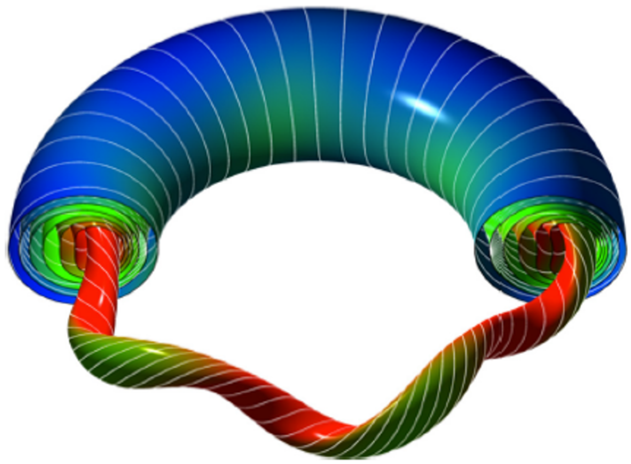

Figure 1.11: Illustration of a single helical axis (SHAx) state in MST. The core flux surface geometry has a helical, rather than axisymmetric, symmetry.

librium [54] referred to as a *single helical axis* (SHAx), illustrated in Figure 1.11. This is in contrast to the standard RFP operating regime in which multiple core-resonant tearing modes have significant amplitude, now called *multiple helicity* (MH). Initially discovered on the Reversed-Field eXperiment (RFX), QSH states can be reliably produced in the MST as well [55].

The physical mechanisms behind the formation and sustainment of the QSH state is still an active area of research. Long-lived, saturated QSH states can be reliably formed in the MST by altering the q profile such that the toroidal magnetic field is exactly zero at the conducting shell ($F = 0$). It is believed that removing the $m = 0$ rational surface from the plasma volume significantly reduces the three-wave coupling between adjacent core-resonant $(1, n)$ and $(1, n+1)$ modes with the $(0, 1)$ mode. This coupling is known to play an important role in the transfer of energy [56, 57, 58, 59]. However, QSH states have also been observed in reversed ($F < 0$) plasmas in both the MST and the RFX. A theory by Terry *et al.* proposes that shear in the magnetic field and/or flow velocity field suppresses the nonlinear coupling between core-resonant modes, thereby reducing energy transfer and sustaining the QSH state [60]. This topic is explored in greater detail in Chapter 6.

During QSH a region of closed flux surfaces is restored within the plasma core, thus reducing stochastic transport. Measurements on the RFX have observed a substantial increase of the electron temperature in the observed helical structure within SHAx plasmas relative to MH plasmas, with sharp gradients indicating the presence of a thermal transport barrier [61]. This indicates that QSH might provide a promising paradigm for a future RFP reactor design [62] which does not suffer the drawbacks encountered with transient current profile control. However, recent experiments indicate that, although thermal ions are well-confined, fast ions are quickly lost from the plasma [63]. Runaway electrons are also known to form within the helical structure, though the resulting HXR emission is not well-localized to the core [64].

The plasma discharges analyzed throughout this thesis will make use of either PPCD or QSH ($F = 0$) improved confinement scenarios. PPCD produces plasmas which are high-temperature and axisymmetric, making them an ideal source of x-rays for the purposes of diagnostic development. PPCD plasmas have also been well-characterized in the literature, making them useful for benchmarking. Non-reversed QSH plasmas are less well-understood, but they do feature long periods with no sawtooth cycle. This permits detectors with low signal (like the ME-SXR detector this thesis focuses on) to integrate over longer time periods without considering rapidly changing plasma dynamics.

1.5 Overview of this thesis

This thesis concerns the continued development of the ME-SXR diagnostic, its synergistic integration with existing diagnostic systems, and the plasma physics problems that can be addressed with these tools. As such, a major objective of this document is to serve as detailed documentation of my own work for future users of ME-SXR diagnostic systems in the hope that they may build upon my progress. The document is organized roughly in the order that the work was accomplished, beginning with diagnostic development before proceeding to analysis techniques, computational modeling, and finally physics results. A more detailed outline is given in the following paragraph, followed by a summary of key results.

Chapter 2 builds upon this introduction to describe the basic concepts behind x-ray plasma diagnostics and details the complementary x-ray diagnostics on the MST which are essential for this work. Chapter 3 discusses the calibration, installation, and configuration of the ME-SXR diagnostic based on the PILATUS3. Chapter 4 describes the development of a quantitative forward model based on the calibrations and underlying physics. Chapter 5 discusses the interpretation of ME-SXR data, including reconstructing T_e profiles, characterizing saturation behavior, and incorporating the ME-SXR into a Bayesian integrated data analysis framework. Finally, Chapter 6 employs the suite of x-ray diagnostics to study the evolution and sustainment of self-organized helical states in the MST RFP.

1.5.1 Summary of key results

The main results of this document are summarized below:

1. **The ME-SXR diagnostic has been calibrated for low-energy ($\sim$ 2 keV) photon detection with a threshold variation of $\Delta E < 100$ eV.** The pixel's sensitivity is well-modeled by an S-curve with a width of $\sigma_E = 300$ eV. This procedure was also performed for the medium energy range, finding $\Delta E < 200$ eV and $\sigma_E = 550$ eV, respectively. A simple model for charge sharing, the phenomenon which occurs when the energy from an absorbed photon is split between two adjacent pixels, was validated for both threshold ranges.

2. **A comprehensive physics-based forward model has been developed for the ME-**

SXR to aid in analysis. This model incorporates information from the spatial and energy calibrations as well as atomic physics modeling from ADAS to simulate the underlying physics, geometry, and detector response. The importance of charge-exchange with the neutral hydrogen population was shown to be important when simulating MST plasmas, and the systematic uncertainty in the model was assessed. The model was quantitatively validated using simultaneous SXR tomography measurements.

3. **The ME-SXR diagnostic has been used to simultaneously extract information about temperature and density profiles for low-to-mid-Z impurities.** Methods were developed for extracting temperature profiles and impurity emission spectra from ME-SXR data. The ME-SXR diagnostic was also incorporated into an integrated data analysis framework with Thomson scattering to produce simultaneous measurements of electron temperature and ion density profiles. The accuracy of these results were further improved by incorporating additional soft x-ray diagnostic measurements. This is an important proof-of-concept for future applications of the diagnostic.

4. **Predator-prey-like dynamics have been directly observed between the dominant mode amplitude, secondary magnetic mode amplitude, and thermal structure of QSH plasmas.** These measurements are consistent with predictions made by a theoretical model of the QSH state proposed by Terry, *et al.* [60] which proposes that shear in the magnetic or flow velocity fields is the mechanism responsible for suppressing the transfer of energy between the dominant and secondary tearing modes. The soft x-ray emissivity profile was seen to oscillate in phase with the dominant mode, corresponding to the behavior of a thermal transport barrier in the model.

5. **Time-resolved 2D electron temperature structure and dynamics have been ob-**

served for the first time in a helical MST plasma. Measurements were made by incorporating soft x-ray tomography and FIR interferometry measurements into an integrated framework based on diagnostic forward models. Temperatures up to $T_e = 700$ eV were observed, well above the normal temperature for similar standard RFP plasmas. Steep temperature gradients up to $|\nabla T_e| \geq 3$ keV m^{-1} were observed, suggesting the presence of a transport barrier. High frequency magnetic fluctuations in the range of 400-800 kHz can be observed only when the helical structure is oriented towards the sensing coil. The steep temperature gradients are suggestive of microtearing modes, but more work will be needed to verify this hypothesis.

6. **Significantly improved confinement is observed during a brief self-organized "quiet phase" at the start of the QSH flattop.** The QSH flattop, the period in which the dominant mode saturates at large amplitude, is divided into two distinct phases. During the "quiet" phase tearing mode activity is almost entirely absent, while small-amplitude tearing fluctuations suddenly resume during the "dynamic" phase. Runaway electrons are well-confined during the quiet phase, but are rapidly lost at the transition to the dynamic phase. A broad hot temperature structure is seen to form during the quiet phase, which rapidly collapses into a hot helical core during the dynamic phase.

Bibliography

[1] M. Arnould and K. Takahashi, "Nuclear astrophysics," *Reports on Progress in Physics*, vol. 62, pp. 395–464, 1999. [Online]. Available: https://doi.org/10.1088/0034-4885/62/3/003

[2] Energy Information Administration, "Monthly Energy Review – May 2020," US Department of Energy, Tech. Rep., 2020. [Online]. Available: www.eia.gov/mer

[3] P. C. Jain, "Greenhouse effect and climate change: scientific basis and overview," *Renewable Energy*, vol. 3, no. 4-5, pp. 403–420, 1993. [Online]. Available: https://doi.org/10.1016/0960-1481(93)90108-S

[4] A. J. Thorpe, "Climate Change Prediction: A challenging scientific problem," Tech. Rep., 2005. [Online]. Available: http://iop.cld.iop.org/publications/iop/archive/page{_}52088.html

[5] D. Wuebbles, D. Fahey, K. Hibbard, D. Dokken, B. Stewart, and T. Maycock, "Climate science special report: Fourth national climate assessment, volume I," *U.S. Global Change Research Program*, vol. 1, p. 470, 2017. [Online]. Available: https://www.globalchange.gov

[6] W. Steffen, J. Rockström, K. Richardson, T. M. Lenton, C. Folke, D. Liverman, C. P. Summerhayes, A. D. Barnosky, S. E. Cornell, M. Crucifix, J. F. Donges, I. Fetzer, S. J. Lade, M. Scheffer, R. Winkelmann, and H. J. Schellnhuber, "Trajectories of the Earth System in the Anthropocene," *Proceedings of the National Academy of Sciences*, vol. 115, no. 33, pp. 8252–8259, 2018. [Online]. Available: www.pnas.org/lookup/suppl/doi:10.1073/pnas.1810141115/-/DCSupplemental.

[7] OECD, "The Costs of Decarbonisation," Nuclear Energy Agency, OECD, Tech. Rep., 2019. [Online]. Available: https://doi.org/10.1787/9789264312180-en

[8] H. Ritchie, "What are the safest sources of energy?" 2010. [Online]. Available: https://ourworldindata.org/safest-sources-of-energy

[9] S. Wheatley, B. K. Sovacool, and D. Sornette, "Reassessing the safety of nuclear power," *Energy Research and Social Science*, vol. 15, pp. 96–100, 2016. [Online]. Available: http://dx.doi.org/10.1016/j.erss.2015.12.026

[10] K. Gi, F. Sano, K. Akimoto, R. Hiwatari, and K. Tobita, "Potential contribution of fusion power generation to low-carbon development under the Paris Agreement and associated uncertainties," *Energy Strategy Reviews*, vol. 27, no. November 2019, p. 100432, 2020. [Online]. Available: https://doi.org/10.1016/j.esr.2019.100432

[11] V. Glukhikh, O. Filatov, and B. Kolbasov, *Fundamentals of Magnetic Thermonuclear Reactor Design*, 1st ed. Elsevier Ltd, 2018.

[12] K. S. Krane, *Introductory Nuclear Physics*, 3rd ed. New York, NY: Wiley, 1988.

[13] V. W. Sidel and B. S. Levy, "Proliferation of Nuclear Weapons: Opportunities for Control and Abolition," *American Journal of Public Health Sidel and Levy | Peer Reviewed | Weapons of Mass Destruction |*, vol. 97, no. 9, p. 1594, 2007. [Online]. Available: https://dx.doi.org/10.2105{%}2FAJPH.2006.100602

[14] R. Hazeltine and J. Meiss, *Plasma Confinement*, 2nd ed. Mineola, NY: Dover Publications, Inc., 2003.

[15] Lawrence Livermore National Laboratory, "How ICF Works." [Online]. Available: https://lasers.llnl.gov/science/icf/how-icf-works

[16] J. Wessen, *Tokamaks*, 4th ed. Oxford, England: Oxford University Press, 2011.

[17] E. A. Azizov, "Tokamaks: from A D Sakharov to the present (the 60-year history of tokamaks)," *Physics-Uspekhi*, vol. 55, no. 2, pp. 190–203, 2012. [Online]. Available: https://doi.org/10.3367/UFNe.0182.201202j.0202

[18] N. Holtkamp and the ITER Project Team, "An overview of the ITER project," *Fusion Engineering and Design*, vol. 82, pp. 427–434, 2007. [Online]. Available: https://doi.org/10.1016/j.fusengdes.2007.03.029

[19] ITER Organization, "What is ITER?" [Online]. Available: https://www.iter.org/proj/inafewlines

[20] G. Federici, W. Biel, M.R. Gilbert, R. Kemp, N. Taylor, and R. Wenninger, "European DEMO design strategy and consequences for materials," *Nuclear Fusion*, vol. 57, no. 092002, pp. 1–26, 2017. [Online]. Available: https://doi.org/10.1088/1741-4326/57/9/092002

[21] B. N. Sorbom, J. Ball, T. R. Palmer, F. J. Mangiarotti, J. M. Sierchio, P. Bonoli, C. Kasten, D. A. Sutherland, H. S. Barnard, C. B. Haakonsen, J. Goh, C. Sung, and D. G. Whyte, "ARC: A compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets," *Fusion Engineering and Design*, vol. 100, pp. 378–405, 2015. [Online]. Available: http://dx.doi.org/10.1016/j.fusengdes.2015.07.008

[22] T. Klinger, T. Adreeva, and The Wendelstein 7-X team, "Overview of first Wendelstein 7-X high-performance operation Recent citations," *Nuclear Fusion*, vol. 59, no. 112004, pp. 1–11, 2019. [Online]. Available: https://doi.org/10.1088/1741-4326/ab03a7

[23] A. H. Boozer, "Why carbon dioxide makes stellarators so important," *Nuclear Fusion*, vol. 60, no. 065001, pp. 1–16, 2020. [Online]. Available: https://doi.org/10.1088/1741-4326/ab87af

[24] I. H. Hutchinson, *Principles of Plasma Diagnostics*, 2nd ed. Cambridge University Press, jul 2002.

[25] L. M. Reusch, M. D. Nornberg, J. A. Goetz, and D. J. Den Hartog, "Using integrated data analysis to extend measurement capability (invited)," *Review of Scientific Instruments*, vol. 89, no. 10, 2018. [Online]. Available: https://doi.org/10.1063/1.5039349

[26] R. Fischer, C. J. Fuchs, B. Kurzan, W. Suttrop, and E. Wolfrum, "Integrated data analysis of profile diagnostics at ASDEX upgrade," *Fusion Science and Technology*, vol. 58, no. 2, pp. 675–684, 2010. [Online]. Available: https://doi.org/10.13182/FST10-110

[27] M. Galante, L. Reusch, D. Den Hartog, P. Franz, J. Johnson, M. McGarry, M. Nornberg, and H. Stephens, "Determination of Z_eff by integrating measurements from x-ray tomography and charge exchange recombination spectroscopy," *Nuclear Fusion*, vol. 55, no. 12, p. 123016, 2015. [Online]. Available: https://doi.org/10.1088/0029-5515/55/12/123016

[28] A. J. H. Donné, A. E. Costley, and A. W. Morris, "Diagnostics for plasma control on DEMO: challenges of implementation," *Nuclear Fusion*, vol. 52, no. 074015, pp. 1–7, 2012. [Online]. Available: http://dx.doi.org/10.1088/0029-5515/52/7/074015

[29] L. Delgado-Aparicio, M. Greenwald, N. Pablant, K. Hill, M. Bitter, J. E. Rice, R. Granetz, A. Hubbard, E. Marmar, K. Tritz, D. Stutman, B. Stratton, and P. Efthimion, "Multi-energy SXR cameras for magnetically confined fusion plasmas (invited)," *Review of Scientific Instruments*, vol. 87, no. 11E204, 2016. [Online]. Available: http://dx.doi.org/10.1063/1.4964807

[30] N. Pablant, L. Delgado-Aparicio, M. Bitter, E. Brandstetter, R. Ellis, K. Hill, P. Hofer, and M. Schneebeli, "Novel energy resolving x-ray pinhole camera on Alcator C-Mod," *Review of Scientific Instruments*, vol. 83, no. 10E526, 2012. [Online]. Available: http://dx.doi.org/10.1063/1.4732177

[31] DECTRIS Ltd., "PILATUS3." [Online]. Available: https://www.dectris.com/products/pilatus3/overview/

[32] T. Barbui, N. Pablant, C. Disch, B. Luethi, N. Pilet, B. Stratton, and P. VanMeter, "Multi-energy calibration of a PILATUS3 CdTe detector for hard x-ray measurements of magnetically confined fusion plasmas (forthcoming)," in *Proceedings of the 23rd Topical Conference on High-Temperature Plasma Diagnostics*. Santa Fe, NM: American Institute of Physics, 2021.

[33] C. Bourdelle, J. Artaud, V. Basiuk, M. Bécoulet, S. Brémond, J. Bucalossi, H. Bufferand, G. Ciraolo, L. Colas, Y. Corre, X. Courtois, J. Decker, L. Delpech, P. Devynck, G. Dif-Pradalier, R. Doerner, D. Douai, R. Dumont, A. Ekedahl, N. Fedorczak, C. Fenzi, M. Firdaouss, J. Garcia, P. Ghendrih, C. Gil, G. Giruzzi, M. Goniche, C. Grisolia, A. Grosman, D. Guilhem, R. Guirlet, J. Gunn, P. Hennequin, J. Hillairet, T. Hoang, F. Imbeaux, I. Ivanova-Stanik, E. Joffrin, A. Kallenbach, J. Linke, T. Loarer, P. Lotte, P. Maget, Y. Marandet, M. Mayoral,

O. Meyer, M. Missirlian, P. Mollard, P. Monier-Garbet, P. Moreau, E. Nardon, B. Pégourí, Y. Peysson, R. Sabot, F. Saint-Laurent, M. Schneider, J. TravèreTrav, E. Tsitrone, S. Vartanian, L. Vermare, M. Yoshida, R. Zagorski, and JET Contributors, "WEST Physics Basis," *Nuclear Fusion*, vol. 55, no. 063017, 2015. [Online]. Available: http://iopscience.iop.org/0029-5515/

[34] L. Delgado-Aparicio, "Active Impurity Control For Maximum Fusion Performance (Early Career Award)," Princeton Plasma Physics Laboratory, Princeton, NJ, Tech. Rep., 2014.

[35] L. Delgado-Aparicio, D. Stutman, K. Tritz, M. Kinenthal, R. Bell, D. Gates, R. Kaita, B. LeBlanc, R. Maingi, H. Yuh, F. Levinton, and W. Heidbrink, "A 'multi-colour' SXR diagnostic for time and space-resolved measurements of electron temperature, MHD activity and particle transport in MCF plasmas," *Plasma Physics and Controlled Fusion*, vol. 49, p. 1257, 2007. [Online]. Available: stacks.iop.org/PPCF/49/1245

[36] L. F. Delgado-Aparicio, K. W. Hill, M. Bitter, B. Stratton, D. Johnson, R. Feder, N. Pablant, J. Klabacha, M. Zarnstorff, and P. Efthimion, "Burning-plasma diagnostics: Photon and particle detector development needs (PPPL-5388)," Princeton Plasma Physics Laboratory, Tech. Rep., 2017. [Online]. Available: https://bp-pub.pppl.gov/pub{_}report/2017/PPPL-5388Report.pdf

[37] H. Bodin and A. Newton, "Reversed-field-pinch research," *Nuclear Fusion*, vol. 20, no. 10, pp. 1255–1324, 1980. [Online]. Available: doi.org/10.1088/0029-5515/20/10/006

[38] A. Gibson, H. Coxell, B. A. Powell, and G. W. Reid, "Plasma confinement during a period of reduced fluctuations in 'ZETA'," Tech. Rep., 1967.

[39] J. B. Taylor, "Relaxation of Toroidal Plasma and Generation of Reverse Magnetic Fields," *Physical Review Letters*, vol. 33, no. 19, pp. 1139–1141, nov 1974. [Online]. Available: https://doi.org/10.1103/PhysRevLett.33.1139

[40] A. H. Boozer, "Physics of magnetically confined plasmas," *Reviews of Modern Physics*, vol. 76, no. 4, pp. 1071–1141, 2004. [Online]. Available: https://doi.org/10.1103/RevModPhys.76.1071

[41] R. N. Dexter, D. W. Kerst, T. W. Lovell, S. C. Prager, and J. C. Sprott, "The Madison Symmetric Torus," *Fusion Technology*, vol. 19, no. 1, pp. 131–139, 1991. [Online]. Available: https://doi.org/10.13182/FST91-A29322

[42] S. D. Terry, D. L. Brower, W. X. Ding, J. K. Anderson, T. M. Biewer, B. E. Chapman, D. Craig, C. B. Forest, R. O'connell, S. C. Prager, and J. S. Sarff, "Measurement of current profile dynamics in the Madison Symmetric Torus," *Physics of Plasmas*, vol. 11, no. 4, pp. 1079–1086, 2004. [Online]. Available: https://doi.org/10.1063/1.1643917

[43] H. D. Stephens, D. J. Den Hartog, C. C. Hegna, and J. A. Reusch, "Electron thermal transport within magnetic islands in the reversed-field pinch," *Physics of Plasmas*, vol. 17, no. 056115, pp. 1–10, 2010. [Online]. Available: https://doi.org/10.1063/1.3388374

[44] M. B. McGarry, "Probing the relationship between magnetic and temperature structures with soft x-rays on the Madison Symmetric Torus," Ph.D. dissertation, University of Wisconsin-Madison, 2013.

[45] T. M. Biewer, C. B. Forest, J. K. Anderson, G. Fiksel, B. Hudson, S. C. Prager, J. S. Sarff, J. C. Wright, D. L. Brower, W. X. Ding, and S. D. Terry, "Electron Heat Transport Measured in a Stochastic Magnetic Field," *Physical Review Letters*, vol. 91, no. 4, pp. 1–4, 2003. [Online]. Available: https://doi.org/10.1103/PhysRevLett.91.045004

[46] L. A. Morton, W. C. Young, C. C. Hegna, E. Parke, J. A. Reusch, and D. J. Den Hartog, "Electron thermal confinement in a partially stochastic magnetic structure," *Physics of Plasmas*, vol. 25, no. 4, 2018. [Online]. Available: https://doi.org/10.1063/1.5021893

[47] L. A. Morton, "Turbulence and transport in magnetic islands in MST and DIII-D," Ph.D. dissertation, University of Wisconsin-Madison, 2016.

[48] J. Sarff, S. Hokin, H. Ji, S. Prager, and C. Sovinec, "Fluctuation and Transport Reduction in a Reversed Field Pinch by Inductive Poloidal Current Drive," *Physical Review Letters*, vol. 72, no. 23, 1994. [Online]. Available: https://doi.org/10.1103/PhysRevLett.72.3670

[49] B. E. Chapman, A. F. Almagri, J. K. Anderson, D. L. Brower, K. J. Caspary, D. J. Clayton, D. Craig, D. J. Den Hartog, W. X. Ding, D. A. Ennis, G. Fiksel, S. Gangadhara, S. Kumar, R. M. Magee, R. O'Connell, E. Parke, S. C. Prager, J. A. Reusch, J. S. Sarff, H. D. Stephens, and Y. M. Yang, "Generation and confinement of hot ions and electrons in a reversed-field pinch plasma," *Plasma Physics and Controlled Fusion*, vol. 52, no. 12, 2010. [Online]. Available: https://doi.org/10.1088/0741-3335/52/12/124048

[50] R. O'Connell, D. J. Den Hartog, C. B. Forest, and R. W. Harvey, "Measurement of fast electron distribution using a flexible, high time resolution hard x-ray spectrometer," *Review of Scientific Instruments*, vol. 74, no. 3 II, pp. 2001–2003, 2003. [Online]. Available: https://doi.org/10.1063/1.1535244

[51] K. J. McCollam, J. K. Anderson, A. P. Blair, D. Craig, D. J. Den Hartog, J. Homepage, F. Ebrahimi, J. A. Reusch, J. S. Sarff, H. D. Stephens, D. R. Stone, D. L. Brower, B. H. Deng, and W. X. Ding, "Equilibrium evolution in oscillating-field current-drive experiments," *Physics of Plasmas*, vol. 17, no. 82506, pp. 1–13, 2010. [Online]. Available: https://doi.org/10.1063/1.3461167

[52] Z. Li, K. J. Mccollam, T. Nishizawa, E. Parke, J. S. Sarff, Z. A. Xing, H. Li, W. Liu, and W. Ding, "Effects of oscillating poloidal current drive on magnetic relaxation in the Madison Symmetric Torus reversed-field pinch Effects of oscillating poloidal current drive on magnetic relaxation in the Madison Symmetric Torus reversed-field pinch," *Plasma Physics and Controlled Fusion*, vol. 61, no. 045004, pp. 1–14, 2019. [Online]. Available: https://doi.org/10.1088/1361-6587/aaf9e0

[53] D. F. Escande, P. Martin, S. Ortolani, A. Buffa, P. Franz, L. Marrelli, E. Martines, G. Spizzo, S. Cappello, A. Murari, R. Pasqualotto, and P. Zanca, "Quasi-single-helicity reversed-field-pinch plasmas," *Physical Review Letters*, vol. 85, no. 8, pp. 1662–1665, 2000. [Online]. Available: https://doi.org/10.1103/PhysRevLett.85.1662

[54] W. F. Bergerson, F. Auriemma, B. E. Chapman, W. X. Ding, P. Zanca, D. L. Brower, P. Innocente, L. Lin, R. Lorenzini, E. Martines, B. Momo, J. S. Sarff, and D. Terranova, "Bifurcation to 3D Helical Magnetic Equilibrium in an Axisymmetric Toroidal Device," *Physical Review Letters*, vol. 107, no. 255001, pp. 1–5, dec 2011. [Online]. Available: https://doi.org/10.1103/PhysRevLett.107.255001

[55] L. Marrelli, P. Martin, G. Spizzo, P. Franz, B. E. Chapman, D. Craig, J. S. Sarff, T. M. Biewer, S. C. Prager, J. C. Reardon, M. Symmetric, T. R. N. Dexter, D. W. Kerst, T. W. Lovell, and J. C. Sprott, "Quasi-single helicity spectra in the Madison Symmetric Torus," *Physics of Plasmas*, vol. 9, no. 7, pp. 2868–2871, aug 2002. [Online]. Available: https://doi.org/10.1063/1.1482766

[56] J. A. Holmes, B. A. Carreras, P. H. Diamond, and V. E. Lynch, "Nonlinear dynamics of tearing modes in the reversed field pinch," *Citation: The Physics of Fluids*, vol. 31, p. 1166, 1988. [Online]. Available: https://doi.org/10.1063/1.866746

[57] Y. L. Ho and G. G. Craddock, "Nonlinear dynamics of field maintenance and quasiperiodic relaxation in reversed-field pinches ARTICLES YOU MAY BE INTERESTED IN," *Physics of Fluids B: Plasma Physics*, vol. 3, p. 721, 1991. [Online]. Available: https://doi.org/10.1063/1.859868

[58] J. S. Sarff, S. Assadi, A. F. Almagri, M. Cekic, D. J. Den Hat-tog, G. Fiksel, S. A. Hokin, H. Ji, S. C. Prager, W. Shen, K. L. Sidikman, and M. R. Stoneking, "Nonlinear coupling of tearing fluctuations in the Madison Symmetric Torus," *Physics of Fluids B*, vol. 5, no. 7, pp. 2540–2545, 1993. [Online]. Available: https://doi.org/10.1063/1.860741

[59] A. K. Hansen, A. F. Almagri, D. Craig, D. J. Den Hartog, C. C. Hegna, S. C. Prager, and J. S. Sarff, "Momentum Transport from Nonlinear Mode Coupling of Magnetic Fluctuations," *Physical Review Letters*, vol. 85, no. 16, pp. 3408–3411, 2000. [Online]. Available: https://doi.org/10.1103/PhysRevLett.85.3408

[60] P. W. Terry and G. G. Whelan, "Time-dependent behavior in a transport-barrier model for the quasi-single helcity state," *Plasma Physics and Controlled Fusion*, vol. 56, no. 9, 2014. [Online]. Available: https://doi.org/10.1088/0741-3335/56/9/094002

[61] P. Franz, M. Gobbin, L. Marrelli, A. Ruzzon, A. Fassina, E. Martines, and G. Spizzo, "Experimental investigation of electron temperature dynamics of helical states in the RFX-Mod reversed field pinch," *Nuclear Fusion*, vol. 53, no. 5, p. 053011, 2013. [Online]. Available: https://doi.org/10.1088/0029-5515/53/5/053011

[62] R. Lorenzini, E. Martines, P. Piovesan, D. Terranova, P. Zanca, M. Zuin, A. Alfier, D. Bonfiglio, F. Bonomo, A. Canton, S. Cappello, L. Carraro, R. Cavazzana, D. F. Escande, A. Fassina, P. Franz, M. Gobbin, P. Innocente, L. Marrelli, R. Pasqualotto, M. E. Puiatti, M. Spolaore, M. Valisa, N. Vianello, and P. Martin, "Self-organized helical equilibria as a new paradigm for ohmically heated fusion plasmas," *Nature Physics*, vol. 5, pp. 570–574, aug 2009. [Online]. Available: https://doi.org/10.1038/nphys1308

[63] P. Bonofiglo, M. Gobbin, D. A. Spong, J. Boguski, E. Parke, J. Kim, and J. Egedal, "Fast ion transport in the quasi-single helical reversed-field pinch," *Physics of Plasmas*, vol. 022502, no. 26, 2019. [Online]. Available: http://dx.doi.org/10.1063/1.5084059

[64] D. J. Clayton, "Fast Electron Transport in Improved-Confinement RFP Plasmas," Ph.D. dissertation, University of Wisconsin-Madison, 2010.

Chapter 2

Characterizing plasmas in the MST

A well-diagnosed plasma is critical to the process of experimental plasma physics. As such, the MST is equipped with a large number of sophisticated diagnostics and analysis methodologies to help the intrepid experimentalist understand what is going on inside the aluminum shell. The goal of this chapter is to provide a "big picture" overview of the diagnostics which are referenced at various other points throughout the rest this thesis. My descriptions are intended to be brief, so references to more thorough sources are also provided.

Section 2.1 provides a general overview of the non-SXR MST diagnostics used throughout this thesis, including the magnetics arrays, FIR, Thomson scattering, and a fast hard x-ray detector. Section 2.2 provides a description of the axisymmetric and helical equilibrium reconstruction codes available for use on the MST. Finally, Section 2.3 discusses the soft x-ray tomography and NICKAL2 Ross spectrometer diagnostics. This section also covers the impact of Al^{+11} and Al^{+12} transition lines on the SXR diagnostic signals. A new multi-energy soft x-ray diagnostic is the subject of Chapters 3, 4, and 5.

2.1 Common MST diagnostics

A large number of sophisticated diagnostics have been developed for the MST during its more than thirty years of operation. Numerous specialized diagnostics and subsystems including spectrometers [1], neutral beams [2], and neutron detectors [3] have been used in scientific studies. A thorough description of the entire MST diagnostic suite is well beyond the scope of this document. Therefore, the goal of this section is to provide a basic overview of a small number of routine diagnostics whose data are used in other chapters. Specifically, we will introduce the magnetics arrays (Section 2.1.1), the Thomson scattering diagnostic (Section 2.1.2), the FIR interferometer (Section 2.1.3), and the fast x-ray camera (Section 2.1.4). References will be provided for further details on each diagnostic.

2.1.1 Magnetics arrays

As with all plasma confinement devices, it is important to be able to characterize the magnetic field configuration of the MST. This is accomplished via multiple arrays of magnetic field-sensing coils along the plasma boundary. A schematic of these arrays is shown in Figure 2.1. Edge measurements of the magnetic field are an important constraint when calculating the global magnetic configuration using an equilibrium reconstruction code (see Section 2.2). Additionally, tearing modes are global instabilities, meaning that edge measurements provide a good indication of magnetic activity in the core.

There are four total magnetics arrays inside the MST vessel, including three poloidal arrays and a toroidal array [4]. Unless otherwise noted, all magnetic field measurements presented in this thesis were taken by the toroidal array. This array consists of 64 evenly-spaced triplets of pickup coils which are used to measure the toroidal and poloidal components of the magnetic field. Typically 64 B_T, 32 B_P, and 32 $\dot{B}_P$ signals are digitized

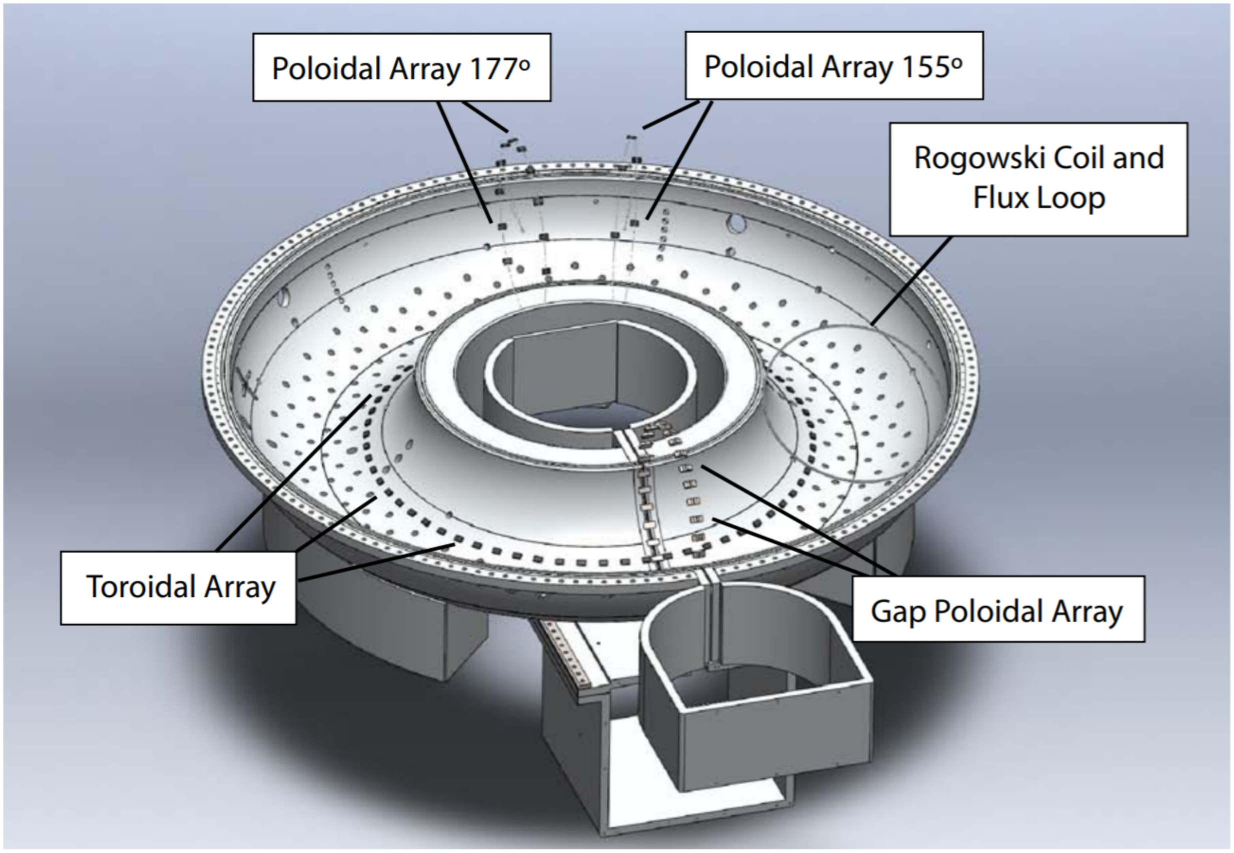

Figure 2.1: Schematic of the various magnetics arrays on the MST. Data shown in this thesis is exclusively from the toroidal array. Reproduced from J. Koliner [4].

at one time. The magnetic field is related to the voltage measured across the coil via

$$B = -\frac{1}{A}\int \varepsilon(t)dt \tag{2.1}$$

where A is the coil area and

$$\varepsilon(t) = -\frac{\partial \Phi}{\partial t} \tag{2.2}$$

is the electromotive force induced across the coil due to the changing magnetic flux $\Phi = \boldsymbol{B} \cdot A\hat{\boldsymbol{n}}$. By measuring the voltage these coils directly measure $\varepsilon \propto \dot{B}$. To convert from $\dot{B}$ to B the signals are passed through an integrator before being digitized at 200 kHz. The $\dot{B}$ signals may also be directly digitized at 3 MHz to preserve sensitivity to high-frequency fluctuations (internally referred to as "fast magnetics" measurements).

Magnetic field fluctuations are measured by the toroidal array and decomposed into Fourier modes according to

$$\tilde{b}_i(\phi, t) = \sum_{n=0}^{N} b_n \cos\left(n\phi - \delta_{i,n}(t)\right), \tag{2.3}$$

where $i = \{\theta, \phi\}$ labels the vector component, ϕ is the toroidal angle, b_n is the amplitude of the n$^{\text{th}}$ harmonic, $\delta_{i,n}$ is the associated phase, and $N = 15$ for $i = \theta$ and $N = 31$ for $i = \phi$. The magnetic field fluctuation amplitudes b_n are typically associated with the tearing mode of the same n. Because the decomposition is in the toroidal direction only, distinct m modes cannot be distinguished. However, given the q profile typical to the RFP (Figure 1.6) it is generally assumed that the core-resonant modes $n = 5, 6, \ldots$ are dominantly $m = 1$, while modes resonant at the reversal surface like $n = 1, 2, \ldots$ are $m = 0$.

The MST also has a Rogowski coil and flux loop to provide measurements on the plasma current I_p and the average toroidal field $\langle B_\phi \rangle$, respectively [4]. These measurements serve as important constraints to equilibrium reconstruction routines.

2.1.2 Thomson Scattering

Routine measurements of the electron temperature are provided by the Thomson scattering diagnostic. Pulses from a high-powered Q-switched Nd:YAG laser are directed into the vacuum vessel where they proceed to scatter off of free electrons in the plasma via the eponymous Thomson scattering process (the low-energy limit of Compton scattering). The scattered light is then collected by the viewing optics along the plasma boundary, as shown in Figure 2.2, allowing localized plasma properties to be inferred. A detailed overview of the physics of Thomson scattering as applied to high-temperature plasma diagnostics is provided by Prunty [5].

The incoherent[1] Thomson scattering process is well-understood in the framework of classical electrodynamics. The incident electromagnetic radiation causes an electron to oscillate and emit dipole radiation, with an intensity proportional to the electron density and dependent on the scattering angle. The spectrum is also broadened due to the Doppler effect, which is dependent on the electron temperature. Relativistic effects are also significant for plasma temperatures on the order of a keV or higher. The spectral density function, which was derived by Selden [7], is given by

$$S(\epsilon, \theta, \alpha) = \frac{c(\alpha)}{A(\epsilon, \theta)} \exp\left[-2\alpha B(\epsilon, \theta) \right], \tag{2.4}$$

where $\alpha = m_e c^2/(2kT_e)$, $\epsilon = \lambda_s/\lambda_i$ is the ratio of the wavelength of the scattered light to that of the incident light, θ is the scattering angle, and

$$A(\epsilon, \theta) = (1+\epsilon)^3 \sqrt{2(1-\cos\theta)(1+\epsilon)+\epsilon^2} \tag{2.5}$$

$$B(\epsilon, \theta) = \sqrt{1+\epsilon^2/[2(1-\cos\theta)(1+\epsilon)} - 1 \tag{2.6}$$

$$C(\alpha) = \sqrt{\frac{\alpha}{\pi}} \left(1 - \frac{15}{16}\alpha^{-1} + \frac{345}{512}\alpha^{-2} + \ldots \right). \tag{2.7}$$

Measurements of the scattered light are taken by the collection optics, which are localized to the intersection of the beam line and the viewing chords with known scattering angle θ. The measurements can then be fit to Equation 2.4 to determine the electron temperature. In principle density measurements can be extracted from the absolute magnitude of the scattered light, but the MST Thomson system is not currently calibrated for that.

The Thomson scattering diagnostic on MST is composed of two Nd:YAG lasers firing

[1]This means that the plasma is sufficiently diffuse that the electromagnetic waves can be considered to scatter off of individual electrons rather than collectively interact with the plasma.

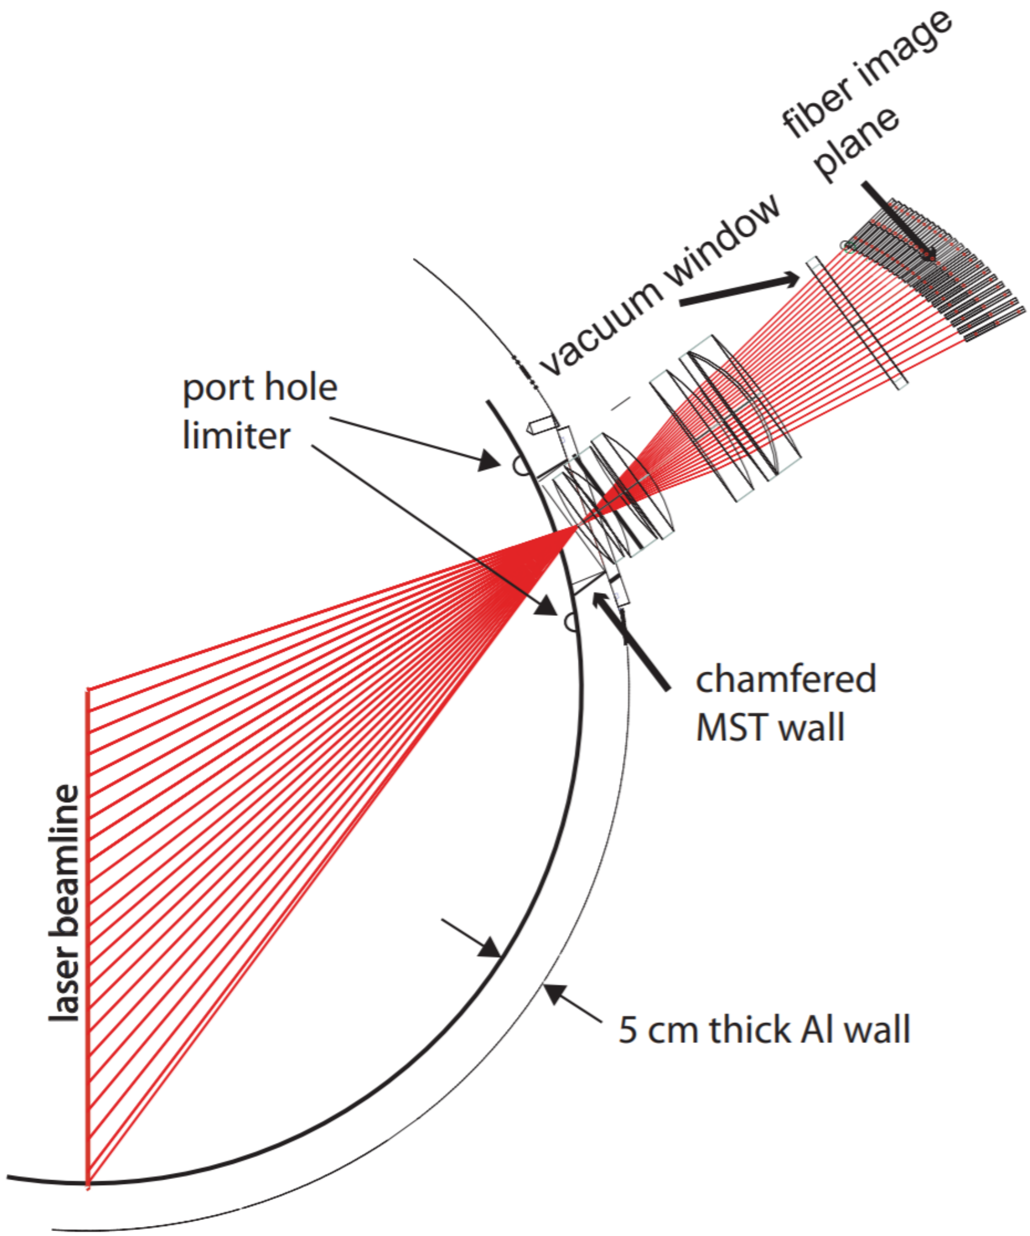

Figure 2.2: Schematic of the Thomson scattering beam-line and collection optics. Electron temperature is measured locally at the points where the beam-line and measurement chords intersect. Reproduced from J. Reusch [6].

at 1 kHz whose timing is staggered to achieve an overall data collection rate of 2 kHz over a 15 ms time window. The diagnostic can also be operated in a pulse-burst mode with resolution of 25 kHz, repeated in 1 kHz bursts [8]. A custom "Fast Thomson" system has also been used in the past to achieve a temporal resolution of up to 250 kHz [9]. However this mode was no longer available during the period in which the research in this thesis was conducted.

2.1.3 FIR interferometry

The far infrared (FIR) interferometry diagnostic on the MST works by measuring the relative phase shift between two $\sim$ 650 GHz beams, a probe beam which passes through the plasma and a reference beam which travels the same distance through the air [10]. The magnitude of the phase shift is a function of the line-integrated plasma index of refraction, itself a function of the electron density,

$$\Delta\phi = \frac{\lambda e^2}{4\pi m_e c^2 \epsilon_0} \int n_e(z)dz \tag{2.8}$$

$$= 2.815 \times 10^{-15} \lambda \int n_e(z)dz, \tag{2.9}$$

where λ is the wavelength of the probe beam and z is the coordinate along a given chord. So by directly measuring the relative shift $\Delta\phi$, the average density $\bar{n}_e = \int n_e(z)dz$ is obtained. Inversion techniques may then be used to obtain the n_e profile.

As shown in Figure 2.3, the FIR system installed on the MST features 11 chords across the plasma volume as well as an external reference channel. The channels are divided into two sets separated toroidally by five degrees (at 250° and 255°). The FIR beams were originally produced by a CO_2 pumped formic acid molecular gas laser. In 2019 this was upgraded to a set of three solid-state FIR diodes, similar to those described in Xie, *et al.* [11].

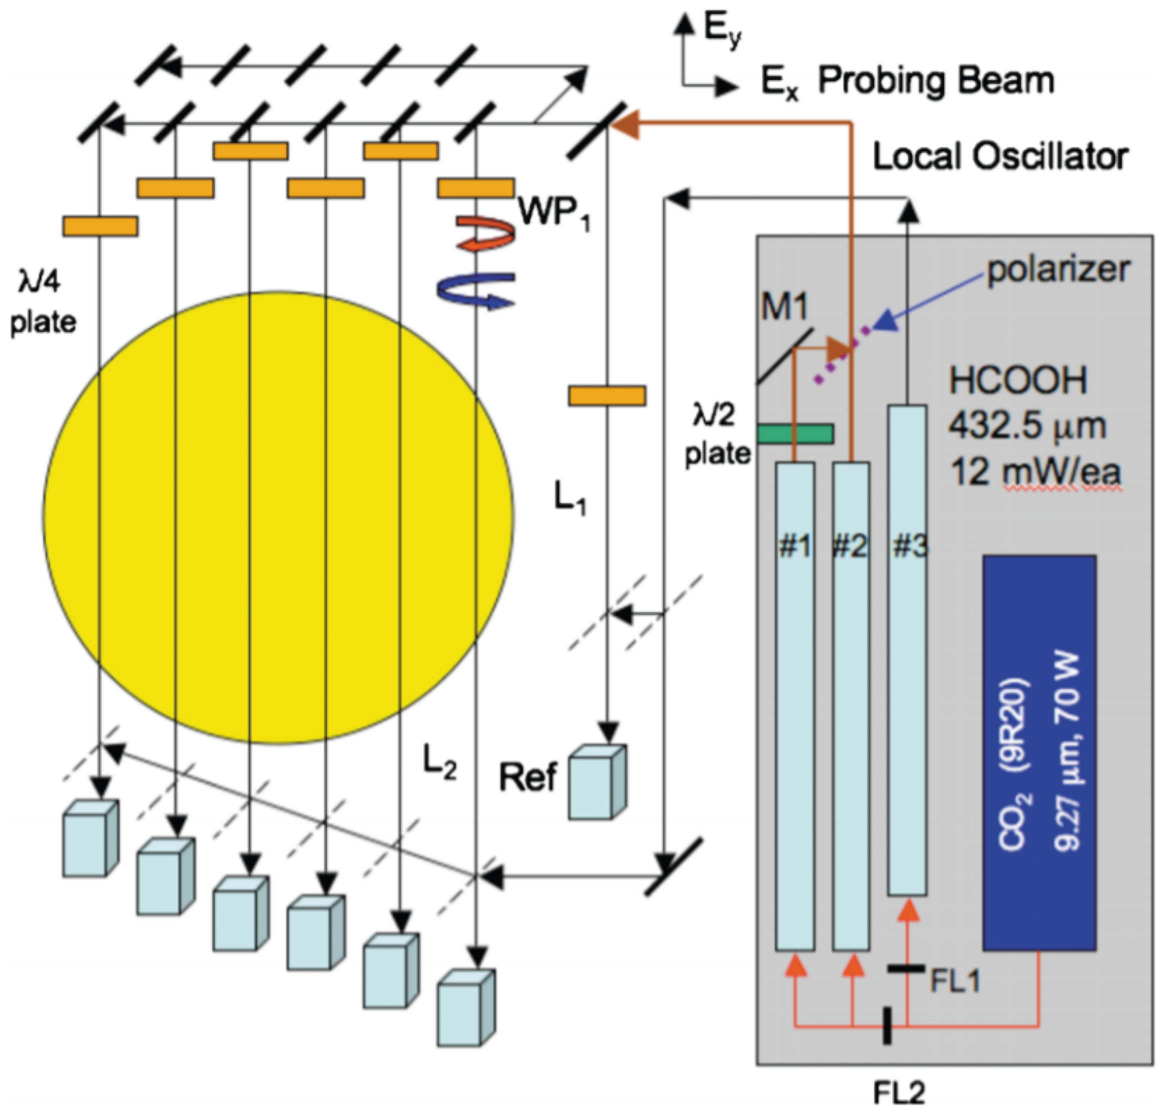

Figure 2.3: Schematic overview of the FIR interferometry system installed on the MST. Note that in 2019 the FIR source was changed from a CO_2 pumped formic acid molecular gas laser (shown in this image) to solid-state diodes. Reproduced from E. Parke, *et al.* [10].

The FIR diagnostic can also be operated as a polarimeter to make measurements of the line-integrated magnetic field [12]. This can be done simultaneously to n_e measurements by operating the diagnostic with three beams. Polarimetry data is not used in this thesis, so a more thorough discussion of this technique is deferred to other sources [13].

2.1.4 The fast x-ray camera

A sixteen-chord hard x-ray spectrometer array was installed in 2003 [14] to measure the radiation produced by fast electrons in the RFP. This capability was enhanced in 2015 by the addition of a single-chord fast x-ray camera which features very high time resolution with minimal dead time [15].

The fast x-ray (FXR) camera is based on a single Si avalanche photodiode (APD) which can detect photons in the range 2-30 keV. The signal output by the detector, which is proportional to the energy of the incident photon, is passed through a Gaussian amplifier with a shaping time of approximately 20 ns before being directly digitized at 500 MHz for later analysis. As a result, relatively high x-ray fluxes can be measured without worrying about significant saturation or pulse pileup. Detected pulses are fit to a characteristic pulse shape derived from an Fe-55 source during the calibration procedure. The R^2 goodness of the fit along with the FWHM of the pulse are used to distinguish real photon pulses from noise. This procedure is shown in Figure 2.4. Once a real photon has been identified, the pulse height is used to determine the photon energy.

There are several physical processes which can result in the presence of fast non-Maxwellian electrons in the RFP. These electrons emit photons in the hard x-ray range via bremsstrahlung, making the FXR camera a very useful tool for diagnosing fast electron populations. For example, the FXR camera was used to show that an anisotropic tail population tends to form during standard plasmas a result of energization during sawtooth (magnetic reconnection) events [16]. The diagnostic has also been used to measure runaway electrons in both PPCD and QSH improved confinement scenarios [15],

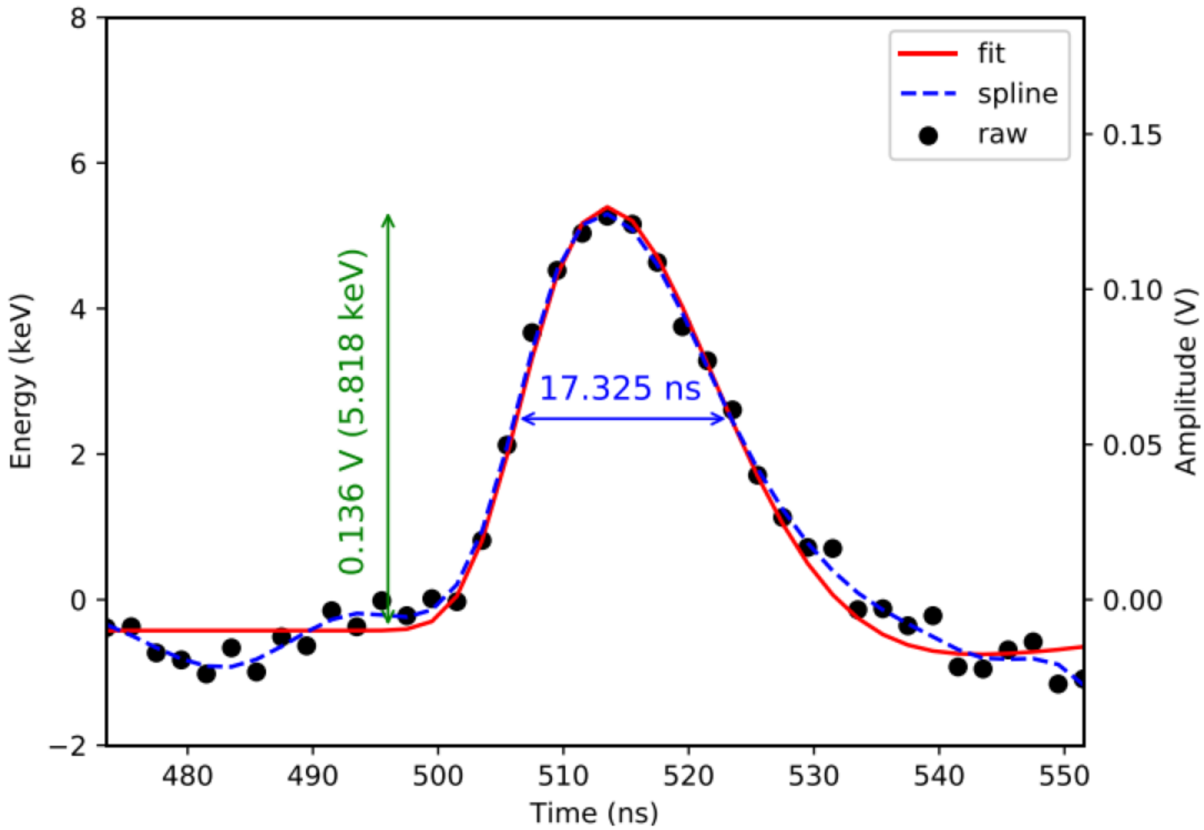

Figure 2.4: A single Fe-55 photon pulse as measured by the FXR detector. The pulse is then fit to a calibrated characteristic pulse and a spline fit to help distinguish real photons from noise and to determine the photon energy. Reproduced from Dubois, *et al.* [15].

which tend to form due to the substantially reduced stochasticity in the core [17].

Figure 2.5 shows an example of FXR data recorded during a QSH plasma discharge. Each pulse represents a single photon detected, and the height of the pulse corresponds to its energy. The pulses are also color-coded by energy in order to help distinguish narrowly-separated pulses. In the example plasma, a large number of high-energy photons are recorded as the magnetic mode grows to large amplitude and begins to saturate. The particular dynamics of fast electrons in the helical RFP is the topic of Section 6.3.

2.2 Equilibrium reconstruction codes

The diagnostics discussed in the previous section each directly measure a plasma's properties for only a small subset of the overall plasma volume (the toroidal array measures B at 32 spots along the boundary, FIR measures $\bar{n}_e$ for its 11 chords, etc.). However, when doing plasma physics it is frequently important to extrapolate these finite measurements into continuous global profiles. This is the purpose of equilibrium reconstruction codes.

An equilibrium reconstruction code solves an equilibrium equation (such as the Grad-Shafranov equation) subject to the constraints of diagnostic measurements. Measurements may be used directly as boundary conditions, or can be incorporated via synthetic diagnostics and solved for iteratively until disagreement is sufficiently low. Reconstruction codes typically output full global profiles for magnetic quantities such as q, ψ, $B_{\theta/\phi}$, $J_{\theta/\phi}$, and sometimes physical parameters like T_e, n_e, etc..

Three separate reconstruction codes were used to some capacity for the work in this thesis. The first, MSTFit, is the focus of Section 2.2.1. This code is known to work well for the cases of axisymmetric RFP plasmas, such as in standard reversed and PPCD scenarios. However, plasmas with strong non-axisymmetric components, such as the helical plasmas which form in the non-reversed QSH scenario, require different tools. Two such tools, SHEq and V3Fit, are discussed in Section 2.2.2. Both helical reconstruction codes

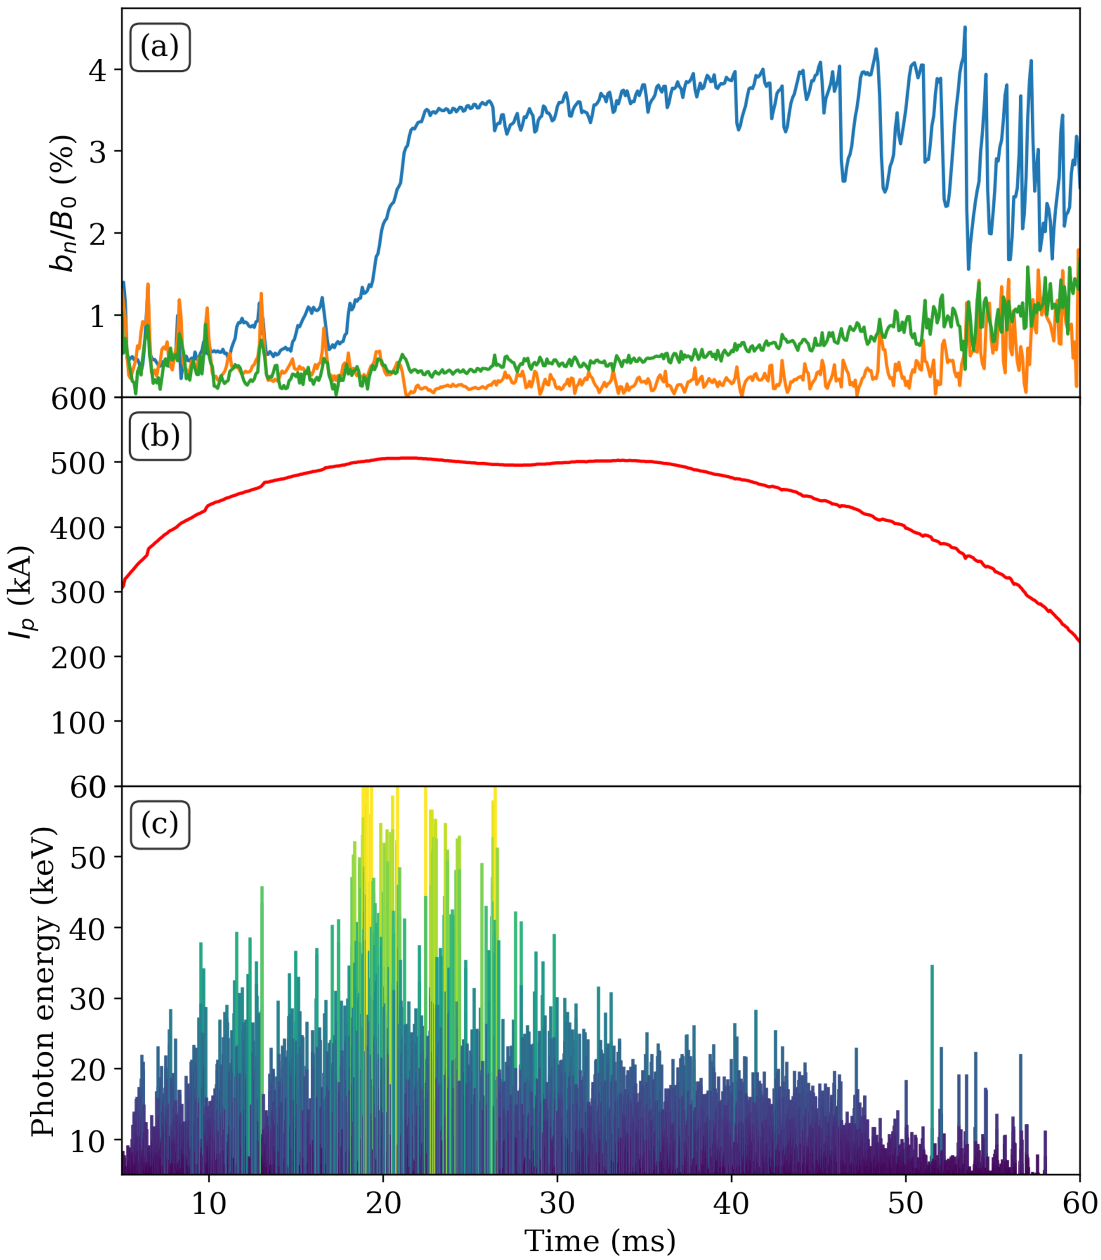

Figure 2.5: Example of photon counts taken with the FXR detector during a 500 kA non-reversed plasma with a quasi single-helicity magnetic spectrum. (a) The Fourier-decomposed magnetic mode amplitudes taken from the toroidal array at the plasma boundary; (b) The plasma current; (c) Each pulse corresponds to a photon detected by the FXR camera. The height of the pulse (and corresponding color coding) gives the energy of the detected photon.

are used in some capacity during the analysis presented in Chapter 6.

2.2.1 MSTFit

MSTFit is a non-linear Grad-Shafranov toroidal equilibrium reconstruction code developed specifically for the RFP [18]. Prior to MSTFit, equilibrium reconstructions of the MST were mostly performed using one-dimensional cylindrical models which did not capture the full geometry of a toroidal device. MSTFit allows users to obtain full 2D axisymmetric toroidal equilibria constrained by numerous diagnostic measurements. MSTFit also correctly accounts for the interaction of the plasma with the close-fitting conducting shell, a detail which is particular to the MST.

The Grad-Shafranov equation assumes an axisymmetric field of the form

$$\boldsymbol{B}(R,Z) = B_\phi(R,Z)\hat{\boldsymbol{\phi}} + \boldsymbol{\nabla}\psi \times \boldsymbol{\nabla}\phi \tag{2.10}$$

which obeys $\boldsymbol{B} \cdot \boldsymbol{\nabla}\psi = 0$. This assumption is generally a good assumption in the tokamak plasmas and for high-performance regimes of the RFP (such as in-between sawteeth or during PPCD). The Grad Shafranov equation, which MSTFit solves, is

$$\Delta^*\psi = -\mu_0 R J_\phi \tag{2.11}$$

$$J_\phi = \frac{2\pi F F'}{\mu_0 R} + 2\pi R p', \tag{2.12}$$

where $\Delta^* = R^2\boldsymbol{\nabla} \cdot (\boldsymbol{\nabla}/R^2)$ is the elliptic operator, ψ is the poloidal magnetic flux, $F = RB_\phi$, and p is the pressure. Both $F = F(\psi)$ and $p = p(\psi)$ are functions of ψ only.

The MSTFit codes solves Equation 2.11 iteratively over an unstructured mesh of 746 elements. At each step in the parameter space synthetic diagnostic measurements are computed and compared to the real data to compute χ^2. Parameters are then varied

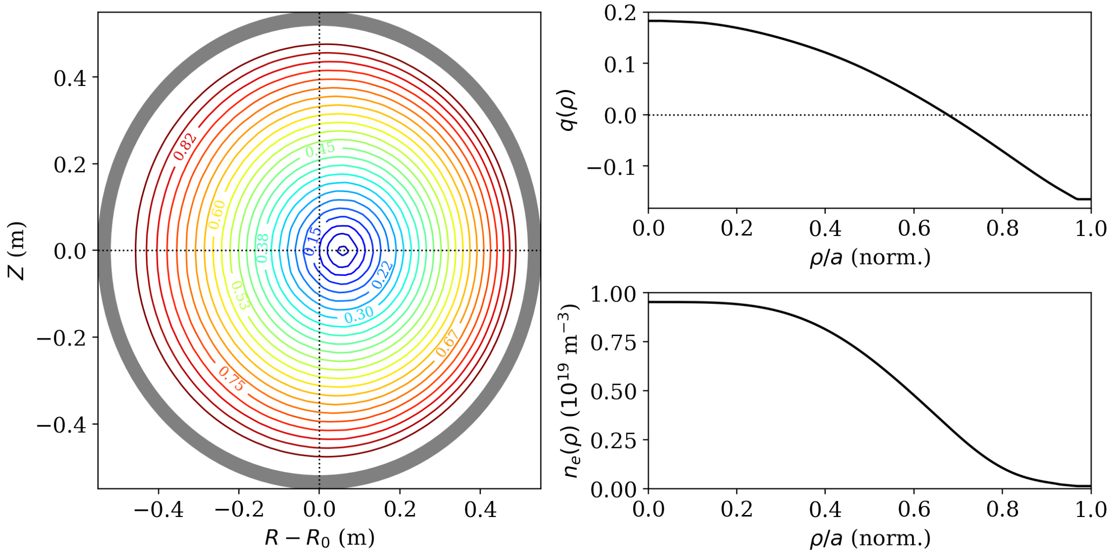

Figure 2.6: Example MSTFit output for plasma discharge #1191204050, $t = 18.5$ ms, 300 kA PPCD. Figure shows nested surfaces of constant ρ/a (left), the safety factor profile (top right), and the reconstructed density profile (bottom right).

to minimize χ^2. This process is iterated a second time using the previous solution as a starting point. The code also provides a convenient radius-like variable $\rho = \rho(\psi)$ which is determined by the volume contained within the flux surface $V(\psi)$,

$$\rho = \sqrt{\frac{V(\psi)}{2\pi^2 R}}, \tag{2.13}$$

which ranges from $\rho = 0$ to $\rho = a$. Figure 2.6 provides an example of an MSTFit reconstruction, showing contours of constant ρ/a.

2.2.2 Helical equilibria

As discussed in Section 1.4.1, high-current non-reversed MST plasmas tend to spontaneously evolve into helical equilibria. This violates MSTFit's assumption of axisymmetry, meaning that a different equilibrium reconstruction method is needed to analyze these plasmas. In Chapter 6 this is mostly accomplished via the NCT/SHEq code. Originally

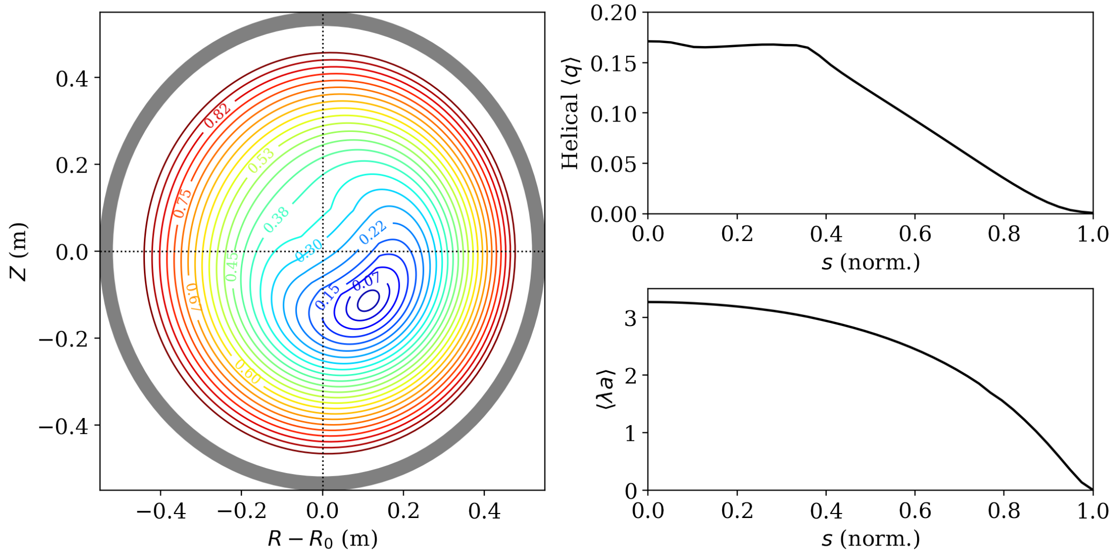

Figure 2.7: Example NCT/SHEq output for plasma discharge #1200701056, $t = 23.5$ ms, 500 kA $F = 0$. Figure shows nested surfaces of constant ρ (left), the helical safety factor profile (top right), and the normalized current density profile (bottom right).

developed for RFX-Mod [19], it was later modified and extended for use on the MST [20].

The NCT/SHEq code assumes that plasma properties such as the toroidal and poloidal magnetic flux can be written as the sum of an axisymmetric component with a helical perturbation,

$$\psi(r,\theta,\phi) = \psi_0(r) + \sum_{m,n} \psi^{m,n}(r) e^{i(m\theta - n\phi)}, \tag{2.14}$$

where r is the toroidal radius, θ is the poloidal (azimuthal) angle, and ϕ is the toroidal (axial) angle. The magnetic field is assumed to have the form

$$B(r,\theta,\phi) = \nabla\psi_T \times \nabla\theta - \nabla\psi_P \times \nabla\phi \tag{2.15}$$

where ψ_P is the poloidal flux and ψ_T is the toroidal flux.

The routine begins by solving for the axisymmetric component of the equilibrium measurements from the edge magnetics arrays as constraints. It also assumes that pressure can be ignored. It then uses NCT to solve a series of toroidal Newcomb-like equations [21, 22] to obtain the perturbations $\psi_T^{m,n}$, $\psi_P^{m,n}$ to the equilibrium. The flux surfaces can then be labeled by the helical flux

$$\chi \equiv m\psi_P - n\psi_T, \tag{2.16}$$

where $n = 5$ is the dominant mode number. The helical flux is usually presented in normalized from,

$$\rho = \sqrt{\frac{\chi - \chi_{\text{min}}}{\chi_{\text{max}} - \chi_{\text{min}}}}, \tag{2.17}$$

which is well-suited for use as a radial variable.

An alternative code, called V3Fit [23], is also available to produce a fully three-dimensional helical equilibrium reconstructions on the MST. Although mostly used with the stellarator community, the code was modified for the helical RFP and has been successfully implemented on the MST [24]. As implemented at the MST, V3Fit uses VMEC [25] as its 3D equilibrium solver. The equilibrium is then used to produce synthetic diagnostic signals, which are compared against the real diagnostic measurements to produce χ^2. The algorithm then adjusts the parameters and repeats the process until a minimum in χ^2 is located. For more details on this process see References [4, 26].

Figure 2.8 shows a comparison between flux surface reconstructions obtained using SHEq and V3Fit. In general the results are similar, though they differ somewhat in shape. Because V3Fit incorporates information from multiple diagnostics, including internal diagnostics like FIR interferometry/polarimetry and SXR tomography, it is generally considered to be better-constrained than NTC/SHEq. VMEC is also fully three-dimensional and incorporates more physics (such as pressure) than NCT/SHEq. However, it also necessitates a large suite of diagnostics when taking data and requires

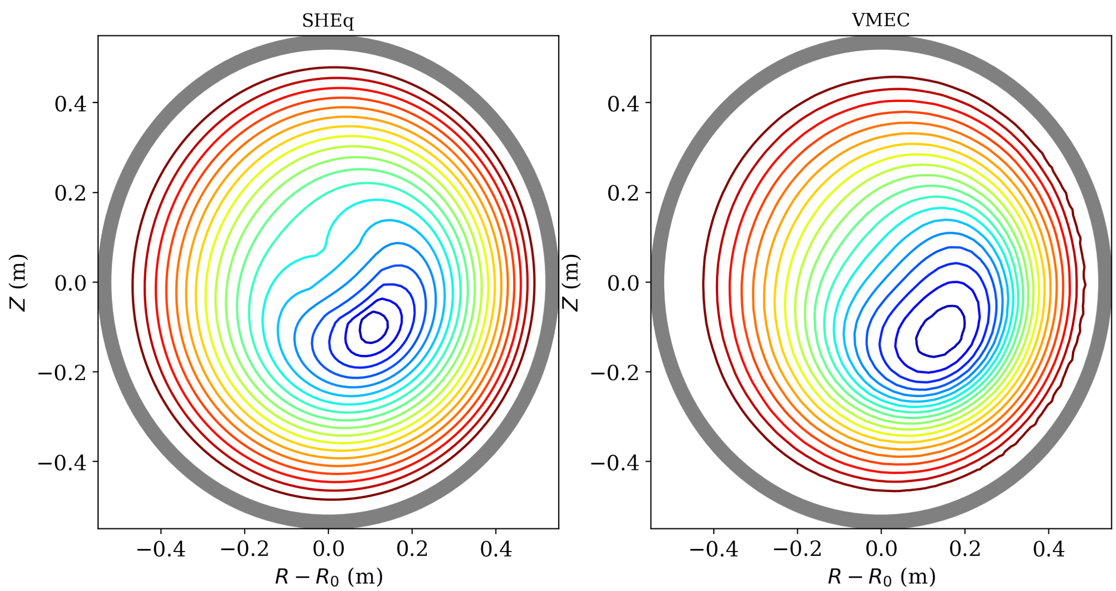

Figure 2.8: Comparison between SHEq (left) and V3Fit (right) reconstructed flux surfaces for the same locking phase.

substantial computational time to produce reconstructions. For that reason I have mostly relied on NTC/SHEq to provide flux surfaces in this thesis.

2.3 Soft x-ray diagnostics

Electromagnetic radiation passively emitted by the plasma provides an excellent opportunity for plasma diagnostics. Spectroscopic diagnostics focusing on the soft x-ray (SXR), portion of the electromagnetic spectrum ($100\,\mathrm{eV} \lesssim h\nu \lesssim 10\,\mathrm{keV}$, or equivalently $1\,\text{Å} \lesssim \lambda \lesssim 100\,\text{Å}$) have been prevalent since the 1960s, when T_e for experimental devices became sufficiently large to stimulate significant SXR emission [27]. The SXR region in particular is of interest because it occupies something of a sweet spot on the electromagnetic spectrum, with a spectrum that is much "cleaner" and easier to interpret than the vacuum ultraviolet ($h\nu \lesssim 100\,\mathrm{eV}$) but still much more intense than hard x-ray ($h\nu \gtrsim 10\,\mathrm{keV}$) emissions at thermal temperatures.

Emission from the SXR spectrum contains a wealth of information about the properties of the emitting plasma, which is reviewed in Section 2.3.1. Section 2.3.2 describes the two-color soft x-ray tomography diagnostic used throughout the rest of this thesis. The methodology used for tomographic inversions is described in Section 2.3.3. Section 2.3.4 explains the difficulties encountered by SXR diagnostics on the MST arising from the presence of mid-Z (Al) impurity transition line emissions. And finally Section 2.3.5 discusses a supplemental Ross spectrometer diagnostic that has recently been deployed to directly measure the Al lines. Chapter 3 extends this discussion to a versatile new multi-energy soft x-ray diagnostic. Together, these soft x-ray diagnostics form a complementary set which can be used to better constrain plasma properties (see Section 5.5).

2.3.1 Elements of the SXR spectrum

The SXR spectrum of a high-temperature magnetically-confined plasma emerges from a combination of several distinct physical processes. If a plasma were composed of 100% ionized hydrogen, its soft x-ray (SXR) spectrum would be simply composed of bremsstrahlung radiation resulting from interactions between freely moving electrons and ions. The electrons move much faster than the ions meaning that they can be treated as stationary, resulting in the simple spectrum given in the equation [28]

$$\varepsilon_{FF}(\nu) = C \cdot n_e n_i Z_i^2 \frac{e^{-h\nu/T_e}}{\sqrt{T_e}} \bar{g}_{ff}(\nu, T_e), \tag{2.18}$$

C is a collection of physical constants, $h\nu$ is the energy of the emitted photon, n_e is the electron density, n_i is the ion density, Z_i is the ionic charge of the ion ($Z = 1$ for hydrogen), T_e is the electron temperature (in energy units), and $\bar{g}_{ff}$ is the Maxwell-averaged Gaunt factor. The Gaunt factor accounts for quantum mechanical corrections to the electron ion scattering process [29]. Note that throughout this discussion we will assume that all particles obey a Maxwell-Boltzmann distribution.

In such a basic plasma diagnosing the electron temperature would be easy. All one would need to do is extract the slope of the log of the spectrum when plotted versus energy. This could readily be accomplished by taking a few measurements with different cutoff thresholds. No real plasmas, however, is this simple. All real plasmas contain low concentrations of higher-Z ions, which we refer to as impurities. Even with densities far less than 1% of the electron density these impurities can dominate the SXR spectrum. They can also add additional features into the spectrum resulting from free-bound interactions. The most significant free-bound effect, commonly called *radiative recombination*, occurs when an ion captures a passing electron and emits its residual kinetic energy as radiation. This results in an enhancement of the free-free spectrum with discrete steps at the ionization energies of the recombined electrons. This effect is described thoroughly in other sources [28, 30] and is illustrated in Figure 2.9. The total continuum emissivity spectrum is therefore a combination of these two effects:

$$\varepsilon_{cont}(\nu) = C \cdot n_e n_i Z_i^2 \frac{e^{-h\nu/T_e}}{\sqrt{T_e}} \left[\bar{g}_{ff}(\nu, T_e) + \bar{g}_{fb}(\nu, T_e) \right]. \tag{2.19}$$

where a new term $\bar{g}_{fb}$ has been introduced to represent the enhancements to the spectrum resulting from radiative recombination. The full equation is given in Equation (5.3.53) of Hutchinson [28]. Note that if $C = 5.03 \times 10^{-54}$ and n_e and n_i are measured in m^{-3}, then the result is measured in units of [W m^{-3} sr^{-1} Hz^{-1}]. The total emissivity spectrum emitted by the plasma is the sum of Equation 2.19 evaluated for each ion species, majority and impurities. In cases where there are no recombination steps in the portion of the spectrum being measured, this sum can be simplified to

$$\varepsilon_{cont}(\nu) \propto n_e^2 Z_{\text{eff}} \frac{e^{-h\nu/T_e}}{\sqrt{T_e}}, \tag{2.20}$$

where $\sum_i n_e n_i Z_i^2 \equiv n_e^2 Z_{\text{eff}}$ is the ion-effective charge.

Although these effects are sufficient to describe the SXR spectrum emitted by most

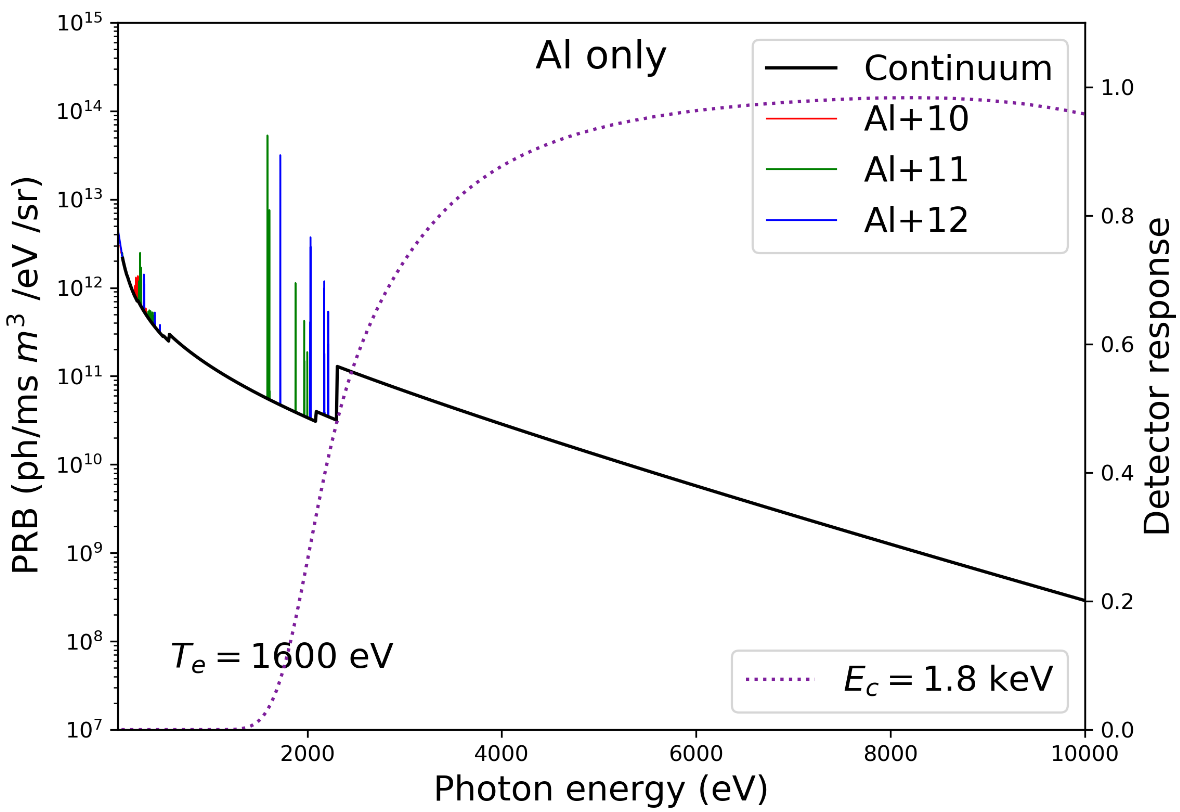

Figure 2.9: The emissivity spectrum of Al with $T_e = 1600eV$. The spectrum is the result of bremsstrahlung (free-free), recombination (free-bound), and excitation lines. The dashed line indicates a typical region of sensitivity for SXR diagnostics.

low-Z impurities (such as C, N, and O) medium- and high-Z ions introduce additional complications[2]. Since these impurities are not fully ionized interactions with passing electrons can excite ground-state electrons into a higher energy level. This electron will then decay back into the ground state, emitting the extra energy as a photon. This process typically leads to the presence of many *excitation lines* in the plasma. Plasmas in MST typically feature very bright excitation lines at $E \sim 2$ keV, which has in the past significantly complicated the interpretation of x-ray measurements [31]. These features are also shown in Figure 2.9. Since the ME-SXR diagnostic is a broad-band detector we can ignore effects such as Doppler and pressure broadening [28], meaning that line emissions are represented well by delta functions:

$$\varepsilon_{exc} = n_e n_i \sum_{\ell \in \mathcal{L}_i} \frac{\mathcal{P}_{i,\ell}(T_e, n_e)}{4\pi} \delta(E - E_\ell). \tag{2.21}$$

Here ℓ indexes the individual transitions, $\mathcal{L}_i$ refers to the set of all transitions for the specified ion, E_ℓ is the line energy, and $\mathcal{P}_{i,\ell}$ is the photon emissivity coefficient, or PEC, which quantifies the strength of the transition. This topic is expanded upon significantly in Section 4.2.3, which is concerned with the development of a quantitative model for the SXR emissivity using the ADAS code [32]. As will be discussed in that section, the values for $\mathcal{P}_{i,\ell}$ depend upon a number of choices about how to model the coupling between the particle population and the emitted radiation.

It is worth noting that there are additional mechanisms which can produce SXR radiation which have not been described here. For example, dielectronic recombination is a three-body interaction which occurs when a free electron is captured and simultaneously excited a bound electron, which then emits a photon as it decays back to the ground state [30]. Preliminary investigations by our collaborator Paolo Franz found that this effect did not contribute significantly to the ME-SXR signal. A similar forward

[2]Partially ionized low-Z impurities in the plasma edge also emit line radiation, but these are outside of SXR range and are typically filtered out by detectors.

model for the NICKAL2 K-edge spectrometer, developed by my colleague Lisa Reusch, does include this effect by way of ADAS.

2.3.2 The Soft X-ray Tomography diagnostic

The Soft X-ray Tomography diagnostic (sometimes abbreviated as SXR tomo. or SXT) was first installed on MST in 2001 [33] and has since served as one of the main diagnostics for examining the plasma's internal structure. The system has undergone several upgrades over the years, with the most significant overhaul occurring around 2010 when the diagnostic was redesigned to allow two-color measurements for each viewing chord, significantly improving the diagnostic's ability to measure T_e [34]. Since that upgrade most advances to the diagnostic have been on the analysis side, such as better understanding of the effect of impurities in the Be foils on filter transmission [35] and validating a synthetic model of the diagnostic to allow absolute comparisons to the data [31].

In its present form, the diagnostic is composed of four detectors at a single toroidal angle ($\phi_{MST} = 90°$) with the viewing geometry chosen to sample the plasma cross-section sufficiently well to enable tomographic inversions (Section 2.3.3). Each detector contains 20 Si photodiodes, each having a thickness of 35 μm and an active area of 4 mm^2. Diodes are paired so that each pair approximately share a line-of-sight, leaving a total of ten chords per detector. This geometry is illustrated in Figure 2.10. The paired diodes look through beryllium foils of differing thicknesses (45 μm and 172 μm for the data in this thesis), often referred to as the "thick" and "thin" channels. As discussed in Section 2.3.1, this allows the diagnostic to sample the slope of SXR continuum, and therefore infer T_e. The signal output from each photodiode is passed through a differential transimpedance amplifier before being digitized at 500 kHz.

Historically T_e measurements have been performed using one of two techniques. The "direct brightness" technique involves taking the ratio of the "thick" and "thin" chan-

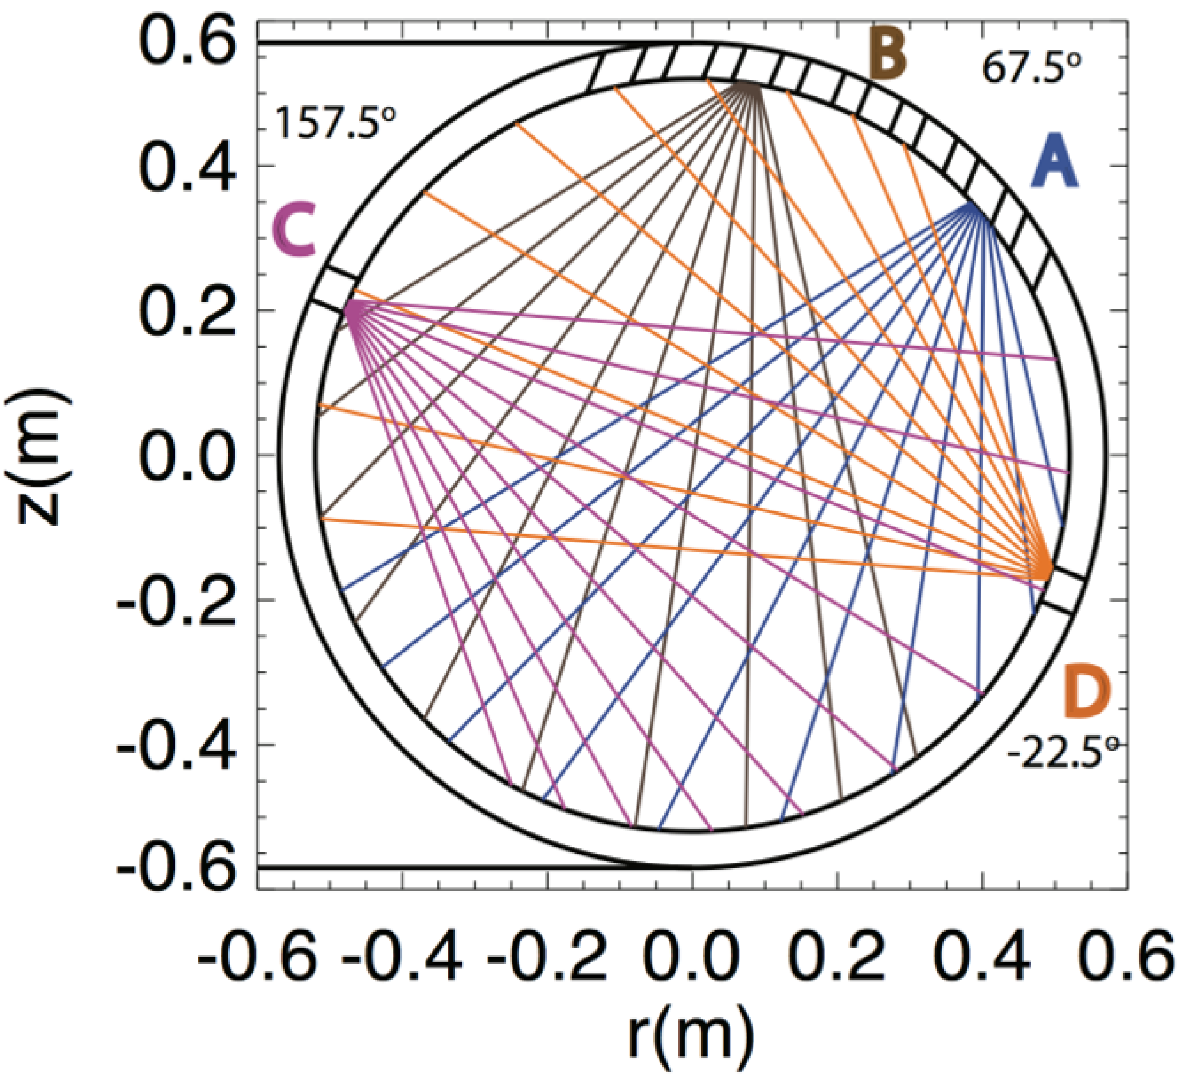

Figure 2.10: The Soft X-ray Tomography diagnostic on MST is composed of four detectors, labeled SXR-A,B,C,D, with ten viewing chords per detector. These chords sample the plasma cross-section from multiple viewing angles, enabling tomographic inversions. There are two Si diodes per chord with Be foils of differing thickness, permitting sensitivity to T_e.

nel chord-integrated measurements and using a ratio curve (computed via synthetic modeling) to map this ratio to the electron temperature [36]. This technique works on the principle that the measured brightness along a chord is mostly a function of the highest temperature (and hence most emissive) point along the chord. As such, it is prone to underestimating the temperature if the geometry of the emission profile was not properly accounted for in the construction of the ratio curve, and can underestimate localized emissive structures [37]. The other method is generate tomographic inversions separately for the thin and thick filter channels, and then take the ratio of the resulting emissivity maps [36]. However, tomographic inversions have some inherent uncertainty which can be amplified when considering the ratios. A third approach, based on modeling, is used throughout this thesis for both the SXR tomography system and the new ME-SXR diagnostic.

The SXR tomography diagnostic has found many applications over the years. Aside from providing T_e measurements, the diagnostic has also been used as a diagnostic of internal magnetic structure. Islands in the emissivity structure have been found to map onto magnetic islands [33], providing a direct measurement of MHD activity. The diagnostic has also been useful for identifying the emergence of helical structures in the magnetic equilibrium [38]. More recently, the diagnostic has become central to the integrated data analysis approach being developed on the MST, which has been used to measure Z_{eff} by synthesizing multiple diagnostics [39].

2.3.3 Tomographic inversion

The soft x-ray tomography diagnostic measures the brightness f along each diode's viewing chord, a line integrated quanitity. However we are typically more interested in the local quantity, the emissivity g, which is the the brightness via the integral:

$$f(p,\phi) = \int_{\mathcal{L}(p,\phi)} g(r,\theta) d\boldsymbol{x}. \tag{2.22}$$

where p and ϕ are the impact parameters specifying the viewing chord (see Section 4.4.1), $\mathcal{L}(p,\phi)$ specifies the viewing chord, and r,θ are the standard polar coordinates.

The class of mathematical techniques for inverting this integral transform given a finite number of measurements of f are referred to as tomographic inversions. The SXR tomography diagnostic on MST employs a technique commonly referred to as the Cormack-Bessel method [40, 41]. This technique expands both the brightness and emissivity as a truncated Fourier series over M terms,

$$f(p,\phi) = \sum_{m=0}^{M} \left[f_m^c(p) \cos(m\phi) + f_m^s(p) \sin(m\phi) \right] \tag{2.23}$$

$$g(r,\theta) = \sum_{m=0}^{M} \left[g_m^c(r) \cos(m\theta) + g_m^s(r) \sin(m\theta) \right]. \tag{2.24}$$

These expressions can be used to derive an expression for the components of emissivity:

$$g_m^{c,s}(r) = -\frac{1}{\pi}\frac{d}{dr}\int_r^1 \frac{r\, f_m^{c,s}(p)\, T_m(p/r)\, dp}{p\sqrt{p^2-r^2}} \tag{2.25}$$

where $T_m(p/r)$ are the Chebyshev polynomials. The Fourier components are further expanded in terms of Bessel functions and recast into a matrix equation of the form

$$\boldsymbol{f} = \boldsymbol{W} \cdot \boldsymbol{g} \tag{2.26}$$

which is computed iteratively until the norm of the residual $||\boldsymbol{f} - \boldsymbol{W} \cdot \boldsymbol{g}||$ falls below a specified tolerance. The result of this process is a two-dimensional map of the emissivity in the poloidal plane. An example of an emissivity map produced with a Cormack-Bessel inversion is shown in Figure 2.11.

The expansions in terms of Bessel functions constrains the symmetry of the resulting reconstructions, and the level of detail is constrained by the finite number of terms retained in the sums. However this method has been found to produce island struc-

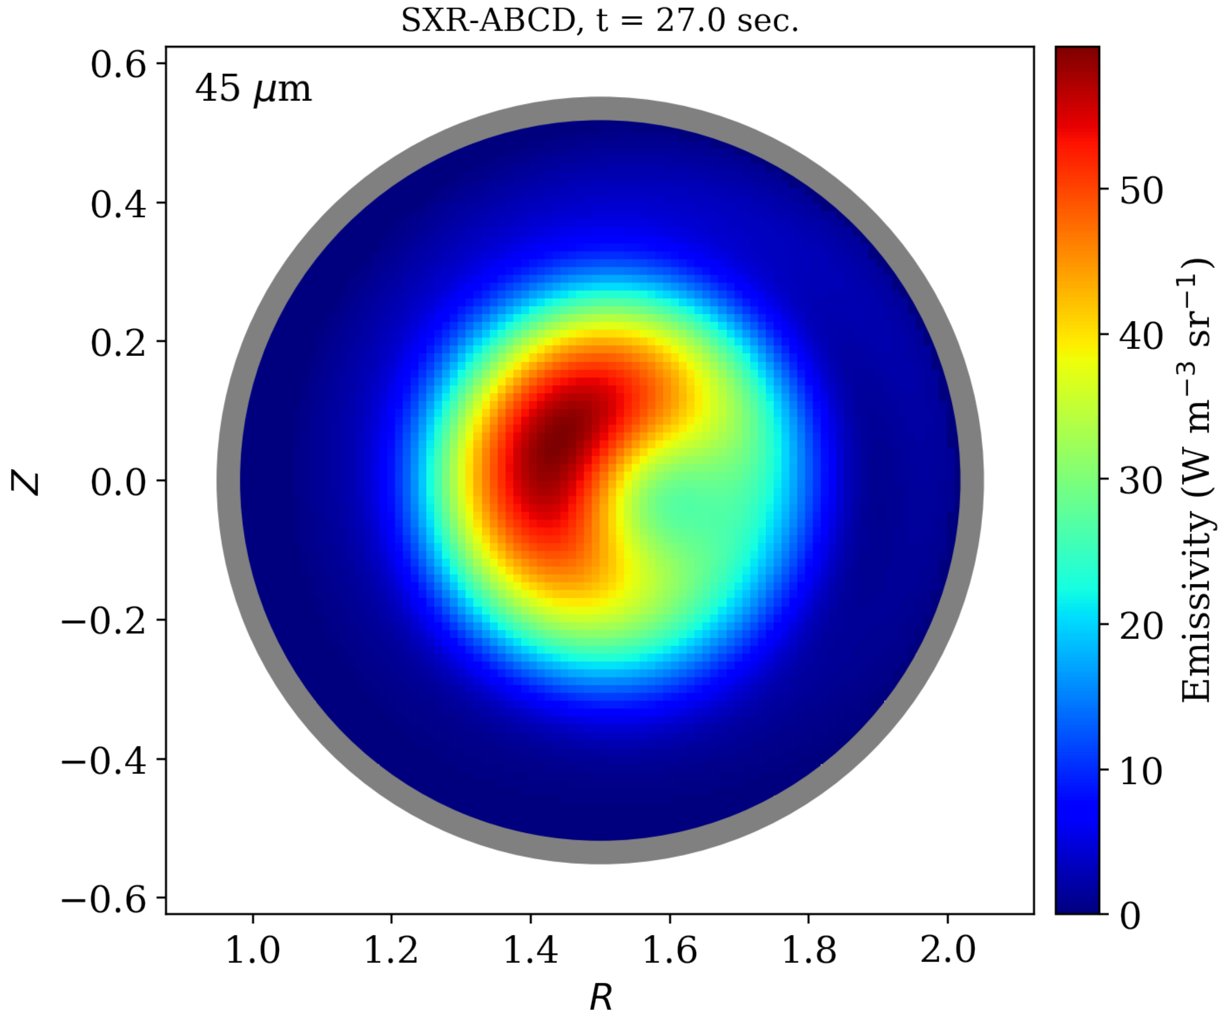

Figure 2.11: Two-dimensional reconstruction of emissivity using the Cormack-Bessel inversion technique. The inversion shown was made using data from the 45 μm channels for a 500 kA non-reversed plasma exhibiting a helical equilibrium. The bright spot corresponding to a “bean-shaped” structure off of the normal magnetic axis is typical of QSH plasmas in the MST.

ture consistent with less-constrained finite element methods [37] while maintaining a significant advantage in computational time. Other modern techniques involving non-stationary Gaussian processes [42] and machine learning [43] have been developed elsewhere but have not yet been tested on the MST.

2.3.4 Effect of Al lines on the SXR spectrum in MST

As a result of long-standing collaborations between the two labs, the two-color soft x-ray tomography diagnostic and analysis techniques described in this section are very similar the system used on RFX-mod in Italy [41]. However, the presence of partially-ionized aluminum ions in MST has introduced significant additional difficulties for the MST system. While RFX-mod commonly uses beryllium filters with thicknesses as low as 25 μm for direct-brightness T_e measurements over a wide range of plasma currents, the strong aluminum lines emissions at ~ 2 keV have historically forced the MST system into using thin/thick filter combinations as thick as 427/805 μm in order to suppress the distortion introduced by these line emissions. The use of such thick filters has limited the application of the diagnostic to the MST's highest performance regime, where electron temperatures $T_e > 1$ keV produce a sufficiently high photon flux above the filter cutoff energies.

This equilibrium has begun to change in recent years. Recent work by L. Reusch and P. Franz [31] has led to the development and validation of a quantitative model of soft x-ray emission in the MST which agrees with measurments over a wide range of plasma temperatures and beryllium filter thicknesses. This was achieved by carefully incorporating the individual mechanisms discussed in Section 2.3.1, and serves as the basis for the multi-energy diagnostic model discussed in Chapter 4. A summary of that work is provided in Figure 2.12.

As shown in the figure, the SXR signal for very thin filters is dominated by Al line emissions, in the most extreme case contributing approximately an order of magnitude

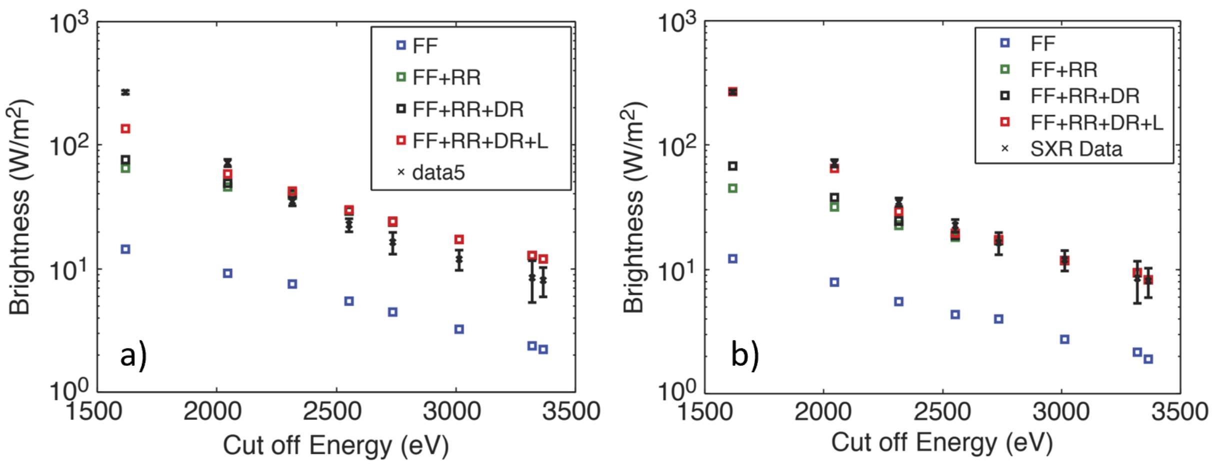

Figure 2.12: These figures show the contribution to the overall SXR signal of the various physical processes discussed in Section 2.3.1. FF refers to "free-free" radiation, RR to "radiative recombination," DR to "dielectronic recombination," and L refers to line emissions. Cutoff energy is increased by increasing the Be filter thickness. Fgiure a) does not consider the impact of charge-exchange with neutral particles on the ionization balance, while figure b) does include this effect. All effects are important for achieving quantitative agreement with experimental data across all filter thicknesses. Figure b) also illustrates the importance of Aluminum emission lines when thin filters are used. Figure reproduced from Reusch, *et al.* [31].

more to the measured signal than other mechanisms. For these thicknesses direct brightness measurements are difficult, as the ratio is dependent not only temperature but also on the line amplitude, which itself depends on n_e, n_{Al}, and the neutral density as well as T_e. As the filter thickness increases these complex dependencies become insignificant and we return to a regime where ratio-based techniques can be used to infer T_e directly. The figure also illustrates the importance of properly characterizing the neutral density when modeling the emission spectrum, a detail which we will return to in Section 4.2.2.

The application of a quantitative model to understand absolute soft x-ray measurements as described above relied upon many years of work and measurements from numerous diagnostics to develop a comprehensive understanding of properties of MST high-performance plasmas. This includes measurements of ion and electron density profiles, an understanding of ion transport phenomena, measurements of neutral density, and Thomson scattering measurements to characterize the typical range and shape of electron temperature profiles. This knowledge was synthesized via a methodology called Integrated Data Analysis [44], discussed more thoroughly in Appendix B and demonstrated in Section 5.5. This analysis has been most-successfully applied to high-current PPCD plasmas.

Due to the complex dependencies of soft x-ray emission on other plasma properties, our understanding of data from the soft x-ray tomography diagnostic can be greatly improved by simultaneous measurements from complementary diagnostics. In the following section I discuss one such complementary diagnostic, a new single-purpose K-edge spectrometer designed to directly measure the aluminum line amplitudes in MST plasmas. In the following chapter I more thoroughly discuss another complementary diagnostic, a highly-customizable solid state pixelated photon counting multi-energy soft x-ray diagnostic which can be used to isolate contributions from emission lines and the continuum.

2.3.5 The NICKAL2 Ross Filter Detector

The NICKAL2[3] is a single-chord Ross spectrometer designed specifically to isolate the signal from the Al^{+11} and Al^{+12} transition lines [45]. The idea for the diagnostic was partially inspired by an previous Ross spectrometer once used on the MST to measure O, C, and Al impurities in standard RFP plasmas [46]. A Ross spectrometer is a radiation detector which uses carefully selected filter pairs to create pass-bands into which incoming electromagnetic radiation can be isolated. Filter thicknesses are chosen so that the transmission curves are nearly identical except for the sharp drops at their respective K-edges[4], which set the bin boundaries. Elements which are consecutive on the periodic table are typically selected in order to minimize the size of the pass band. The result can be seen in Figure 2.13, which shows the transmission curves for the three filters installed in the NICKAL2 detector.

The NICKAL2 detector consists of three columnated photodiodes, of the same type used in the tomography diagnostic, each separated from the plasma by a different x-ray filter. The filters were designed using a simple optimization routine to simultaneously minimize transmission outside of the passbands while maximizing it inside the bands (see N. Lauersdorf's senior thesis [45] for a thorough description of the diagnostic and the optimization process). The initial design called for a mix of aluminum (1.56 keV), silicon (1.84 keV), and phosphorus (2.14 keV) filters in order to effectively isolate the Al lines in the 1.5-2 keV region of the spectrum. However a phosphorous-based filter was deemed impractical due to its high combustibility, so it was replaced with zirconium which has an L-edge around (2.23 keV). Because L-edges are weaker than K-edges, an even thinner filter was required. An additional thin film (either beryllium or Mylar) was also deposited onto each filter in order to manipulate out the lower-energy part of the

[3] It is so named because it is the second Ross spectrometer used on the MST to be created by someone named Nick (Lanier and Lauersdorf), and it is designed to measure Al line radiation.

[4] A filter's K-edge is a sudden increase in x-ray absorption, or equivalently a drop in transmission, which occurs at the binding energy of the innermost (K-shell) electrons.

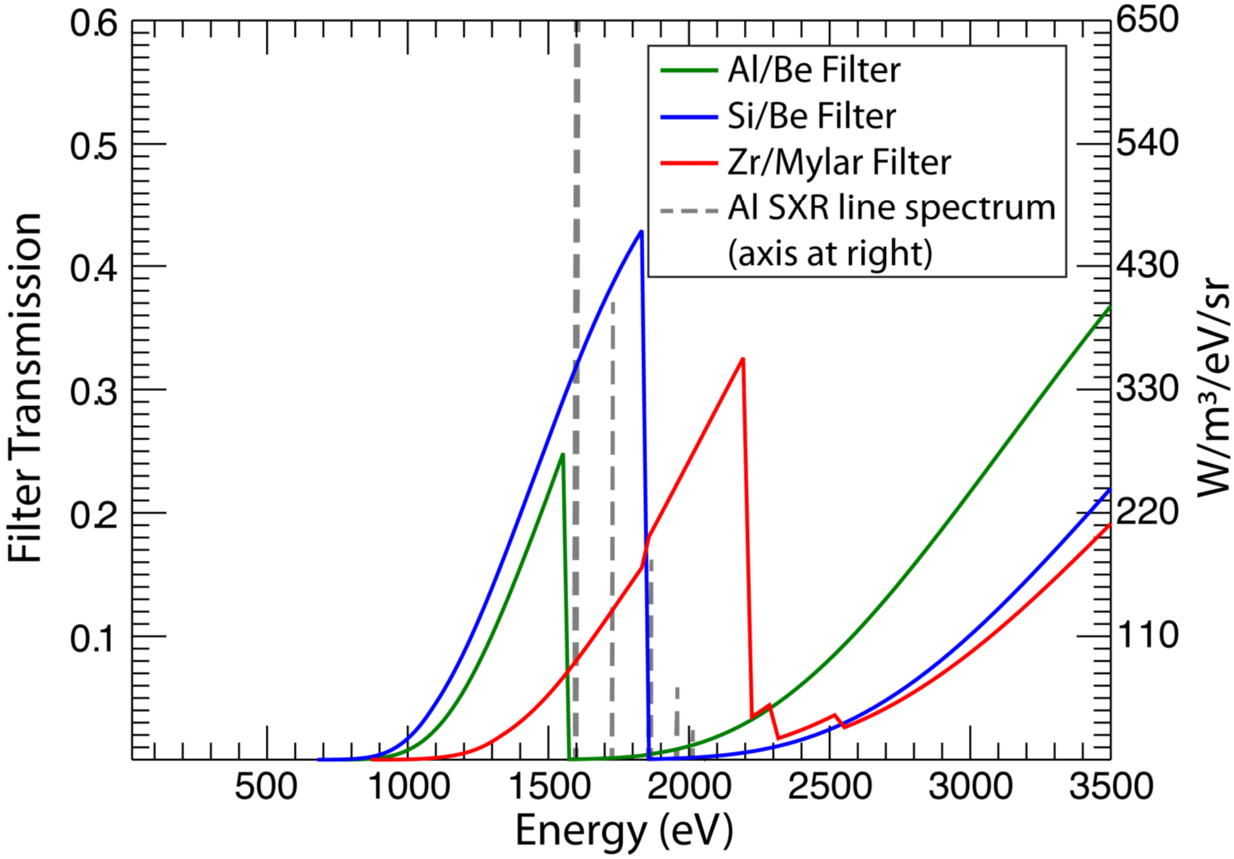

Figure 2.13: Illustration of the transmission bands for the three core-viewing NICKAL2 channels. The sudden drops transmission are used to independently measure the line amplitudes. The dashed lines correspond to the brightest Al emission lines in typical MST plasma discharges, as modeled by ADAS. Figure reproduced from N. Lauersdorf's undergraduate thesis [45].

transmission curve and to provide additional structural support to the very thin filters. The final filter design was:

1. 2.0 μm Zr with 8.0 μm Mylar
2. 6.8 μm Al with 31.2 μm Be
3. 10.0 μm Si with 12.7 μm Be.

Bibliography

[1] T. Nishizawa, M. D. Nornberg, D. J. Den Hartog, and D. Craig, "Upgrading a high-throughput spectrometer for high-frequency (<400 kHz) measurements," *Review of Scientific Instruments*, vol. 87, no. 11E530, 2016. [Online]. Available: http://dx.doi.org/10.1063/1.4960073

[2] J. Waksman, J. K. Anderson, M. D. Nornberg, E. Parke, J. A. Reusch, J. Homepage, D. Liu, G. Fiksel, V. I. Davydenko, A. A. Ivanov, N. Stupishin, P. P. Deichuli, and H. Sakakita, "Neutral beam heating of a RFP plasma in MST," *Physics of Plasmas*, vol. 19, no. 122505, 2012. [Online]. Available: http://dx.doi.org/10.1063/1.4772763

[3] W. J. Capecchi, J. K. Anderson, P. J. Bonofiglo, J. Kim, and S. Sears, "A collimated neutron detector for RFP plasmas in MST," *Review of Scientific Instruments*, vol. 87, no. 11D826, 2016. [Online]. Available: http://dx.doi.org/10.1063/1.4961304

[4] J. J. Koliner, "Neutral Beam Excitation of Alfvén Continua in the Madison Symmetric Torus Reversed Field Pinch," Ph.D. dissertation, University of Wisconsin-Madison, Madison, oct 2013.

[5] S. L. Prunty, "A primer on the theory of Thomson scattering for high-temperature fusion plasmas," *Physica Scripta*, vol. 89, no. 128001, 2014. [Online]. Available: https://doi.org/10.1088/0031-8949/89/12/128001

[6] J. A. Reusch, "Measured and Simulated Electron Thermal Transport in the Madison Symmetric Torus Reversed Field Pinch," Ph.D. dissertation, University of Wisconsin-Madison, Madison, 2011.

[7] A. C. Selden, "Simple analytic form of the relativistic Thomson scattering spectrum," *Physics Letters*, vol. 79A, no. 5-6, pp. 405–406, oct 1980. [Online]. Available: https://doi.org/10.1016/0375-9601(80)90276-5

[8] E. Parke, "Diagnosis of equilibrium magnetic profiles, current transport, and internal structures in a reversed-field pinch using electron temperature fluctuations," Ph.D. dissertation, University of Wisconsin-Madison, 2014.

[9] W. Young and D. Den Hartog, "Thomson scattering at 250 kHz," *Journal of Instrumentation*, vol. 10, no. C12021, dec 2015. [Online]. Available: https://doi.org/10.1088/1748-0221/10/12/C12021

[10] E. Parke, W. X. Ding, J. Duff, and D. L. Brower, "An upgraded interferometer-polarimeter system for broadband fluctuation measurements," *Review of Scientific Instruments*, vol. 87, pp. 11–115, 2016. [Online]. Available: http://dx.doi.org/10.1063/1.4960731

[11] J. Xie, H. Wang, W. Ding, H. Li, T. Lan, A. Liu, W. Liu, and C. Yu, "Design of interferometer system for Keda Torus eXperiment using terahertz solid-state diode sources," *Review of Scientific Instruments*, vol. 85, pp. 11–828, 2014. [Online]. Available: https://doi.org/10.1063/1.4890249

[12] T. F. Yates, W. X. Ding, T. A. Carter, and D. L. Brower, "Simultaneous density and magnetic field fluctuation measurements by far-infrared interferometry and polarimetry in MST," *Review of Scientific Instruments*, vol. 79, no. 10E714, 2008. [Online]. Available: https://doi.org/10.1063/1.2966377

[13] J. R. Duff, "Observation of trapped-electron mode microturbulence in improved confinement reversed-field pinch plasmas," Ph.D. dissertation, University of Wisconsin-Madison, 2018.

[14] R. O'Connell, D. J. Den Hartog, C. B. Forest, J. K. Anderson, T. M. Biewer, B. E. Chapman, D. Craig, G. Fiksel, S. C. Prager, J. S. Sarff, S. D. Terry, and R. W. Harvey, "Observation of Velocity-Independent Electron Transport in the Reversed Field Pinch," *Physical Review Letters*, vol. 91, no. 4, pp. 8–11, 2003. [Online]. Available: https://doi.org/10.1103/PhysRevLett.91.045002

[15] A. M. Dubois, J. D. Lee, and A. F. Almagri, "A high time resolution x-ray diagnostic on the Madison Symmetric Torus," *Review of Scientific Instruments*, vol. 86, no. 073512, pp. 1–6, 2015. [Online]. Available: https://doi.org/10.1063/1.4927454

[16] A. M. DuBois, A. Scherer, A. F. Almagri, J. K. Anderson, M. D. Pandya, and J. S. Sarff, "Turbulence-driven anisotropic electron tail generation during magnetic reconnection," *Physics of Plasmas*, vol. 25, no. 055705, 2018. [Online]. Available: http://dx.doi.org/10.1063/1.5016240

[17] D. J. Clayton, "Fast Electron Transport in Improved-Confinement RFP Plasmas," Ph.D. dissertation, University of Wisconsin-Madison, 2010.

[18] J. K. Anderson, C. B. Forest, T. M. Biewer, J. S. Sarff, and J. C. Wright, "Equilibrium reconstruction in the Madison Symmetric Torus reversed field pinch," *Nuclear Fusion*, vol. 44, no. 1, pp. 162–171, 2004. [Online]. Available: https://doi.org/10.1088/0029-5515/44/1/018

[19] P. Zanca and D. Terranova, "Reconstruction of the magnetic perturbation in a toroidal reversed field pinch," *Plasma Physics and Controlled Fusion*, vol. 46, no. 7, pp. 1115–1141, 2004. [Online]. Available: https://doi.org/10.1088/0741-3335/46/7/011

[20] F. Auriemma, P. Zanca, W. F. Bergerson, B. E. Chapman, W. X. Ding, D. L. Brower, P. Franz, P. Innocente, R. Lorenzini, B. Momo, and D. Terranova, "Magnetic reconstruction of nonaxisymmetric quasi-single-helicity configurations in the Madison Symmetric Torus," *Plasma Physics and Controlled Fusion*, vol. 53, no. 10, 2011. [Online]. Available: https://doi.org/10.1088/0741-3335/53/10/105006

[21] W. A. N. Lawrence, "Hydromagnetic Stability of a Diffuse Linear Pinch*," *Annals of Physics*, vol. 10, pp. 232–267, 1960. [Online]. Available: https://doi.org/10.1016/0003-4916(60)90023-3

[22] A. H. Glasser, "The direct criterion of Newcomb for the ideal MHD stability of an axisymmetric toroidal plasma," *Physics of Plasmas*, vol. 23, no. 72505, 2016. [Online]. Available: https://doi.org/10.1063/1.4958328

[23] J. D. Hanson, S. P. Hirshman, S. F. Knowlton, L. L. Lao, E. A. Lazarus, and J. M. Shields, "V3FIT: a code for three-dimensional equilibrium reconstruction," *Nuclear Fusion*, vol. 49, no. 075031, 2009. [Online]. Available: https://doi.org/10.1088/0029-5515/49/7/075031

[24] J. J. Koliner, M. R. Cianciosa, J. Boguski, J. K. Anderson, J. D. Hanson, B. E. Chapman, D. L. Brower, D. J. Den Hartog, W. X. Ding, J. R. Duff, J. A. Goetz, M. Mcgarry, L. A. Morton, and E. Parke, "Three dimensional equilibrium solutions for a current-carrying reversed-field pinch plasma with a close-fitting conducting shell," *Phys. Plasmas*, vol. 23, p. 32508, 2016. [Online]. Available: http://dx.doi.org/10.1063/1.4944670]

[25] S. P. Hirshman, W. I. Van Rij, and P. Merkel, "Three-dimensional free boundary calculations using a spectral Green's function method," *Computer Physics Communications*, vol. 43, pp. 143–155, 1986. [Online]. Available: https://doi.org/10.1016/0010-4655(86)90058-5

[26] J. Boguski, "Local Ion Velocity Measurements in the MST Saturated Single Helical Axis State," Ph.D. dissertation, University of Wisconsin-Madison, 2019.

[27] C. De Michelis and M. Mattioli, "Soft-X-ray spectroscopic diagnostics of laboratory plasmas," *Nuclear Fusion*, vol. 21, pp. 677–754, 1981. [Online]. Available: https://doi.org/10.1088/0029-5515/21/6/007

[28] I. H. Hutchinson, *Principles of Plasma Diagnostics*, 2nd ed. Cambridge University Press, jul 2002.

[29] P. A. van Hoof, R. J. Williams, K. Volk, M. Chatzikos, G. J. Ferland, M. Lykins, R. L. Porter, and Y. Wang, "Accurate determination of the free-free gaunt factor - I. Non-relativistic gaunt factors," *Monthly Notices of the Royal Astronomical Society*, vol. 444, no. 1, pp. 420–428, 2014. [Online]. Available: https://doi.org/10.1093/mnras/stu1438

[30] H. R. Griem, *Principles of Plasma Spectroscopy*. Cambridge, UK: Cambridge University Press, 1997.

[31] L. M. Reusch, P. Franz, D. J. Den Hartog, J. A. Goetz, M. D. Nornberg, and P. VanMeter, "Model validation for quantitative X-ray measurements," *Fusion Science and Technology*, vol. 74, no. 1-2, pp. 167–176, 2018. [Online]. Available: https://doi.org/10.1080/15361055.2017.1404340

[32] H. P. Summers and M. G. O'Mullane, "Atomic data and modelling for fusion: The ADAS Project," *AIP Conference Proceedings*, vol. 1344, no. May 2011, pp. 179–187, 2011. [Online]. Available: https://doi.org/10.1063/1.3585817

[33] P. Franz, L. Marrelli, P. Piovesan, I. Predebon, F. Bonomo, L. Frassinetti, P. Martin, G. Spizzo, B. E. Chapman, D. Craig, and J. S. Sarff, "Tomographic imaging of resistive mode dynamics in the Madison Symmetric Torus reversed-field

pinch," *Physics of Plasmas*, vol. 13, no. 1, pp. 1–15, 2006. [Online]. Available: https://doi.org/10.1063/1.2160519

[34] M. B. McGarry, P. Franz, D. J. Den Hartog, J. A. Goetz, M. A. Thomas, M. Reyfman, and S. T. A. Kumar, "High-performance double-filter soft x-ray diagnostic for measurement of electron temperature structure and dynamics," *Review of Scientific Instruments*, vol. 83, no. 10, 2012. [Online]. Available: https://doi.org/10.1063/1.4740274

[35] M. B. McGarry, P. Franz, D. J. Den Hartog, and J. A. Goetz, "Effect of beryllium filter purity on x-ray emission measurements," *Plasma Physics and Controlled Fusion*, vol. 56, no. 125018, 2014. [Online]. Available: https://doi.org/10.1088/0741-3335/56/12/125018

[36] P. Franz, F. Bonomo, L. Marrelli, P. Martin, P. Piovesan, G. Spizzo, B. E. Chapman, D. Craig, D. J. Den Hartog, J. A. Goetz, R. O'connell, S. C. Prager, M. Reyfman, and J. S. Sarff, "Two-dimensional time resolved measurements of the electron temperature in MST," *Review of Scientific Instruments*, vol. 77, no. 10F318, 2006. [Online]. Available: https://doi.org/10.1063/1.2229192

[37] M. B. McGarry, "Probing the relationship between magnetic and temperature structures with soft x-rays on the Madison Symmetric Torus," Ph.D. dissertation, University of Wisconsin-Madison, 2013.

[38] F. Bonomo, B. E. Chapman, P. Franz, L. Marrelli, P. Martin, P. Piovesan, I. Predebon, G. Spizzo, and R. B. White, "Imaging of a Double Helical Structure in the Reversed Field Pinch," *IEEE Transactions on Plasma Science*, vol. 33, no. 2, 2005. [Online]. Available: https://doi.org/10.1109/TPS.2005.845334

[39] M. Galante, L. Reusch, D. Den Hartog, P. Franz, J. Johnson, M. McGarry, M. Nornberg, and H. Stephens, "Determination of Z_eff by integrating measurements from x-ray tomography and charge exchange recombination spectroscopy," *Nuclear Fusion*, vol. 55, no. 12, p. 123016, 2015. [Online]. Available: https://doi.org/10.1088/0029-5515/55/12/123016

[40] A. M. Cormack, "Representation of a function by its line integrals, with some radiological applications. II," *Journal of Applied Physics*, vol. 35, no. 10, pp. 2908–2913, 1964. [Online]. Available: https://doi.org/10.1063/1.1713127

[41] P. Franz, L. Marrelli, A. Murari, G. Spizzo, and P. Martin, "Soft X ray tomographic imaging in the RFX reversed field pinch," *Nuclear Fusion*, vol. 41, no. 6, pp. 695–709, 2001. [Online]. Available: https://doi.org/10.1088/0029-5515/41/6/304

[42] D. Li, J. Svensson, H. Thomsen, F. Medina, A. Werner, and R. Wolf, "Bayesian soft X-ray tomography using non-stationary Gaussian Processes," *Review of Scientific Instruments*, vol. 84, no. 8, 2013. [Online]. Available: https://doi.org/10.1063/1.4817591

[43] D. D. Carvalho, D. R. Ferreira, P. J. Carvalho, M. Imrisek, J. Mlynar, H. Fernandes, and J. Contributors, "Deep neural networks for plasma tomography with applications to JET and COMPASS," *Journal of Instrumentation*, vol. 14, no. C09011, sep 2019. [Online]. Available: https://doi.org/10.1088/1748-0221/14/09/C09011

[44] L. M. Reusch, M. D. Nornberg, J. A. Goetz, and D. J. Den Hartog, "Using integrated data analysis to extend measurement capability (invited)," *Review of Scientific Instruments*, vol. 89, no. 10, 2018. [Online]. Available: https://doi.org/10.1063/1.5039349

[45] N. J. Lauersdorf, "Development of a Ross Filter Based Aluminum Line Radiation (NICKAL2) Detector in Madison Symmetric Torus," University of Wisconsin-Madison, Madison, Tech. Rep., 2018.

[46] N. E. Lanier, S. P. Gerhardt, and D. J. Den Hartog, "Low-cost, robust, filtered spectrometer for absolute intensity measurements in the soft x-ray region," *Review of Scientific Instruments*, vol. 72, no. 1 II, pp. 1188–1191, 2001.

Chapter 3

The Multi-Energy Soft X-Ray diagnostic

The multi-energy soft X-ray (ME-SXR) diagnostic at the Madison Symmetric Torus (MST) has been developed as a collaboration between the University of Wisconsin-Madison and the Princeton Plasma Physics Lab (PPPL). The diagnostic is based around a novel implementation of the PILATUS3 100K detector which has been calibrated to simultaneously sample the plasma emission at multiple x-ray energy ranges. This provides sensitivity to a variety of important plasma properties such as core T_e and n_e, as well as impurity species content [1].

The PILATUS3 detector was calibrated for multi-energy operation following a procedure developed for the PILATUS2 detector at Alcator C-Mod [2] and later extended to the PILATUS3 [3]. This chapter builds upon this prior work by applying this procedure to a new system and then using the results to analyze pixel to pixel variation across the detector. Of particular interest was the resolution to which a specific photon energy threshold could be set due to uncertainty in the calibration procedure and the discrete nature of the PILATUS3 threshold settings. This resolution was found to be $\Delta E < 100$ eV for a 1.6-to-6 keV calibration and $\Delta E < 200$ eV for a 4-to-14 keV calibration. These results are discussed in Section 3.2.

An introduction to the diagnostic hardware, including the PILATUS3 module itself, is provided in Section 3.1. Section 3.3 describes the in-vessel spatial calibration technique used to determine each pixel's line-of-sight through the plasma. Section 3.4 discusses the phenomenon of charge sharing, in which the charge generated by an absorbed photon is partially collected by two adjacent pixels. Finally, Section 3.5 demonstrates multiple ways the detector can be configured and shows corresponding data.

3.1 ME-SXR diagnostic hardware

The ME-SXR diagnostic essentially consists to two parts: the physical hardware (the PILATUS3 detector and associated housing) and the configurations based on our custom energy calibration. This section discusses the former. Section 3.1.1 describes the PILATUS3 detector itself, including its general description, fundamental operating principles, and associated auxiliary systems. Section 3.1.2 covers the detector installation, including the housing, insertion mechanism, vacuum configuration, geometry, and grounding scheme.

3.1.1 The PILATUS3 Detector

The PILATUS3 100K-M detector is a pixelated hybrid photon-counting x-ray detector produced commercially by DECTRIS Ltd. and specifically modified for PPPL. The device is composed of a single 450 μm Si sensor which absorbs incident photons, producing a cloud of photoelectrons with a total charge proportional to the photon energy [4]. This charge is then transferred via a bump-bonded indium connection to one of the many charge-sensitive preamplifiers (CSA) located on one of the 16 application-specific integrated circuits (ASIC) that compose the detector. The charge is converted to a pulse which is discriminated against a threshold by a comparator, rejecting photons with energies below the threshold. The threshold is controlled by a global comparator V_{cmp} setting, but can be further adjusted, or trimmed, on an individual level by an additional setting stored in a 6-bit register called the "trimbit" setting. Pulses which pass this threshold are recorded into a 20-bit counter and read out at pre-set intervals. These ASICs are arranged in an 8x2 grid, each contain an array of 60x97 individual pixels (each with its own CSA, comparator, trimbit setting, and counter), leading to a total of 480x194 = 93,120 pixels (often referred to as 100k). Individual pixels have an effective area of $172 \times 172\ \mu\text{m}^2$. Pixels along the edge of modules are 50% larger in order to span the gap. When using factory settings, the counts from two adjacent "over-sized" pixels are reallocated into an synthetic pixel along the gap in order to avoid discontinuities. This reallocation technique is inapplicable to our custom multi-energy settings, so the synthetic pixels are ignored.

The individual trimbit settings exist to permit the detector to compensate for inhomogeneities resulting from the manufacturing process and achieve a uniform photon energy threshold, which is the intent of the standard factory calibration. The ME-SXR concept, however, use a custom calibration to take advantage of the trimbit settings in order to intentionally set different energy thresholds for pixels across the detector. This allows the implementation of custom configurations which combine spatial and spectral resolution into a single diagnostic which can be quickly and easily configured for a specific scientific goal.

Global settings determine the minimum energy threshold (V_{cmp}), gain (V_{rf}), and the extent of individual trimbit increments on the energy threshold (V_{trm}). Appropriate global settings were determined which permitted sensitivity to multiple energy ranges of interest. Then, a custom calibration procedure (Section 3.2) was applied to determine the mapping between trimbit setting and the energy threshold E_c for each individual pixel. Two specific energy-range calibrations are discussed later in this chapter: a 1.6-to-6 keV calibration (Section 3.2.1) and a 4-to-14 keV calibration (Section 3.2.4). A higher-energy calibration has not been used in this thesis, but will be explored in future devices (i.e. WEST).

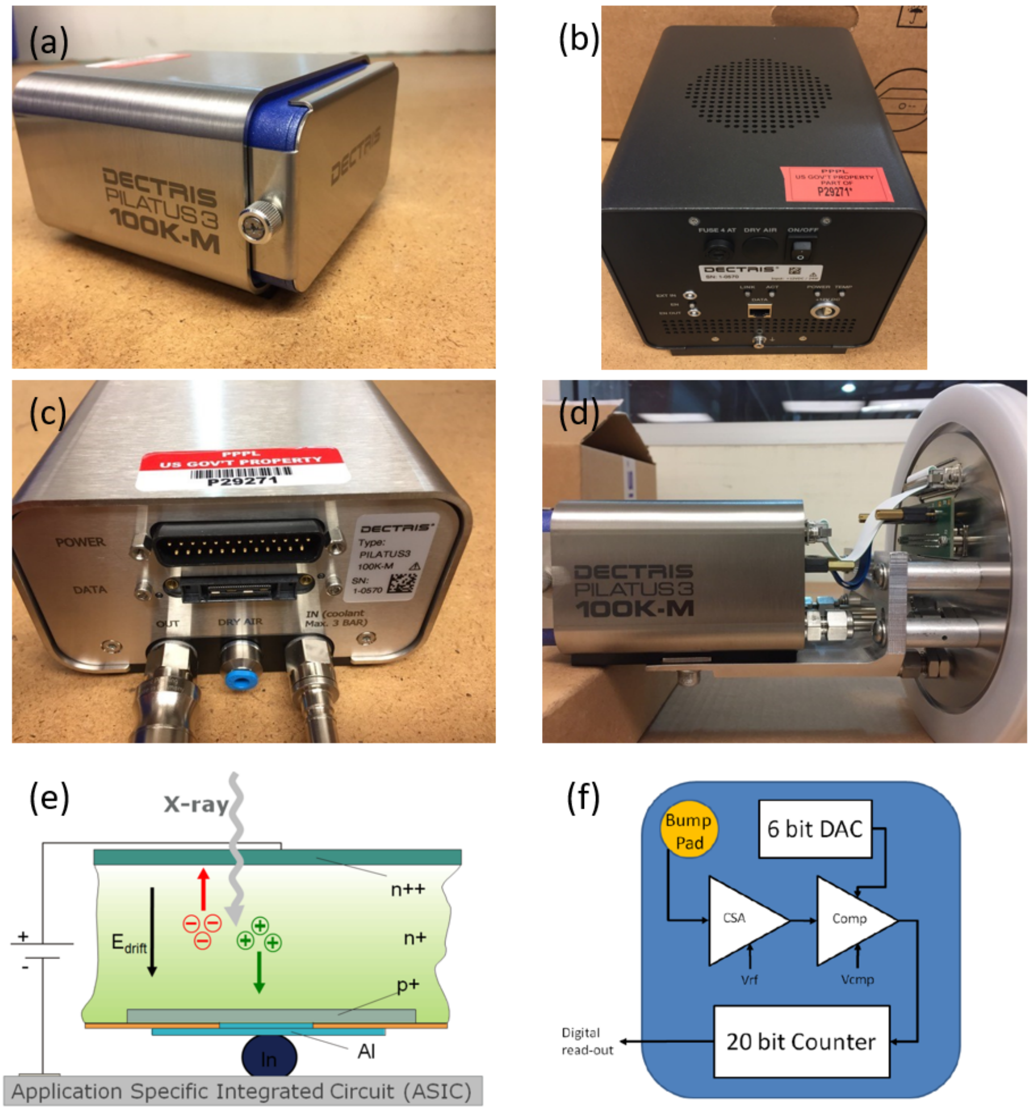

Figure 3.1: (a) PILATUS3 100K-M frontend with a cover over the detector screen; (b) PILATUS3 backend; (c)various connection ports on the frontend; (d) detector mounted onto the flange with vacuum feedthroughs for connectors; (e) Photons are absorbed into the Si semiconductor layer, producing a cloud of electrons and holes; (b) the Si absorber is bump-bonded directly to the detector electrons via small Indium balls; and (f) schematic of the pulse-counting circuit for each pixel.

During a run day, the PILATUS3 detector is configured using these custom settings prior to the plasma discharge. The detector then receives a trigger and begins recording just as the MST capacitor banks begin to discharge, allowing the resulting image files to be synced up with other diagnostic signals. The timing and duration of each image depends on the chosen settings, but the most commonly used configuration is a cycle of a one millisecond exposure followed by a one millisecond readout period (dead time), for an effective time resolution of 500 Hz. Fluctuations within the exposure window cannot be discerned, making the diagnostic best-suited for tracking the evolution of the equilibrium. The settings can be quickly adjusted between plasma discharges.

The PILATUS3 100K-M hardware is composed of three distinct parts: the frontend, the backend, and the control unit (server). The frontend is a small module which contains the pixel grid and the electronics discussed in the preceding paragraphs. The backend is a somewhat larger box which operates the detector, performs the readout, and communicates with the control unit. For standard PILATUS models the pixel grid and readout electronics are combined into a single unit, but for the 100K-M (modified by request for PPPL) they are separate. This allows the more-sensitive electronics in the backend to be placed further away from the vacuum vessel to avoid magnetic interference. The two ends are connected by 2 meter long two cables, a high-speed data ribbon for operation/readout and a parallel cable which powers the frontend and sets the voltage levels. The control unit is a Dell PowerEdge server which fits into a standard server rack. The backend is connected to the server via Ethernet cable (though we convert that signal to fiber, for grounding reasons explained in the following Section). The PILTUS3 frontend, backed, connection ports, vacuum interface, and circuit design are all depicted in Figure 3.1.

3.1.2 Detector installation

The ME-SXR diagnostic on the MST consists of the PILATUS3, the detector housing, and the pinhole through which the plasma is viewed. The detector is mounted inside a "can" and physically separated from the plasma by a 25 μm Be window over the pinhole. This window filters low-energy photons and also separates the low-vacuum region of the detector can from the high-vacuum of the MST. An additional Mylar filter was later added to further suppress low-energy photons, which is discussed in more detail in Section 5.4. Vacuum feedthroughs were designed for the PILATUS3 power and data cables. The detector mount assembly was connected to a drive train (powered by a hand drill) to allow it to be retracted and sealed behind a gate valve when not in use. An exploded view of the components of the detector mount assembly is shown in Figure 3.2.

The "pinhole" is actually a small rectangular slit which allows each pixel to focus on a specific region of the plasma. The actual slit can be removed from the flange and replace with one of several available options (2×1, 2×2, and 4×1, all in mm). In each case, the slit was oriented so that the longer dimension extends in the toroidal direction, taking advantage of the plasma's axisymmetry to increase signal levels. For the data presented in this thesis the 2×1 slit was used. Above the pinhole slit is a curved opening over which the Be filter is mounted. The curve ensures that photons travel through the same distance of Be before impacting a pixel, thus allowing consistent energy resolution. This is especially important because the detector is only 30.5 mm from the pinhole, resulting in a wide viewing angle $2\theta > 90^{\circ}$. The flange, window, and pinhole are shown in Figure 3.3.

The detector was installed with a radial viewing geometry, as shown in Figure 3.4. This viewing geometry allows the detector to measure emissivity within a poloidal cross-section, which can be inverted into one-dimensional radial profiles [5]. Due to the fact that pixels are arranged in a two-dimensional grid (with most not perfectly aligned to the

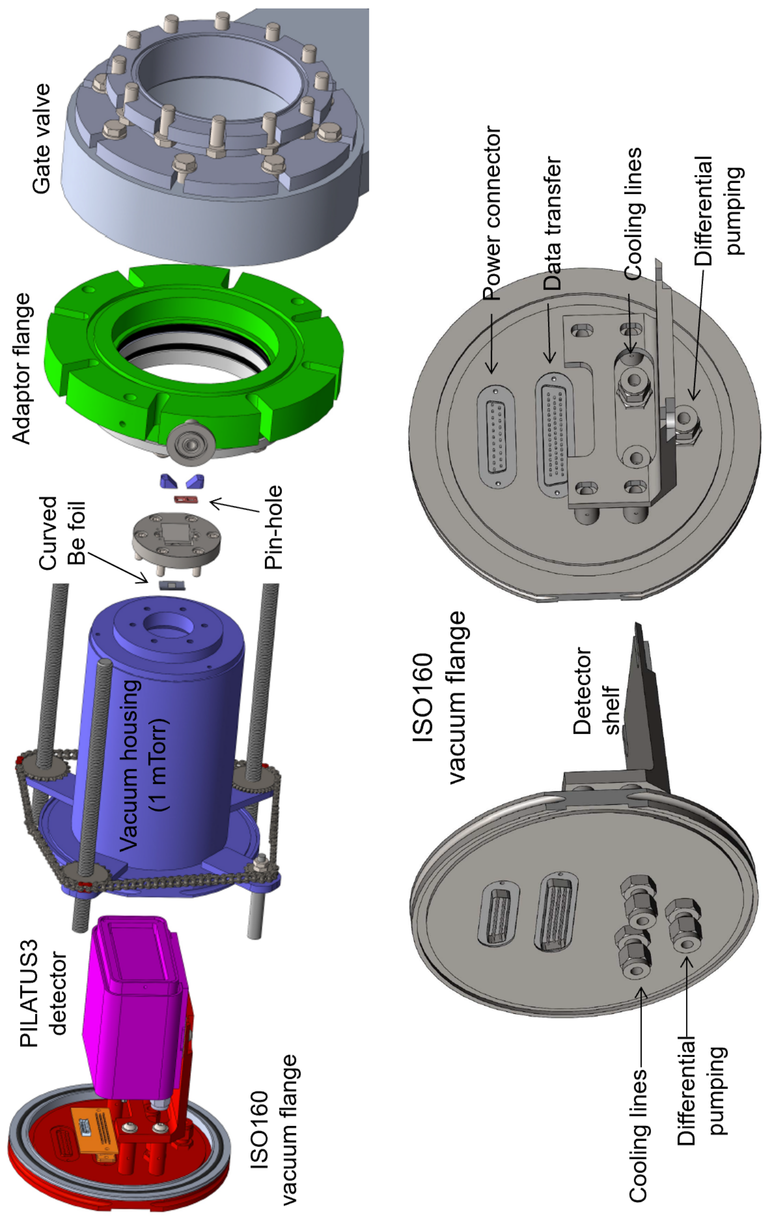

Figure 3.2: Exploded view of the components of the ME-SXR detector mount assembly.

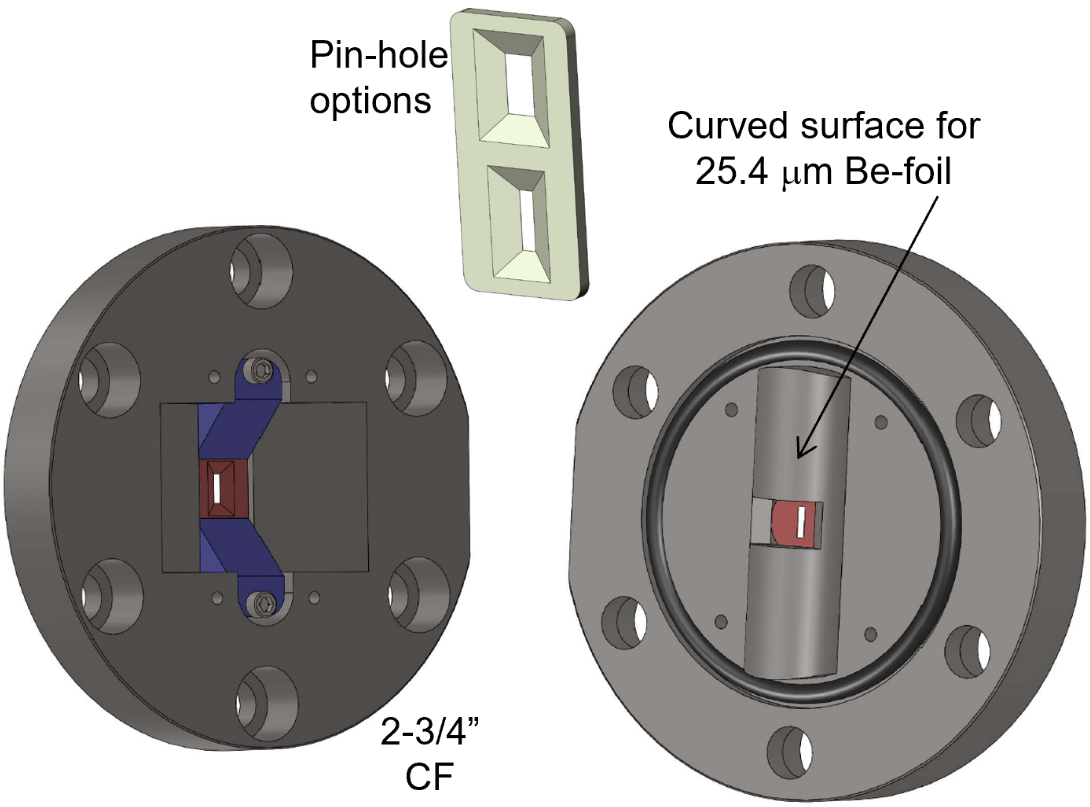

Figure 3.3: Close-up schematic view of the ME-SXR pinhole and 25 μm Be window.

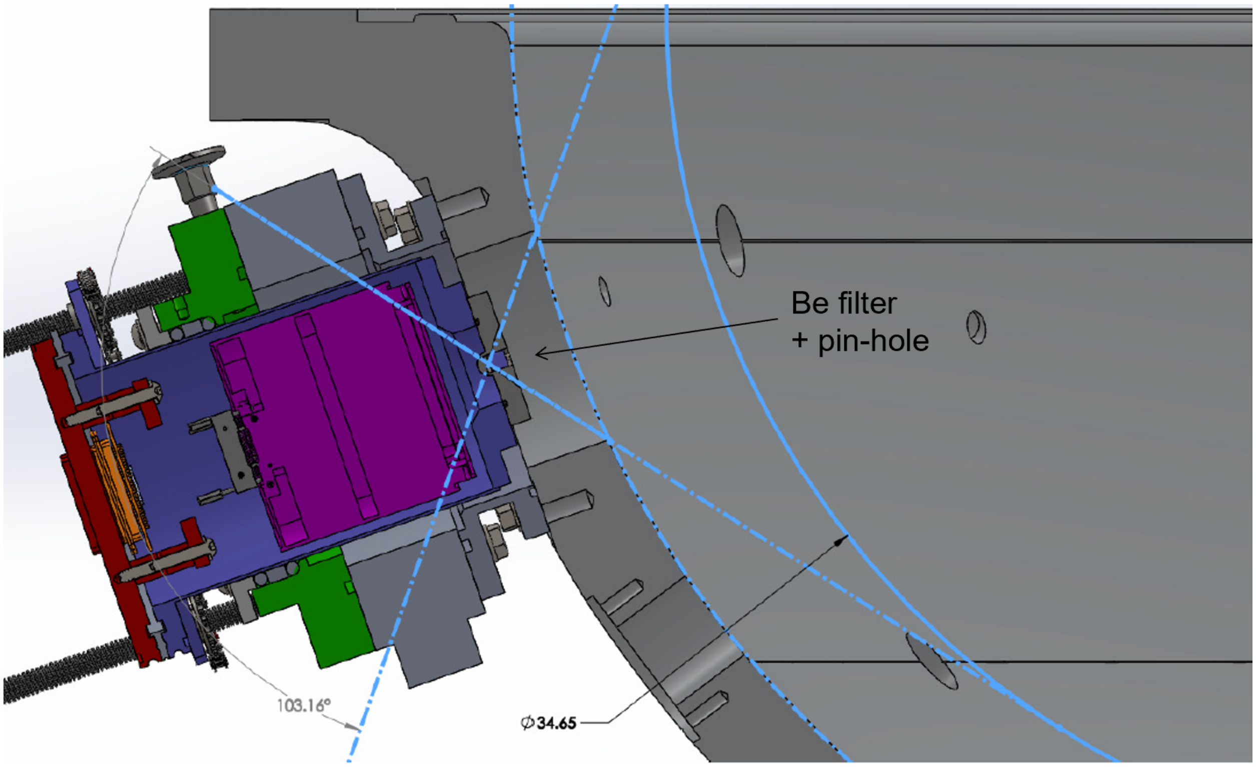

Figure 3.4: The fully assembled detector in its inserted position, looking through a gate valve with a radial view of the vacuum vessel. The dashed blue lines illustrate the wide pinhole viewing angle of $2\theta = 103.16°$.

pinhole), there is a spread of about $\pm 30°$ in the toroidal direction. However, since most MST plasmas are nearly axisymmetric this variation is not considered to be a problem. A tangential view was considered but not implemented on the MST.

The grounding scheme for the detector also required special consideration. The PILATUS3 unit was in direct electrical contact with the detector housing, which was in direct physical contact with the MST conducting shell. This meant that the detector must be on vacuum vessel (VCV) ground. Significant care was therefore required to avoid a ground loop between VCV and building ground, which includes the power outlets that the detector second stage draws its electricity from. This is illustrated in Figure 3.5, which shows a schematic I implemented when tracking down grounding issues. When a ground loop was present, it frequently resulted in data being lost due to a readout

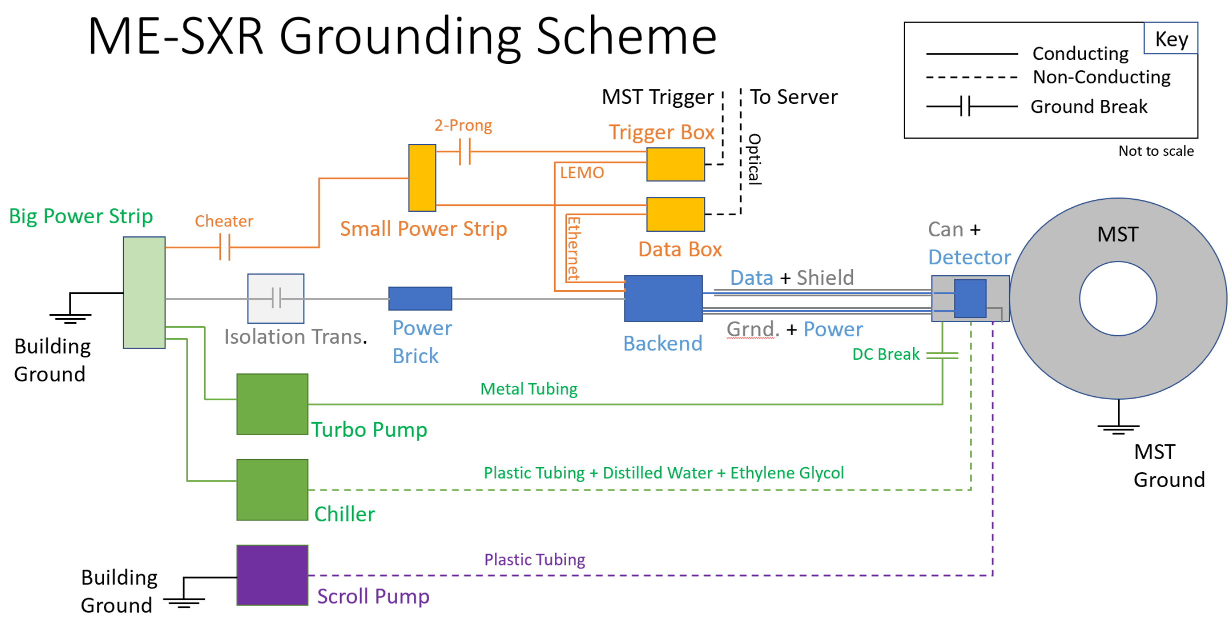

Figure 3.5: Schematic illustration of the grounding scheme used with the ME-SXR diagnostic. Solid lines indicate connections via mediums which conduct electricity (i.e., copper wires) while dashed lines indicate non-conducting connections (rubber hosing). Breaks indicate conducting connections where the ground is left floating.

error.

A fiber optic connection was used to connect the ME-SXR to the data server without creating a ground loop. Data is carried out of the PILATUS3 backend via an Ethernet cable an connected to a 10 gigabit Ethernet-to-fiber converter box where it is converted to an optical cable which is run to the server. It is important that the converter box operate at 10 gigabits; we initially installed the system with a slower converter and, as a result, the backend failed to establish a connection to the server.

The detector frontend is also connected to a chiller unit via non-conductive tubing. The chiller uses a mixture of distilled water and ethylene glycol to keep the unit at a stable operating temperature of 10° C. A vacuum feedthrough interface for the chiller lines was included on the can. The PILATUS3 frontend can also be connected to a dry nitrogen line in order to keep humidity low during up-to-air operations (such as bench testing).

During plasma operations data collection is initiated by a trigger. The PILATUS3 backend accepts a connection for a LEMO Type 00 trigger cable, which was taken from a nearby data crate in order to ensure the ME-SXR data was well-synchronized with the rest of the MST data tree. On run days the PILATUS3 detector was appropriately configured and placed into "trigger" mode, which instructs the detector to take multiple simultaneous images when the trigger is received. These images are then stored directly on the server hard disk and were then copied into the MST MDSPlus database. The detector then automatically re-arms itself and waits for the next trigger. The full run data Python routine, which can be run remotely from the control room, is included in Appendix A.

A table of various technical specifications for the diagnostic, and their associated uncertainties, is given in Table 3.1. The tolerances are derived from measurements, provided specifications, or error propagation. Entries in this table will be referenced in Chapter 4 when developing a forward model.

Parameter	**Symbol**	**Value**	**Tolerance**
Pixel size	s	$172\mu m$	$\pm 1\mu\text{m}$
Pixel Area	A_{pix}	$2.96 \times 10^{-2}\,\text{mm}^2$	$\pm 3.44 \times 10^{-4}\text{mm}^2$
Distance to pinhole	d	$30.5mm$	$\pm 1mm$
Pinhole area	A_{pin}	$2mm^2$	$\pm 0.224\,\text{mm}^2$
Be filter thickness	t_{Be}	$25\mu m$	$\pm 0.1\mu m$
Mylar filter thickness	t_{Mylar}	$96\mu m$	$\pm 5\mu m$

Table 3.1: Measured values and tolerances for various ME-SXR geometric parameters.

3.2 Energy calibration of the PILATUS3 detector

The goal of the energy calibration is to determine for each pixel the mapping between the trimbit register setting and its corresponding photon cutoff energy threshold, E_c [3]. Data for the energy calibration was collected at the DECTRIS facility in Switzerland. The detector module was exposed to a nearly-uniform x-ray source generated by fluorescence. Once the appropriate global settings were determined and exposure was taken with each pixel's trimbit value set to $\hat{t} = 0$. This exposure was then repeated with all trimbits set to $\hat{t} = 1, 2, ...63$. The data of this trimbit scan has a characteristic S-curve shape [2], as shown in Figure 3.6. This curve is well-described by the equation

$$N(\hat{t}) = \frac{1}{2}\left[\mathrm{erf}(-\frac{\hat{t}-a_0}{a_1\sqrt{2}}) + 1\right]\Big(a_2 + a_3(\hat{t}-a_0)\Big) + a_4 + a_5(\hat{t}-a_0),$$

where $\hat{t}$ is the trimbit setting, here allowed to assume non-integer values, a_0 is the location of the S-curve inflection point, a_1 is the width of the error function (corresponding to the standard deviation of the integrated Gaussian), a_2 is the signal level, a_3 is the slope of a linear distortion due to charge-sharing (CS) by adjacent pixels [6], and a_4 and a_5 describe a linear offset due to the background signal (BG). The relation between these fit terms and the S-curve shape is illustrated by FIG. 3.6.

The term a_0 describes the trimbit value which sets E_c for that specific pixel to the energy of the source photons. This value was obtained for each pixel by performing a nonlinear fit of the trimbit scan data for each pixel to Equation (3.1) using the `scipy.optimize.curve_fit` [7] function which implements a Levenberg-Marquardt based nonlinear least-squares algorithm [1]. The algorithm also provides an estimate for the uncertainty σ_{a_0} based on the diagonal component of the estimated covariance matrix. This estimate was generally found to be reliable as the covariance between the fit parameters $\{a_i\}$ was generally found to be weak.

[1] https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.curve_fit.html.

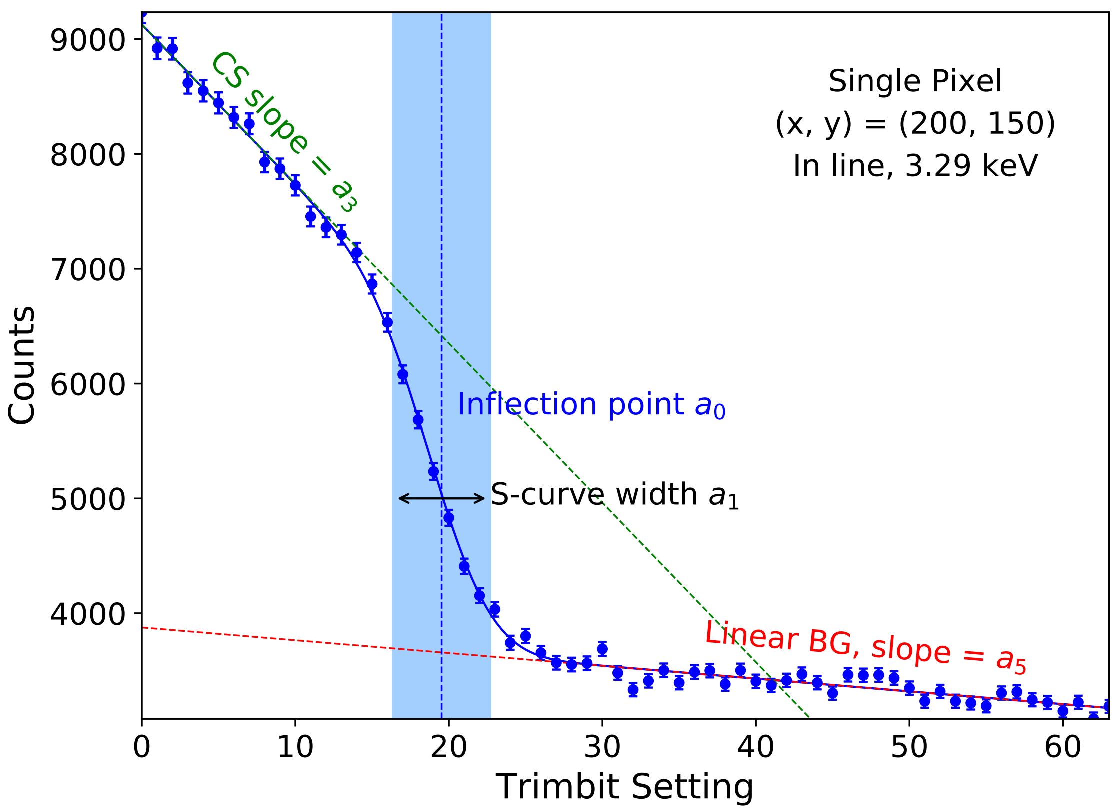

Figure 3.6: The trimbit scan calibration S-curve for a single pixel exposed to the Indium line at 3.29 keV. Key features of the curve are annotated as they relate to Equation 3.1.

A mapping between trimbit setting and cutoff energy $f : E_c \mapsto \hat{t}$ such that $\hat{t} = f(E_c)$ was generated for each pixel by repeating this procedure for multiple target sources with emission lines at different energies within the desired calibration range. This mapping is well-described by a quadratic polynomial, as given by

$$f(E_c) = c_0 + c_1 E_c + c_2 E_c^2. \tag{3.1}$$

For each pixel a fit was performed in order to allow interpolation between calibration energies. The fit was performed for each pixel using the `numpy.polyfit` function [8] which implements a standard linear least-squares technique to fit Equation 3.1 to the data $\hat{\boldsymbol{t}} = (a_0^{\mathrm{Zr}}, a_0^{\mathrm{Mo}}, a_0^{\mathrm{Ag}}, a_0^{\mathrm{In}}, a_0^{\mathrm{Ti}}, a_0^{\mathrm{V}})$ with a standard deviation given by $\boldsymbol{\sigma}_{\hat{t}}^2 = (\sigma_{a_0}^{\mathrm{Zr}}, \sigma_{a_0}^{\mathrm{Mo}}, \sigma_{a_0}^{\mathrm{Ag}}, \sigma_{a_0}^{\mathrm{In}}, \sigma_{a_0}^{\mathrm{Ti}}, \sigma_{a_0}^{\mathrm{V}})$.

The `numpy.polyfit` also returns an estimate of the covariance matrix of the fit parameters. Unlike with the previous case the parameters $\boldsymbol{c}$ are strongly correlated, meaning that an appropriate estimate of the uncertainty in the corresponding trimbit value must properly account for the covariance. It is worth taking a moment to consider this more carefully, because it turns out that naively relying on just the variance to estimate the fit uncertainty will result in substantially overestimating the uncertainty in the trimbit setting.

The covariance matrix of a set of random variables is a symmetric matrix whose diagonal elements describe the variance of each parameter and whose off-diagonal elements describe the covariance between parameters [9]. For a set of three random variables, the covariance matrix looks like

$$\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_0^2 & \sigma_0\sigma_1 & \sigma_0\sigma_2 \\ \sigma_1\sigma_0 & \sigma_1^2 & \sigma_1\sigma_2 \\ \sigma_2\sigma_0 & \sigma_2\sigma_1 & \sigma_2^2. \end{bmatrix}. \tag{3.2}$$

For our particular case, the random variables in question are the fit parameters $\boldsymbol{c} =$

(c_0, c_1, c_2). Given the inherent noisiness of the data $\hat{\boldsymbol{t}}$, as encoded in the standard deviations $\boldsymbol{\sigma}_{\hat{t}}$, if we were repeat the calibration procedure many times by taking new data and performing the fit again we would find a distribution of the values of $\boldsymbol{c}$, which we are now assuming to be normal. Ideally, the distribution of the coefficients of $\boldsymbol{c}$ each have a mean given by $\boldsymbol{\mu} = (\bar{c}_0, \bar{c}_1, \bar{c}_2)$ and are distributed according to a multivariate normal distribution,

$$p(\boldsymbol{c}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^3 \det(\boldsymbol{\Sigma})}} \exp\left(-\frac{1}{2}(\boldsymbol{c} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{c} - \boldsymbol{\mu}) \right). \tag{3.3}$$

Technically, we cannot exactly measure $\boldsymbol{\mu}$ or $\boldsymbol{\Sigma}$ from a finite set of samples. However, in practice there are a number of common methods for estimating these matrix elements from a set of samples [7]. The ultimate goal of this exercise is to characterize the variation on the mapping f due to the variation in coefficients $\boldsymbol{c}$. This variation is given by the standard error propagation formula [9],

$$\sigma_f^2(E_c) = \sum_{i=0}^{2} \sum_{j=0}^{2} \frac{\partial f}{\partial c_i} \frac{\partial f}{\partial c_j} \sigma_i \sigma_j \tag{3.4}$$

$$= S_{\text{var}}^2(E_c) + S_{\text{cov}}^2(E_c), \tag{3.5}$$

which is exact for the case of a linear fit and a first-order approximation in the case of a nonlinear fit. Notice that in Equation 3.5 we have organized the result into two terms, one depending on the variance,

$$S_{\text{var}}^2(E_c) = \left(\frac{\partial f}{\partial c_0} \sigma_0 \right)^2 + \left(\frac{\partial f}{\partial c_1} \sigma_1 \right)^2 + \left(\frac{\partial f}{\partial c_2} \sigma_2 \right)^2 \tag{3.6}$$

$$= \sigma_0^2 + E_c^2 \sigma_1^2 + E_C^4 \sigma_2^2, \tag{3.7}$$

and the other on the covariance,

$$S^2_{\text{cov}}(E_c) = 2\frac{\partial f}{\partial c_0}\frac{\partial f}{\partial c_1}\sigma_0\sigma_1 + 2\frac{\partial f}{\partial c_0}\frac{\partial f}{\partial c_2}\sigma_0\sigma_2 + 2\frac{\partial f}{\partial c_1}\frac{\partial f}{\partial c_2}\sigma_1\sigma_2 \tag{3.8}$$

$$= 2E_c\sigma_0\sigma_1 + 2E_c^2\sigma_0\sigma_2 + 2E_C^3\sigma_1\sigma_2. \tag{3.9}$$

This is significant because for this application it is always the case that $S_{var} > 0$ and $S_{cov} < 0$, meaning that ignoring the off-diagonal terms in the covariance matrix will result in a significant over-estimate of the uncertainty on the chosen trimbit. Considering this, the appropriate trimbit setting in order to achieve of threshold of E_c is determined by $\hat{t} = f(E_c) \pm \sigma_f(E_c)$.

An inverse mapping, $g : \hat{t} \mapsto E_c$, is also useful. This can be readily obtained by applying the quadratic formula to Equation 3.1, giving

$$g(\hat{t}) = \frac{1}{2c_2}\left[\sqrt{c_1^2 - 4c_0(c_2 - \hat{t})} - c_1\right]. \tag{3.10}$$

One could repeat the process of Equation 3.7 and compute the various partial derivatives of $g(\hat{t})$ in order to determine $\sigma_g(\hat{t})$. However, we can use the requirement that this result should be consistent with the forward mapping f to simplify the calculation. We assert that, by definition, the uncertainty of the output σ_g is related to the uncertainty on the input $\sigma_{\hat{t}}$, which we can identify with $\sigma_{\hat{t}} = \sigma_f\big(g(\hat{t})\big)$, giving

$$g(\hat{t} + \sigma_{\hat{t}}) = g(\hat{t}) + \sigma_g(\hat{t}). \tag{3.11}$$

As we can see in Figure 3.8 it is generally safe to assume that $\sigma_{\hat{t}} << 1$. So we can expand Equation 3.11 to first order and simplify,

$$\sigma_g(\hat{t}) = g(\hat{t}) + \sigma_{\hat{t}}\frac{dg}{d\hat{t}} + \mathcal{O}(\sigma_{\hat{t}}^2) - g(\hat{t}) \tag{3.12}$$

$$\approx \frac{dg}{d\hat{t}}\sigma_{\hat{t}}, \tag{3.13}$$

where $\partial g/\partial\hat{t}$ is readily found by differentiating Equation 3.1 with $E_c = g(\hat{t})$:

$$\frac{dg}{d\hat{t}} = \frac{1}{c_1 + 2c_2 E_c(\hat{t})}. \tag{3.14}$$

This gives a final expression for computing σ_g for arbitrary $\hat{t}$:

$$\sigma_g(\hat{t}) = \frac{\sigma_f^2\big(g(\hat{t})\big)}{c_1 + 2c_2 E_c(\hat{t})}. \tag{3.15}$$

Since a pixel's trimbit may only take on an integer value, when setting a desired threshold value the result of Equation 3.1 must be rounded to the nearest integer, $\hat{T} = [f(E_c)]$, where $[x]$ denotes rounding to the nearest integer. This will cause the actual set threshold to deviate somewhat from the desired value. Making use of both mappings f and g, we can calculate the magnitude of this deviation from the chosen threshold E_c:

$$\Delta E(E_c) = g\big([f(E_c)]\big). \tag{3.16}$$

The statistical analysis of this deviation over the entire detector is the focus of Section 3.2.3.

3.2.1 High-gain calibration results: 2-to-7 keV (lowE)

The detector global settings were configured with a threshold under 2 keV up to just more than 6 keV, and a calibration was performed following the procedure outlined above using fluorescent Zr, Mo, Ag, In, Ti, and V targets with energies ranging between 2-5 keV (see Figure 3.7). This range is of interest on the MST in order to diagnose the strong Al^{+11} and Al^{+12} lines which are observed between 1.6 and 2 keV [10], as well as to provide continuum T_e measurements.

The calibration procedure described in Section 3.2 was performed individually for each of the $\sim$ 100k pixels on the detector. The resulting S-curves for an example pixel are shown in Figure 3.7, with the linear background removed for the purpose of illustration. The detector counts have also been normalized so that the response is equal to one when threshold is at half of the incident photon energy. The fit values a_0 are then used to generate a trimbit-E_c curve as shown in Figure 3.8. These mappings allow for the implementation of custom E_c configurations within this sensitivity range.

The calibration data was sufficient to well-characterize the mapping, as demonstrated by the small region of uncertainty surrounding the fit values in Figure 3.8. It is notable that the uncertainty in the trimbit-E_c mapping for any given individual pixel is much smaller than the variation between pixels across the detector.

3.2.2 Pixel-to-pixel variation of the calibration

Substantial correlated variation in the results of the energy calibration was observed across the detector's $\sim$ 100k pixels. This variation can be seen in Figure 3.9, which shows the trimbit-E_c mappings for 500 randomly-selected pixels, demonstrating a variation of the order of 10 trimbits to achieve the same energy threshold. The standard deviation of the inflection point a_0 was found to be significantly larger for the Zr target, whose emission comes from a pair of narrowly-spaced lines at approximately 2.04 keV. This is because the S-curve threshold for this target was near the lower-limit of the detectors sensitivity, meaning that the algorithm sometimes struggled to determine where the S-curve flattened out at the top. This uncertainty was accounted for in the fit with larger error-bars.

The quality of the trimbit-E_c fit was explored by calculating $\chi^2 = \sum_s (a_{0,s} - \hat{t}(E_s))^2/\sigma_s^2$ for each pixel, where s labels each x-ray source, E_s is the characteristic line energy of that

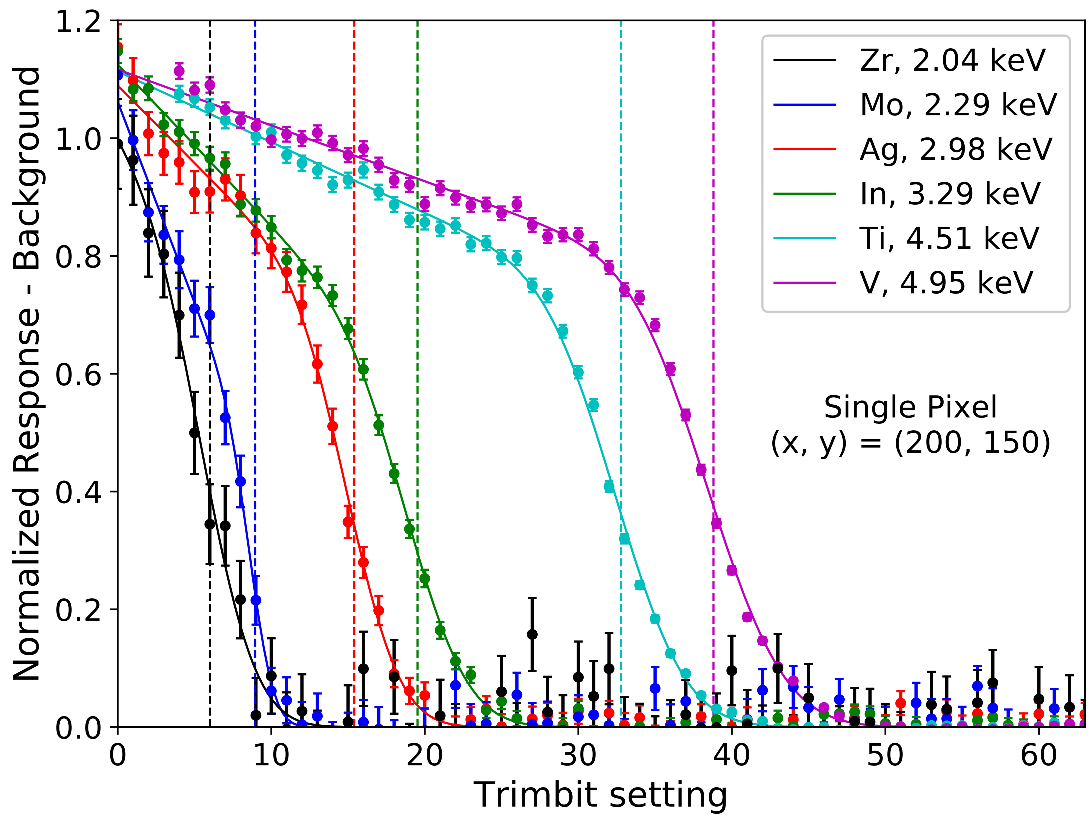

Figure 3.7: All trimbit scan calibration S-curves for a single pixel. For this plot the linear background was subtracted off and the signal level was normalized. The dashed lines indicate the location of the inflection points a_0. Uncertainties in the individual counts were assumed to follow Poisson statistics.

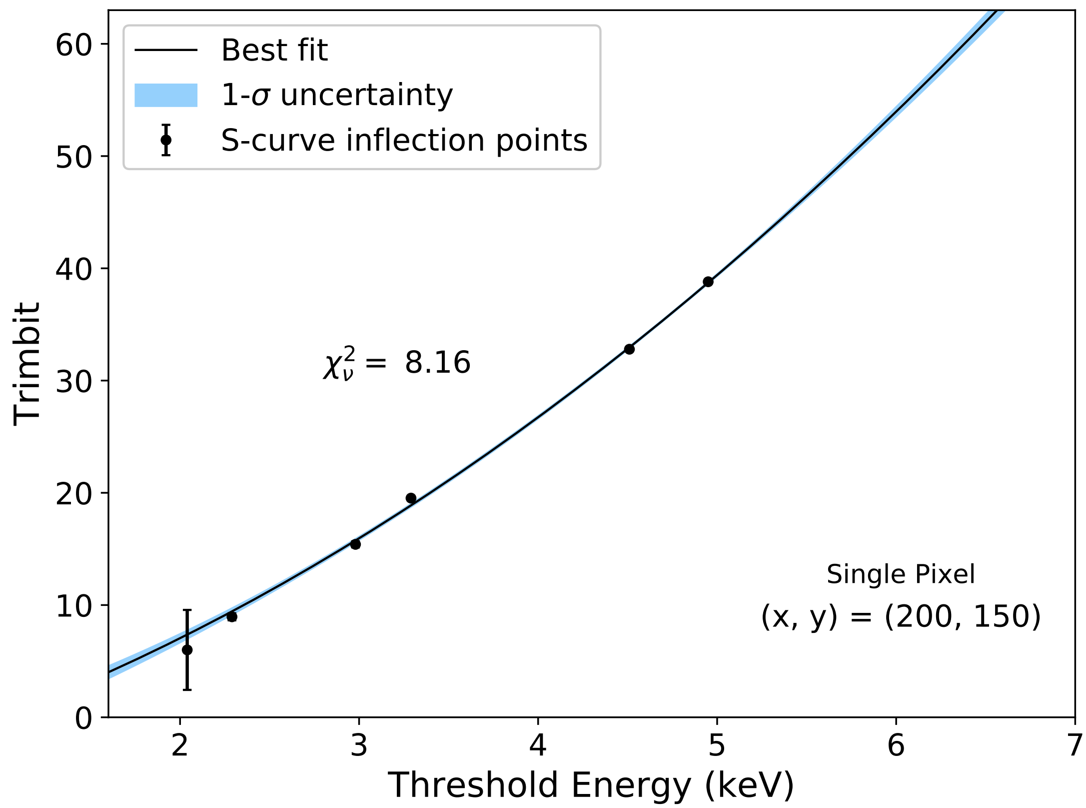

Figure 3.8: Single pixel inflection point a_0 data fit to a quadratic curve with the resulting 1σ uncertainty region highlighted. The uncertainty in the a_0 values, indicated by the black error bars, was taken from the variance in the S-curve fits. The shape of the fit and level of uncertainty is typical of all pixels across the detector.

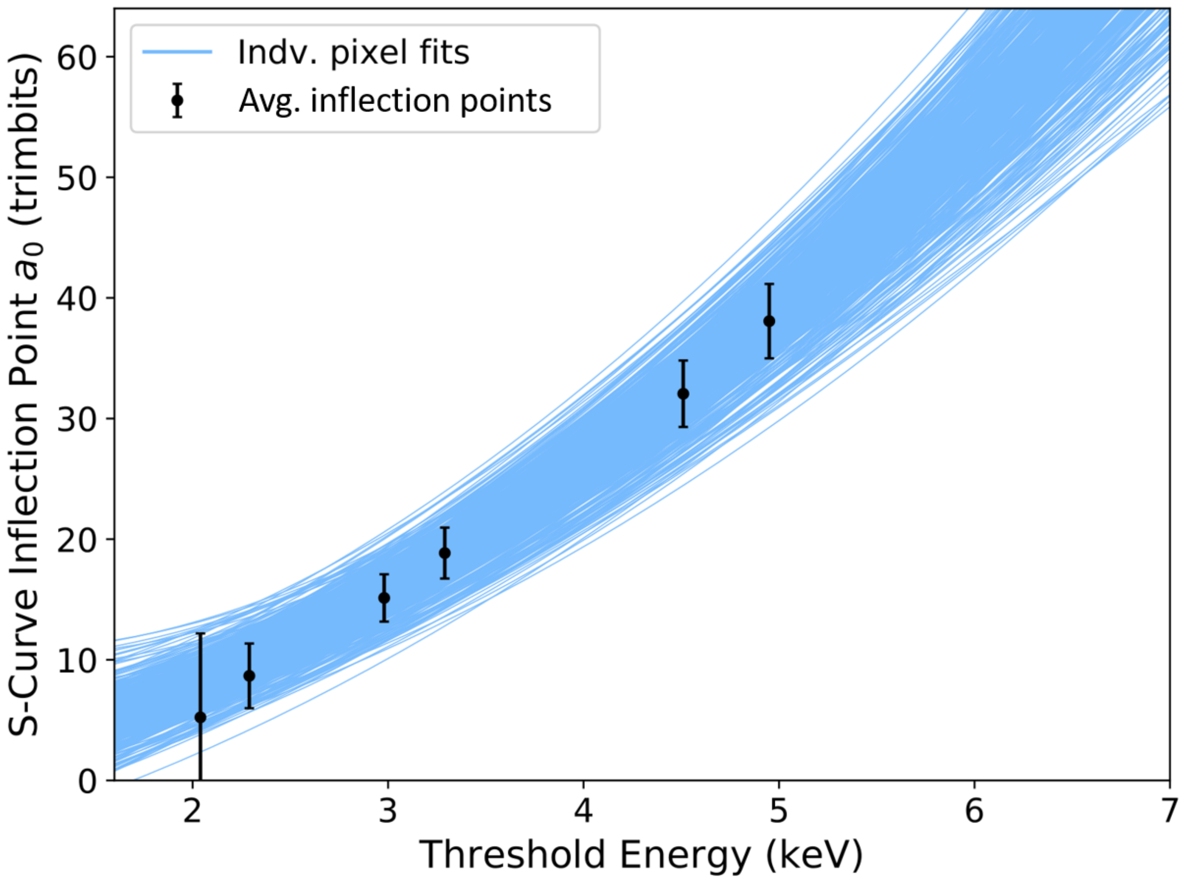

Figure 3.9: The blue lines show the trimbit-E_c mapping for 500 randomly-selected pixels, showing a spread of $\sim$ 10 trimbits for a given threshold. The black points show the average trimbit setting of the S-curve inflection point for each x-ray source, and the error bars show the standard deviation.

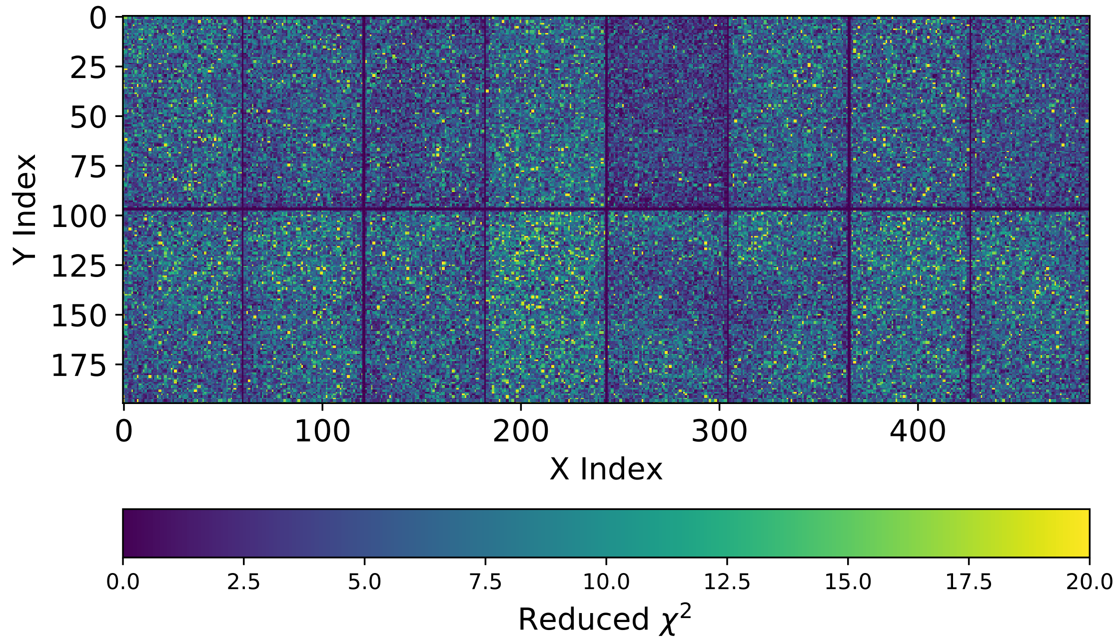

Figure 3.10: Map of the reduced χ^2 for the quadratic trimbit-E_c fit for each pixel. The columns and rows between the rectangular chips do not collect data and have thus been zeroed out.

source, and $\hat{t}(E_c)$ is the best-fit quadratic trimbit-E_c mapping. As shown in Figure 3.10, this was found to be relatively uniform across the detector, though some correlation between pixels on the same ASIC can be observed.

A series of trimbit configurations was produced which each set a uniform threshold across the detector ranging from 1.8 keV to 6.3 keV in order to investigate how the variation between pixels affects energy resolution. For each energy the trimbit-E_c calibration was used to determine the exact trimbit value which would be required to set each pixel to that particular energy threshold, permitting non-integer values. The distribution of

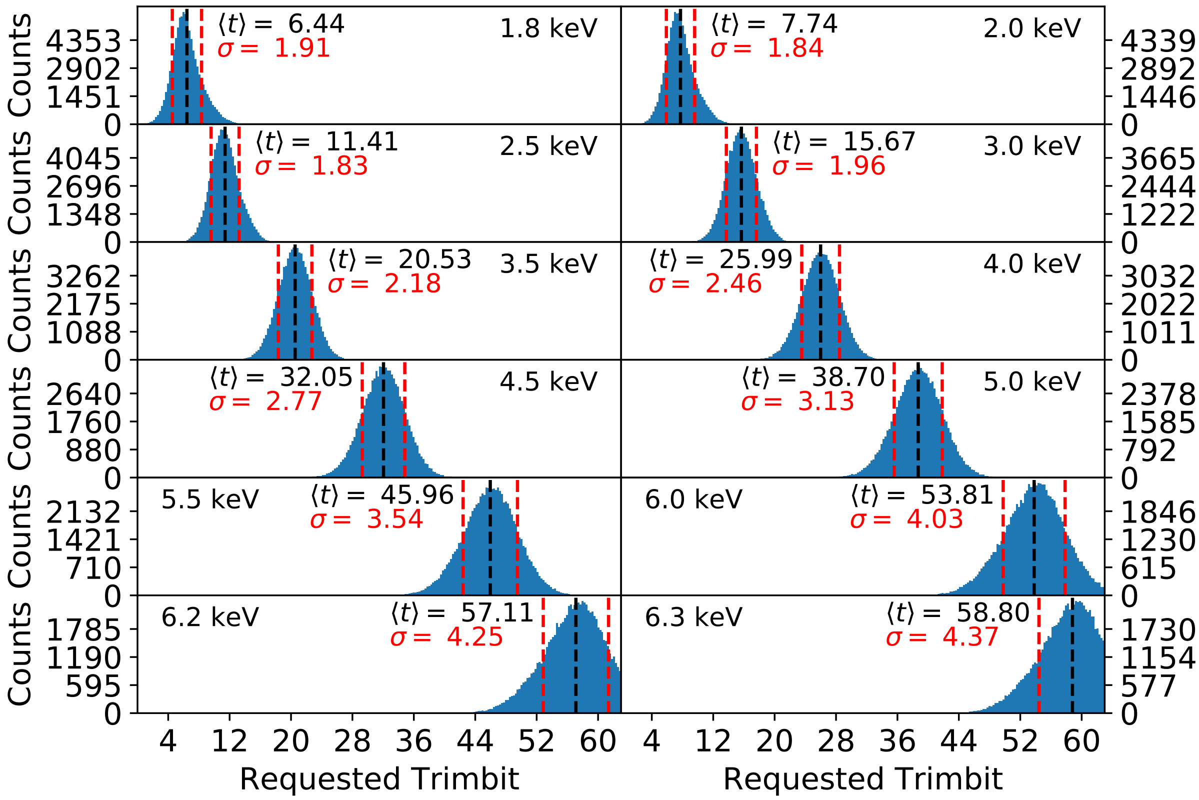

Figure 3.11: Distribution of trimbit settings required to set the detector to the specified uniform threshold. These values must be rounded to the nearest integer before the detector can be configured.

these trimbit settings, shown in Figure 3.11, is well-described as normal. The distribution is seen to widen as the desired threshold energy is increased, with a standard deviation varying from less than 2 trimbits to more than 4. This plot also shows the limits of this calibration's energy range, as threshold settings above 6.0 keV result in an appreciable number of pixels requiring a trimbit setting above the hardware limit of 63.

3.2.3 Effect of pixel-to-pixel variation on energy resolution

For an appropriate ME-configuration of the detector the trimbit settings obtained from the calibration must be rounded to the nearest integer value. As a result of this rounding all pixels in a particular row or column will actually be set to slightly different thresholds within the range $E_c \pm \Delta E$. The value ΔE therefore serves as the limitation on the resolution of the detector under this calibration.

The value of ΔE was determined by taking the pixel trimbit settings for each of the uniform configurations as shown in Figure 3.11, rounding them to the nearest integer value, and then mapping the resultant integer back to its corresponding energy threshold using the calibrated trimbit-E_c mapping, of the type shown in FIG. 3.8. The resulting distributions, shown in Figure 3.12 are nearly uniform with sloped edges. This uniformity is a result of the fact that the variation between pixels is arbitrary and no particular trimbit-E_c mapping is preferred. The slopes at the edge of the distribution are a result of uncertainty in the calibration procedure. The largest values of ΔE are seen at low specified E_c, where the trimbit-E_c mapping is the least steep. For all energies within the range of this calibration the threshold can be set with a resolution of less than 100 eV.

Since this analysis depends on the assumption that the trimbit-E_c mappings produced by the calibration procedure are essentially correct, it is worthwhile to quantify the level of uncertainty in the calibration results at each considered energy threshold. This was determined by making a histogram of the calibration uncertainty associated with each pixel. For the pixel shown in Figure 3.8, for instance, this is the size of the blue shaded region evaluated at the appropriate trimbit. The results, displayed in Figure 3.13, show an uncertainty of about the same size as ΔE for low threshold settings but that becomes minimal by a threshold of about 3.0 keV. This explains the sloping feature on the edges of the low-threshold uniform distributions in Figure 3.12.

Another consideration for describing the resolution of the detector is the width of the S-curve, described by the parameter a_1 in Equation 3.1. This parameter characterizes

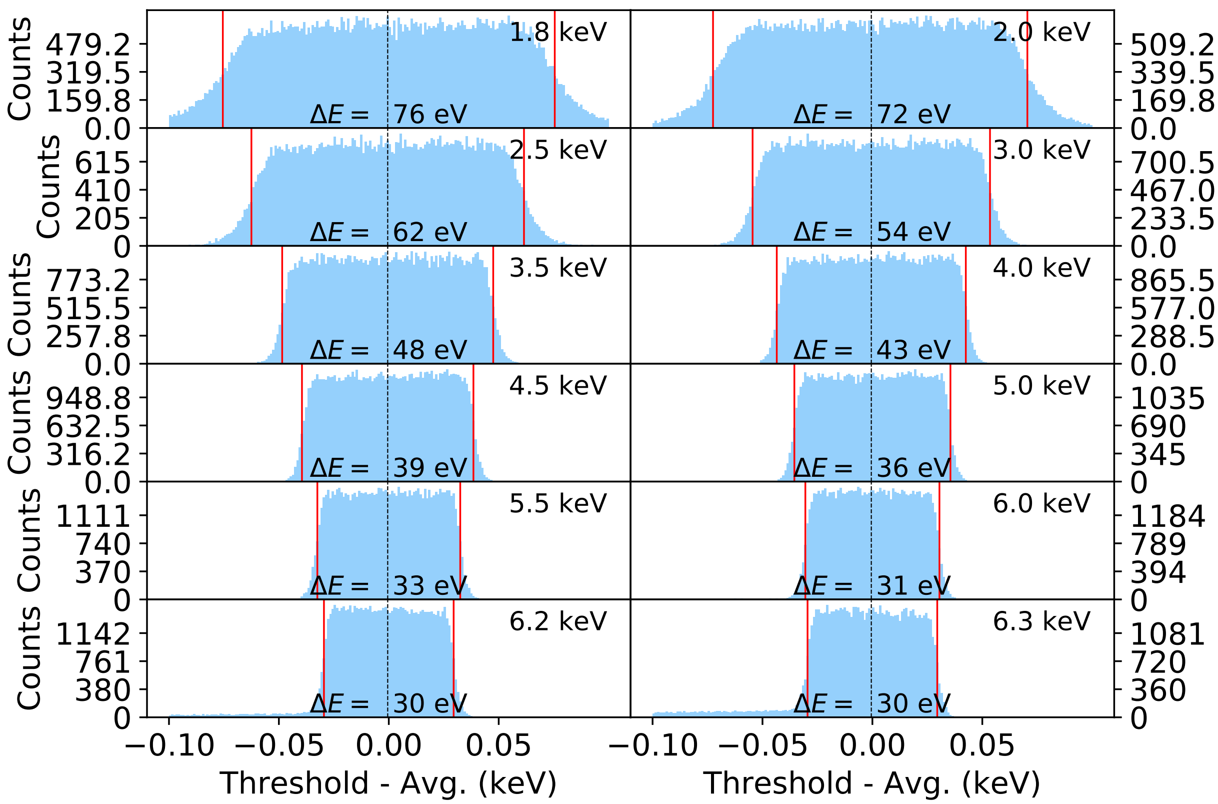

Figure 3.12: Distribution of energy thresholds $\Delta E - \langle \Delta E \rangle$ across all pixels given a common target threshold resulting from discrete trimbit settings. The distribution is observed to be nearly uniform, with ΔE ranging from 30-76 eV depending on the requested threshold. This represents a variation of $\sim 1 - 2\%$ of the threshold value.

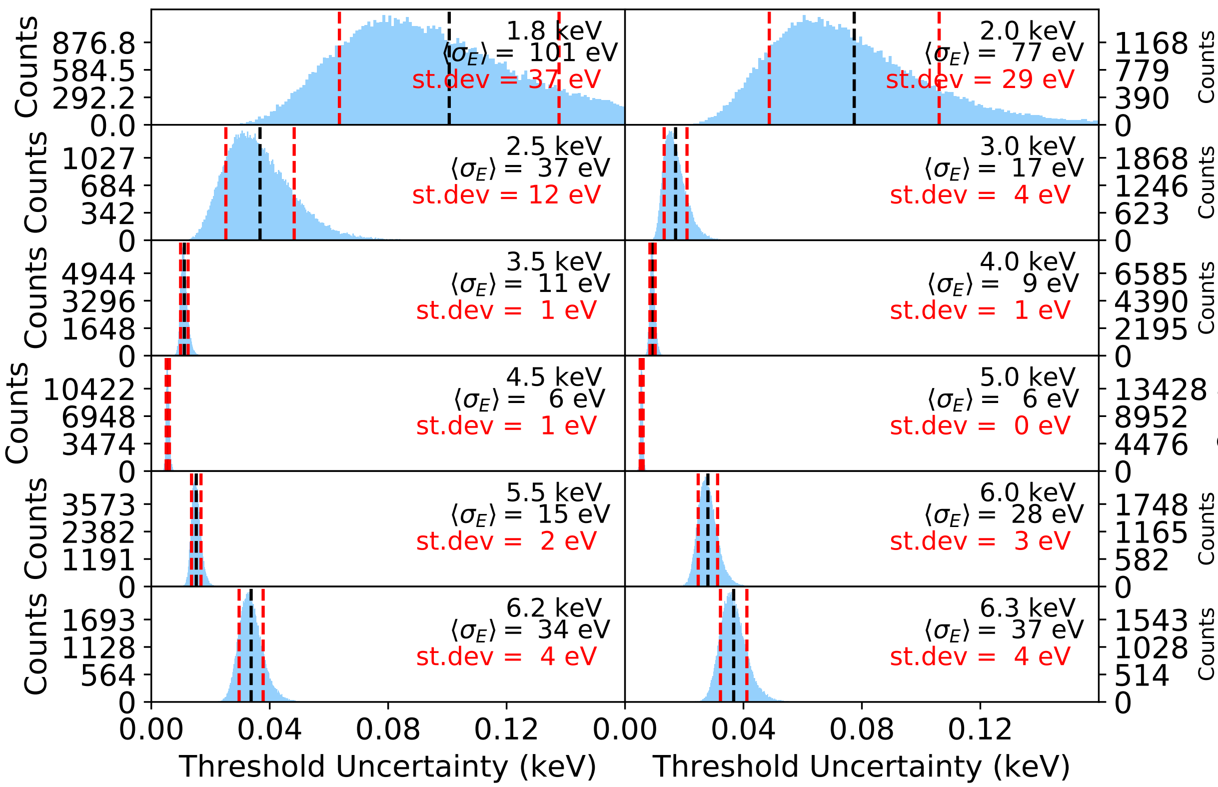

Figure 3.13: Uncertainty in the threshold energy resulting from Poisson uncertainty propagating through the calibration procedure. This uncertainty is seen to be substantially smaller than the threshold variation ΔE due to rounding for threshold settings of ~ 3 keV and larger. Note that since the distribution is somewhat skewed toward high uncertainty the mean and mode of a given distribution does not exactly agree.

the width of the region between which the detector transmission, also described by an S-curve [1], increases from 0 to 1 so that any photons within this energy range have a fractional chance of being counted. Figure 3.14 shows the average width for all pixels in the data set at each of the calibration line energies. Here the width is presented in terms of a full width at half maximum, $FWHM = 2\sqrt{2\ln 2} \cdot a_1 \cdot \frac{\partial E_c}{\partial \hat{t}}$. This plot also shows the S-curve width for 1000 randomly selected pixels. A higher variance in the FWHM is seen in the lower-energy datasets, attributable to the difficulty in performing the S-curve fits when the inflection point is near the detector's lower-energy limit. The S-curve width was found to be independent of the energy threshold on average, with a $FWHM \approx 0.7$ keV. This does vary between pixels, and it was observed that a_0 and a_1 tend to be positively correlated.

3.2.4 Medium-gain calibration results: 4-to-14 keV (midE)

This calibration procedure was also performed with global settings chosen for threshold ranges from 4 keV to above 12 keV. This calibration used emission lines from fluoresced Cr, Fe, Cu, Ge, and Br sources at 5.41, 6.40, 8.05, 9.89, and 11.92 keV respectively.

The level of pixel-to-pixel variation in the trimbit-E_c mapping was found to be consistent with that seen in the 1.6-to-6 keV calibration. ΔE ranges from approximately 70-200 eV, an increase of about 2.5 times over the 1.6-to-6 keV calibration. This is consistent with the overall increase energy range covered by the calibration. The uncertainty in the threshold energy resulting from propagated counting error was found to be smaller than ΔE by a factor of 3 or more (depending on the threshold chosen). The scalings of ΔE with threshold setting for both calibrations are shown in Figure 3.15. Quadratic fits are also provided for interpolation. S-curves were found to have a width $FWHM \approx 1.3$ keV which was independent of threshold energy.

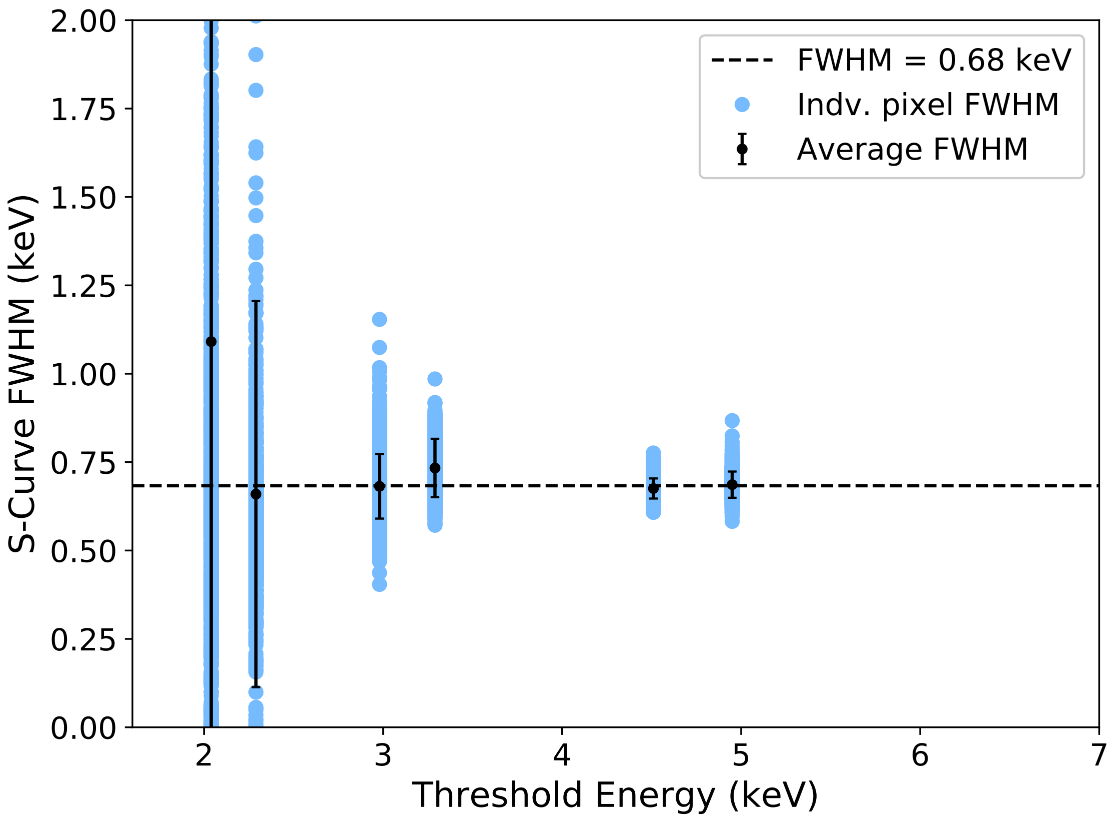

Figure 3.14: On average, the width of the S-curves for the 1.6-to-6 keV calibration is well described as a constant FWHM ~ 0.7 keV. The light blue points represent the width for 1000 randomly selected pixels. The black points are the average of the entire detector, and the error bars are the standard deviations.

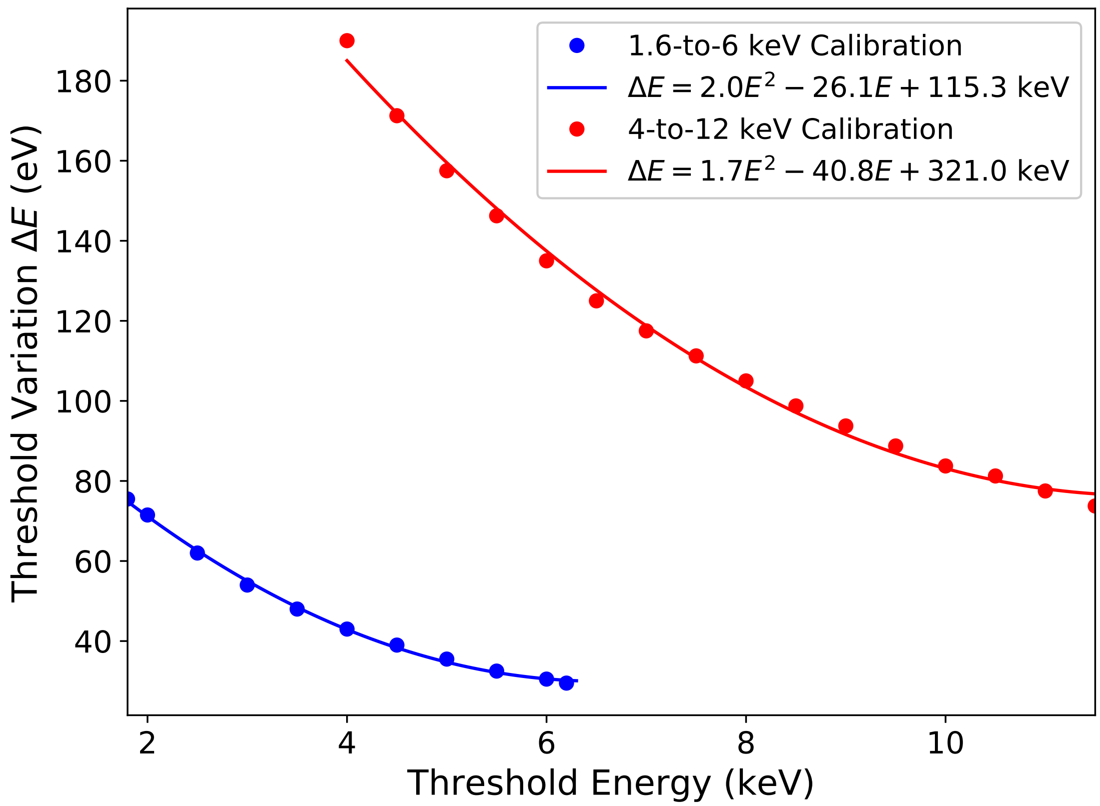

Figure 3.15: Threshold variation ΔE vs threshold energy for both the 1.6-to-6 and 4-to-14 keV calibrations. Points were calculated using uniform threshold configurations (as in FIG. 3.12) and interpolated via quadratic fit.

3.3 MST *in-situ* spatial calibration

The goal of the spatial calibration presented in this section is to understand where the PILATUS3 detector is in physical space, and what portion of the plasma volume each pixel can see through the pinhole. This is just as essential to understanding the resulting data as the energy calibration presented earlier in this section. This is done by imaging a small radiation source at multiple known vertical positions. Because MST plasmas are nearly toroidally-symmetric, the calibration was only performed in the vertical direction.

The calibration was performed *in-situ* by placing an Fe-55 source into a custom holder on the end of an insertable probe. This probe was then lowered into the MST by small increments, taking a detector image at each position. The insertion depth was measured outside of the vacuum vessel using a measuring stick mounted onto the probe housing, making sure to account for known offsets including the thickness of the MST wall, the port height, and the size of holder. The insertion depth is given as the quantity Z and is equal to zero when the probe is fully retracted to the wall. This geometry is illustrated in Figure 3.16.

The calibration was performed by taking images with the PILATUS3 for 24 different source insertion depths. The scan began with the probe fully extracted and moved downward in steps of 10 cm. Due to the relatively low activity of the Fe-55 source, exposure periods ranged from 20-60 minutes, depending on the insertion depth, in order to gather sufficient numbers of photons. At the bottom of the MST the position was shifted by 5 cm and and subsequent measurements were taken in steps of 10 cm upwards. This interleaving of measurements helps to reduce systematic errors resulting from the initial probe positioning. This results in an overall spatial resolution of 5 cm, although additional measurements were taken in the core. Images taken from three Z positions are shown in Figure 3.17. Signal from the Fe-55 source is well-localized to a small bar moving up and down within the center of the image (marked in the figures by a red region). All signal outside of this region is due to background noise (such as

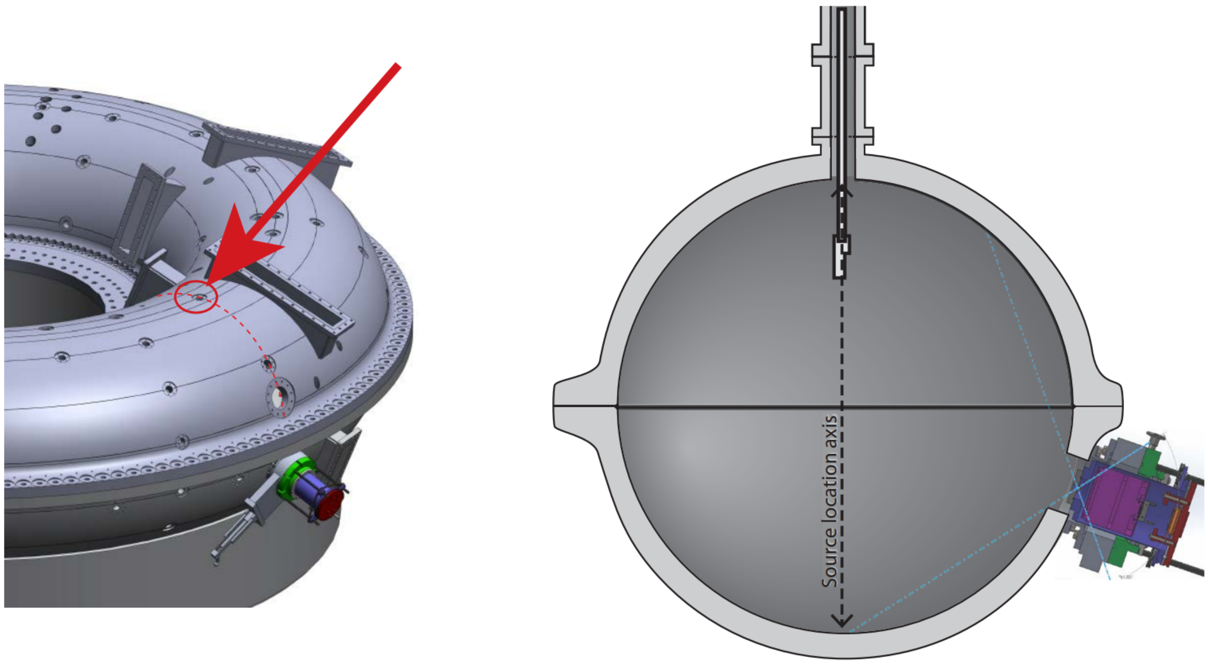

Figure 3.16: Illustration of the spatial calibration probe probe housing (left) and insertion geometry (right). Figure courtesy of L. Reusch.

cosmic rays) resulting from the long exposure times.

For each Z, the image files were used to determine which pixels could directly see the source through the pinhole. Since only the vertical displacement is of interest, the signal was summed over the horizontal direction to produce a 1D pulse. Signal outside of the core "track" (the red region in Figure 3.17) were discarded. The resulting pulse was then fit to a "tophat" function with rounded corners, as shown in Figure 3.18 (a). The "tophat" fitting function is given by a sum of logistic functions of the form

$$N(X) = N_0 + A \cdot \left[f_k(X - X_L) - f_k(X - X_R) \right], \tag{3.17}$$

where x is the pixel index and

$$f_k(X - X_{L/R}) = \frac{1}{1 + e^{-k(X - X_{L/R})}}, \tag{3.18}$$

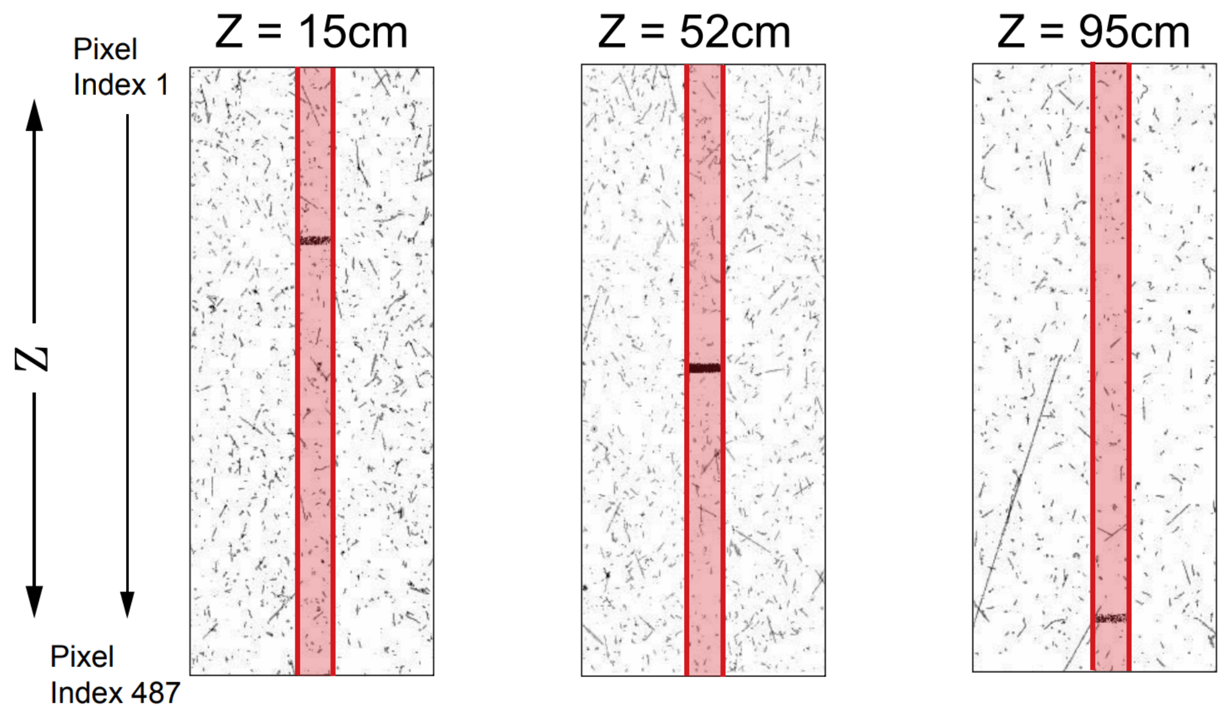

Figure 3.17: Raw spatial calibration data taken at three separate Z locations. The calibration signal moves vertically within the red highlighted region as Z is adjusted. Signal outside of the red region is assumed to be background noise. Figure courtesy of Dr. L. Reusch.

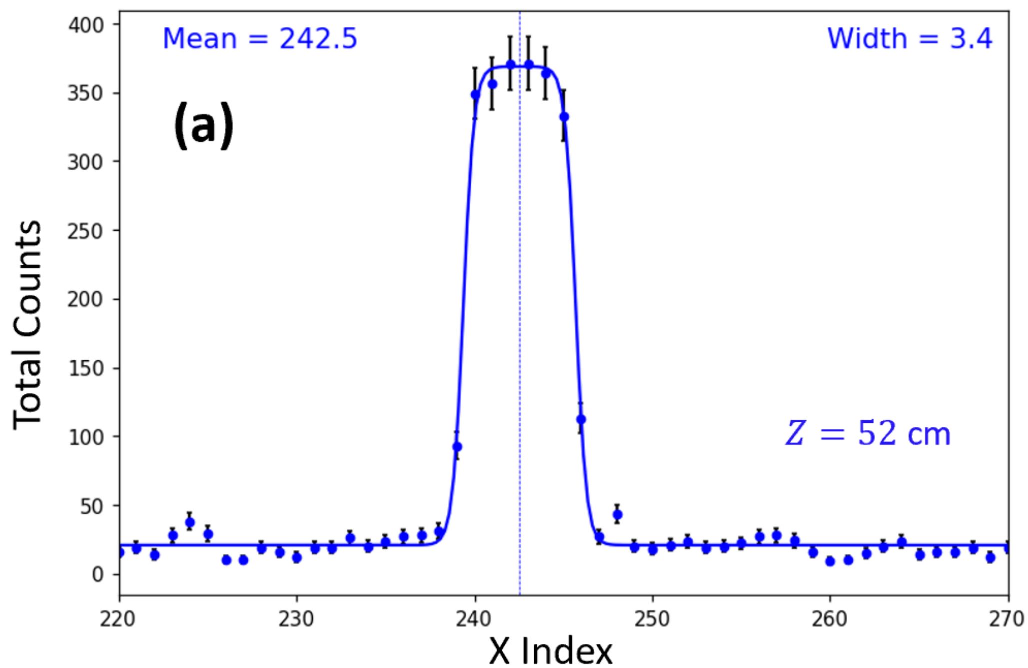

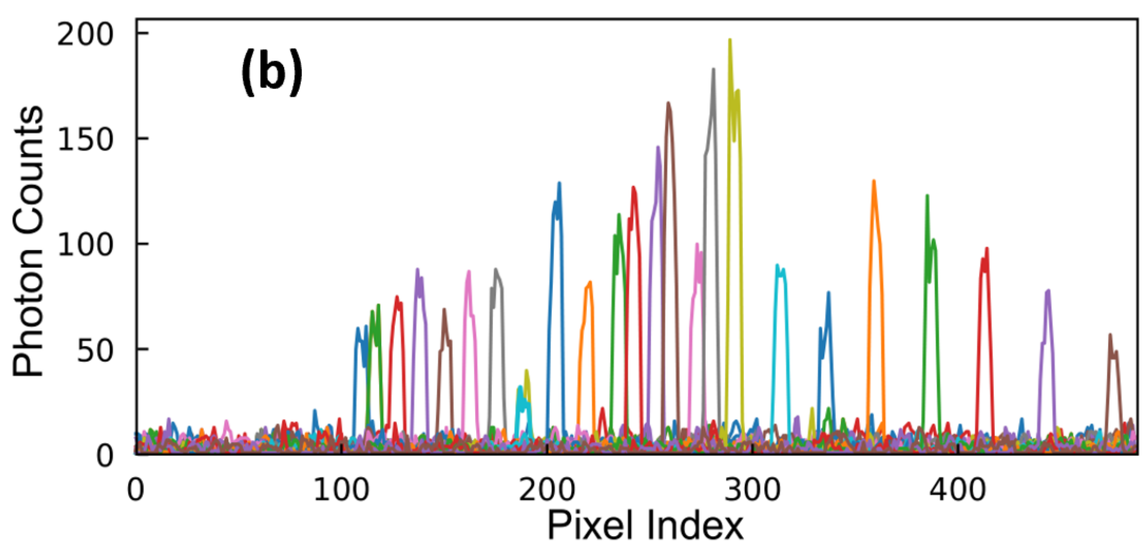

Figure 3.18: (a) Summed at a from a single Z position is fit to a tophat function. (b) Summed data for all Z locations used in the calibration.

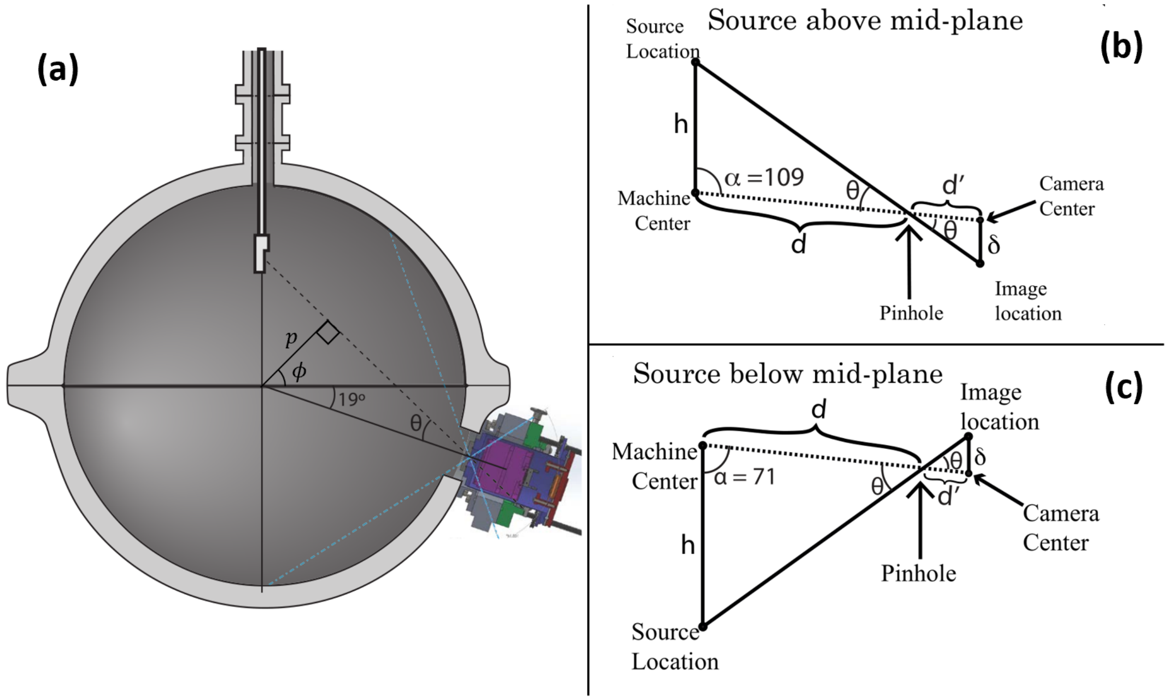

Figure 3.19: Spatial calibration impact parameter geometry: (a) relation between the source angle θ and impact angle ϕ; (b) geometry when the source is above mid-plane; (c) geometry when the source is below mid-plane. Figure courtesy of L. Reusch.

where k is the logistic "width" (or steepness) of the edge in number of pixels and $X_{L/R}$ are respectively the left and right midpoints of the step. The midpoint of the pulse is taken to be the average of X_L and X_R, $\bar{X} = (X_L + X_R)/2$. The combined set of pulses for every Z location in the scan is shown in Figure 3.18 (b), with each color corresponding to a different camera exposure/Z location.

By considering the geometry of the source-detector system we can determine the impact parameter and angle characterizing each line of sight. The basic geometry is illustrated in Figure 3.19, considering separately the cases when the source is above or below mid-plane. The image location on the face of the detector is given by

$$\delta = \frac{d' h \sin\alpha}{d - h\cos\alpha}, \tag{3.19}$$

where h is the vertical location of the source relative to MST, d is the distance from the machine center to the pinhole, d' is the distance from the pinhole to the detector face, and α is the interior angle between the vertical diameter and the detector normal. Both h and α are a set by the source height Z:

$$h = Z - 52.0\,\mathrm{cm} \tag{3.20}$$

$$\alpha = \begin{cases} (90^\circ + 19^\circ = 109^\circ) & h \geq 0 \\ (90^\circ - 19^\circ = 71^\circ) & h < 0. \end{cases} \tag{3.21}$$

We can convert from δ (measured in cm) to pixel index X by

$$X = 243.5 - \frac{\delta}{0.0172\,\mathrm{cm}}. \tag{3.22}$$

The goal of this exercise is to determine the best-fit values for the unknown parameters d and d'. From there it is straightforward to characterize the viewing chord of each pixel via (p, φ).

The impact parameter p is the distance from the geometric axis to the viewing chord along a line which is normal to the viewing chord. The impact angle φ is the angle between the line p and horizontal axis, with $\varphi = 0$ at the outboard mid-plane. These two parameters, which are also shown in Figure 3.19, are sufficient to parameterize an arbitrary viewing chord. These parameters are related to the fit parameters via an intermediary angle θ, given by

$$\theta = \tan^{-1}(d/d') \tag{3.23}$$

$$p = d\sin(\theta) \tag{3.24}$$

$$\varphi = \alpha - \theta. \tag{3.25}$$

Using the data collected during the Z scan, an optimal fit was found with parameters given by $d = 55.2$ cm and $d' = 3.47$ cm. This fit is shown in Figure 3.20. As seen in the figure, the data is well-described by this geometric model. The results of this fit were used to calculate p and θ (and hence φ) for each pixel, as shown in Figure 3.21. The use of a geometrical model allows reliable extrapolation beyond the region that was directly sampled sampled during the spatial scan. This mapping between X and p is used to characterize nearly all ME-SXR data presented in this thesis.

The results of the spatial calibration are illustrated in a different form in Figure 3.22, which shows ME-SXR chords color-coded by impact parameter. As can be seen, positive p looks toward the top of the vacuum vessel and slightly outboard, while negative p looks toward the bottom of the vacuum vessel and slightly inboard. In order to ensure reproducibility of the detector insertion position (and therefore the validity of the spatial calibration), markings were placed on the detector housing.

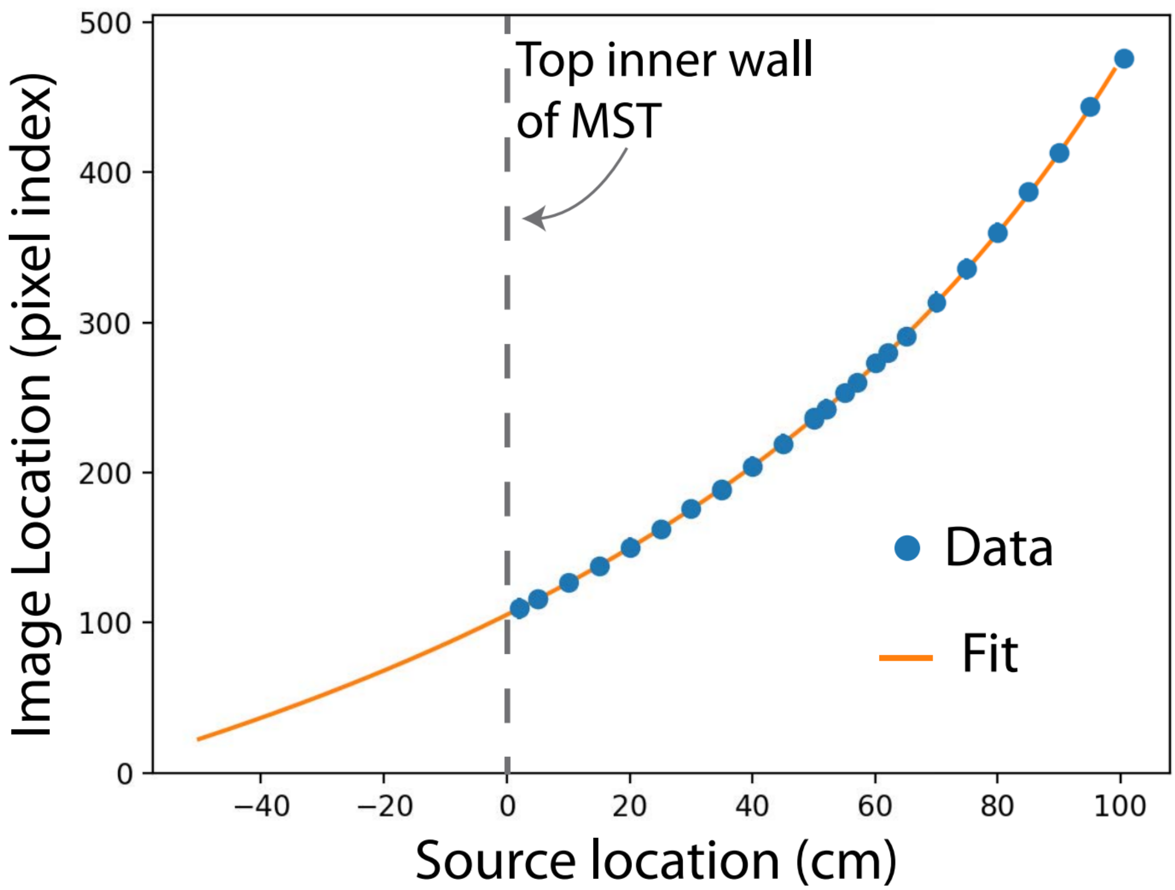

Figure 3.20: Fit between Equation 3.19 and spatial calibration data. The fit has optimal parameters $d = 55.2$ cm and $d' = 3.47$ cm. Figure courtesy of L. Reusch.

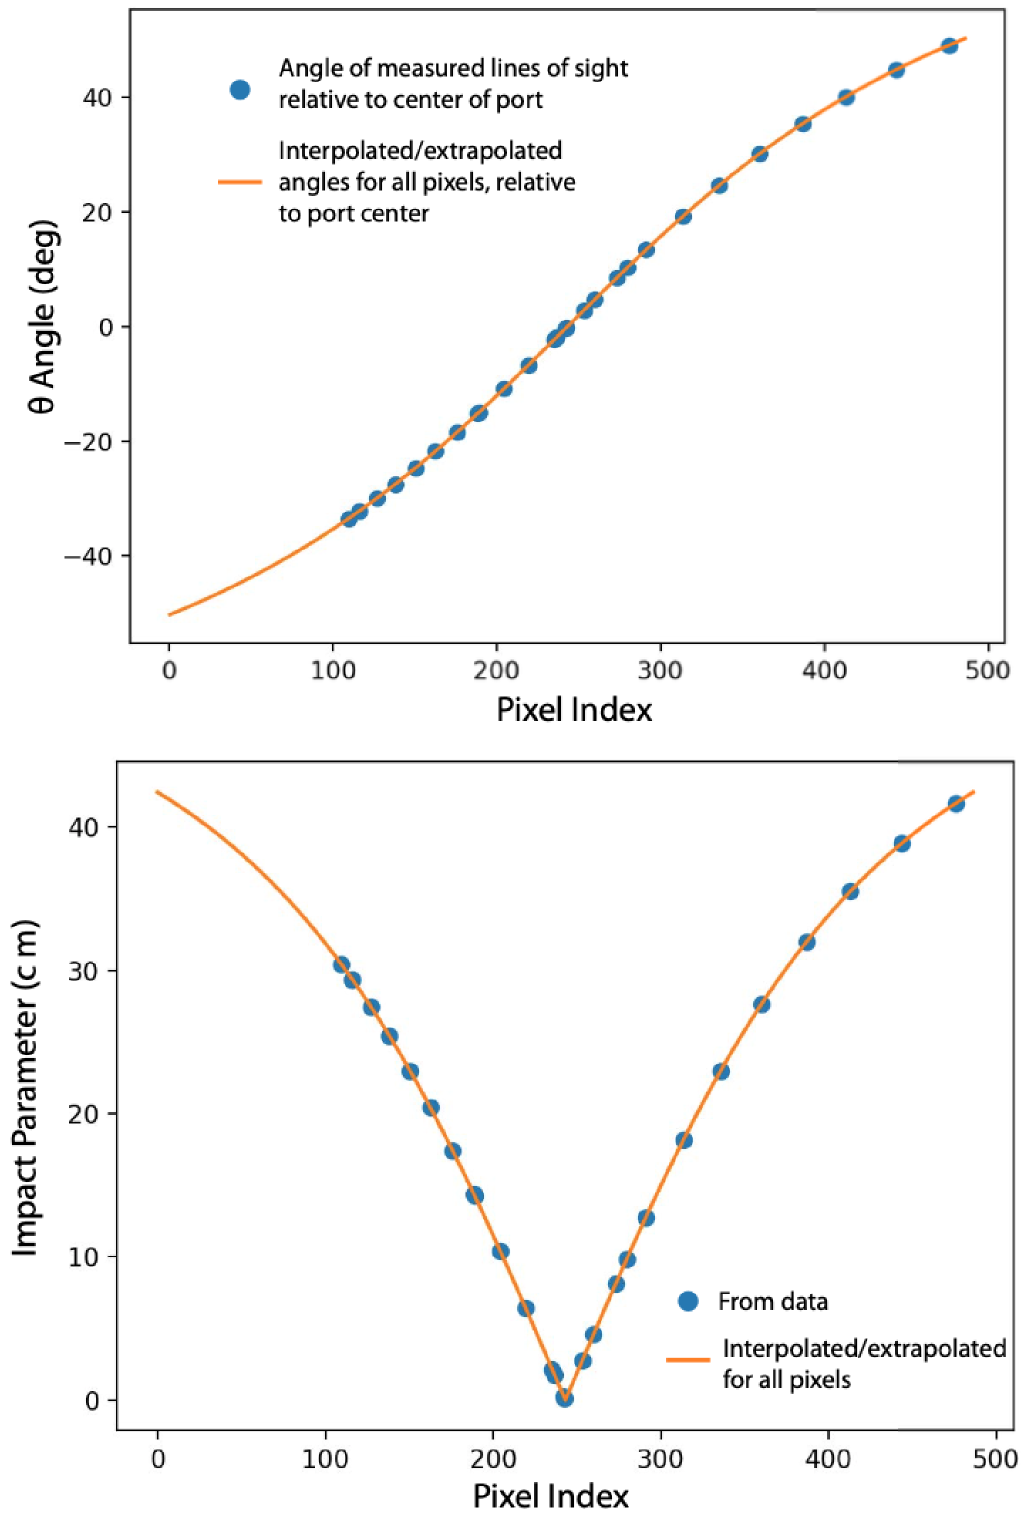

Figure 3.21: Results of the spatial calibration showing the resulting angle θ and impact parameter p as a function of pixel index. Figure courtesy of L. Reusch.

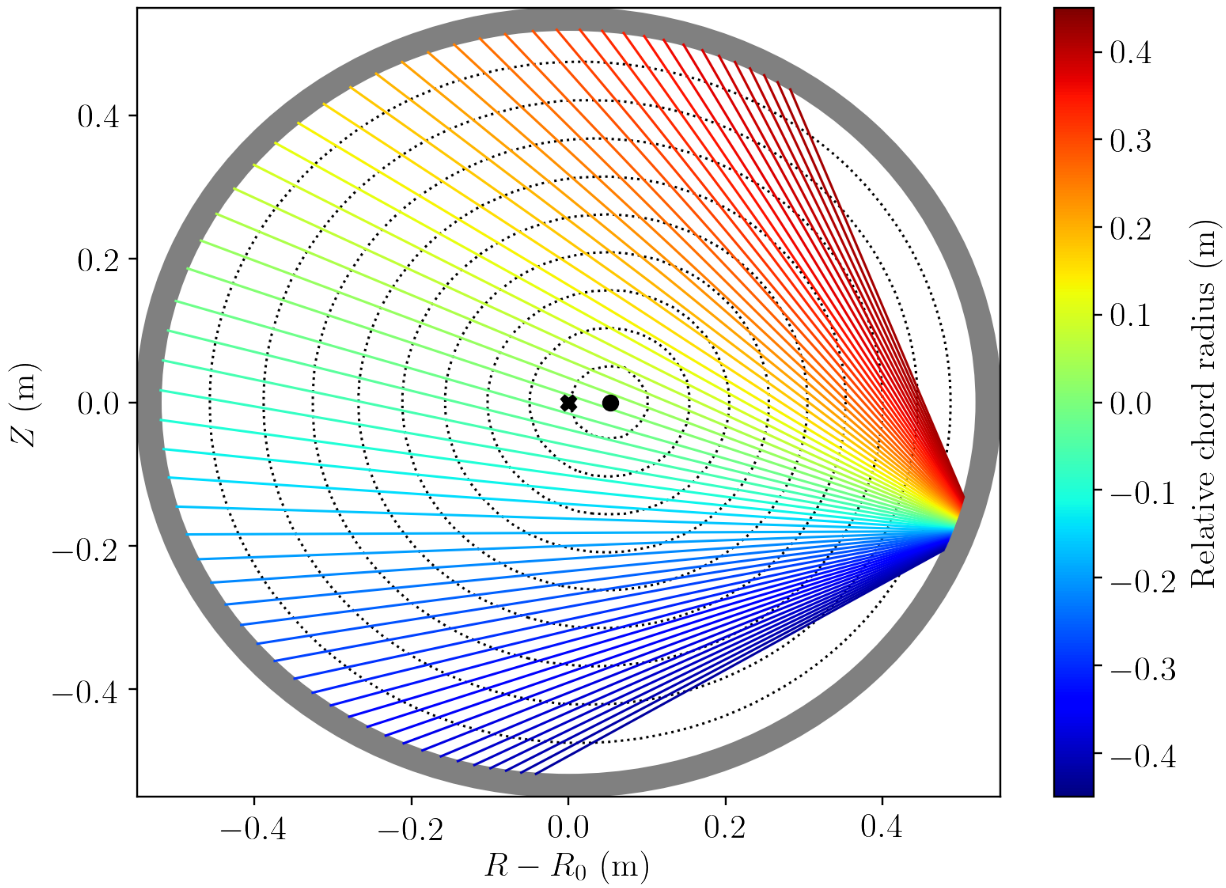

Figure 3.22: Impact parameter for ME-SXR lines of sight within a poloidal plane, as determined by the spatial calibration. Due to space considerations only $1/8^{th}$ of the total 487 vertical chords are shown. Positive impact parameter increases as the line of sight moves upward in the MST vessel. The "X" designates the geometric axis of the MST, and the circular point designates the magnetic axis assuming a 6 cm Shafranov shift. Typical flux surfaces are also shown.

3.4 Characterization of charge-sharing between adjacent pixels

In order to properly interpret the data produced by the ME-SXR diagnostic, it is essential to understand the impact of charge-sharing between adjacent pixels. This was briefly mentioned in Section 3.2 to explain the presence of a linear distortion in the trimbit scan calibration data, described by the parameter a_3. This section expands upon this explanation and paves the way for an implementation of the effect into the ME-SXR diagnostic forward model, described in Section 4.3.2.

The phenomenon of charge-sharing is common to all types of detectors which rely on a monolithic absorbing material to convert photons into charged particles (such as electrons/holes) before transporting these charges to discrete detectors. As the generated charges drift towards the discrete detector elements, there is a finite time during which those charges will undergo expansion due to Coulomb repulsion [11]. As a result, some fraction of the original charge cloud may be collected by an adjacent pixel, thereby appearing to be two photons of lower energies. Whether one, two, or zero photons are detected depends upon the threshold settings of the two pixels. This process is illustrated by a simplified cartoon in Figure 3.23.

As explained in Section 3.1.1, the PILATUS series of detectors were designed to be operated with monochromatic x-ray sources. In this scenario the impact of charge-sharing can be mitigated by setting the global detector threshold to half of the incident photon energy. This ensures that all photons are counted exactly once and that over-counting and under-counting are equally likely (and therefore cancel out on average). However high-temperature plasmas emit a broad spectrum of x-rays, rendering this strategy inapplicable. Instead we must sufficiently characterize the effect and incorporate it into our diagnostic model.

A simplified model for interpreting charge-sharing slope was described in [4]. That

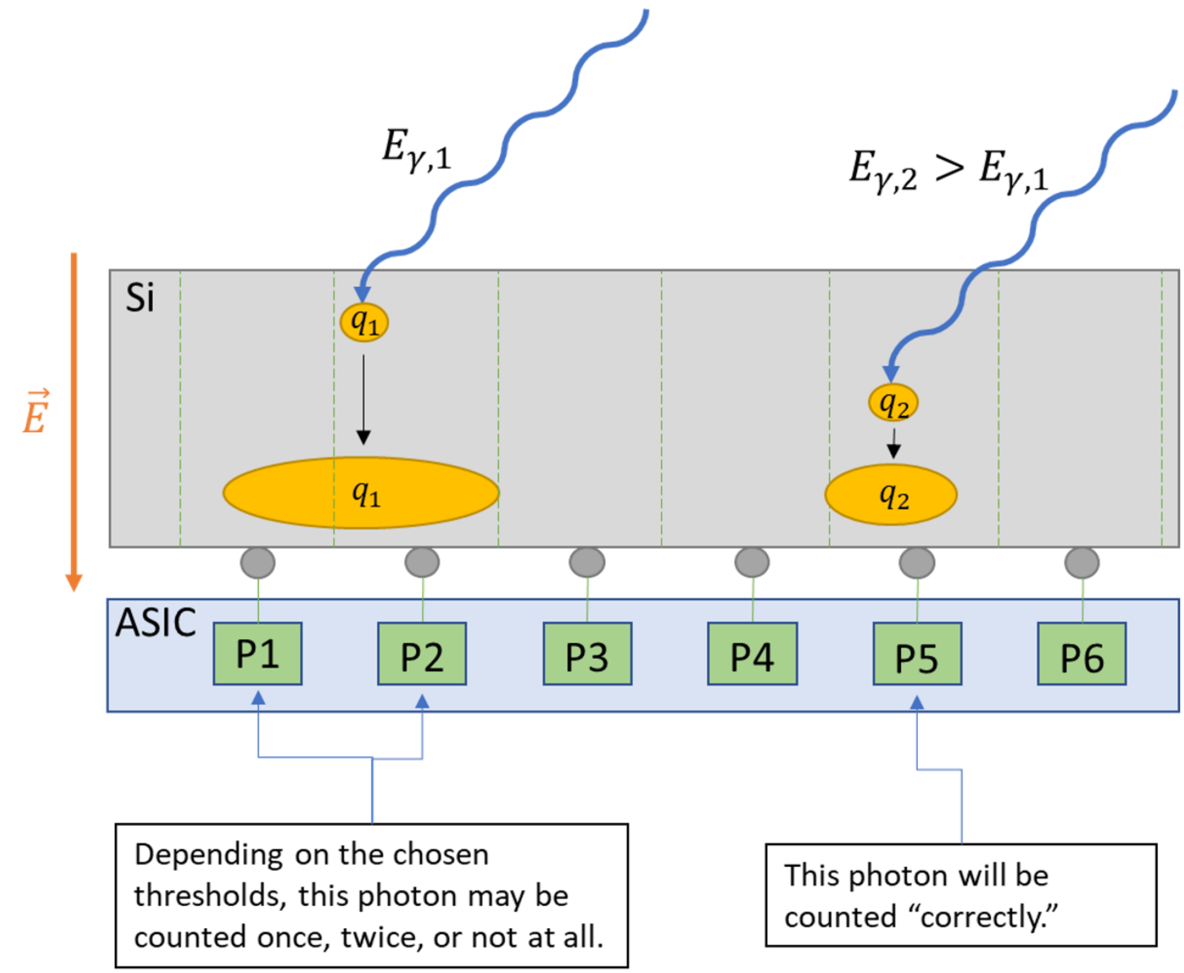

Figure 3.23: Simplified cartoon illustrating the basic concepts of charge-sharing due to Coulomb repulsion in the depleted Si absorber. The effect is more severe for lower-energy photons as they have a smaller penetration depth, thus providing more time for the charge cloud to expand. Note that in reality absorbed photons generate electron-hole pairs and that it is the expansion of the holes which leads to charge sharing in the PILATUS3 detector.

paper is about the characteristics of the original PILATUS detector, which has the same physical dimensions as the PILATUS3. The key insight is that the impact of charge-sharing on registered photon counts is principally determined by the ratio of the area around the pixel border for which charge-sharing is significant to the total pixel area, a quantity I will refer to as $f = A_{CS}/A_{pix}$. This quantity is in turn proportional to the incident photon energy, E_0:

$$-k\frac{E_0}{2} = f \tag{3.26}$$

where k is the charge-sharing slope in units of reciprocal energy. This slope is therefore inversely proportional to photon energy, aligning with our intuition that the effect is most significant for lower-energy photons.

Although the charge-sharing slope was characterized by the parameter a_3 earlier in this chapter, this was done as a scan over trimbits. In order to align with previous literature and to better-characterize the energy response for the purpose of modeling the calibration data was re-processed as a function of the threshold energy. This was done by using Equation 3.1 to convert the trimbit scan to a threshold scan for each pixel and then fitting the results to Equation 3.27. Prior to the fit the background signal (described by a_4 and a_5) was subtracted in order to simplify the analysis.

$$N(E_c) = \frac{1}{2}\left[\mathrm{erf}(-\frac{E_c - b_0}{b_1\sqrt{2}}) + 1\right]\left(b_2 + b_3\big(E_c - b_0\big)\right) \tag{3.27}$$

The charge-sharing slope is then defined to be $k = b_3/N_{50}$, where N_{50} is the number of counts registered when the threshold is set to be half of the photon energy (and hence is the true number of photons). Note that as defined this quantity will always be negative.

The value of k was determined for a subset[2] of the calibration scans for both lowE and

[2]Elements such as Zr, Mo, and Ag were exluded because the photon energy is too near the lower end of the range of thresholds which can be set. In such cases the charge-sharing slope is not well-resolved,

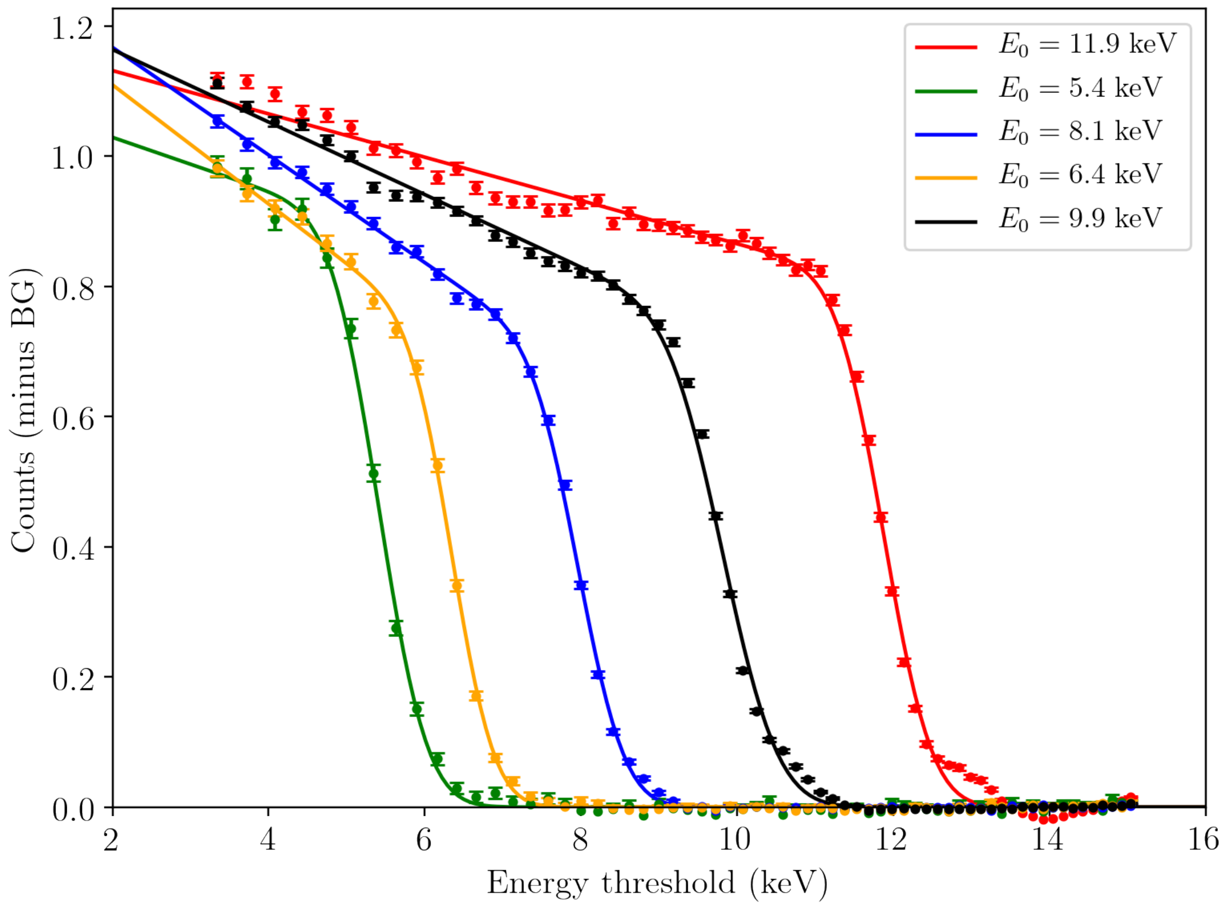

Figure 3.24: Using the results of the energy calibration described in Section 3.1.1, the calibration data was converted from a scan over trimbit settings to a scan over energy thresholds. This was then used to characterize the charge-sharing slope as a function of photon energy. The points are measured counts for a single pixel and the solid line is a fit to Equation 3.27. The background signal was subtracted prior to the fit.

midE settings. Since the effect is based on the physical drift of electrons in the detector, the slope is (approximately) independent of detector gain. These values for k were then fit to Equation 3.26 in order to determine the ratio of the ratio of areas:

$$f = 0.266 \pm 0.008. \tag{3.28}$$

This fit is illustrated by the black data points in Figure 3.24. The value for f obtained via this calibration is in good agreement with the value of 0.272 obtained by Kraft [4]. As an added note, it was also found that $b_1 \approx 0.3$ keV for lowE settings and $b_1 \approx 0.55$ keV for midE, in good agreement with the FWHM values reported in Section 3.2.

The applicability of this model can be seen by extrapolating it to larger pixels. As previously described, the pixels which span the gaps between ASICs are 50% larger than standard pixels. If we consider the charge-sharing area A_{CS} to be defined by a rectangular border around the outer edge of the pixel, its width is found to be 12.3 μm. If we assume that the width of this strip is the same for the larger pixels, we can recalculate the area ratio to be $f = 0.205$. This can then be used to predict k for these larger pixels. As shown in Figure 3.25, these predictions are in good agreement with the data.

3.5 Detector configurations and first data

First data has been collected on MST using three different configurations, discussed in this section. All data shown here were taken using MST's improved confinement mode, PPCD, with $I_p = 400$ kA. ME-SXR data was taken using custom configurations with high-gain global settings and 1 ms exposure times.

Figure 3.26 a) shows data taken in Config. #1. In this configuration, pixels in the same column (toroidal direction) are set to have the same threshold while allowing variation

and N_{50} cannot be reliably determined.

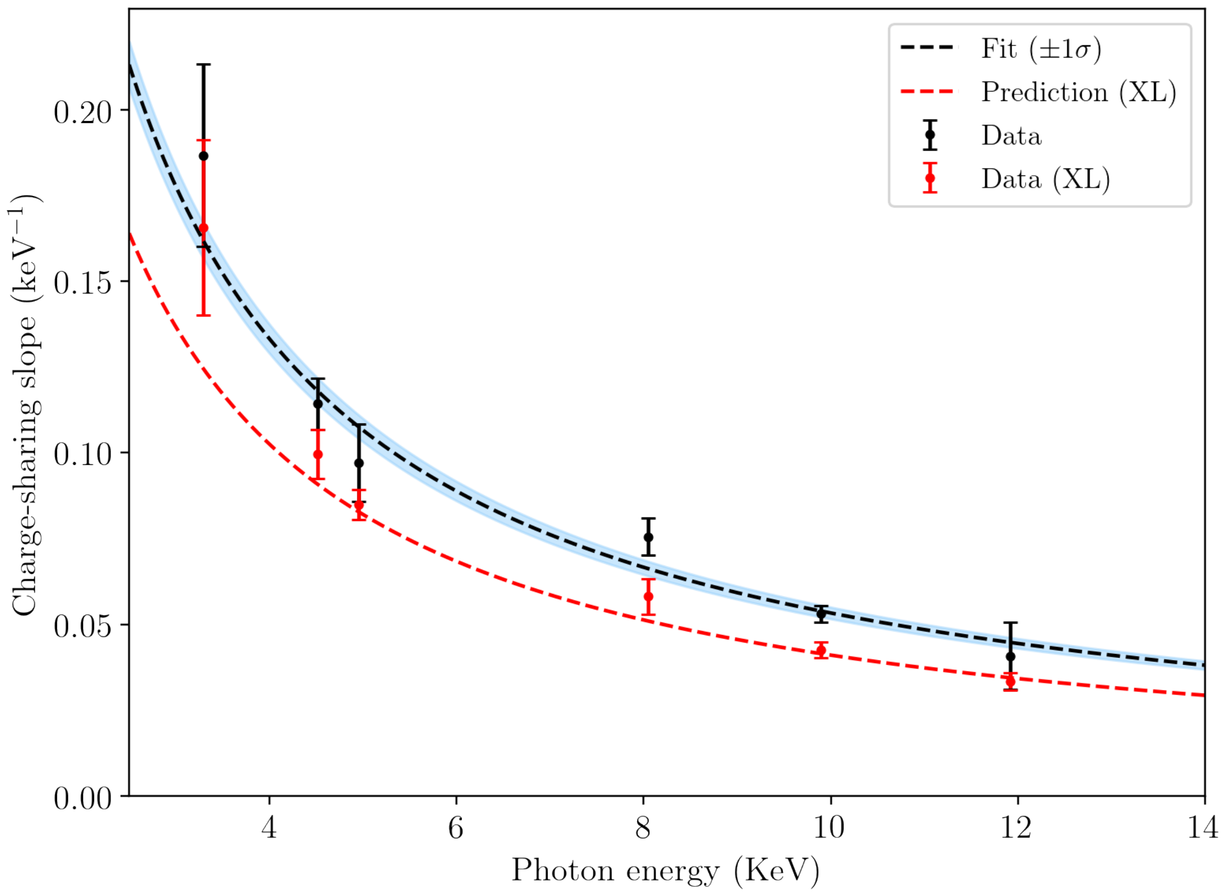

Figure 3.25: Comparison of charge-sharing slope obtained from the threshold scan data to Equation 3.26. The black points are the values of k for different incident photon energies averaged over all regular pixels on the detector and the black dashed line is the best fit with $f = 0.266$. The blue region represents the 1σ uncertainty of the fit. The red points are the values of k average over all extra-large (XL) pixels. The red dashed line is the prediction based on the regular pixel fit. Data from both lowE (In, Ti, V) and midE (Cu, Ge, Br) calibration scans were used in this analysis.

from column-to-column. Since most MST plasmas can be assumed to be symmetric in the toroidal direction, this configuration allows us to use this symmetry to increase the signal level. In this configuration a small number of thresholds (eight for the data shown) are repeated radially in order to achieve the desired spatial resolution. Each cluster of eight columns can be thought of as a single unit, roughly sharing a line of sight but with different spectral sensitivities. This data can be further processed by summing all of the photon counts within a column and separating the columns by threshold energy. The result is eight separate 1D radial emission profiles taken simultaneously from the same plasma, as shown in Figure 3.26 b).

An alternative technique, Config. #2, is shown in Figures 3.27 a) and b). Here constant thresholds are set for each row but allowed to vary within a column. Pixels are separated out by threshold and then summed within the same column to form 1D profiles. The resulting profiles have much higher spatial resolution than Config. #1 (480 points vs 60) at the cost of significantly lower total counts (and thus higher statistical noise). For the installation at MST this configuration significantly over-samples the plasma, but for a larger vessel such as WEST the increased spatial sensitivity might be an important feature.

The third and final configuration that was tested (Config. #3), which we refer to as a "metapixel" configuration. Small blocks of adjacent pixels (2x2, 3x3,...) are set to different thresholds. These blocks are repeated regularly across the face of the detector. The resulting image can be split into multiple images by threshold. This configuration provides a means of simultaneous energy-resolved 2D imaging. The installation geometry on MST is not ideal for this configuration, but it is expected to be useful for future applications with a pinhole geometry and a toroidal viewing angle. An example of this configuration with 2x2 threshold blocks is shown in Figures 3.28 a) and b).

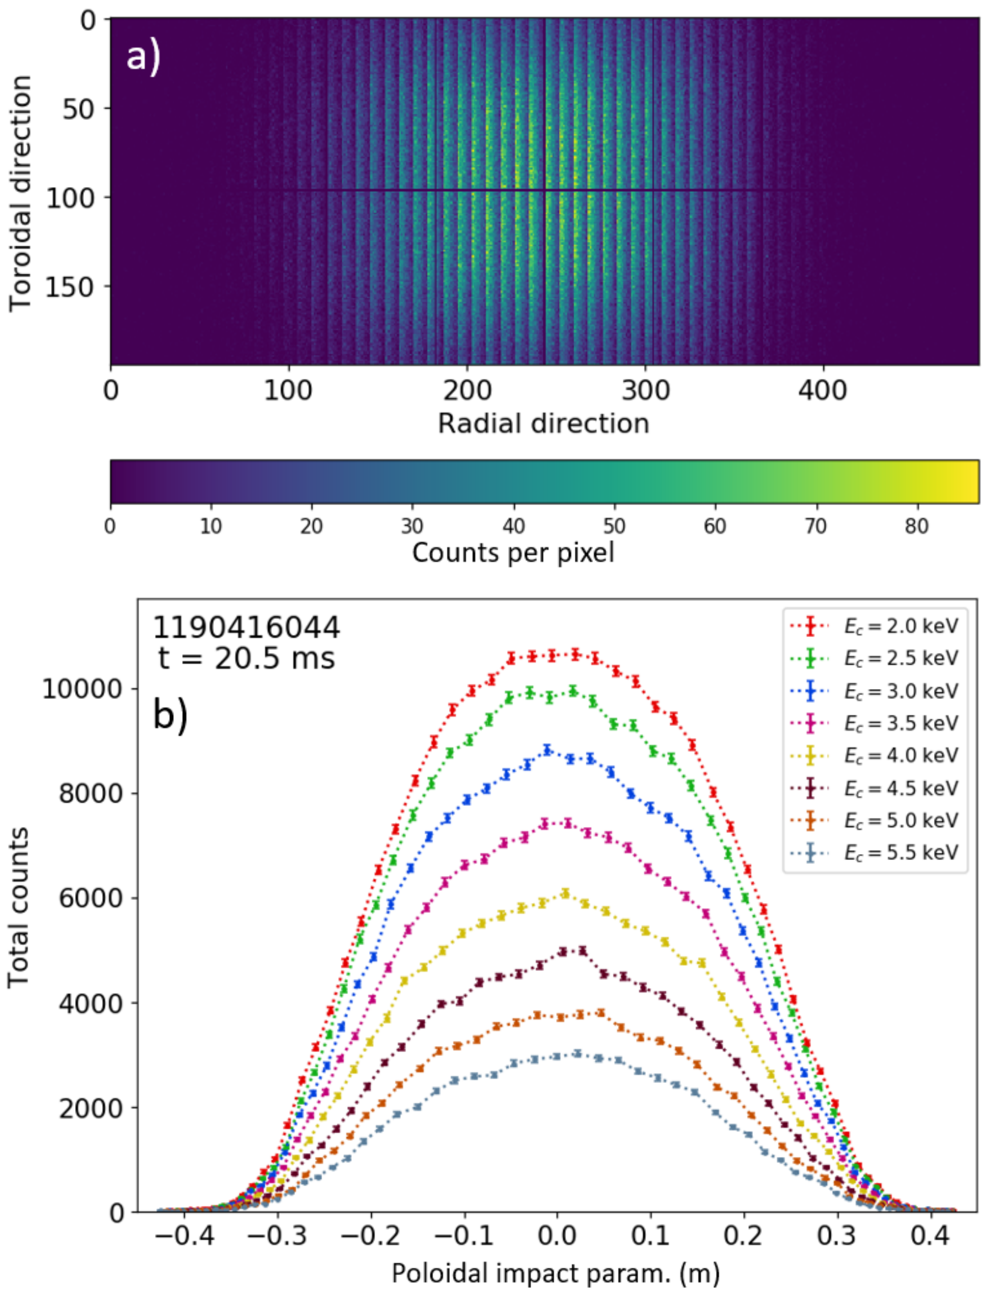

Figure 3.26: (a) Detector image for a 400 kA PPCD plasma using Config. #1. (b) The integrated 1D profile. There are 60 effective chords, one for each cluster of eight thresholds. The profiles are mapped to lines of sight (defined by impact parameter) via the results of a radial calibration. Such a spatial calibration has not been performed for the toroidal direction.

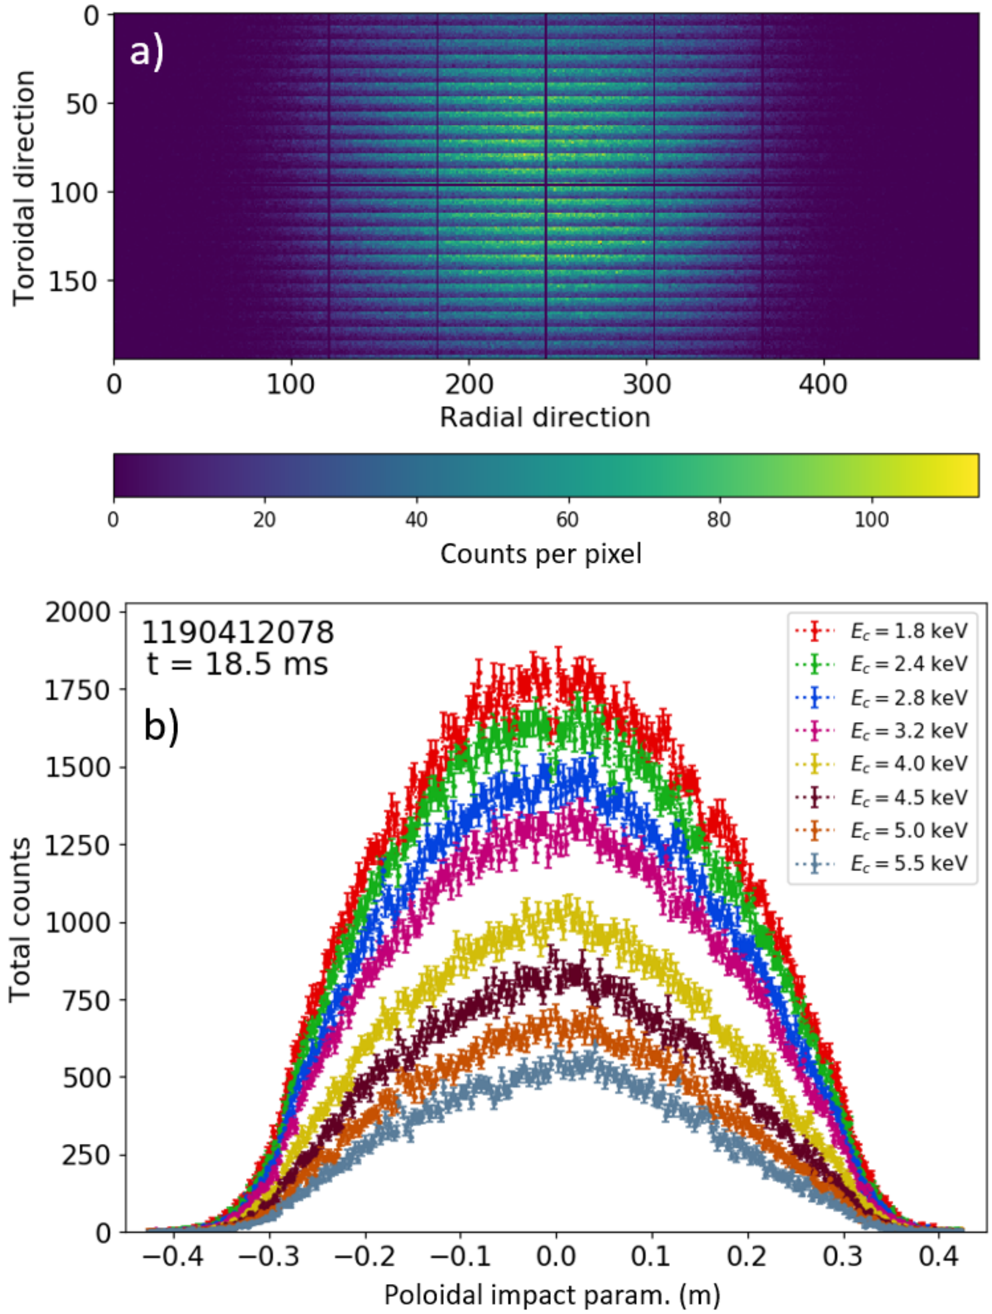

Figure 3.27: (a) Detector image for a 400 kA PPCD plasma using Config. #2. (b) The integrated 1D profile. Spatial resolution is greatly increased to compared to Config. #1 with 480 chords, at the cost of fewer total counts.

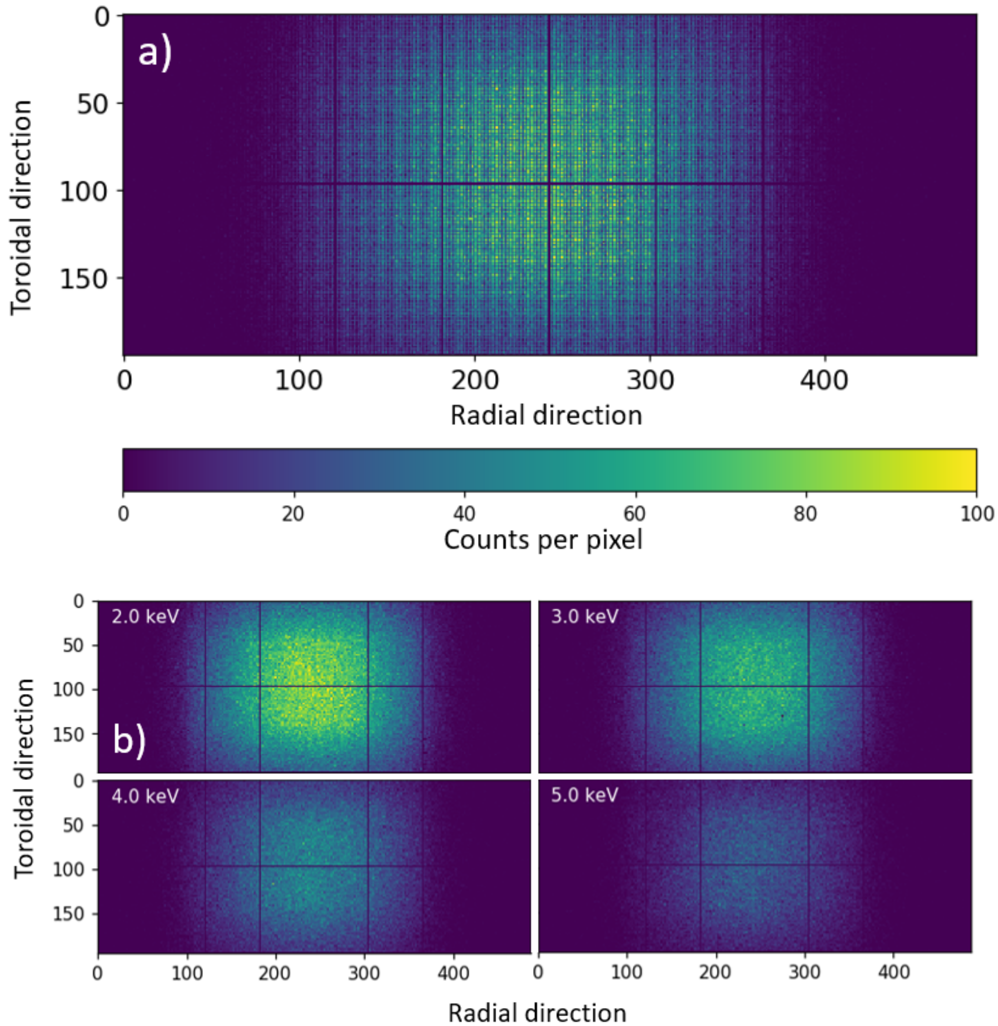

Figure 3.28: (a) Detector image for a 400 kA PPCD plasma using 2x2 metapixels using Config. #3. (b) The resulting energy-resolved re-sampled images.

Bibliography

[1] L. Delgado-Aparicio, M. Greenwald, N. Pablant, K. Hill, M. Bitter, J. E. Rice, R. Granetz, A. Hubbard, E. Marmar, K. Tritz, D. Stutman, B. Stratton, and P. Efthimion, "Multi-energy SXR cameras for magnetically confined fusion plasmas (invited)," *Review of Scientific Instruments*, vol. 87, no. 11E204, 2016. [Online]. Available: http://dx.doi.org/10.1063/1.4964807

[2] N. Pablant, L. Delgado-Aparicio, M. Bitter, E. Brandstetter, R. Ellis, K. Hill, P. Hofer, and M. Schneebeli, "Novel energy resolving x-ray pinhole camera on Alcator C-Mod," *Review of Scientific Instruments*, vol. 83, no. 10E526, 2012. [Online]. Available: http://dx.doi.org/10.1063/1.4732177

[3] J. Maddox, N. Pablant, P. Efthimion, L. Delgado-Aparicio, K. Hill, M. Bitter, M. Reinke, M. Rissi, T. Donath, B. Luethi, and B. Stratton, "Multi-energy x-ray detector calibration for T e and impurity density (nZ) measurements of MCF plasmas," *Review of Scientific Instruments*, vol. 87, no. 11E320, 2016. [Online]. Available: http://dx.doi.org/10.1063/1.4960602

[4] P. Kraft, A. Bergamaschi, C. Bronnimann, R. Dinapoli, E. F. Eikenberry, H. Graafsma, B. Henrich, I. Johnson, M. Kobas, A. Mozzanica, C. M. Schleputz, and B. Schmitt, "Characterization and calibration of PILATUS detectors," *IEEE Transactions on Nuclear Science*, vol. 56, no. 3, pp. 758–764, 2009. [Online]. Available: https://doi.org/10.1109/TNS.2008.2009448

[5] R. E. Bell, "Inversion technique to obtain an emissivity profile from tangential line-integrated hard x-ray measurements," *Review of Scientific Instruments*, vol. 66, no. 1, pp. 558–560, 1995. [Online]. Available: https://doi.org/10.1063/1.1146350

[6] P. Kraft, A. Bergamaschi, C. Broennimann, R. Dinapoli, E. F. Eikenberry, B. Henrich, I. Johnson, A. Mozzanica, C. M. Schlepütz, P. R. Willmott, and B. Schmitt, "Performance of single-photon-counting PILATUS detector modules," *Journal of Synchrotron Radiation*, vol. 16, no. 3, pp. 368–375, 2009. [Online]. Available: https://dx.doi.org/10.1107/S0909049509009911

[7] E. Jones, T. Oliphant, P. Peterson, and Others, "SciPy: Open source scientific tools for Python," 2001. [Online]. Available: http://www.scipy.org/

[8] S. Van Der Walt, S. C. Colbert, and G. Varoquaux, "The NumPy array: A structure for efficient numerical computation," *Computing in Science and Engineering*, vol. 13, no. 2, pp. 22–30, 2011. [Online]. Available: https://numpy.org

[9] P. Bevington and D. Robinson, *Data Reduction and Error Analysis for the Physical Sciences*, 3rd ed. McGraw-Hill Education, 2002.

[10] L. M. Reusch, P. Franz, D. J. Den Hartog, J. A. Goetz, M. D. Nornberg, and P. VanMeter, "Model validation for quantitative X-ray measurements," *Fusion*

Science and Technology, vol. 74, no. 1-2, pp. 167–176, 2018. [Online]. Available: https://doi.org/10.1080/15361055.2017.1404340

[11] L. Tlustos, M. Campbell, E. Heijne, and X. Llopart, "Signal variations in high granularity Si pixel detectors," *IEEE Nuclear Science Symposium Conference Record*, vol. 3, no. 6, pp. 1588–1593, 2003. [Online]. Available: https://doi.org/10.1109/TNS.2004.839095

Chapter 4

The ME-SXR forward model

A *diagnostic forward model* is essential to the quantitative interpretation of plasma x-ray measurements. A diagnostic forward model is a computational framework (either an analytic model or a computational model) which takes as its inputs the properties of the system to be diagnosed (in this case, the plasma) and outputs a realistic prediction of the diagnostic signals. This typically requires a detailed understanding of both the physical system being measured and the diagnostic response. This chapter discusses the development of a computational diagnostic forward model for the ME-SXR diagnostic on MST.

In practice this forward model has had two practical applications. One is predictive: as new pinhole geometries, detector settings, and plasma conditions were tested, the forward model was used to set expectations about the resulting detector signal. In this application the model also provided evidence that the fundamental nature of the detector's response to the plasma was well-understood. The other application of the model is used to classify how well an experimentally obtained measurement is explained by a given set of plasma properties. In this application the model is used to infer the most likely plasma properties based on the data. This type of analysis, using a Bayesian framework, is the focus of Section 5.5. Diagnostic forward models are a key component of integrated data analysis [1].

The importance of forward models for soft x-ray diagnostics is especially important on the MST. As discussed in Chapter 2, the soft x-ray spectrum contains numerous "non-continuum" features such as transition lines and recombination steps around 2 keV due to the trace presence of partially ionized aluminum ions. These features limit the applicability of traditional SXR measurement techniques (i.e., continuum ratios) and instead require detailed atomic physics modeling in order to properly interpret the measured signals [2]. This requires a significant time investment but also brings significant benefits: the resulting non-linear model can provide useful constraints on many plasma parameters (T_e, n_e, n_Z, ...) and can be used to significantly extend measurement capabilities [2].

Section 4.1 provides an overview of the forward model framework and the sources of information it draws upon. Section 4.2 describes the plasma component of the model, which uses the ADAS code to generate synthetic spectra for a given set of plasma parameters. Section 4.3 describes the hardware response of the diagnostic to photons, including both electronic and filter effects. Section 4.4 considers the geometry of the detector-plasma system. Section 4.5 combines all of the elements to produce synthetic measurements which compare favorably with real data. Finally, Section 4.6 estimates the uncertainty inherent in the model.

4.1 Overview of the model

The ME-SXR forward model simulates and interfaces the three essential components of the plasma-diagnostic system. The system can be broadly divided into three categories:

1. The soft x-ray spectrum emitted by the plasma (Section 4.2)

2. The detector's spectral response (Section 4.3)

3. The geometry of the system (Section 4.4)

The interfacing between these components is discussed in Section 4.5.

A graphical overview of the model is shown in Figure 4.1. This chart shows how the user input (some number of parameters specifying the plasma properties and the detector configuration) propagate through the different sections of the model to yield the output (synthetic measurements). It also illustrates that the model uses information from four "external" sources: the energy calibration (Section 3.2), the spatial calibration (Section 3.3), the filter specifications (Section 3.1.2), and ADAS (Section 4.2.3). Arrows entering a box represent inputs, arrows leaving represent outputs, and the circles represent mathematical operations on the inputs.

This chapter focuses on the theoretical foundations of the forward model and covers the essential details which must be considered when implementing the model in code. My main goal with this chapter is to provide sufficient details to aid the development of a forward model for future iterations of the ME-SXR diagnostic on other devices. This description is agnostic to the specific programming language or computing style. My implementation of the model was written in Python.

4.2 Modeling the emissivity spectrum

This section details the portion of the model which maps localized plasma properties (specifically T_e, n_e, n_0, and $\{n_Z\}$) to an emissivity spectrum which will observed by the detector in Section 4.4.3. How exactly these plasma properties are defined will be delayed until later. As discussed in Section 2.3.1, soft x-ray emissions in the MST are due to several different atomic processes. Even though they account for a very small fraction of the overall plasma density, the SXR spectrum is dominated by emissions from the impurity ions. As such, this problem is fundamentally one of atomic physics.

In Section 4.2.1 we will cover the basic types of thermodynamic equilibria and the

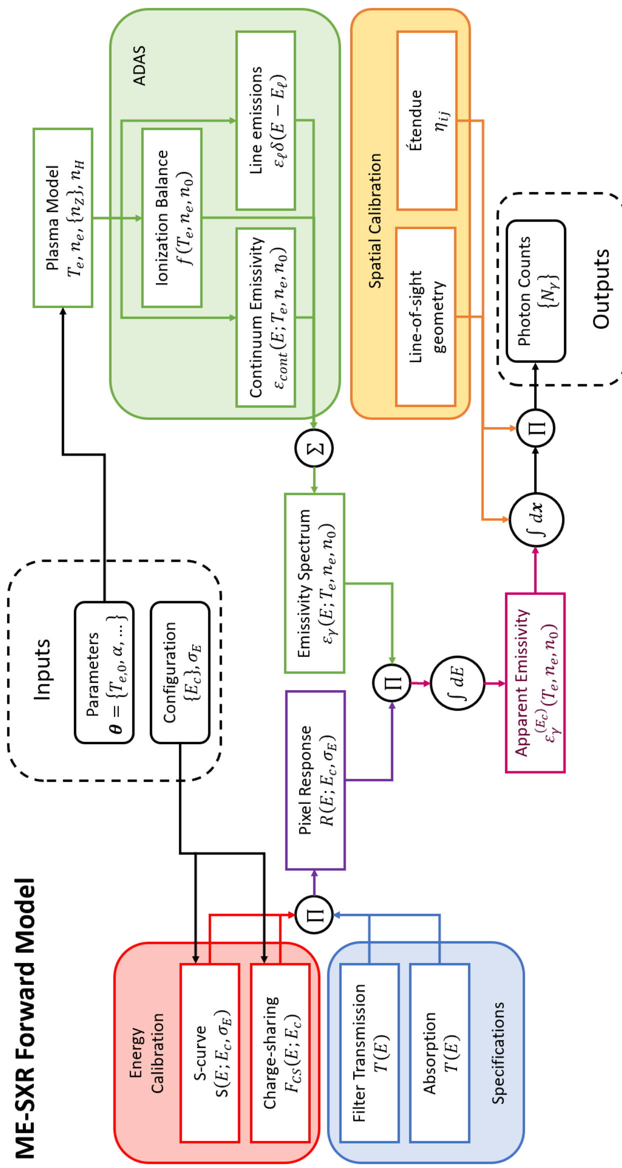

Figure 4.1: Flow chart illustrating the connection between the components of the ME-SXR forward model. The colored boxes represent sources of information.

basic assumptions that go into our chosen atomic physics model. Section 4.2.2 considers the importance of including the neutral density in the ionization fraction calculation, an often overlooked detail which is found to be important for the MST. Section 4.2.3 walks through the detailed calculations required to model the SXR spectrum. Finally, Section 2.3.1 includes the additional consideration for quasi-neutrality and the calculation of the ion-effective charge, Z_{eff}.

4.2.1 Collisional radiative modeling

In order to quantitatively model the plasma's x-ray emissions, we must first decide on a set of assumptions that are appropriate for plasmas in the MST. In many situations, this means determining the most appropriate type of thermodynamic equilibrium [3]. These are:

1. *(Global) Thermodynamic equilibrium*: The particle distribution for the entire system (electrons, ions, and photons) is described by a distribution specified by a single temperature parameter. The principle of detailed balance holds; i.e., forward (emission) and reverse (absorption) processes are equally likely. Emission follows the Plank Law for blackbody radiation. This model is rarely applicable to laboratory plasmas but may be used for some astrophysical systems like stellar interiors.

2. *Local thermodynamic equilibrium (LTE)*: Thermodynamic equilibrium approximately holds for each point within the system (or a subset of the system), described by a locally varying temperature parameter. This assumes that collisions are the primary means of de-excitation so that particles equilibriate more quickly than they diffuse, allowing local equilibration to occur. The plasma is effectively de-coupled from the radiation field. As a result, this model requires the plasma be sufficiently dense. The regime in which this model applies is approximately given by

$$\left(\frac{n_e}{10^{19}\mathrm{m}^{-3}}\right) \gg \left(\frac{T_e}{\mathrm{eV}}\right)^{1/2}\left(\frac{\Delta E}{\mathrm{eV}}\right)^3, \tag{4.1}$$

where ΔE is the excitation energy and T_e is the electron temperature in energy units [4]. This may apply to some laboratory systems as well as many astrophysical systems such as stellar atmospheres.

3. *Coronal equilibrium*: The opposite extreme of LTE. Ions are excited by collisions and de-excitied via spontaneous emission. The plasma is assumed to be optically thin. This requires a relatively low plasma density. This model is applicable to diffuse plasmas like the stellar corona.

It is easy to see that none of these models perfectly describe conditions of a high-temperature magnetically-confined plasma like that produced by the MST. One of the principal goals of magnetic confinement is to produce a plasma with a core much hotter than the edge, so global equilibrium is inapplicable. Given MST's relatively high temperature ($T_e \sim 1000$ eV) and the ionization energy of hydrogen ($\Delta E \sim 13.6$ eV), it can be seen that typical densities of $n_e \sim 10^{19}\,\mathrm{m}^{-3}$ are far below the LTE limit in the core. Likewise, the MST plasmas are insufficiently diffuse to apply coronal equilibrium. Instead, a non-equilibrium model is needed, which will account for both local collisions and non-local coupling to the radiation field by directly solving the set of differential equations which describe state populations. These equations account for processes such as ionization, recombination, free-bound transitions, and three-body interactions. This is called a *collisional-radiative model*.

The forward model developed throughout this section will rely on the *Atomic Data and Analysis Structure*, or ADAS, code [5] to provide the collisional radiative modeling. Unlike other popular CR codes, ADAS was developed specifically for magnetically-confined fusion plasmas. The code can be used (among many other things) to calculate the ionization balance for a given impurity species as well as its x-ray emission spectrum.

ADAS makes a few assumptions which are worth noting here:

- Electrons obey a Maxwellian distribution.
- Spontaneous emission is fast relative to collisions, so that most ions are in the ground or meta-stable states.
- Local equilibration is fast relative to the time scale of collisional transport, allowing ADAS calculations to decouple from impurity transport modeling.
- ADAS relies on an extensive database for atomic rate coefficients $\langle \sigma v \rangle$.

The ADAS code is composed of numerous sets of subroutines, termed "series" [6]. We will most frequently make use of the 400 series, which models ionization, more specifically the *ADAS 405* routine. This routine takes local plasma properties as inputs (T_e, n_e, n_0) and returns the ionization balance for an ion of charge number Z. Transition line amplitudes can be accessed via processed ADF files.

4.2.2 The impact of neutral density on the ionization balance

Previous work [7] [8] has found that the neutral density in MST plasmas (that is, the density of un-ionized hydrogen) is high enough to significantly alter the ionization-fraction balance. This is due to the impact of charge exchange between these neutrals and the ionized population. This interaction tends to push the population towards lower-Z charge states; i.e., there is relatively less Al^{+13} and more Al^{+12} and Al^{+11}. This is illustrated in Figures 4.2 and 4.3, which show the output of ADAS 405 ignoring and accounting for this effect, respectively.

Clearly, an accurate model of the SXR emission in MST must account for this effect. I have included this in the ME-SXR forward model by using neutral density profiles obtained by my colleague A. Xing [9], which uses H-α measurements and modeling with

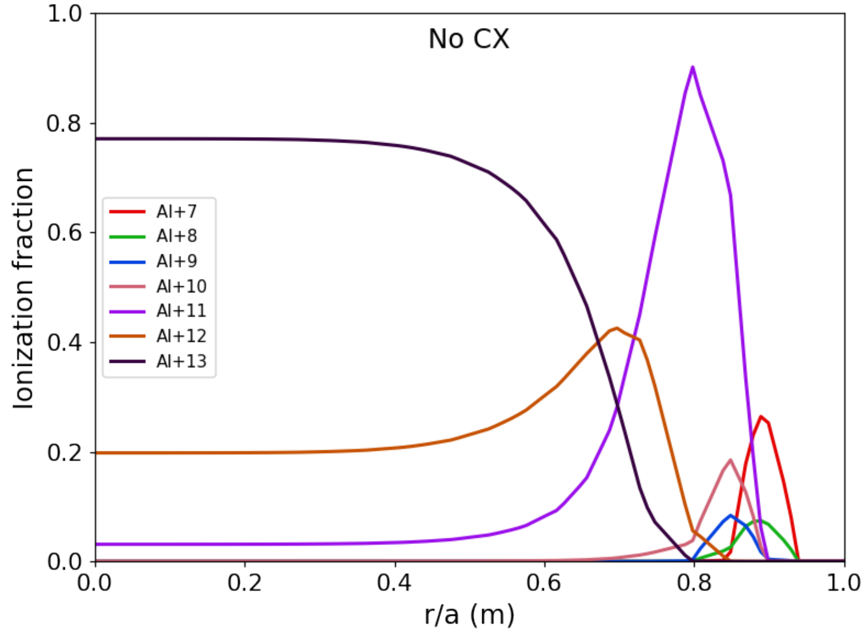

Figure 4.2: The ionization-fraction balance according to ADAS 405 not accounting for charge exchange with neutral hydrogen. Model $T_e = 1500$ eV.

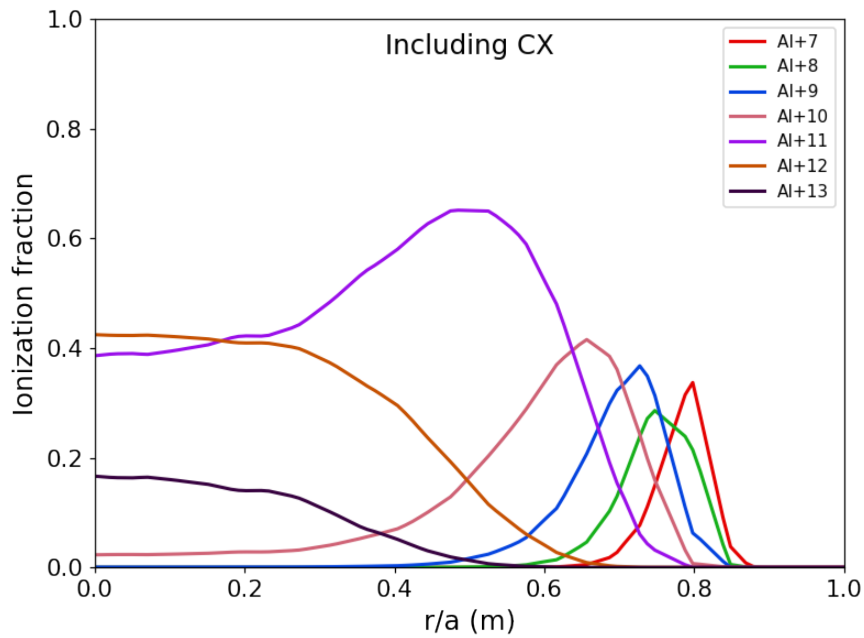

Figure 4.3: The ionization-fraction balance according to ADAS 405 accounting for charge exchange with neutral hydrogen using typical values for n_0 in PPCD. Model $T_e = 1500$ eV.

DEGAS2 to constrain the average neutral profile for high-I_p PPCD plasma conditions. For QSH plasmas, I also used similar measurements obtained by R. Norval [10].

These measurements are then inputted into ADAS 405 along with T_e and n_e in order to obtain the fraction $f_{Z,i}$ of an ion species Z (such as Al, O, N, etc.) which will be found in particular ionization state i (such as Al^{+11} or Al^{+12}). This is defined explicitly in Equation 4.2.

$$f_{Z,i}(T_e, n_e, n_0) \equiv n_i / n_Z \tag{4.2}$$

It is worth noting that if we could ignore the effects of charge-exchange then this quantity would simply reduce to a function of electron temperature only, $f_{Z,i}(T_e)$, which could significantly simplify the spectrum model. This is the convention used by some codes such as FLYCHK [11]. For other devices such as tokamaks with a divertor (NSTX-U, WEST, ITER, etc.) this assumption might be valid. However, as we can see in Figures 4.2 and 4.3, this assumption is not valid on the MST.

Throughout the rest of this chapter I will typically omit the arguments to $f_{Z,i}$ in order to avoid needless clutter, but they should always be taken as implied. Also when it is clear I will sometimes shorten the symbol down to f_i, with the indexing by Z being implied. It is also worth noting that these fractions can also be interpreted as the probability of an ion being in the state i when chosen at random. That means that they follow the typical properties of a complete set of probabilities such as $\sum_i f_i = 1$. We will also encounter terms which are most readily understood as being averaged over the ionization states. For example, the ion-averaged value of the charge Z_i of a given ion (e.g., $Z_i = 11$ for Al^{+11}) is given by

$$\langle Z_i \rangle_Z \equiv \sum_{i \in Z} f_i Z_i. \tag{4.3}$$

This notation will turn out to be useful in constructing the final spectrum for the ME-

SXR forward model. As stated above, this average quantity has an implied dependence on (T_e, n_e, n_0).

Given that the neutral density significantly impacts the emissivity coming from the plasma, it is important to quantify how sensitive the model we will build in this chapter will be to uncertainty in the n_0 profile. This will put a limit on the accuracy of the model based on the accuracy of the available neutral density measurements. This study was performed using ADAS 405 to vary the neutral density by $+10\%$ at various nominal neutral density and electron temperatures and quantifying the percent change according to

$$\Delta f_i(n_0, T_e)\,[\%] = \frac{|f_i(n_0, T_e) - f_i(1.1 \cdot n_0, T_e)|}{f_i(n_0, T_e)} \times 100\%. \tag{4.4}$$

The results for Al are shown in Figure 4.4, focusing only on the higher ionization stages with a non-negligible abundance in the MST. Two features stand out: the magnitude of the effect, sometimes accounting for a change in f_i on the order of 40%, and the highly non-linear dependence on T_e. This does not mean that we should expect the total number of emitted photons to vary this substantially, however, as in some cases the large percentage changes in abundance $\Delta f_i / f_i$ for a given ionization state are due to low absolute abundance $f_i << 1$. Instead what we want to determine is how an uncertainty in n_0 impacts the total emissivity ε. This is shown in Figure 4.5, which shows

$$\left(\frac{\delta\varepsilon}{\varepsilon}\right)_{n_0}[\%] = \frac{|\varepsilon(T_e, n_e, n_0) - \varepsilon(T_e, n_e, 1.1 \cdot n_0)|}{\varepsilon(T_e, n_e, n_0)} \times 100\% \tag{4.5}$$

evaluated over a grid of (T_e, n_0) points, where $n_e = 1 \times 10^{19}\ \mathrm{m}^{-3}$ has been held constant.

We can see that the impact of a change in n_0 on ε is much more modest than the change in f_i suggested, with the emissivity shifting by as little as 1% or as much as 5%, depending on the other parameters. This provides a baseline which we will use when estimating the uncertainty in the model later in Section 4.6. It should also be noted

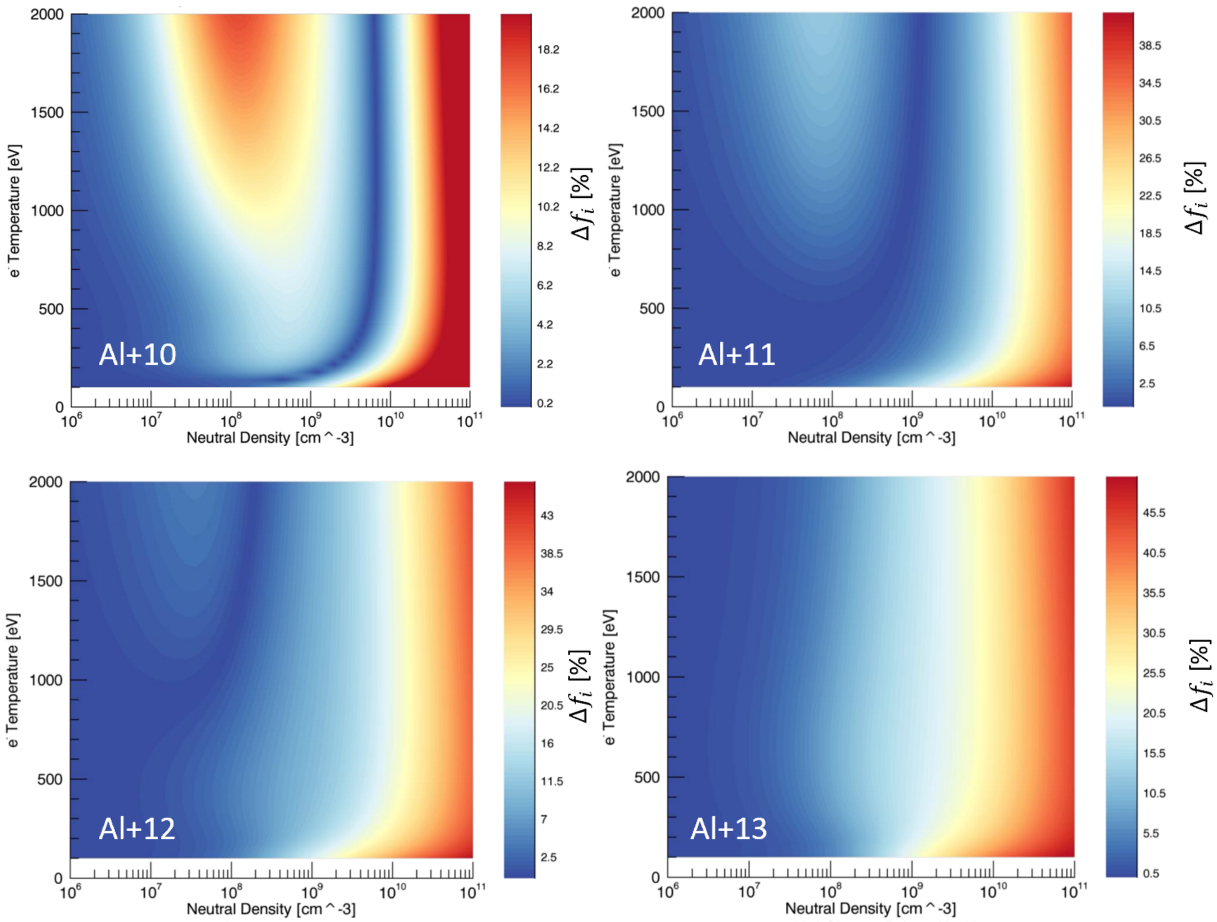

Figure 4.4: The balance of ionization states $f_{Z,i}$ changes significantly in response to variations in the neutral density n_0. These plots depict the percentage change in emissivity that results from a 10% increase in n_0 relative to the value depicted on the x-axis. Notice also that T_e significantly impacts these results.

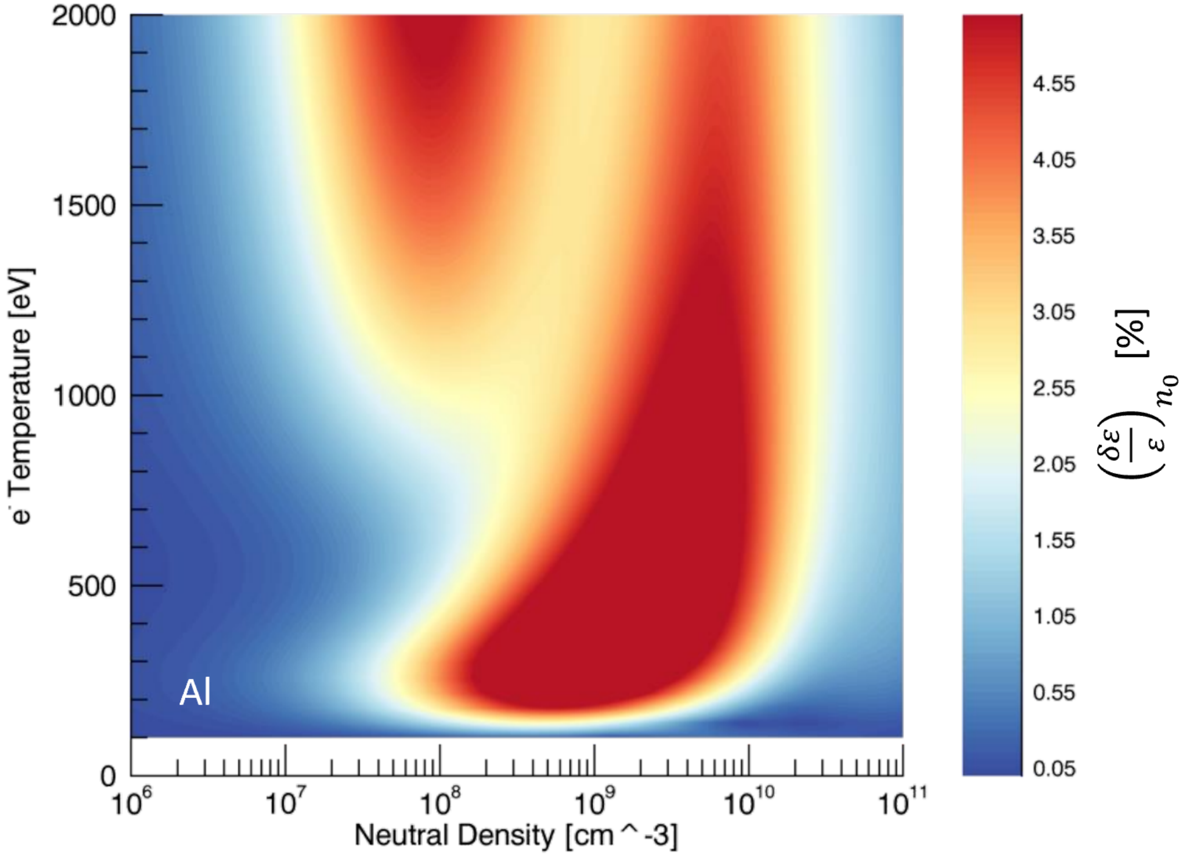

Figure 4.5: The impact of a 10% underestimation of n_0 on the total Al emissivity. This results in a $\leq 5\%$ change in the number of photons emitted by the plasma.

that, although we only considered the impact on various Al ionization stages, other ion species (like C, O, etc.) are almost completely ionized in the MST and thus subject to very minimal variation. Therefore, when quantifying the uncertainty only Al is expected to matter.

4.2.3 Modeling the spectrum with ADAS

The continuum spectrum was constructed by a custom-modified version of the ADAS *continuo* sub-routine. The routine uses calculations from Burgess, Hummer, and Tully [12] to calculate the effect of free-bound transitions. I modified the routine to use a more modern calculation of the free-free Gaunt factor [13] and to use the current NIST values for ionization energies [14] rather than a Rydberg approximation. For the remainder of this section I will represent the output of this code for the i^{th} ion state of the species Z as $\mathcal{F}_{Z,i}(E, T_e)$. It is scaled by the ion and impurity densities and has units of [ph ms^{-1} m^3 sr^{-1}]. Note that this quantity is equivalent to Equation 2.19 divided by the photon energy E. The total continuum emission is given by

$$\varepsilon_{cont,Z}(E, \mathbf{x}) = n_e \sum_{i \in Z} n_i \mathcal{F}_{Z,i}(E, T_e). \tag{4.6}$$

As discussed above, the ion density n_i is related to the total impurity species density n_Z via the ionization fraction calculated by ADAS 405, as in Equation 4.2. It follows that

$$\varepsilon_{cont,Z}(E, \mathbf{x}) = n_e n_Z \sum_{i \in Z} f_i \mathcal{F}_{Z,i}(E, T_e) = n_e n_Z \langle \mathcal{F}(E) \rangle_Z. \tag{4.7}$$

This is useful since it means that we can work in terms of the ion-averaged spectrum of a species rather than dealing with all the individual contributions from each ion. By summing over all impurity species we get the total continuum spectrum:

$$\varepsilon_{cont}(E, \boldsymbol{x}) = n_e \sum_{Z} n_Z \langle \mathcal{F}(E) \rangle_Z. \tag{4.8}$$

This is all we need to correctly model the continuum portion of the plasma SXR spectrum. Unfortunately this process turns out to be relatively slow, requiring multiple seconds to build $\mathcal{F}_{Z,i}(E, T_e)$ for sufficiently many points in E for all species $\{Z\}$ in a typical MST plasma. This process can be sped up significantly by instead interpolating over a grid of pre-evaluated points at run time. However, storing $\langle \mathcal{F}(E) \rangle$ as a four-dimensional array in (T_e, n_e, n_0, E) with sufficiently many points in every dimension turns out to be prohibitively large. Therefore, we will need to carry the processing out further before building lookup tables.

The ME-SXR is a broadband SXR diagnostic, meaning that it measures the integral of the spectrum multiplied by some response function. This response function accounts for the transmission of x-rays through filters, absorption into the detector, electronic thresholds, etc. For now we will think about a generic family of response functions $R_k(E)$. The details of this function for the ME-SXR model will be discussed in Section 4.3.

What the detector actually sees is the *apparent emissivity* $\varepsilon_{cont}^{(k)}$ due to the response indexed by k. This is a local quantity defined by the local plasma properties:

$$\begin{aligned} \varepsilon_{cont}^{(k)}(T_e, n_e, n_0) &= \int_0^\infty dE \, \varepsilon_{cont}(E) \, R_k(E) \\ &= n_e \sum_Z n_Z \int_0^\infty dE \, \left\langle \mathcal{F}(E, T_e) \right\rangle_Z R_k(E). \end{aligned} \tag{4.9}$$

This integral is what we actually want to use to build our lookup tables. We will define it in a way which is independent of the ion density:

$$\epsilon_{cont,Z}^{(k)}(T_e, n_e, n_0) \equiv \int_0^\infty dE \, \left\langle \mathcal{F}(E, T_e) \right\rangle_Z R_k(E). \tag{4.10}$$

Now we finally have an expression for the total apparent emission for a detector with response function $R_k(E)$:

$$\varepsilon_{cont}^{(k)}(T_e, n_e, n_0) = n_e \sum_Z n_{Z_0} \epsilon_{cont,Z}^{k}(T_e, n_e, n_0). \tag{4.11}$$

As discussed in Section 2.3.1, a model of the SXR spectrum in the MST must also consider emission due to excitation lines. This is included in the ADAS code via calculation of the *photon emissivity coefficients*, or PECs, which relates the radiated photons for a given transition to to the electron and ion densities [15]. The values for these PECs are stored in a file format called ADF15 and can be interpolated for a specified T_e and n_e using the *read_adf15* routine [5]. I will refer to the output of this routine as $\mathcal{P}_{i,\ell}(T_e, n_e)$, which I have converted to units of [ph ms^{-1} m^3]. Following this convention the emissivity from a given ion is calculated as in Equation 2.21. We can again invoke Equation 4.2 to represent this in terms of the ion fractions:

$$\varepsilon_{exc} = n_e n_Z \sum_{i \in Z} f_i \left[\sum_{\ell \in \mathcal{L}_i} \frac{\mathcal{P}_{i,\ell}(T_e, n_e)}{4\pi} \delta(E - E_\ell) \right]. \tag{4.12}$$

The interpretation of this term is less intuitive than Equation 4.7 due to the inclusion of the delta function terms. Therefore, we will instead go straight to the apparent emission due to excitation lines:

$$\varepsilon_{exc,Z}^{(k)} = \int_0^\infty dE \, \varepsilon_{exc,Z}(E) \, R_k(E). \tag{4.13}$$

We can manipulate this expression in order to write the answer in terms of a single ion-averaged quantity:

$$\begin{aligned}
\varepsilon_{exc,Z}^{(k)} &= n_e n_Z \int_0^\infty dE \sum_{i=0}^{Z} f_{Z,i} \left[\sum_{\ell \in \mathcal{L}_{Z,i}} \frac{\mathcal{P}_{Z,i,\ell}(T_e, n_e)}{4\pi} \delta(E - E_\ell) \right] R_k(E) \\
&= n_e n_Z \sum_{i=0}^{Z} f_{Z,i} \left[\sum_{\ell \in \mathcal{L}_{Z,i}} \int_0^\infty dE \, \frac{\mathcal{P}_{Z,i,\ell}(T_e, n_e)}{4\pi} \delta(E - E_\ell) R_k(E) \right] \\
&= n_e n_Z \sum_{i=0}^{Z} f_{Z,i} \left[\sum_{\ell \in \mathcal{L}_{Z,i}} \frac{\mathcal{P}_{Z,i,\ell}(T_e, n_e)}{4\pi} R_k(E_\ell) \right] \\
&= n_e n_Z \left\langle \sum_{\ell} \frac{\mathcal{P}_\ell(T_e, n_e)}{4\pi} R_k(E_\ell) \right\rangle_Z \\
&\equiv n_e n_Z \epsilon_{exc,Z}^{(k)}(T_e, n_e, n_0).
\end{aligned} \tag{4.14}$$

.

This is a term which can be readily indexed and interpolated that depends on the chosen response function and the parameters (T_e, n_e, n_0). The total emission due to excitation lines is given by the sum

$$\varepsilon_{exc}^{(k)}(T_e, n_e, n_0) = n_e \sum_Z n_Z \epsilon_{exc,Z}^{(k)}(T_e, n_e, n_0). \tag{4.15}$$

Now that we have methods for calculating both the total apparent continuum emissivity and the total apparent excitation line emissivity, we can combine these terms into a full total apparent emissivity. We first define the normalized terms:

$$\epsilon_Z^{(k)}(T_e, n_e, n_0) \equiv \epsilon_{cont,Z}^{(k)}(T_e, n_e, n_0) + \epsilon_{exc,Z}^{(k)}(T_e, n_e, n_0). \tag{4.16}$$

The total apparent emissivity at a point $\boldsymbol{x}$ in the plasma, which will be integrated along a given line of sight, is therefore given by Equation 4.17.

$$\varepsilon^{(k)}(\boldsymbol{x}) = n_e \sum_Z n_Z \epsilon_Z^{(k)} \big(T_e(\boldsymbol{x}), n_e(\boldsymbol{x}), n_0(\boldsymbol{x}) \big) \tag{4.17}$$

As I have previously indicated, the main reason that I took this approach was so I could construct tables of $\epsilon_Z^{(k)}(T_e, n_e, n_0)$ which are then interpolated at run time using the Scipy RegularGridInterpolator class [16]. Such a grid must be constructed for each ion species and each response, but for seven ion species (D, C, O, N, B, Al, and Ar) and eight thresholds (responses) this does not require a prohibitive amount of time or an unreasonable amount of storage space. All of this effort allows the model to perform rapidly at run time.

4.2.4 Quasi-neutrality and ion-effective charge

We want to ensure that the model plasma obeys the principle of quasi-neutrality, that the total negative charge due to electrons in the plasma is exactly negated by the positive charge from ions. Since the n_e and n_Z are set as inputs to the model (and are typically constrained by measurements), we should therefore set the density of the majority species (deuterium in the MST) in order to account for this principle. The relation for this balance can be worked out from the mathematical expression of quasi-neutrality as follows:

$$\begin{aligned} n_e &= \sum_j n_j Z_j \\ &= n_D + \sum_Z \sum_i^Z n_{Z,i} Z_i \\ &= n_D + \sum_Z n_Z \sum_i^Z f_{Z,i} Z_i \\ &= n_D + \sum_Z n_Z \langle Z \rangle_Z. \end{aligned} \tag{4.18}$$

In the first term the index j refers generically to all ions, while the next line separates out the deuterium density and relabels the remaining ions using the notation introduced in Section 4.2.3. We can then solve directly for the deuterium density:

$$\frac{n_D}{n_e} = 1 - \sum_Z \frac{n_Z}{n_e} \langle Z \rangle_Z. \tag{4.19}$$

Finally, we will comment on another term which is of common interest for plasma physicists: the ion-effective charge Z_{eff}. The value of Z_{eff} is useful not only as a summary of the overall impurity content of the plasma but also for calculating properties such as resistivity and the related Lundquist number. In fact the desire to calculate Z_{eff} was one of the principal original motivations for the application of IDA techniques at the MST [17]. The value of Z_{eff} for a modeled plasma can be readily extracted from the input profiles. We begin with the definition

$$n_e^2 Z_{eff} \equiv \sum_j n_e n_j Z_j^2. \tag{4.20}$$

We can rewrite this to explicitly separate out the impurity species and ions:

$$\begin{aligned} Z_{eff} &= \frac{1}{n_e^2} \sum_Z \sum_i^Z n_e n_{Z,i} Z_i^2 \\ &= \frac{1}{n_e} \sum_Z n_z \sum_i^Z f_{Z,i} Z_i^2. \end{aligned} \tag{4.21}$$

Finally, we rewrite this in terms of Equation 4.3:

$$Z_{eff} = \frac{n_D}{n_e} + \sum_Z \frac{n_Z}{n_e} \langle Z^2 \rangle_Z. \tag{4.22}$$

One might see these results and conclude that we also need to construct databases of both $\langle Z \rangle_Z (T_e, n_E, n_0)$ and $\langle Z^2 \rangle_Z (T_e, n_E, n_0)$. However, in practice, it tends to be the case that even for heavier-Z impurities like Al the distribution of ion-states tends to be very narrowly peaked. This means that it is a reasonable assumption to take $\langle Z^2 \rangle \approx \langle Z \rangle^2$ for the purpose of the Z_{eff} calculation. I have generally found the error introduced by this assumption to be $< 1\%$.

4.3 Modeling the detector response

This section will discuss the particular details of the detector response functions, referred to in the previous section as generic functions $R_k(E)$. In general, for each pixel in the ME-SXR model, this term is composed of five separate effects:

$$R(E) = T_{Be}(E) \cdot T_{Mylar}(E) \cdot A_{Si}(E) \cdot S(E; E_c, \sigma_E) \cdot F_{CS}(E). \tag{4.23}$$

$T_{Be}(E)$ is the proportion of the incoming photons of energy E which are transmitted through the 25 μm beryllium window; $T_{Mylar}(E)$ is the proportion which are transmitted through the Mylar filter (see Section 5.4); $A_{Si}(E)$ is the proportion which are absorbed by the silicon. The term $S(E)$ is called the *S-curve* and models the effect of the electronics in either counting or rejecting a photon. The final term $F_{CS}(E)$ is an empirical adjustment to the response function in order to account for the effect of charge-sharing (Section 3.4). The remainder of this section will go into more detail on how these terms are calculated.

4.3.1 Filter response

The filter transmission is based on the table of linear attenuation coefficients $\mu(E)$ published by Henke *et al.* [18], as made available on the Lawrence Berkeley National Lab Center for X-ray Optics website [1]. The proportion of incident photons transmitted through a material of density ρ and thickness t is then given by

$$T(E) = \exp\Big(- \mu(E)\rho t \Big). \tag{4.24}$$

For these materials, reflection is minimal, meaning that photons are either transmitted or absorbed. This gives a simple relation for the absorption factor $A(E)$ given that $\mu(E)$ is known for the material.

[1] https://henke.lbl.gov/optical_constants/

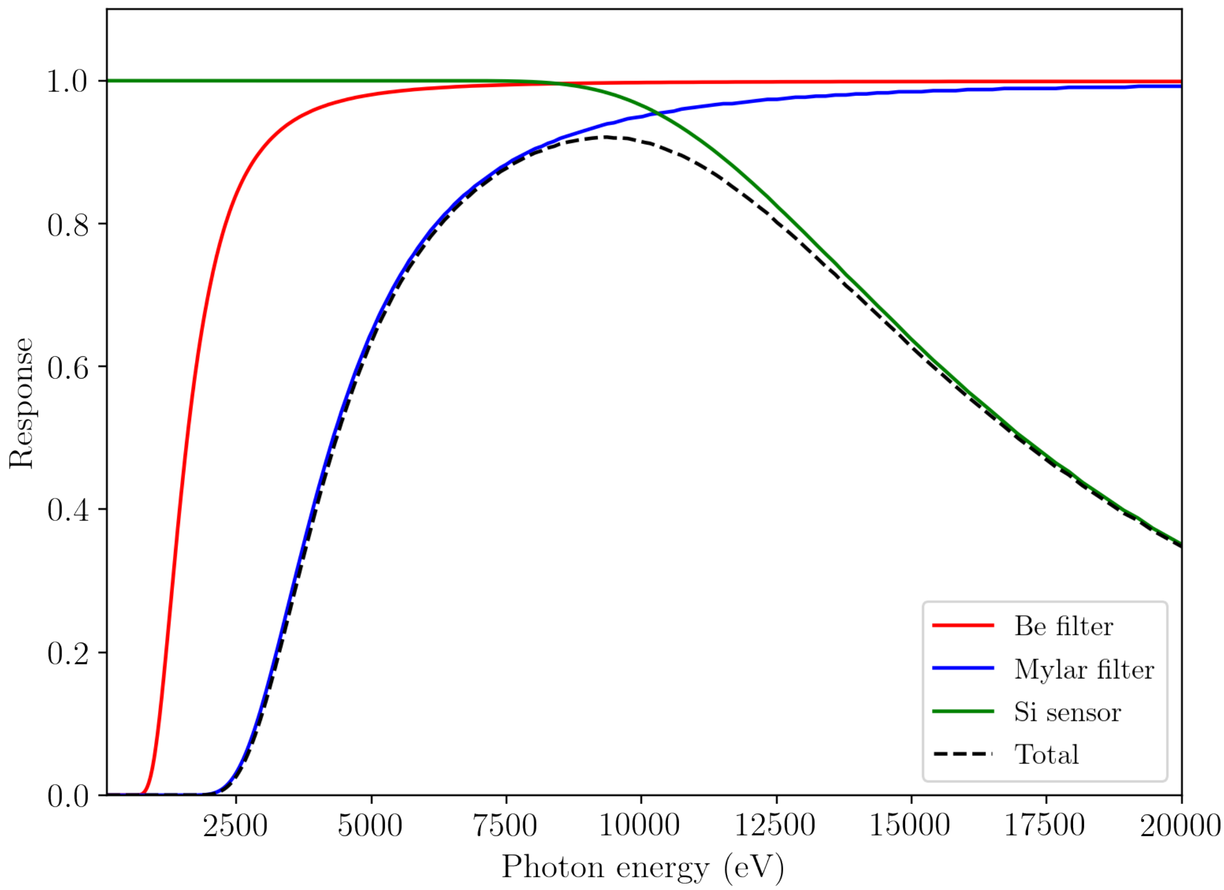

Figure 4.6: Illustration of photon response to physical layers in the ME-SXR model. This includes transmission through the 25 μm Be and 100 μm Mylar filters as well as absorption into the 450 μm Si sensor. The model stops at 20 keV because emissivity past this point (for thermal electrons) is negligible.

$$A(E) = 1 - T(E) \tag{4.25}$$

The transmission curves for 25 μm Be and 100 μm Mylar, the absorption curve for 450 μm Si, and their composite effect are all shown in Figure 4.6. In the current configuration the low-energy response is dominated by the Mylar filter while the response trails off at high energies due to the silicon. Absorption is minimal past 30 keV.

The forward model developed for the SXR tomography diagnostic [2] also accounts for the effect of non-normal photon angle-of-incidence on the effective Si thickness. When a photon strikes the detector at an angle different from 90°, it will travel through a greater depth of absorber material, increasing the likelihood of being absorbed, especially for higher-energy photons. This effect was considered for the ME-SXR model but was ultimately found to be unnecessary. This is because the PILATUS3's 450 μm Si sensor was already sufficiently thick to ensure that most photons are absorbed.

4.3.2 Response of the PILATUS3 electronics

The next element to consider is the response of the detector electronics. As discussed in Section 4.3, once a photon is absorbed it generates a cloud of charge which in turn creates a pulse in the preamplifier. This pulse is then compared against a reference threshold in a comparator. Ideally this would be modeled like a step function. However, we saw previously that in reality this transition is smooth, taking on the form of an S-curve. This was previously modeled in the energy calibration as Equation 3.1. This is a straightforward effect to include in the model response function:

$$S(E; E_c, \sigma_E) = \frac{1}{2}\mathrm{erfc}\left(-\frac{E - E_c}{\sqrt{2}\,\sigma_E} \right). \tag{4.26}$$

E_c represents the chosen 50% threshold, or cutoff, energy. σ_E is the width of the S-curve, taken from the calibration data in presented in Section 3.2. It was found that $\sigma_E \approx 300$ eV for lowE settings and $\sigma_E \approx 550$ eV for midE settings. These widths are roughly constant for any chosen threshold value.

As discussed in Section 3.4, properly modeling the effect of charge-sharing is also essential to developing a quantitative understanding of ME-SXR data when applied to a broad-spectrum source like a high-temperature plasma. It was found that charge-sharing leads to a linear distortion in the observed photon counts for a given energy E as the threshold E_c is varied. By way of comparison with Equation 3.27 we can define a charge-sharing response term,

$$F_{CS}(E; E_c) = c \cdot \big(1 - k\,(E - E_c)\big), \tag{4.27}$$

where c is a constant to be determined.

The slope k is given by Equation 3.26, with $f = 0.266$ for the PILATUS3 as previously determined. The constant c is set by the requirement that for a constant photon energy $E = E_0$, $F_{CS} = 1$ when $E_c = E_0/2$. This gives $c = (1+f)^{-1}$, meaning that the full equation is given by

$$F_{CS}(E; E_c) = \frac{1 + 2f \cdot \left(1 - \frac{E_c}{E}\right)}{1+f}. \tag{4.28}$$

The impact of charge-sharing on the resulting pixel response function is illustrated in Figure 4.7. The dashed lines illustrate the basic S-curve model as given by Equation 4.26, and the solid lines illustrate the product of Equations 4.26 and 4.28. In most cases, the inclusion of charge-sharing will reduce the transmission, thereby resulting in fewer counts.

This charge-sharing model can be directly validated by reproducing the threshold scan calibration data, like that shown in Figure 3.24. As in that case, the background signal was removed, and the total counts were normalized to N_{50}, the number of photons detected when the threshold was set to 50% of the incident photon energy (and hence the true number of photons). This data can be directly compared to a simple model in which the detector response function is multiplied by a "narrow-band" source spectrum which integrates to unity, here represented by a highly peaked Gaussian centered upon the known source energy E_0:

$$N(E_c; E_0) = \int_0^\infty S(E; E_c, \sigma_E)\, F_{CS}(E; E_c)\, \frac{1}{\sqrt{2\pi}\delta_E} \exp\left[-\frac{1}{2}\left(\frac{E - E_0}{\delta_E}\right)^2\right] dE, \tag{4.29}$$

where the source width δ_E is typically set to a few eV, effectively mimicking a delta function. Note that no filters were used in the calibration procedure, so none are included in this model. The result of this comparison is shown in Figure 4.8 for three well-resolved calibration scans with lowE settings. The results agree closely with the measured data, lending confidence to the model. Similar agreement was also obtained for midE calibration data.

4.4 Modeling the detector geometry

An accurate model must include knowledge of the geometry of the diagnostic-plasma system. This is considered in three parts: the mapping between spatial coordinates in the plasma and ME-SXR pixel viewing chords (Section 4.4.1), the impact of the viewing angle on signal intensity (Section 4.4.2), and the definition of plasma profiles (Section 4.4.3).

4.4.1 Lines of sight

The line-of-sight geometry of the detector was modeled based on the spatial calibration discussed in Section 3.3, wherein an impact parameter and impact angle were determined for each pixel in the plane of the poloidal cross-section. The detector sees x-rays emitted from the plasma within a cone extending from the pixel into the plasma volume, subtended by the pinhole. In practice, such cones are small and nearly cylindrical, so it suffices to integrate along the center of the cone (called the "line of sight") and scale the result by an étendue factor which accounts for the viewing volume. This means we are ultimately interested in integrating long each pixel's line of sight,

$$\int_{LoS_i} f(\boldsymbol{x}) d|\boldsymbol{x}| = \int_{-L/2}^{L/2} f(z) dz, \tag{4.30}$$

where i indexes each pixel, $f(\boldsymbol{x})$ is an arbitrary scalar field describing the emissivity, $\boldsymbol{x}$

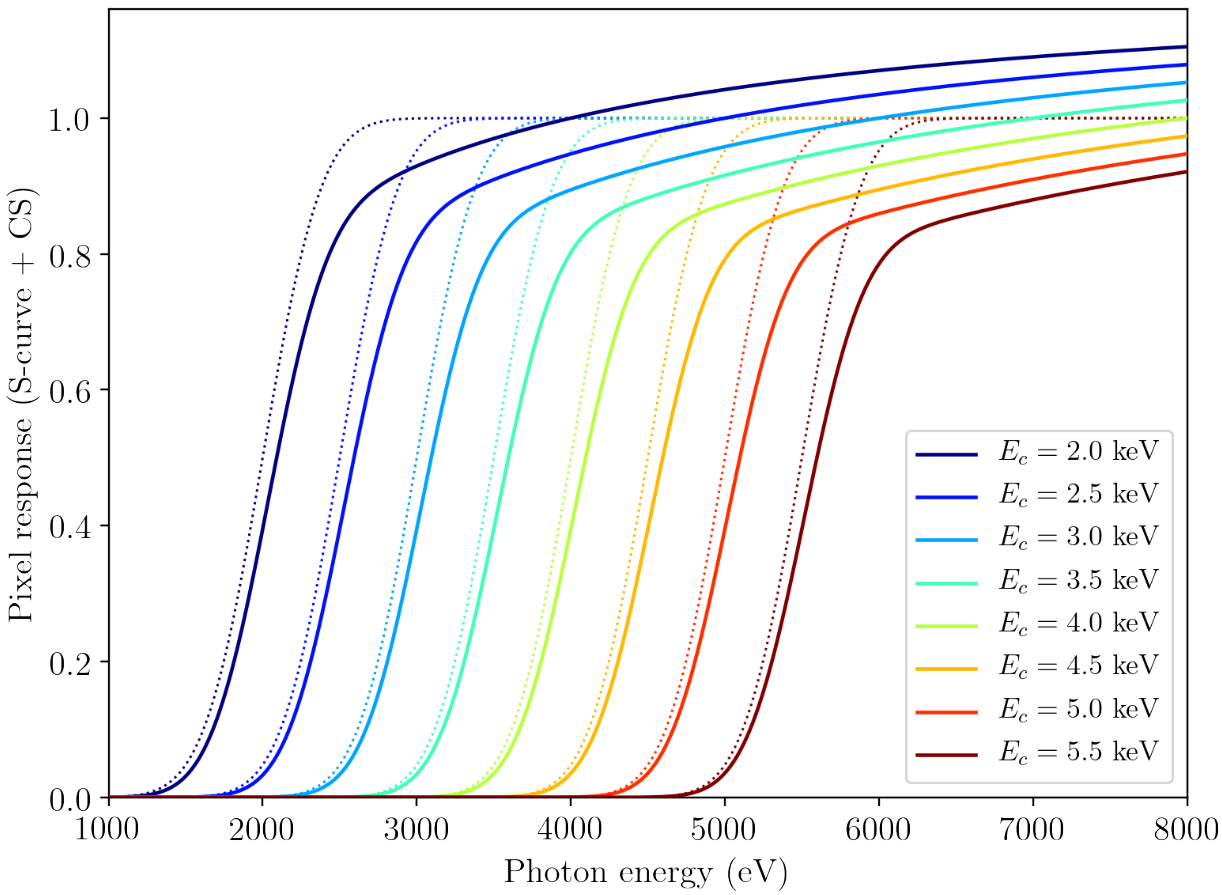

Figure 4.7: Model of the S-curve response with (solid lines) and without (dashed lines) the charge-sharing term given by Equation 4.28 for a range of energy thresholds.

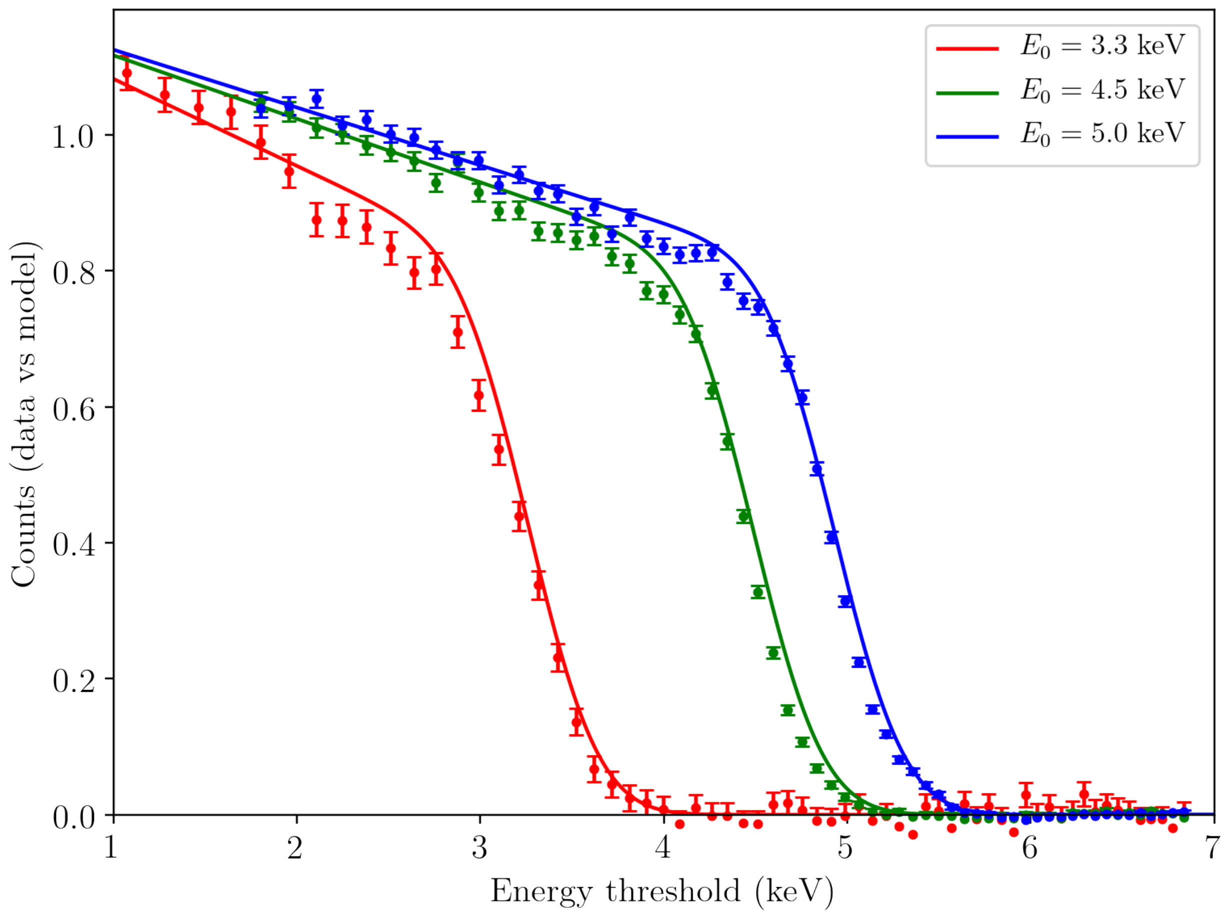

Figure 4.8: Comparison between calibration data (lowE) for a representative pixel and the synthetic threshold scan.

is the position vector, and z is a new one-dimensional coordinate along the line of sight. The calculation of the étendue for each pixel is the focus of Section 4.4.2.

We now want to define a consistent geometry characterizing the transformation between the two-dimensional (x, y) Cartesian coordinate system and a one-dimensional coordinate z defined for an arbitrary line-of-sight. This transformation will be characterized by the impact parameter p and impact angle φ, as defined in Section 3.3. The geometry is depicted in Figure 4.9. In practice we will only need to determine the inverse transformation $A(p, \varphi) : z \mapsto (x, y)$, which will allow us to perform the line integrals. This transformation is given by Equation 4.31. Notice also that if $z = 0$ is defined to be the point of intersection between the perpendicular line p and the chord defining the line-of-sight, the chord always intersects the wall at $\pm\sqrt{a^2 - p^2}$, where $a = 0.52\,\mathrm{m}$ is the minor radius of the MST, meaning these points are the bounds of the integral.

$$\begin{bmatrix} x(z) \\ y(z) \end{bmatrix} = \begin{bmatrix} \cos\varphi & \sin\varphi \\ \sin\varphi & -\cos\varphi \end{bmatrix} \cdot \begin{bmatrix} p \\ z \end{bmatrix} \tag{4.31}$$

Under this scheme we will introduce a simplified notation when evaluating functions along a specified line-of-sight,

$$f(z) \equiv f\big(x(z), y(z)\big). \tag{4.32}$$

Notice that this transformation is simply a rotation of the coordinate system, meaning that no additional scaling factors are required in the integral (i.e., $d|\boldsymbol{x}| = dz$). Putting all of this together, we get the final expression for an integral of a 2D scalar field $f(\boldsymbol{x})$ along a line of sight specified by $(p, \varphi)_i$:

$$\int_{V_i} f(x, y) d|\boldsymbol{x}| = \int_{-\sqrt{a^2 - p_i^2}}^{+\sqrt{a^2 - p_i^2}} f(z) dz. \tag{4.33}$$

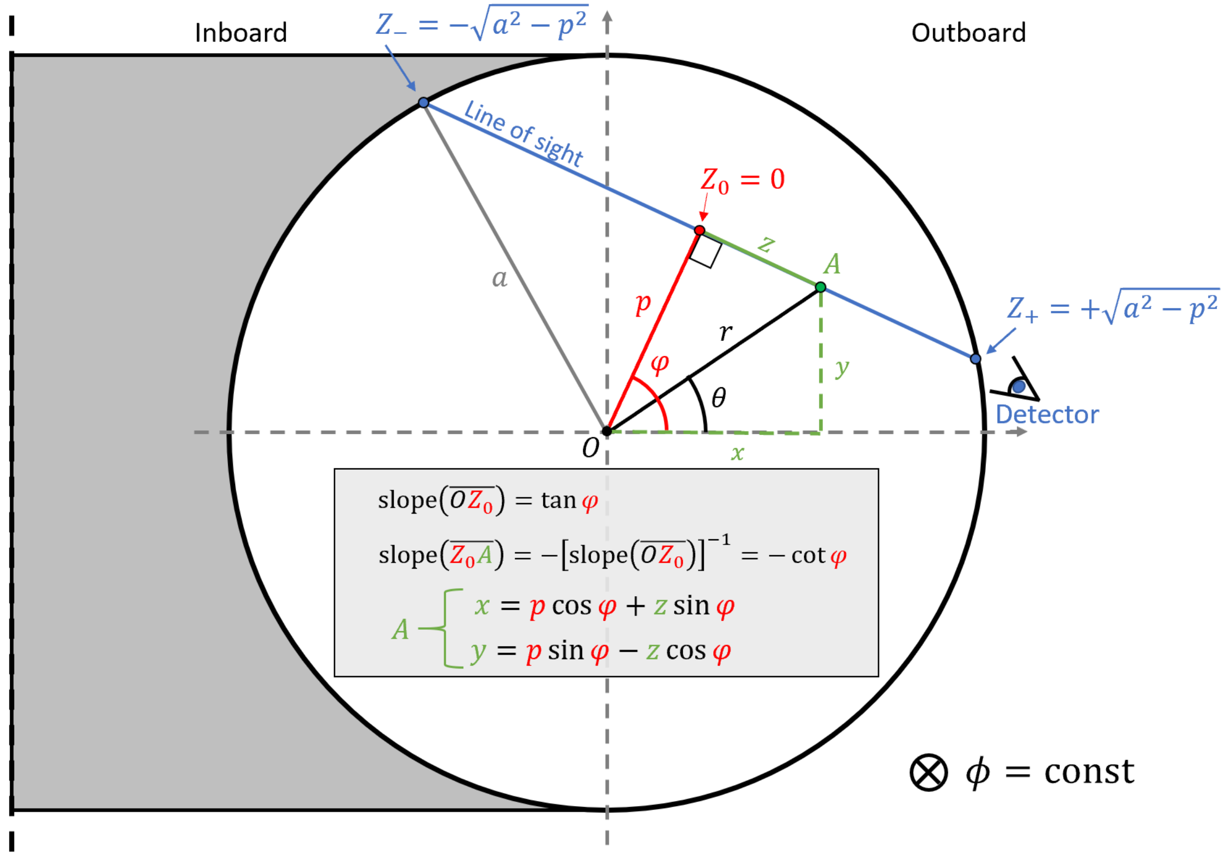

Figure 4.9: Illustration of the line-of-sight geometry for an arbitrary detector within a poloidal cross-section of the MST. Machine coordinates are colored black, the impact parameter and angle are red, the detector and line-of-sight are blue, and the coordinates associated with an arbitrary point A along the line of sight are green. Note that the toroidal angle ϕ increases into the page.

4.4.2 Calculating the pixel étendue

It is also important to determine the detector's étendue, which characterizes how much of the source light reaches a given pixel. Mathematically the étendue dG of a differential surface element relates the radiance L of an emissive source within line-of-sight of the surface element to the power density dP incident on the surface:

$$\begin{aligned} dP &= LdG \\ &= \frac{dG}{4\pi}\varepsilon(z)dz \end{aligned} \tag{4.34}$$

The radiance has been written in terms of the emissivity ε defined in Section 4.2.3. Note that this works equally well whether measured in terms of number of photons or radiated power. The étendue G of an optical detector exposed to a source in a vacuum is given in differential form by the equation:

$$dG = dA_{det} \cos\theta d\Omega \tag{4.35}$$

where dA_{det} is a differential element of the detector, $d\Omega$ is the solid angle which subtends the source, and θ is the angle between the detector element's normal vector and the line of sight to the source. This geometry is illustrated for the ME-SXR diagnostic in Figure 4.10.

This calculation is straightforward for a pixelated pinhole detector like the ME-SXR diagnostic on the MST. We can treat each pixel independently and assume that the photon flux on each pixel is approximately constant such that $dA_{det} \approx A_{pix}$. Furthermore we can compute the pixel solid angle by treating the pixel as a single point. We can then calculate Ω by considering the solid angle from a pixel through the pinhole:

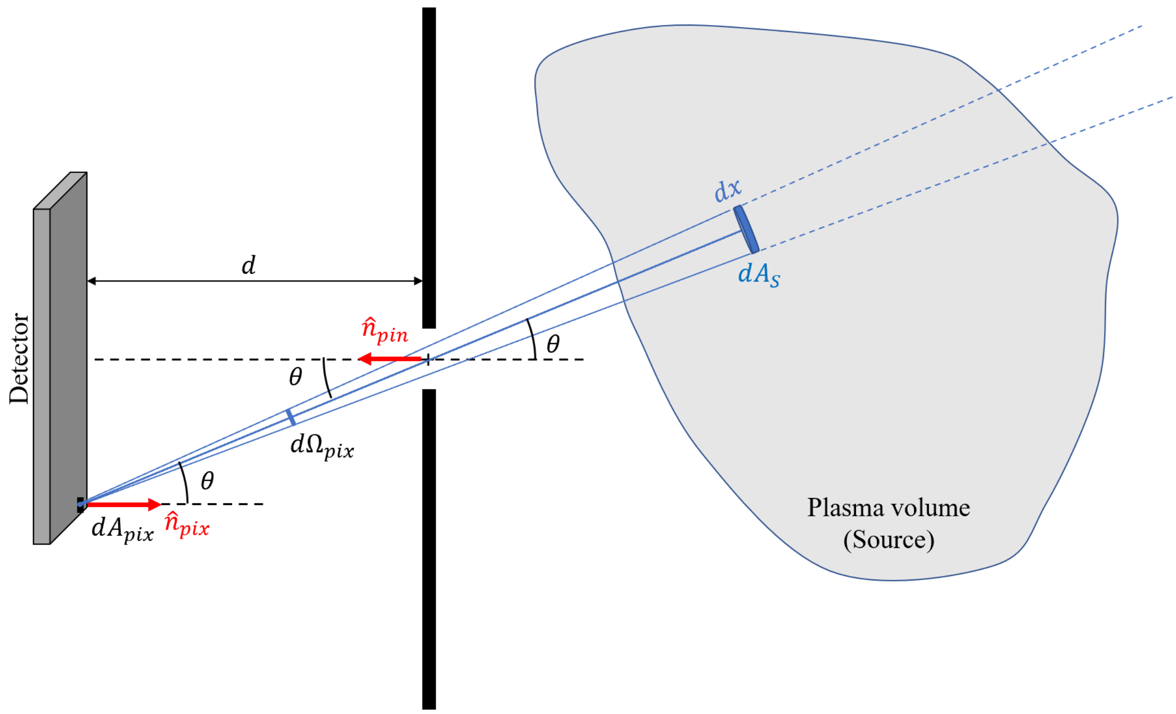

Figure 4.10: Illustration of the pixel-pinhole system solid angle geometry.

$$\begin{aligned}\Omega &= \frac{(\hat{\boldsymbol{r}} \cdot \hat{\boldsymbol{n}}) A_{pin}}{r^2} \\ &= \frac{\cos\theta A_{pin}}{(d/\cos\theta)^2} \\ &= \frac{A_{pin}}{d^2}\cos^3\theta,\end{aligned} \tag{4.36}$$

where A_{pin} is the pinhole area, d is the distance from the detector to the pinhole, and θ is the angle between the pixel's viewing chord and the vector normal to the pinhole. This gives each pixel an étendue of:

$$G = \frac{A_{pin}A_{pix}}{d^2} cos^4\theta. \tag{4.37}$$

We will combine this with the 4π factor from Equation 4.34 to get $\eta = G/4\pi$, also called étendue by some sources. This notation is in line with that used in previous publications on the ME-SXR detector [19]. This term must be calculated independently for each pixel in the 2D array, indexed by (i, j):

$$\eta_{ij} = \frac{A_{pin}A_{pix}}{4\pi d^2}\cos^4\theta_{ij}, \tag{4.38}$$

where $A_{pix} = s^2$ is the area of a pixel with side s. The geometric quantities are labeled in Figure 4.10. The measured values for the distances and areas are given in Table 3.1. The angle θ_{ij} is given by the equation

$$\theta_{i,j} = \arctan\left(\frac{s}{d}\sqrt{g_{ij}}\right). \tag{4.39}$$

The term g_{ij} is a geometric factor which depends only on the pixel indices (i, j) and the total number of pixels in each dimension of the detector screen, $N_x = 195$ and $N_y = 487$. As illustrated in Figure 4.11, given the way I have defined my pixel indices, this factor is given by:

$$g_{ij} = \left(i + \frac{1}{2} - N_x/2\right)^2 + \left(j + \frac{1}{2} - N_y/2\right)^2. \tag{4.40}$$

We can combine this with Equation 4.38 and make use of the relation $cos(\arctan x) = (x^2+1)^{-1/2}$ to rewrite the expression into a more direct form,

$$\eta_{ij} = \frac{A_{pin}A_{pix}}{4\pi d^2}\omega_{ij}, \tag{4.41}$$

where for convenience I have defined that

$$\omega_{ij} \equiv \cos^4\theta_{ij} \tag{4.42}$$

$$= \frac{1}{(\frac{s^2}{d^2}g_{ij}+1)^2}. \tag{4.43}$$

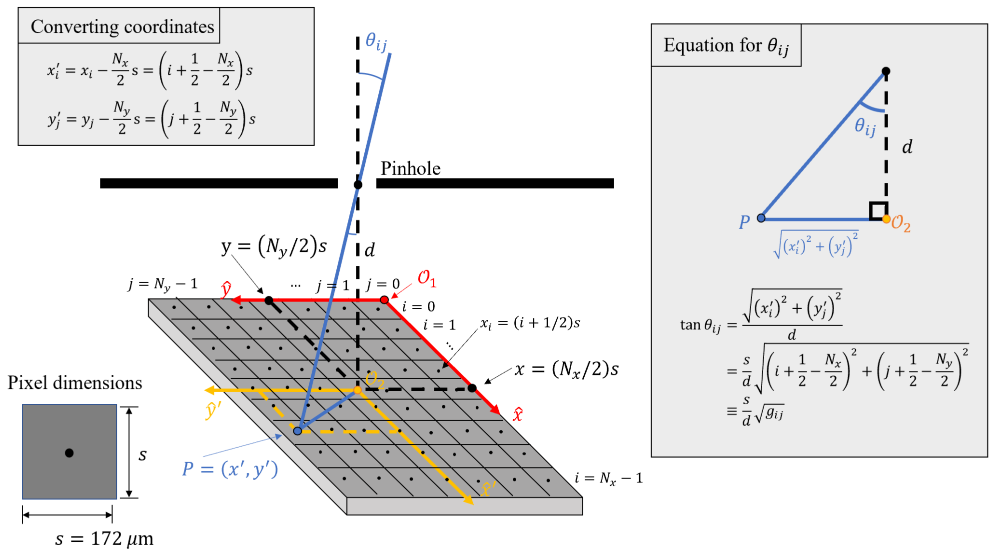

Figure 4.11: Geometric derivation relating the pixel index to the geometric factor g_{ij}, needed to compute the pixel étendue η_{ij}.

The model étendue, as calculated with Equation 4.41, is shown in Figure 4.12. Due to the distance from the pinhole, the étendue near the edges of the detector is very low. This is consistent with what has been observed in the data.

Since we are interested in modeling energy-resolved 1D profiles of the type shown in Figure 3.26, we need to model the detector response for each pixel and then sum over each column. Given the assumption that pixels of the same column effectively see the same plasma (due to toroidal symmetry), this equates to summing the étendues:

$$\eta_i = \sum_{j}^{N_y} \eta_{ij}. \tag{4.44}$$

4.4.3 Plasma profiles and flux coordinates

One final detail to be considered is the mapping between the spatial coordinates $\boldsymbol{x} = (x, y)^T$, typically Cartesian coordinates measured with respect to the MST geometric axis and the plasma's properties, i.e., $T_e(\boldsymbol{x})$, $n_e(\boldsymbol{x})$, etc. This will typically be accomplished with parameterized profiles based on the physics knowledge of the symmetry of the

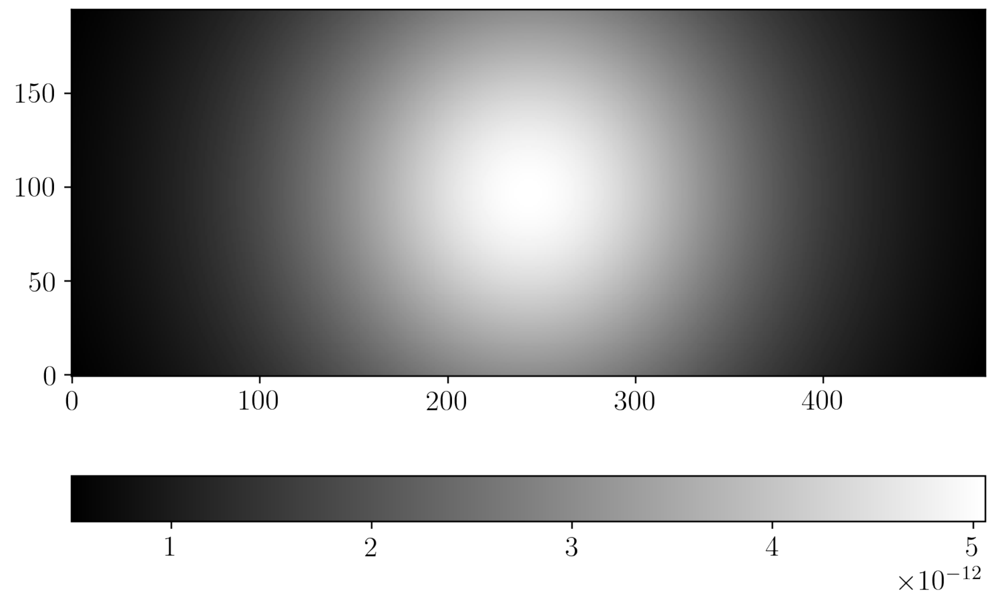

Figure 4.12: Étendue η_{ij} for each pixel on the detector. The x-axis corresponds to the i index and the y-axis to the j index.

profiles.

The simplest possible scenario is axisymmetry, $f(\boldsymbol{x}) = f(r)$. Many physical profiles are well-described by a two-parameter α-β model. The parameters α and β jointly describe the relative flatness of the core as well as and the steepness of the edge,

$$f(r) = \left(1 - (r/a)^{\alpha}\right)^{\beta}, \tag{4.45}$$

where $r = \sqrt{x^2 + y^2}$ is the radial coordinate and $a = 0.52$ m is the MST minor radius.

Such models can often be improved by making use of flux-surface reconstructions to more accurately describe the profile symmetry. For example, standard RFP and PPCD plasma profile symmetry is well-described by a radial-like ρ coordinate which lies on surfaces of constant flux. This can be computed using the MSTfit code [20] and is the basis for the analysis considered in Section 5.5.

4.5 The full forward model

The elements described in the previous section must now be combined into a full forward model of the ME-SXR detector's response to a plasma. The goal is to combine the plasma emissivity, detector geometry, and pixel response into a realistic prediction of total number of photons the detector should measure for a given set of plasma conditions. Mathematically, we want to integrate over both space and energy:

$$n_i^{(k)} = \eta_i \int_{LoS_i} dz \int dE\, \varepsilon(E, z) R_k(E) \tag{4.46}$$

$$= \eta_i \int_{-\sqrt{a^2-p_i^2}}^{+\sqrt{a^2-p_i^2}} dz\, \varepsilon^{(k)}(z), \tag{4.47}$$

where $n_i^{(k)} = N_i^{(k)}/\Delta t$ is the photon count rate for the i^{th} line of sight, k indexes the response functions by threshold, z is the one-dimensional spatial coordinate along the given line of sight, and $\varepsilon^{(k)}(z)$ is the relative emissivity for the k^{th} response function evaluated along that line of sight (as described in Section 4.2.3). Also note that we have used the summed version of η_i, as defined in Equation 4.44, because we are interested in simulating the sum of a row of pixels.

Due to the complex nature of the observed emissivity $\varepsilon^{(k)}$, a general closed-form solution to this spatial integral is not possible. Instead, numerical integration must be used. It is straightforward enough to iterate over a series of sums to estimate the value of the line integral for each detector chord i:

$$n_i^{(k)} \approx \eta_i \Delta Z \sum_j \varepsilon^{(k)}(z_{ij}), \tag{4.48}$$

where z_{ij} is the j^{th} point along the i^{th} detector chord using the transformation described in Section 4.4.

The output of the ME-SXR model can be compared directly to actual data to confirm its similarity. Figure 4.13 shows model temperature, electron density, neutral density, impurity density, and Z_{eff} profiles which are considered to be typical for high-current (400-500 kA) PPCD plasmas (see Section 5.5.1). These profiles will be used to generate a synthetic set of profiles comparable to those shown in Figure 3.26.

The plasma profiles are used to calculate the plasma emissivity $\varepsilon^{(k)}(z_{ij})$ for a set of points along each chord, given that chord's lower threshold setting $E_c^{(k)}$. An example of the resulting plasma emissivity is shown in Figure 4.14. This profile is then integrated spatially along each chord to produce realistic synthetic measurements. This output is shown in Figure 4.15.

It is immediately clear that the ME-SXR model output is both qualitatively and quantitatively similar to the real measurements shown in Figure 3.26. This similarity, when coupled with our previous experience modeling and validating similar SXR diagnostics [21], provides a great deal of confidence in the model. This model is further bolstered by the results discussed in Section 5.5, in which the model is found to be consistent with simultaneous measurements obtained by the SXR tomography diagnostic.

It is also worth taking a moment to consider what the model is not and possible ways it could be extended in future work. By using ADAS, the model assumes that that both the electrons and ions obey strict Maxwellian distributions. This could be extended to support other equilibria such as the κ distribution, which accounts for a small runaway population using calculations such as those in [22], though this would require substantial work. The model also presently does not include any time-dependent phenomena (like transport) other than that implied by chosen profile shapes. It is intended only to simulate ME-SXR measurements at a particular instance in time given some instantaneous plasma profiles.

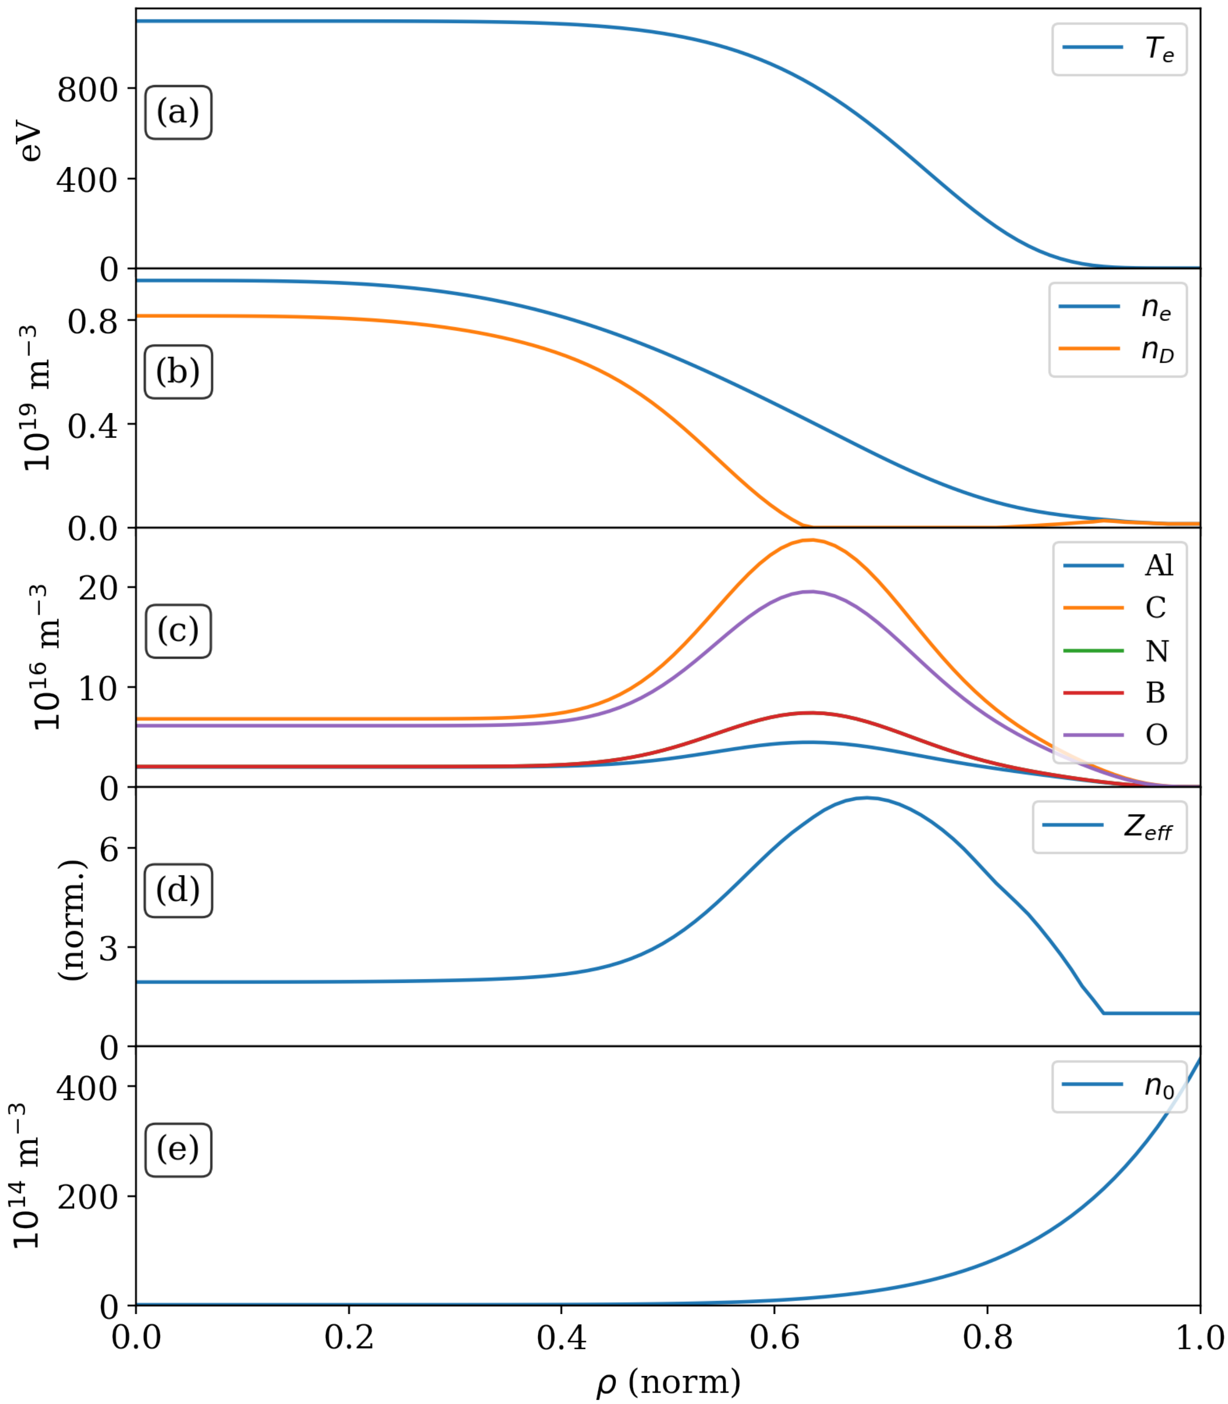

Figure 4.13: Example plasma profiles as a function of a Shafranov-shifted radial coordinate ρ. These profiles were chosen based on typical profile values and shapes for high-current (400-500 kA) PPCD plasmas in the MST.

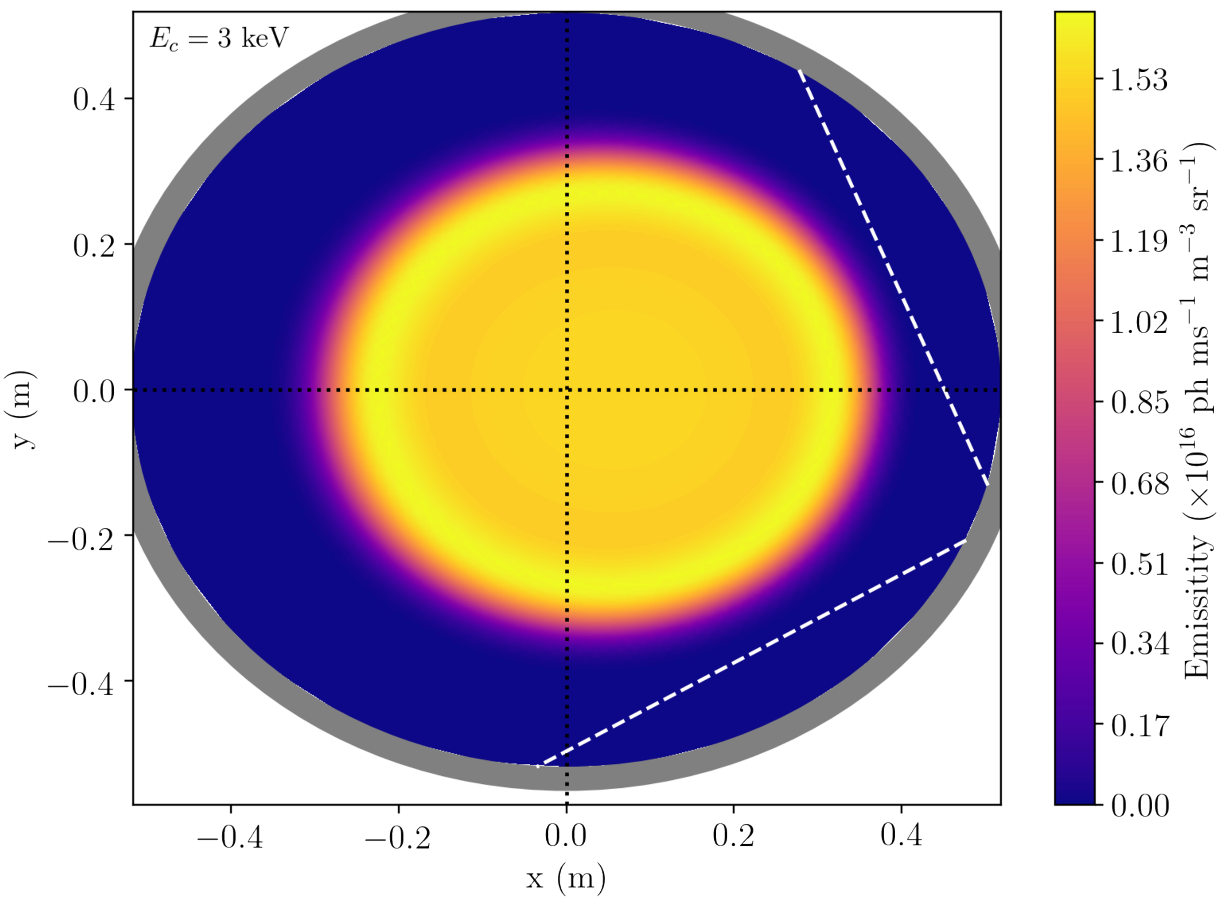

Figure 4.14: Local emissivity map $\varepsilon^{(3.0)}$ for the ME-SXR detector with a lower threshold of $E_c = 3.0$ keV given the plasma profiles shown in Figure 4.13. The core is bright and very nearly flat, but emissivity drops off rapidly with temperature and density near the mid-radius. The dashed line represents the boundaries of the ME-SXR diagnostic's viewing angle.

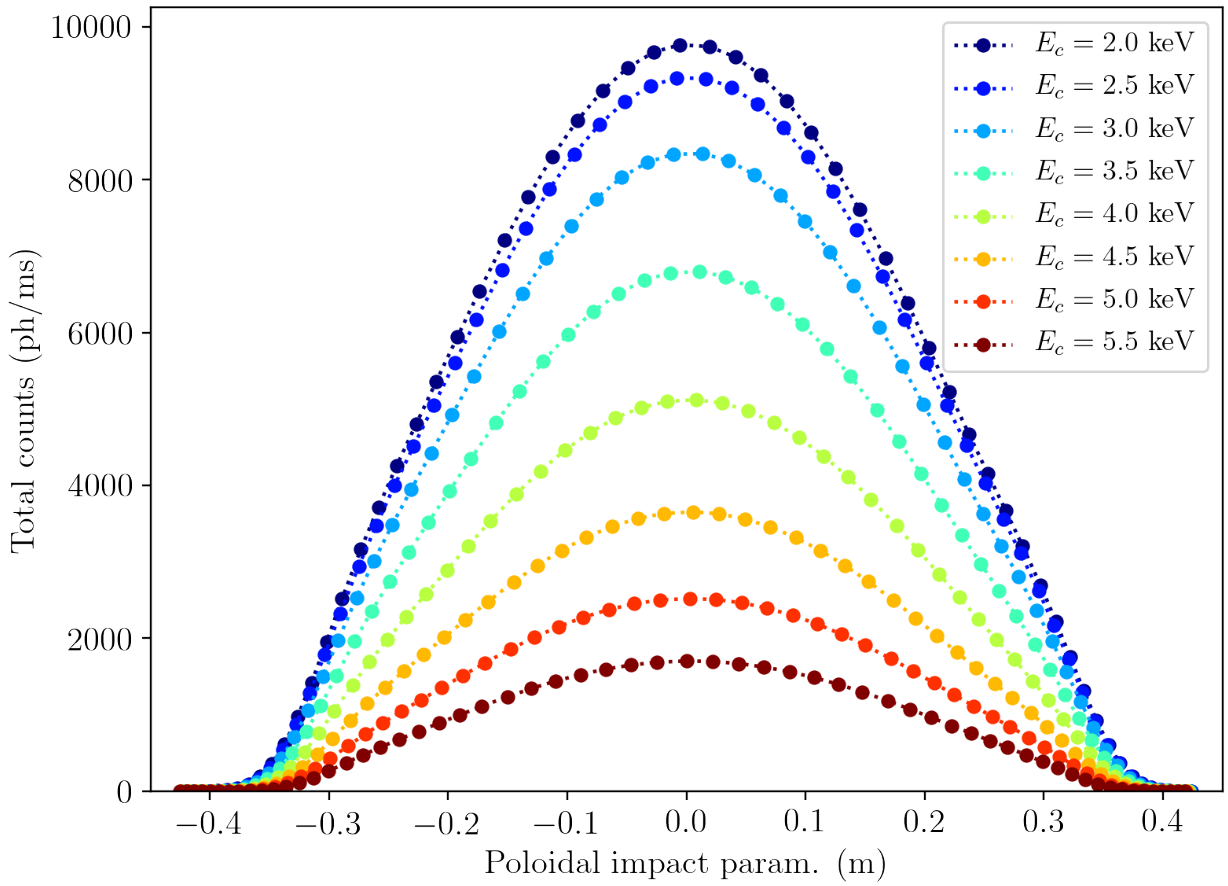

Figure 4.15: Full synthetic ME-SXR output for the the plasma profiles shown in Figure 4.13. The resulting output is both qualitatively and quantitatively similar to Figure 3.26.

4.6 Uncertainty analysis

In order to make meaningful comparisons between the ME-SXR model and collected data, we must first understand the model's limitations. Numerous parameters that are based on real-world measurements are subject to measurement uncertainty, and these uncertainties propagate to the final model output. This section considers three significant sources of uncertainty in the model: the étendue η_i, the Mylar filter thickness t_{My}, and the neutral density n_0.

The calculation of the pixel étendue, as discussed in Section 4.4.2, depends on three parameters which are subject to measurement uncertainty: the pinhole area A_{pin}, the distance from the detector screen to the pinhole d, and the area of each pixel A_{pix}. Recall from Section 4.4.2 that the étendue can be represented as $\eta_i = \eta_0 \sum_j \omega_{ij}$ where

$$\eta_0 = \frac{A_{pix} A_{pin}}{4\pi d^2}. \tag{4.49}$$

We can assume that the uncertainty is mostly due to the η_0 factor. The remaining factor is summed over each pixel, so assuming uncorrelated errors we expect overestimates and underestimates to approximately cancel out. This leaves

$$\begin{aligned}\sigma_{\eta_0}/\eta_0 &= \frac{1}{\eta_0}\left[\left(\frac{\partial \eta_0}{\partial A_{pix}}\sigma_{pix}\right)^2 + \left(\frac{\partial \eta_0}{\partial A_{pin}}\sigma_{pin}\right)^2 + \left(\frac{\partial \eta_0}{\partial d}\sigma_d\right)^2\right]^{1/2} \\ &= \left[\left(\frac{\sigma_{pix}}{A_{pix}}\right)^2 + \left(\frac{\sigma_{pin}}{A_{pin}}\right)^2 + \left(\frac{2\sigma_d}{d}\right)^2\right]^{1/2}.\end{aligned} \tag{4.50}$$

Given that $s = 172$ μm, $A_{pix} = s^2$, $A_{pin} = 2$ mm^2, $d = 30.5$ mm, $\sigma_d = 1$ mm^2, $\sigma_{pin} = 0.224$ mm^2, and $\sigma_{pix} = 3.44 \times 10^{-4}$ mm^2 (as reported in Table 3.1), then $\sigma_{\eta_0}/\eta_0 = 0.13$. This uncertainty is dominated by the uncertainty in the pinhole area (around $\sim 11\%$). A more complex calculation which treats each pixel in the sum individually was also considered; however, the result was found to be nearly identical. This is again due to the fact that the pinhole area, which is independent of pixel label, is the dominant term.

Another important consideration is the uncertainty in the thickness of the beryllium and Mylar filters. The filter transmission model discussed in Section 4.3.1 was built using the nominal thickness values as listed in Table 3.1. However, we need to account for the possibility that these nominal values might vary within the specified tolerances and that this could result in a significant variation in the final photon count rate. This variation can be investigated using the model itself. A 0.1μm variation in the beryllium thickness was found to be inconsequential. However, the ± 10 μm tolerance on the Mylar filter was found to be important. The effect of this variation on the response function is shown in Figure 4.16.

The impact of this unknown variation on the detector signal was investigated by taking a basic Monte Carlo approach. Assuming a mean Mylar thickness of 100 μm and a standard deviation equal to the stated tolerance (10 μm), I generated a set of filter thicknesses $t_{My} \sim \mathcal{N}(100\,\mu\text{m}, 10\,\mu\text{m})$ and calculated the resulting emissivity according to Equations 4.10 and 4.13. This procedure was repeated over a range of electron temperatures, holding other parameters fixed ($n_e = 1 \times 10^{19}\text{m}^{-3}$ and $n_0 = 2 \times 10^{14}\text{m}^{-3}$). The results of this scan are shown in Figure 4.17. This shows that the relative uncertainty ranges from about $10 - 20\%$, decreasing as the temperature increases.

We see that for a plasma with core temperature $T_e \sim 1000$ eV, we would estimate the contribution from the Mylar filter as $(\delta\varepsilon/\varepsilon)_{Mylar} \approx 0.12$.

As discussed in Section 4.2.2, uncertainty in the neutral density input into the model is also a potential source of uncertainty in the model output. The n_0 profiles used in subsequent chapters will be based off of previous work using ensemble measurements and therefore are not expected to perfectly represent any given plasma discharge. We will therefore assume an inherent variation on the order of 10%, which as seen in Figure 4.5 leads to an uncertainty in the emissivity of roughly 5%.

Now we need to combine these uncertainties in a consistent way. We can make use

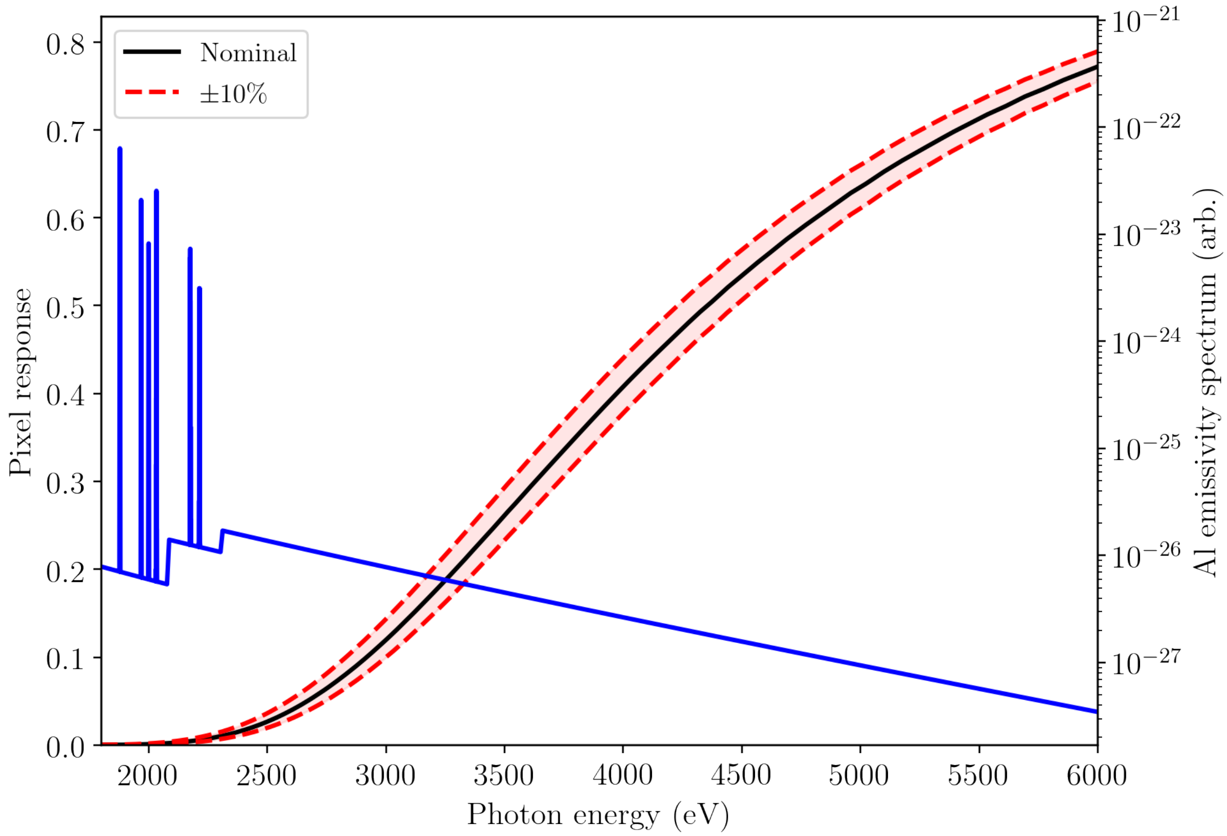

Figure 4.16: ME-SXR response function including a Mylar filter of thickness of 100 μm, with dashed lines illustrating the effect of varying the thickness by $\pm10\%$.

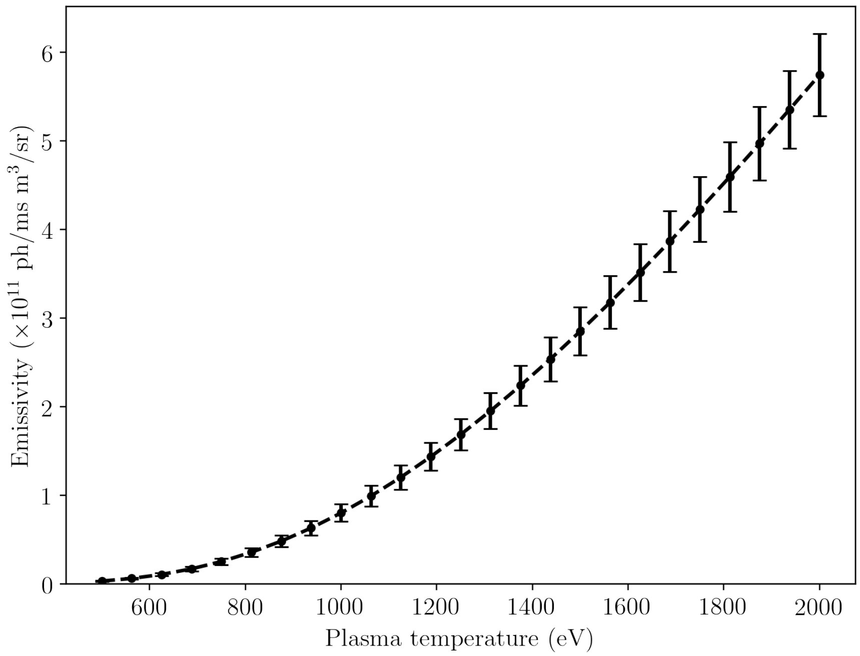

Figure 4.17: Mean and standard deviation of the local emissivity obtained by sampling over the Mylar filter thickness $t_{My} \sim \mathcal{N}(100\,\mu\text{m}, 10\,\mu\text{m})$. The results are plotted against electron temperature, holding other parameters fixed at typical values.

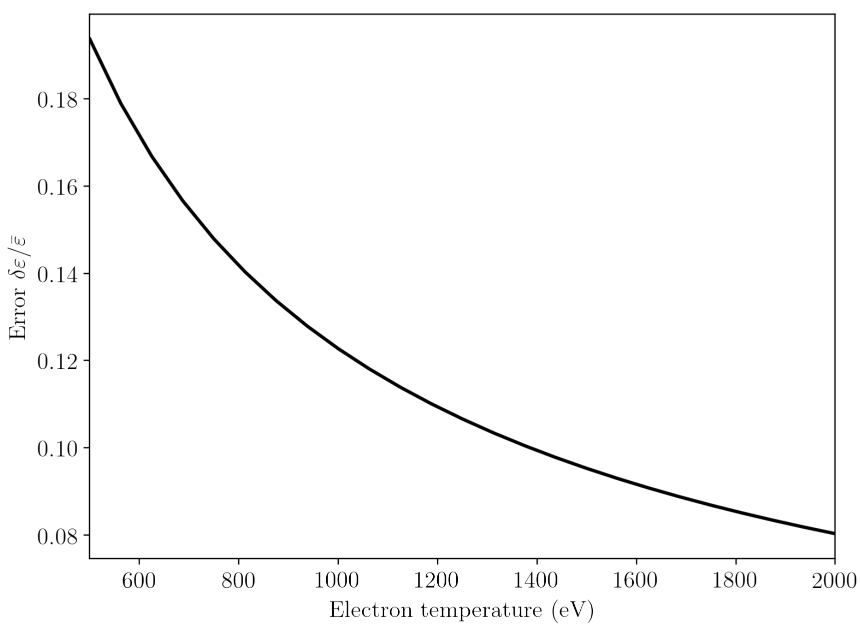

Figure 4.18: Variation in local emissivity normalized to the nominal emissivity, as a function of electron temperature.

of the fact that plasma emissivity profiles tend to be fairly flat in the core such that

$$n = \eta \int dz\, \varepsilon(z) dz \approx \eta \varepsilon L, \tag{4.51}$$

where ε is the emissivity in the core and L is the chord length through the core. The uncertainty is then given by

$$\begin{aligned} \frac{\delta n}{n} &\approx \left[\left(\frac{\delta \eta}{\eta} \right)^2 + \left(\frac{\delta \varepsilon}{\varepsilon} \right)^2_{Mylar} + \left(\frac{\delta \varepsilon}{\varepsilon} \right)^2_{n_0} \right]^{1/2} \\ &\approx \left[(0.13)^2 + (0.12)^2 + (0.05)^2 \right]^{1/2} \\ &\approx 0.18. \end{aligned} \tag{4.52}$$

In the remaining chapters we will make use of this as an estimate for the inherent uncertainty in the model, denoted by $\sigma_m \approx 18\%$. The magnitude of the uncertainty is mostly due to two factors: the necessary addition of the Mylar filter in order to combat saturation from low-energy Al photons and the relatively large uncertainty in pinhole area.

Finally, it is worth saying a few words about a source of uncertainty that was not considered: the atomic physics calculations performed by ADAS. ADAS makes use of numerous databases of atomic physics factors such as interaction cross-sections $\langle \sigma v \rangle$, themselves drawn from a range of experimental measurements and theoretical computations. In most cases the uncertainties on these factors have not been assessed (or reported), and it is therefore impossible to incorporate that source of uncertainty into σ_M. For the remainder of this thesis we will operate under the assumption that this source of uncertainty is small relative to the sources already discussed.

Bibliography

[1] R. Fischer, C. J. Fuchs, B. Kurzan, W. Suttrop, and E. Wolfrum, "Integrated data analysis of profile diagnostics at ASDEX upgrade," *Fusion Science and Technology*, vol. 58, no. 2, pp. 675–684, 2010. [Online]. Available: https://doi.org/10.13182/FST10-110

[2] L. M. Reusch, M. D. Nornberg, J. A. Goetz, and D. J. Den Hartog, "Using integrated data analysis to extend measurement capability (invited)," *Review of Scientific Instruments*, vol. 89, no. 10, 2018. [Online]. Available: https://doi.org/10.1063/1.5039349

[3] I. H. Hutchinson, *Principles of Plasma Diagnostics*, 2nd ed. Cambridge University Press, jul 2002.

[4] H. R. Griem, *Principles of Plasma Spectroscopy*. Cambridge, UK: Cambridge University Press, 1997.

[5] H. P. Summers and M. G. O'Mullane, "Atomic data and modelling for fusion: The ADAS Project," *AIP Conference Proceedings*, vol. 1344, no. May 2011, pp. 179–187, 2011. [Online]. Available: https://doi.org/10.1063/1.3585817

[6] "ADAS: Docmentation." [Online]. Available: http://www.adas.ac.uk/manual.php

[7] S. T. Kumar, D. J. Den Hartog, B. E. Chapman, M. O'Mullane, M. Nornberg, D. Craig, S. Eilerman, G. Fiksel, E. Parke, and J. Reusch, "High resolution charge-exchange spectroscopic measurements of aluminum impurity ions in a high temperature plasma," *Plasma Physics and Controlled Fusion*, vol. 54, no. 1, 2012. [Online]. Available: https://doi.org/10.1088/0741-3335/54/1/012002

[8] T. Barbui, L. Carraro, D. J. Den Hartog, S. T. Kumar, and M. Nornberg, "Impurity transport studies in the Madison Symmetric Torus reversed-field pinch during standard and pulsed poloidal current drive regimes," *Plasma Physics and Controlled Fusion*, vol. 56, no. 7, 2014. [Online]. Available: https://doi.org/10.1088/0741-3335/56/7/075012

[9] Z. A. Xing, "Ion Thermal Transport and Heating in Reduced Tearing RFP," Ph.D. dissertation, University of Wisconsin-Madison, 2019.

[10] R. J. Norval, "Plasma-neutral interactions as an energy sink in the edge of the Madison Symmetric Trous," Ph.D. dissertation, University of Wisconsin-Madison, 2019.

[11] H. Yamazaki, L. F. Delgado-Aparicio, R. Groebner, B. Grierson, K. Hill, N. Pablant, B. Stratton, P. Efthimion, A. Ejiri, Y. Takase, and M. Ono, "A computational tool for simulation and design of tangential multi-energy soft x-ray pin-hole cameras for tokamak plasmas," *Review of Scientific Instruments*, vol. 89, no. 10, 2018. [Online]. Available: http://dx.doi.org/10.1063/1.5038788

[12] A. Burgess, D. G. Hummer, and J. A. Tully, "Electron impact excitation of positive ions," *Philosophical Transactions of the Royal Society of London*, vol. 266, no. 1175, pp. 225–279, apr 1970. [Online]. Available: https://doi.org/10.1098/rsta.1970.0007

[13] P. A. van Hoof, R. J. Williams, K. Volk, M. Chatzikos, G. J. Ferland, M. Lykins, R. L. Porter, and Y. Wang, "Accurate determination of the free-free gaunt factor - I. Non-relativistic gaunt factors," *Monthly Notices of the Royal Astronomical Society*, vol. 444, no. 1, pp. 420–428, 2014. [Online]. Available: https://doi.org/10.1093/mnras/stu1438

[14] A. Kramida, Y. Ralchenko, J. Reader, and The NIST ASD Team (2018), "NIST Atomic Spectra Database (ver. 5.6.1)," 2019. [Online]. Available: http://physics.nist.gov/asd

[15] H. P. Summers, W. J. Dickson, M. G. O'Mullane, N. R. Badnell, A. D. Whiteford, D. H. Brooks, J. Lang, S. D. Loch, and D. C. Griffin, "Ionization state, excited populations and emission of impurities in dynamic finite density plasmas: I. The generalized collisional-radiative model for light elements," *Plasma Physics and Controlled Fusion*, vol. 48, no. 2, pp. 263–293, 2006. [Online]. Available: https://doi.org/10.1088/0741-3335/48/2/007

[16] E. Jones, T. Oliphant, P. Peterson, and Others, "SciPy: Open source scientific tools for Python," 2001. [Online]. Available: http://www.scipy.org/

[17] M. Galante, L. Reusch, D. Den Hartog, P. Franz, J. Johnson, M. McGarry, M. Nornberg, and H. Stephens, "Determination of Z_eff by integrating measurements from x-ray tomography and charge exchange recombination spectroscopy," *Nuclear Fusion*, vol. 55, no. 12, p. 123016, 2015. [Online]. Available: https://doi.org/10.1088/0029-5515/55/12/123016

[18] B. L. Henke, E. M. Gullikson, and J. C. Davis, "X-ray interactions: Photoabsorption, scattering, transmission, and reflection at E = 50-30, 000 eV, Z = 1-92," pp. 181–342, 1993. [Online]. Available: https://henke.lbl.gov/optical{_}constants/

[19] L. Delgado-Aparicio, J. Wallace, H. Yamazaki, P. Vanmeter, L. Reusch, M. Nornberg, A. Almagari, J. Maddox, B. Luethi, M. Rissi, T. Donath, D. Den Hartog, J. Sarff, P. Weix, J. Goetz, N. Pablant, K. Hill, B. Stratton, P. Efthimion, Y. Takase, A. Ejiri, and M. Ono, "Simulation, design, and first test of a multi-energy soft x-ray (SXR) pinhole camera in the Madison Symmetric Torus (MST)," *Review of Scientific Instruments*, vol. 89, no. 10, 2018. [Online]. Available: https://doi.org/10.1063/1.5038798

[20] J. K. Anderson, C. B. Forest, T. M. Biewer, J. S. Sarff, and J. C. Wright, "Equilibrium reconstruction in the Madison Symmetric Torus reversed field pinch," *Nuclear Fusion*, vol. 44, no. 1, pp. 162–171, 2004. [Online]. Available: https://doi.org/10.1088/0029-5515/44/1/018

[21] L. M. Reusch, P. Franz, D. J. Den Hartog, J. A. Goetz, M. D. Nornberg, and P. VanMeter, "Model validation for quantitative X-ray measurements," *Fusion Science and Technology*, vol. 74, no. 1-2, pp. 167–176, 2018. [Online]. Available: https://doi.org/10.1080/15361055.2017.1404340

[22] J. Dudík, J. Kašparová, E. Dzifčáková, M. Karlický, and S. MacKovjak, "The non-Maxwellian continuum in the X-ray, UV, and radio range," *Astronomy and Astrophysics*, vol. 539, pp. 1–12, 2012. [Online]. Available: https://doi.org/10.1051/0004-6361/201118345

Chapter 5

Interpreting ME-SXR data

Now that the ME-SXR diagnostic has been calibrated, configured, and installed onto the MST, we need to develop methods to interpret the resulting data and extract useful physical information. That is the focus of this chapter. The diagnostic forward model, developed in Chapter 4, will be of critical importance to this endeavor. Additionally, this chapter will also detail particular challenges encountered when operating the diagnostic, and how these were overcome.

Section 5.1 demonstrates how subsequent ME-SXR images can be combined to observe the temporal evolution of the plasma's soft x-ray emission profile under various conditions. Section 5.2 uses argon doping to demonstrate the diagnostic's sensitivity to mid-Z impurities. Section 5.3 discusses techniques for extracting the electron temperature profile from ME-SXR images. Section 5.4 discusses early difficulties related to detector saturation and the modifications required to suppress these distortions. And finally, Section 5.5 discusses an integrated data analysis framework which can be used to synthesize ME-SXR measurements with other complementary diagnostics to extract not just temperature but also ion density profiles.

5.1 Time evolution of plasma soft x-ray emission

During normal operations, the ME-SXR records a series of images at the selected cycle rate (usually 500 Hz). Once an image is recorded, the data can be separated into one-dimensional profiles in the poloidal plane according to threshold as shown in Figure 3.26. The evolution of these profiles over time can provide significant insight into the heating and profile evolution of a plasma. An example is shown in Figure 5.1, which depicts the evolution of a single PPCD plasma (see Section 1.4) during the good confinement period with eight distinct thresholds. PPCD has been well-studied and thoroughly characterized over the last two decades, and as such it makes a good test case for a new diagnostic. Initially the SXR emissivity is very low, corresponding to a low T_e. However once the PPCD banks begin to discharge, driving $J_{||}$ and thereby suppressing core tearing mode activity and reducing thermal transport, the plasma emissivity begins to increase rapidly. This emissive structure gradually broadens as the plasma bulk continues to heat. This comes to a sudden end at approximately 21 ms, once the PPCD banks have discharged all of their available energy. The "improved confinement" period rapidly collapses as the return of tearing modes allows the stored thermal energy to rapidly transport out of the core. This frequently manifests as a very bright flash on a single frame.

The diagnostic has also been applied to high-current non-reversed ($F = 0$) quasi-single helicity (QSH) plasmas (also described in Section 1.4). The plasma transitions from multi-helicity to a quasi-single helicity state during the "rise phase" from about 15-20 ms, during which the rotating structure locks to the wall and the SXR emissivity reaches a maximum. As the mode saturates the SXR emissivity begins to decrease at an approximately linear rate, even though the electron density is held constant (Section 6.4). This has to do with how re-emergence of secondary tearing mode activity affects the thermal confinement and production of runaway electrons, which the PILATUS3's 450 μm Si sensor can detect. The phenomenology of QSH evolution, and how this manifests

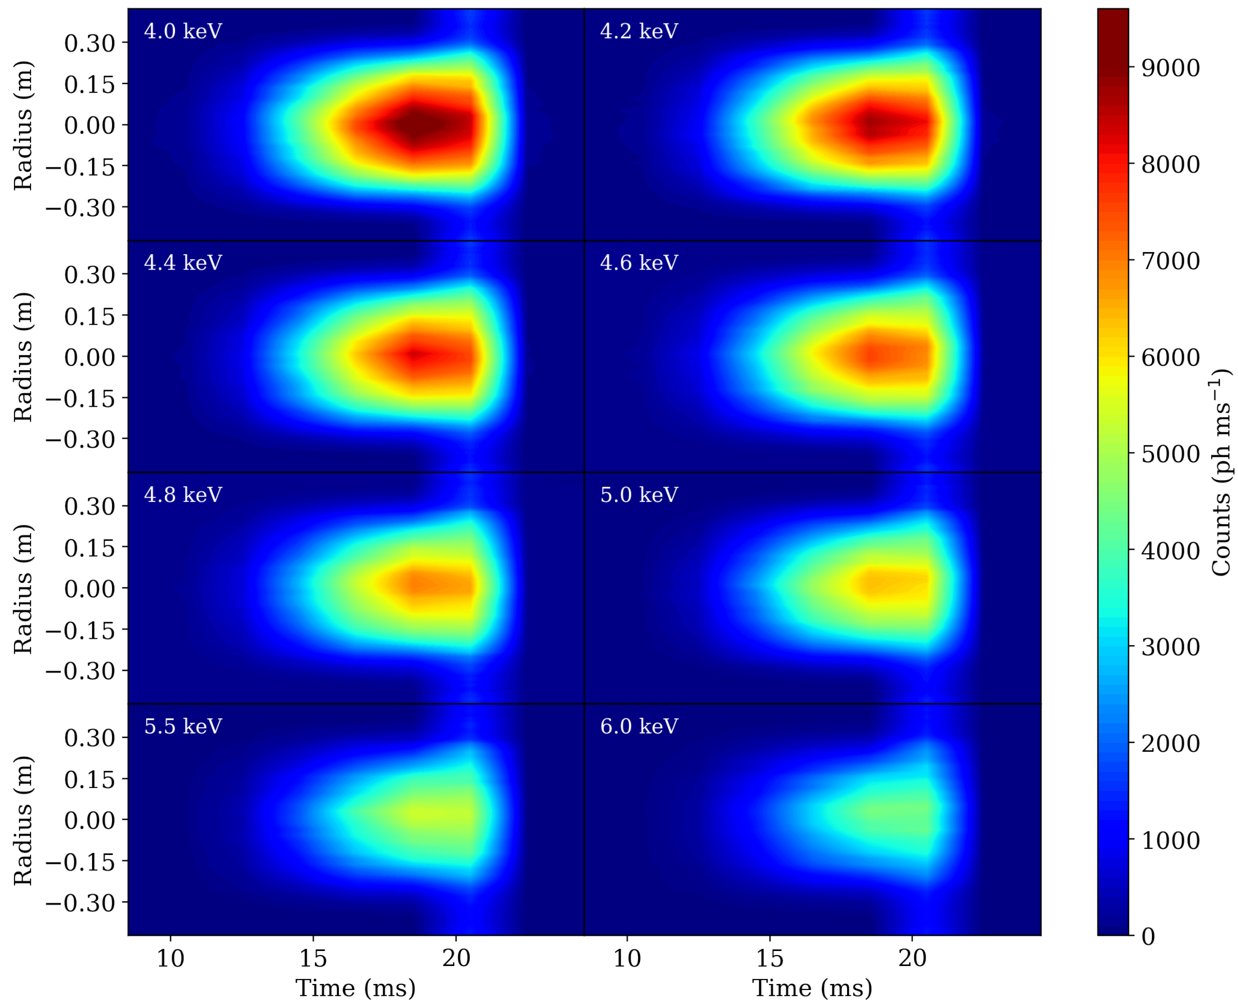

Figure 5.1: Evolution of observed soft x-ray emission during a single PPCD discharge. Each panel corresponds to a different energy threshold. The plasma heats up and becomes increasingly emissive as tearing modes are suppressed and thermal transport is reduced. This ends suddenly around 21 ms as the improved confinement period ends.

in ME-SXR measurements, is discussed at length in Chapter 6.

A related phenomenon can be observed during single frames of many QSH discharges. As shown in Figure 5.3, a "ring-like structure" can be observed. The count rate in the core of the plasma remains relatively unchanged, while the substantially more photons are seen near the edge. Furthermore, these photons seem to increase the count rate for all thresholds, implying they are high energy. These images are known to be connected to the generation of runaway electrons. The unusual geometry is due to the relatively large size of the diagnostic's porthole, causing a break in the magnetic symmetry and resulting in error fields penetrating through the porthole and into the detector housing or the porthole edge. Whenever there is an interruption in the core confinement, the energized electrons stream along the field lines through the porthole before colliding with the housing and emitting target emission. We will return to this discussion in Section 6.3, which focuses on the observation and diagnosis of runaway electrons during QSH states.

5.2 Detection of mid-Z impurities

The ME-SXR diagnostic can also be used to identify the presence of mid-Z[1] impurities in the plasma. Such impurities will be partially ionized and typically feature strong emission lines comfortably within the detector's energy range, which can easily be discerned by a proper selection of thresholds. This was tested by injecting small amounts of argon gas into non-reversed plasmas as a pre-fill, allowing the gas to diffuse throughout the vacuum vessel before the plasma is formed. As illustrated in Figure 5.4, for typical MST parameters Ar features strong emission lines around 3 keV. The ME-SXR diagnostic allows us to set thresholds both above and below this energy allowing for a direct detection of these photons.

[1]The term "mid-Z," as used here, roughly means "atomic number Z of approximately one to a few dozen," though there is not a hard cutoff.

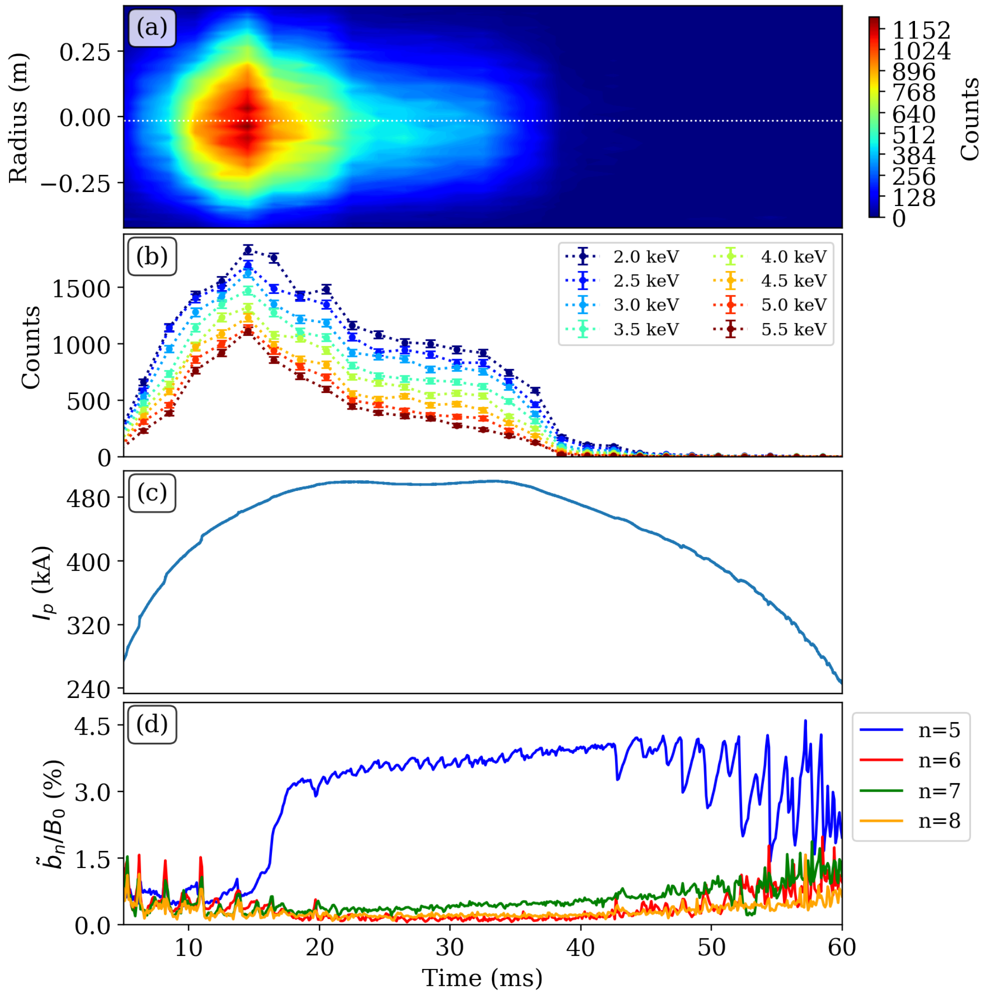

Figure 5.2: Evolution of observed x-ray emission during the course of a single non-reversed ($F = 0$) plasma discharge in which a QSH state forms. Panel (a) shows the the evolution of a single emissive profile (with a lower threshold of 2 keV), (b) shows the temporal evolution of the core-most chord for each threshold, (c) is the plasma current, and (d) shows the evolution of the dominant (blue) and secondary (others) mode magnetic fluctuation amplitudes.

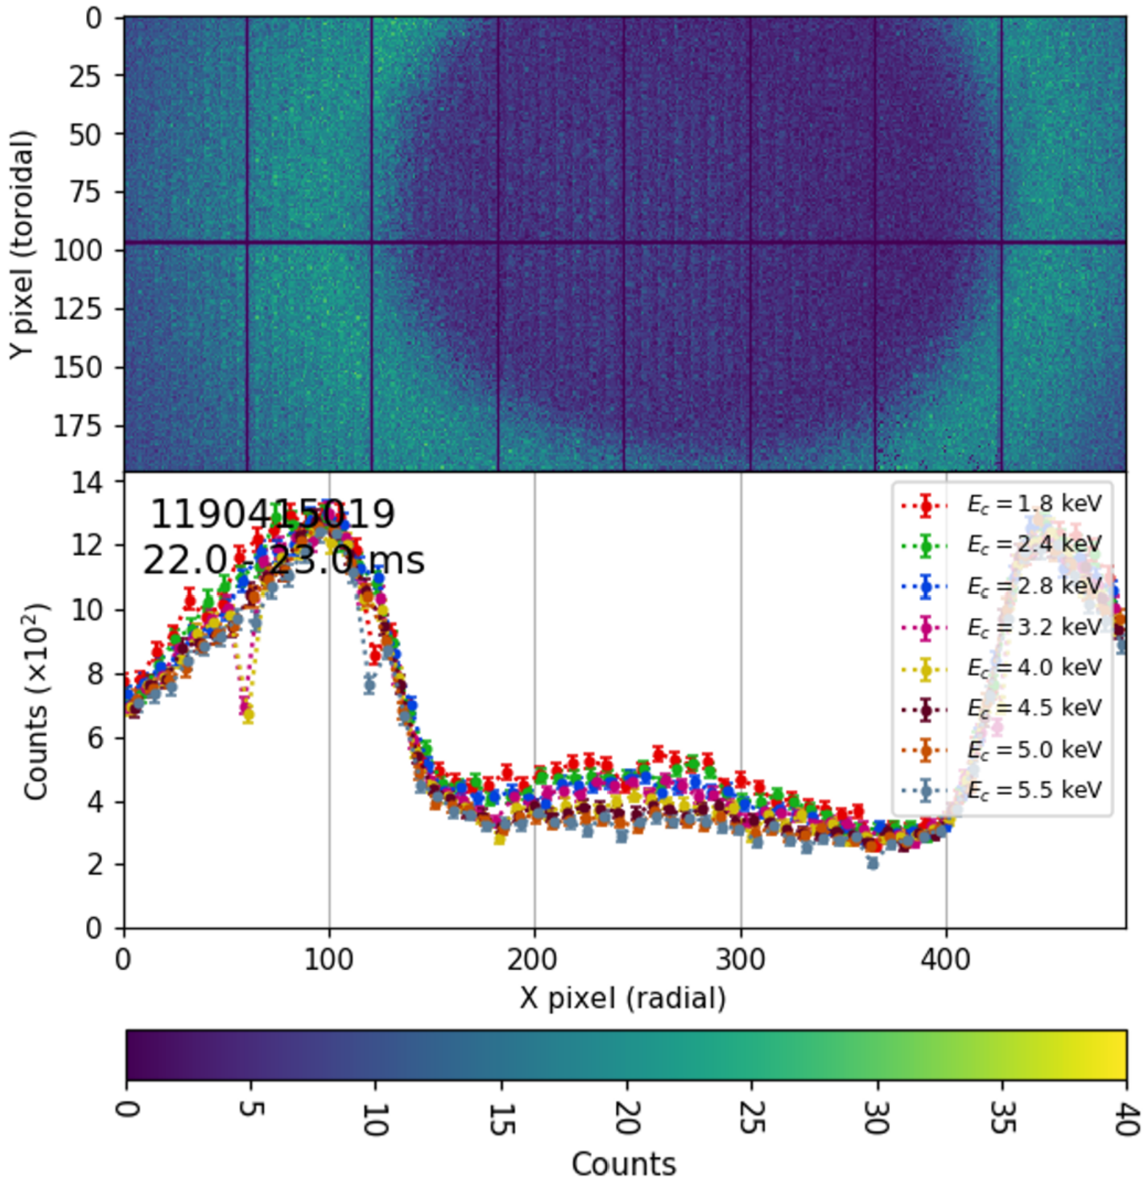

Figure 5.3: An emissivity "ring structure" observed during a single frame in a QSH plasma. Such phenomena occur frequently when operating at low density, providing evidence for a connection with runaway electron generation.

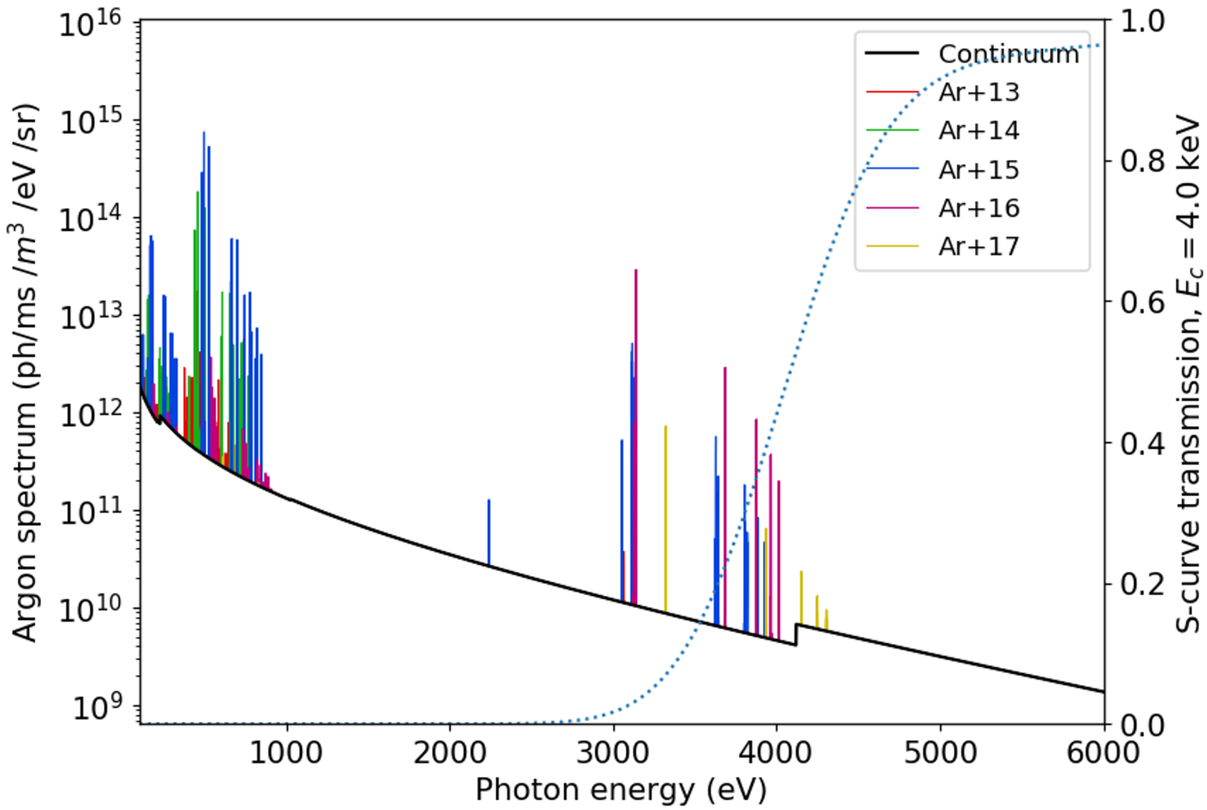

Figure 5.4: Spectrum for argon impurities as modeled by ADAS, given typical MST parameters. The dashed line illustrates the response function of a pixel with $E_c = 4$ keV.

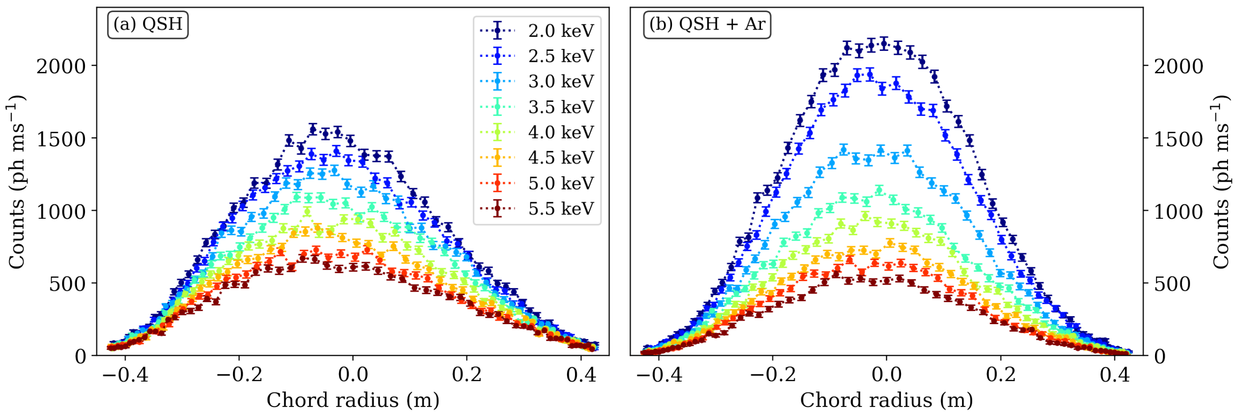

Figure 5.5: Frames from two non-reversed plasmas (a) without and (b) with an injected argon dopant, with otherwise similar plasma characteristics. Increased emission is for thresholds below the 3 keV Ar emission lines.

This capability was tested in 500 kA QSH plasmas by alternating between clean and Ar-doped plasmas. Figure 5.5 shows single-frame eight-color 1D measurements for two similar plasmas, except that (b) features the argon dopant and (a) is argon-free. The counts for $E_c \leq 3$ keV increase substantially when argon is present, while for higher energies the measurements are comparable. This comparison is shown more clearly in Figure 5.6, which shows the photon counts (averaged over five central chords) vs the threshold energy for each case. This clearly demonstrates that the increased signal is entirely due to Ar emission lines, meaning that the concentration of argon is low enough to not substantially change Z_{eff} (which would affect all thresholds). This demonstrates that the ME-SXR diagnostic can be used as an *ad-hoc* spectrometer to diagnose the presence of mid-Z like Ar impurities in the plasma.

Figure 5.7 extends this analysis over multiple time points, illustrating how the intensity of the Ar line amplitudes varies as the plasma evolves. We see that the increase in signal due to argon lines is greatest during the middle of the plasma lifetime, when the plasma temperature is at its peak. This serves as a proof-of-concept that the diagnostic can be used to characterize the emission spectrum over time, which has potential

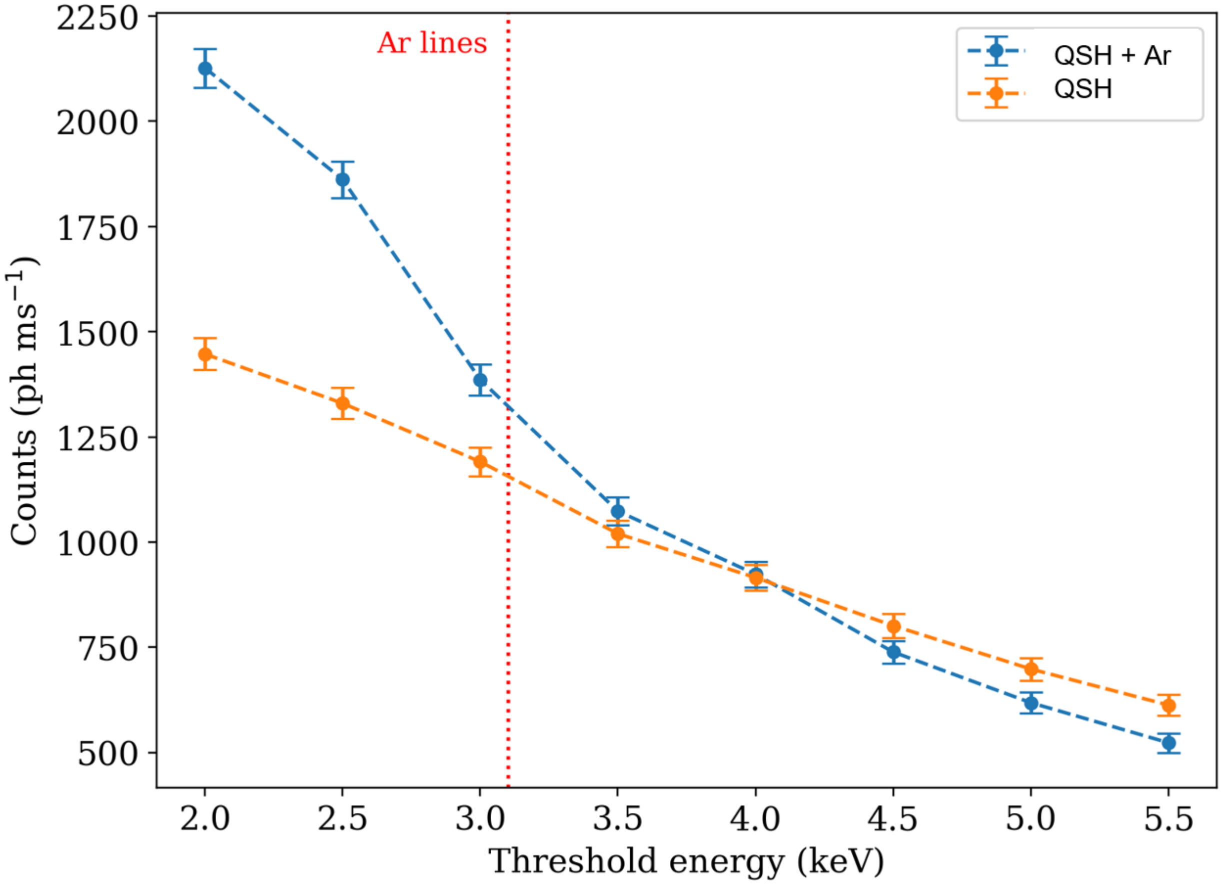

Figure 5.6: Measured single frame spectrum (counts vs lower threshold) for similar non-reversed plasmas with and without an argon pre-fill. Counts are averaged over central chords. The dominant emission line energy for argon is denoted by a vertical red line.

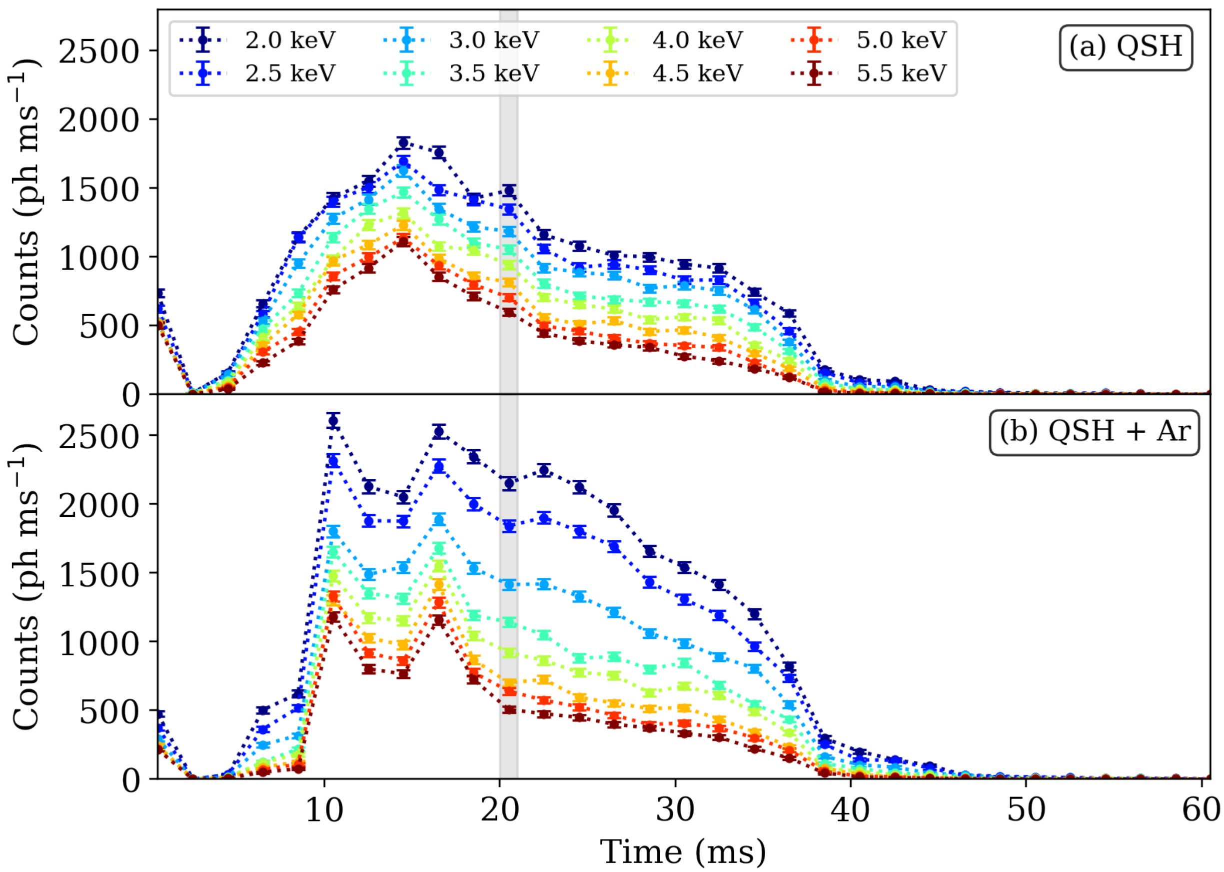

Figure 5.7: Temporal evolution of the central chord count rates for similar QSH plasma discharges (a) without and (b) with an argon pre-fill. The frame analyzed in Figures 5.5 and 5.6 is highlighted in gray.

applications in mid-Z impurity transport studies for long-pulse devices like tokamaks.

Finally, we can use the difference in the measured spectra to positively identify argon as the source of the increased emissivity. Since the two spectra are very similar for $E_c > 3$ keV, we will assume that the entire difference in signal is due to the presence of an additional source of photons of a single characteristic energy, E_0. We will denote the two spectra shown in Figure 5.6 as $\boldsymbol{y}_{Ar}$ and $\boldsymbol{y}_0$ for the Ar-doped and clean measurements, respectively. Then, the difference, $\boldsymbol{d} = \boldsymbol{y}_{Ar} - \boldsymbol{y}_0$, approximately forms an S-curve, just like those encountered during the energy calibration procedure. This is shown as the data points in Figure 5.8. Points where $\boldsymbol{y}_0 > \boldsymbol{y}_{Ar}$ were set to zero, since this is presumably the result of normal variation between the two plasma discharges.

The S-curve traced out by $\boldsymbol{d}$ can be fit directly to Equation 3.27, the same model we used in the energy calibration. We will include the results of the charge-sharing analysis of Section 3.4 and set the S-curve width to $\sigma_E = 0.3$ keV. That leaves a model with two free parameters, the amplitude N_{50} and the source energy E_0, given by

$$f(E_c; N_{50}, E_0) = \frac{N_{50}}{2}\left[\mathrm{erf}\left(-\frac{E_c - E_0}{\sigma_E\sqrt{2}}\right) + 1\right]\left(1 + k \cdot \left(E_c - b_0\right)\right) \tag{5.1}$$

where k is related to E_0 by Equation 3.26.

In order to properly account for counting statistics, we will do the fit using Bayesian methods (explained in Appendix B) with a Poisson likelihood function, given in Equation B.9, where the parameter $\lambda(N_{50}, E_0)$ is set to Equation 5.1. We will assume independent uniform priors, $N_{50} \sim \mathcal{U}(300, 700)$ and $E_0[keV] \sim \mathcal{U}(2, 5.5)$, denoted by $\pi(N_{50})$ and $\pi(E_0)$. Then, the posterior distribution is given by Bayes' Rule (Equation B.18) as $p(N_{50}, E_0|\boldsymbol{d}) \propto \mathcal{L}(N_{50}, E_0)\, \pi(N_{50})\, \pi(E_0)$. Since we are only interested in the source energy, we will marginalize N_{50} out as a nuisance parameter to obtain

$$p(E_0|\boldsymbol{d}) = \int p(N_{50}, E_0|\boldsymbol{d})\, dN_{50}. \tag{5.2}$$

This distribution is shown by the red curve in Figure 5.8.

The resulting posterior distribution is very narrowly peaked, with a 1σ credible interval of $E_0 = 3.05 \pm 0.02$ keV. This E_0 is close to the energy of the brightest line of Ar^{+16} as illustrated in Figure 5.4. Since the He-like state is often the most common ionization state of mid-Z impurities in the plasma core, this measurement would be sufficient to identify the impurity (or at least narrow the range of candidates). Figure 5.8 also shows the fit between the model (Equation 5.1) and the measured data $\boldsymbol{d}$, with error bars representing the 95% confidence level. Agreement is not perfect, but some discrepancy is expected given that the reference signal $\boldsymbol{y}_0$ was taken from a different plasma discharge

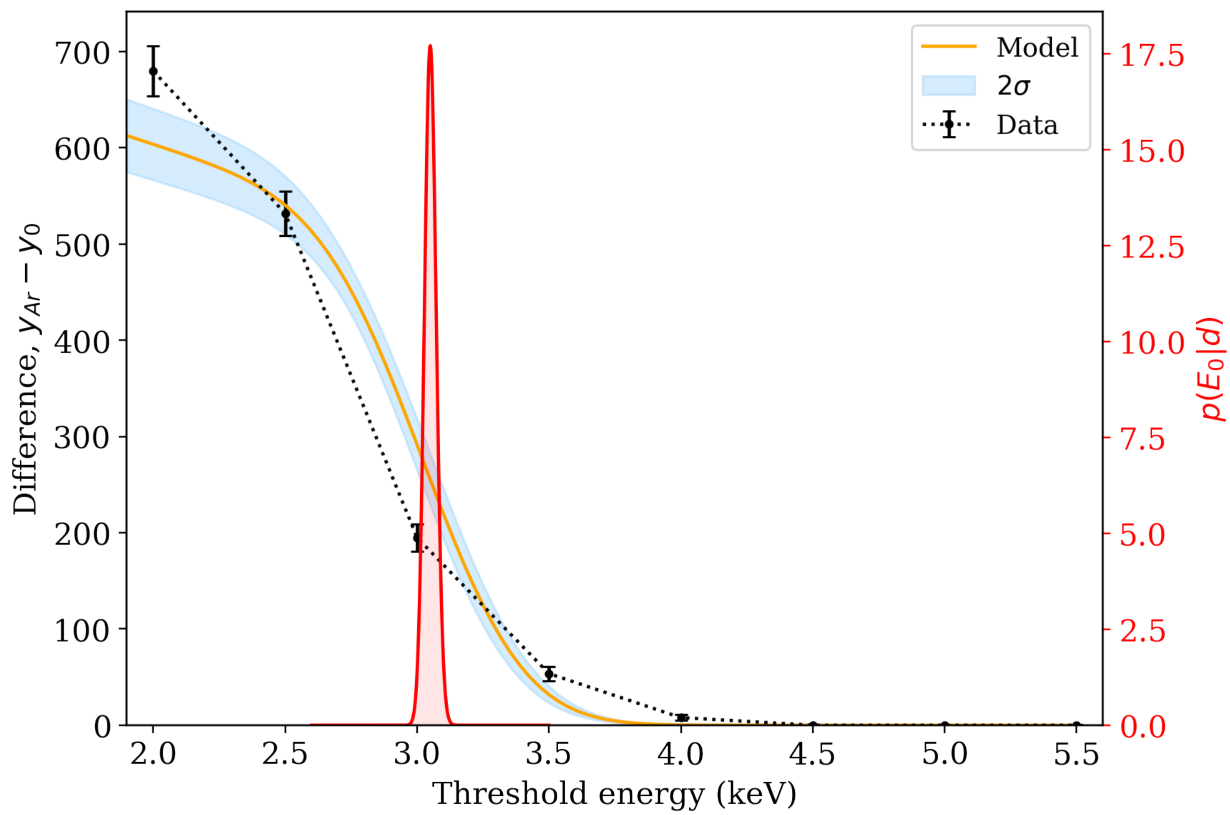

Figure 5.8: The difference between the two spectra shown in Figure 5.6 form an S-curve which can be used to determine the source energy E_0 to within 40 eV. The black points are the measured data $\boldsymbol{d}$, the best-fit model is shown with a 95% confidence region, and the posterior distribution over E_0 is shown in red. The narrowness of this distribution shows a high confidence in the results.

than $\boldsymbol{y}_{Ar}$. Even still, it was possible to accurately determine the energy of the emission lines to within 40 eV. In situations where the variability between plasma discharges can be further reduced, accuracy can presumably be improved.

The analysis reported in this section demonstrates that, in addition to characterizing thermal properties, the ME-SXR diagnostic can be used to measure the spectrum of mid-Z impurities, identify their characteristic line energies, and track this information over time. This analysis was repeated for a series of PPCD plasmas, however the resulting increase SXR emission was so bright that the detector began to saturate. This is due to the much higher T_e which is typical during PPCD, increasing the abundance of the bright Ar^{+15} and Ar^{+16} lines. However for higher-Z impurities a detector configuration with lower gain (Section 3.2.4) could be used, eliminating the saturation concern. In principle, the PILATUS3 can be calibrated for thresholds up to $E_c > 20$ keV, making this technique applicable to a wide range of impurities.

5.3 Temperature profile analysis

There are multiple techniques that could be employed to infer the electron temperature profile from ME-SXR data. In situations where circular symmetry can be assumed a straightforward Abel inversion can be used to extract the emissivity profile [1] for each threshold, which are then related to T_e directly via ratios. The method presented here is similar to this approach, but employs a Bayesian methodology to systematically propagate uncertainty and accounts for non-cylindrical flux surface geometries. However it relies upon prior knowledge of the profile shape, which can be a major drawback in some situations.

In general, the count rate detected by a cluster of pixels indexed by i on the ME-SXR with a shared threshold E_c and plasma volume is related to the local photon emissivity rate via line integral

$$N_{\gamma,i}(p,\phi) = \eta_i \int_{\mathcal{L}(p,\phi)} \varepsilon_\gamma(z)dz, \tag{5.3}$$

where p and ϕ are the tangency radius and angle which parameterize the chord, η_i is the étendue of the pixel(s), and z parameterizes the distance along the chord.

This equation can be recast as a matrix equation which maps the emissivity as a function of flux radius $\varepsilon(\rho)$ to the measured photon count rates for each cluster of pixels $N_{\gamma,i}$. This linear map is represented as the geometry matrix $\underline{\underline{\boldsymbol{R}}}$, which can be determined numerically using the results of the ME-SXR spatial calibration (Section 3.3) and the output of a flux surface reconstruction code (such as MSTfit [2]). In matrix form,

$$\underline{\boldsymbol{n}} = \underline{\underline{\boldsymbol{R}}} \cdot \underline{\boldsymbol{\varepsilon}}_\gamma, \tag{5.4}$$

where $\underline{\boldsymbol{n}} = (N_1, N_2, \ldots, N_{\gamma,N})^T$, $\underline{\boldsymbol{\varepsilon}}_\gamma = (\varepsilon_\gamma(\rho_1), \varepsilon_\gamma(\rho_2), \ldots, \varepsilon_\gamma(\rho_M))^T$, and $\rho \in [0,1]$ is some normalized 1D radial coordinate. Note that because $\underline{\underline{\boldsymbol{R}}}$ includes an ètendue factor it is unique to each threshold.

The next step is to choose a representation for the emissivity profile. We will select a 3-parameter $\alpha - \beta$ profile which is commonly used for the RFP [3]. This profile shape assumes a relatively flat core which then decreases monotonically around the mid-radius before going to zero at the edge.:

$$\varepsilon(\rho) = \varepsilon_0(1 - \rho^\alpha)^\beta. \tag{5.5}$$

This shape has the advantage that it forces the core of the profile to be flat, stabilizing the ratio. However parameterizations of other forms could be used instead, or possibly even nonparametric models based on Gaussian processes (this is beyond the scope of this dissertation, but see [4] for an explanation of this kind of analysis).

Next, the task is to constrain the profile parameters given the measured data. This could be done with a number of nonlinear fitting techniques, but I have chosen to use

Bayesian inference bases on Markov chain Monte Carlo sampling via the code *emcee* using a Poisson likelihood function and uniform priors. This type of methodology is described in more detail in Section 5.5 and Appendices B and C, but for now the important detail is that it produces a set of samples from the probability distribution for each parameter which are most likely to have generated the measured data, given the assumption that the model is correct. These samples can then be extrapolated to produce samples from derived quantities. This procedure was applied to a single frame recorded during of 300 kA PPCD operations, and the resulting emissivity profiles are shown in Figure 5.9. The resulting emissivity profile is flat in the core out to $r/a \approx 0.4$ before decreasing to near-zero at around $r/a \approx 0.8$.

These results are illustrated in more detail by frames a) and b) of Figure 5.10, which shows the agreement between the fit profiles and the measured data (a) as well as a two-dimensional poloidal cross-section of the resulting emissivity structure (b).

These emissivity profiles, and their associated uncertainties, can be used to produce estimates of the electron temperature by considering the ratio of the observed emissivity at a given point in the radial profile relative to a chosen reference threshold. If we assume that the threshold energy is sufficiently high to justify ignoring emission lines and recombination steps then the emissivity (up to a multiplicative constant) is given by the energy integral $\mathcal{I}$,

$$\begin{aligned} \mathcal{I}(T_e, E_c) &\equiv \int_0^\infty \frac{e^{-E/T_e}}{E\sqrt{T_e}} R(E; E_c) dE \\ &\propto \varepsilon_\gamma(T_e; E_c), \end{aligned} \tag{5.6}$$

where $R(E; E_c)$ is the total pixel response function which accounts for transmission through filters, absorption into Si, and energy discrimination due to detector electronics as a function of the chosen threshold (or "cutoff energy") E_c. When taking the ratio of two emissivity profiles the multiplicative constant cancels out, leaving a ratio of the

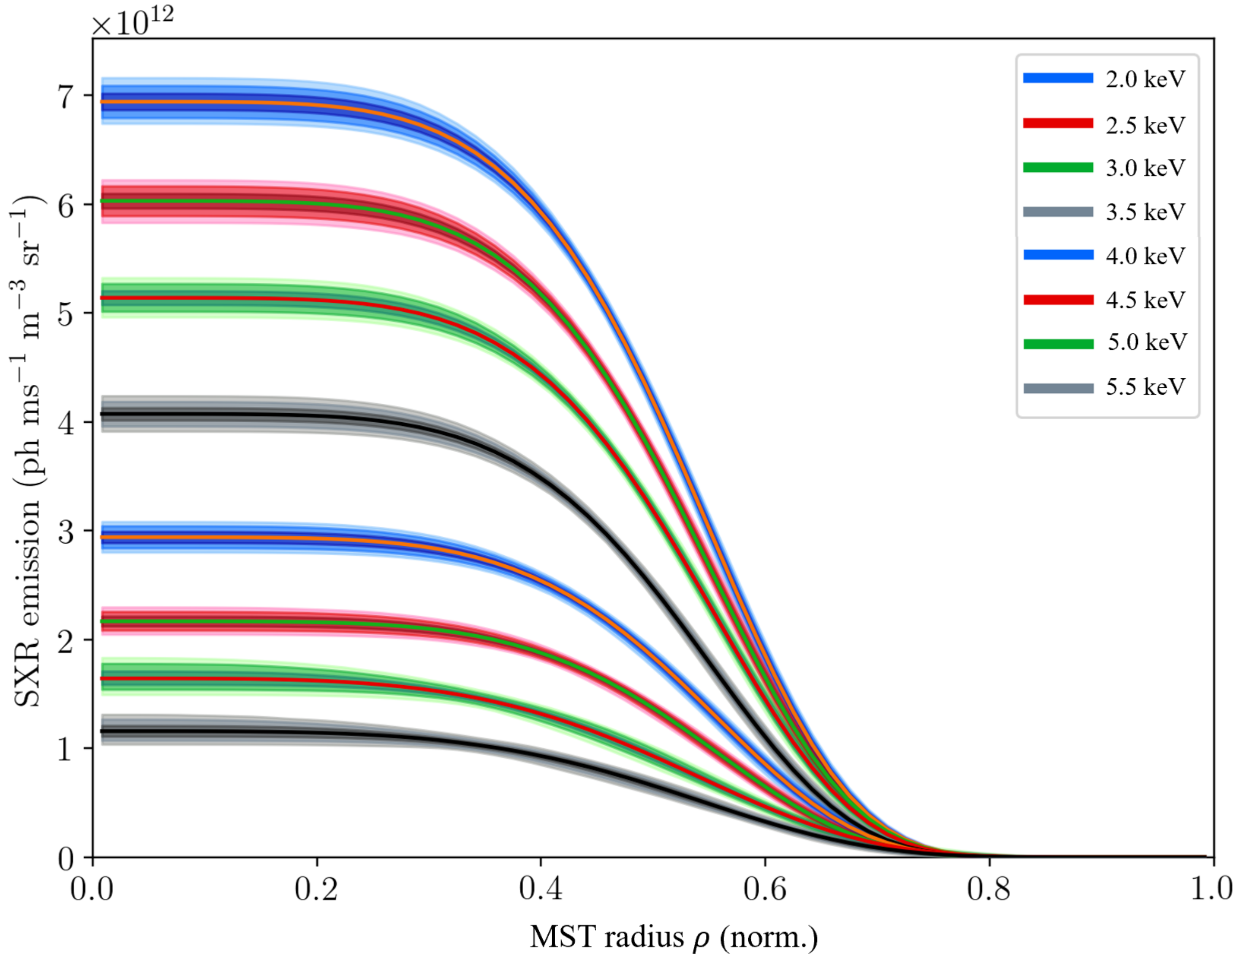

Figure 5.9: Reconstructed emissivity profile as a function of ρ, the MSTfit radius-like normalized flux surface label. MSTfit assumes a geometry of nested circular flux surfaces with a Shafranov fit.

energy integrals,

$$
\begin{aligned}
R(T_e, E_c) &= \frac{\varepsilon_\gamma(T_e, E_c)}{\varepsilon_\gamma^{(}T_e, E_{ref})} \\
&= \frac{\mathcal{I}(T_e, E_c)}{\mathcal{I}(T_e, E_{ref})},
\end{aligned}
\tag{5.7}
$$

where E_{ref} is threshold corresponding to the reference profile selected to serve as the denominator. Ratio curves produced by this method are shown in Figure 5.10 c).

Using these curves the electron temperature may be directly inferred. Figure 5.10 d) shows the inferred temperature for each of the remaining thresholds relative to $E_{ref} = 3$ keV. Notice that $E_c = 2$ and 2.5 keV significantly underestimate T_e, a consequence of ignoring the contribution of emission lines and recombination steps at energies where these features are not negligible. These values were not included in the final T_e estimation. The remaining thresholds > 3 keV were averaged to produce an overall estimation of T_e, and this process was repeated at each radial point. The resulting profile is shown in Figure 5.11 along with corresponding Thomson scattering points for reference. Comparison in the core to mid-radius is generally favorable, but $r/a = 0.6$ the profile begins to increase unphysically. However soft x-ray emission from this far out in the radial profile tends to be very low due to the declining electron density, so this portion of the profile is mostly constrained by low count rates and is therefore subject to significant noise. As a result, values for T_e for $\rho > 0.6$ are be omitted.

This methodology was used to individually analyze an ensemble of 35 individual 300 kA PPCD shots, each including three time points. The core T_e for each point was then compared against the corresponding Thomson scattering measurement, with the results shown in Figure 5.12. As suggested by the individual profile inversion shown in Figure 5.11, the temperature inferred from ME-SXR data is systematically higher than the Thomson measurements by $\Delta T_e \approx 180$ eV. The additional variation in the data is well

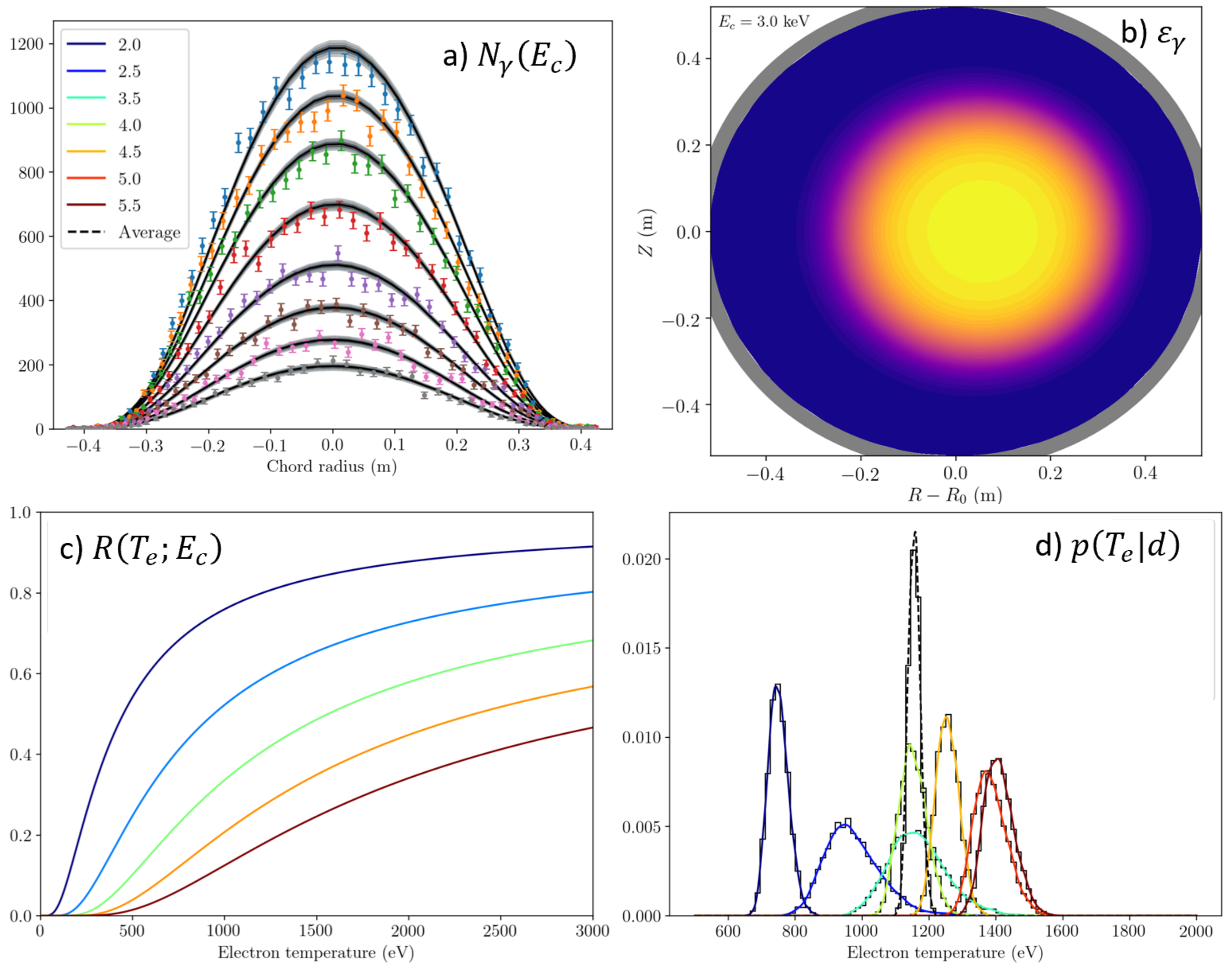

Figure 5.10: Summary of the methodology for obtaining T_e from ME-SXR measurements: a) Comparison between input data and best fit profiles; b) reconstructed 2D apparent emissivity profile for $E_c = 3$ keV; c) ratio curves relating emissivity to T_e, normalized to 3 keV; and d) posterior distribution for T_e in the core for each threshold and the average (black dashed line). Note that $E_c = 2.0$ and 2.5 keV were not used to compute the average and are included to illustrate the impact of Al lines on T_e inference.

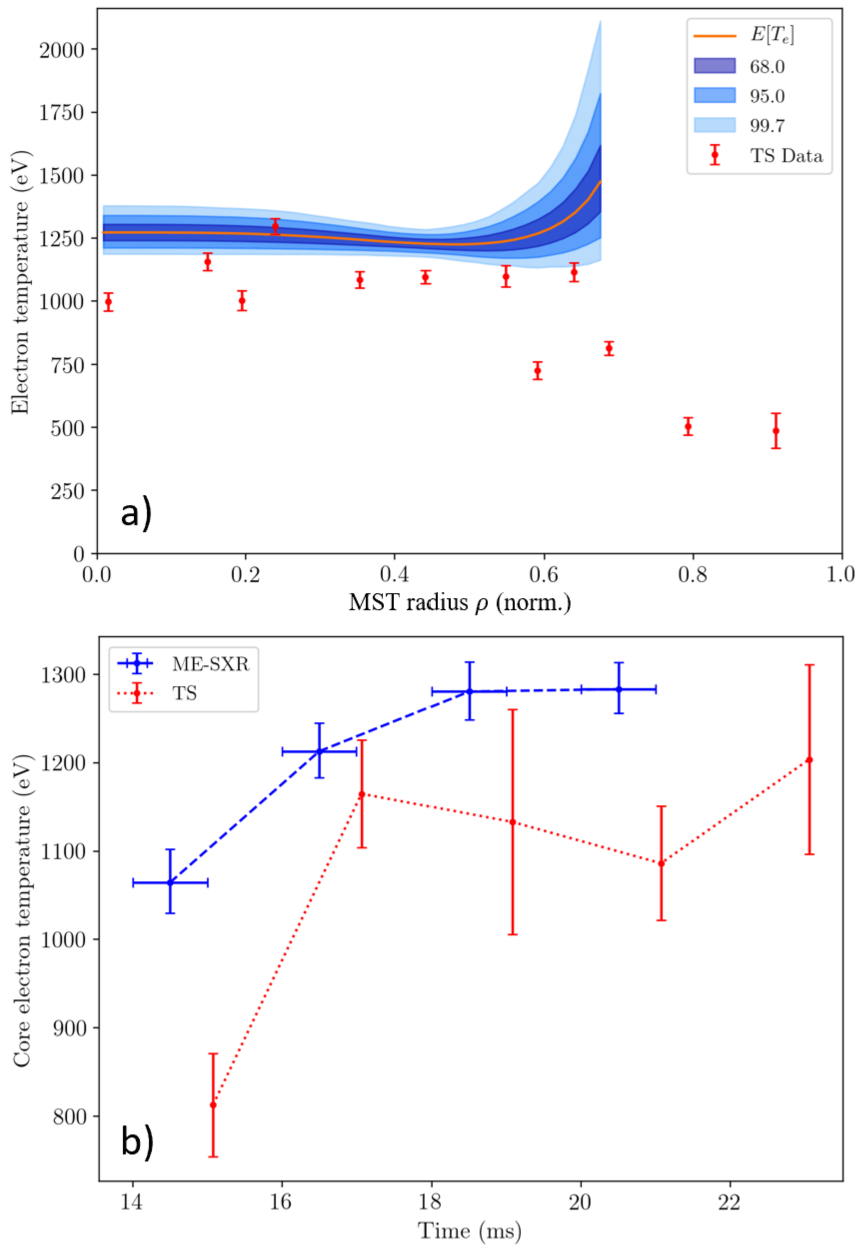

Figure 5.11: a) Reconstructed electron temperature profile as a function of ρ compared to simultaneous Thomson scattering data. Profiles are similar from the core to mid-radius, though the ME-SXR technique tends to produce poor results nearer the edge. Notice that the ME-SXR profile is somewhat higher than the TS data suggests. b) Evolution of the average core temperature over time for this plasma discharge.

explained by the variance in the Thomson scattering data, which features significant error bars when operated at densities this low. The source of the 180 eV offset has not been confirmed, though the effect of high-energy photons originating from runaway electrons (known to be present in PPCD [5]) is a possible candidate. Variation in the Mylar thickness may also be relevant. Regardless of the source, because the offset is constant (and therefore independent of the incident photon rate) it is not a result of detector saturation (as discussed in Section 5.4). Therefore, we can treat ΔT_e as a calibration factor. Using this methodology, ME-SXR measurements can be used to effectively determine T_e.

5.4 Troubleshooting detector saturation and pulse pileup

After the initial installation of the ME-SXR diagnostic, the resulting output images featured persistent non-uniformities which varied on a chip-by-chip basis. A typical example of this behavior can be seen in Figure 5.13, using integration and readout times of 1 ms. One of the ASIC chips, which we identify as Chip #2, records noticeably fewer counts than would be expected when compared to its neighbors. Chip #12 also shows signs of under-counting, though this is not uniform across the chip. These artifacts are not the result of plasma structure, as the features were consistently observed across a wide range of plasma operating conditions.

Several tests were performed in order to identify the source of this unexpected behavior. We ruled out electrical interference, problems with the grounding scheme, and issues with specific detector software settings. We performed bench tests with an Fe-55 source, and noted that the issue was not present in this data. Building off of this hint, we developed the hypothesis that the anomalous behavior was the result of saturation by photons with enough energy to pass through the Be filter and be absorbed into the Si, but below the threshold to be counted. There are many such photons from the Al excitation lines ≤ 2 keV. At a high enough flux they generate substantial pulse-pileup (saturation)

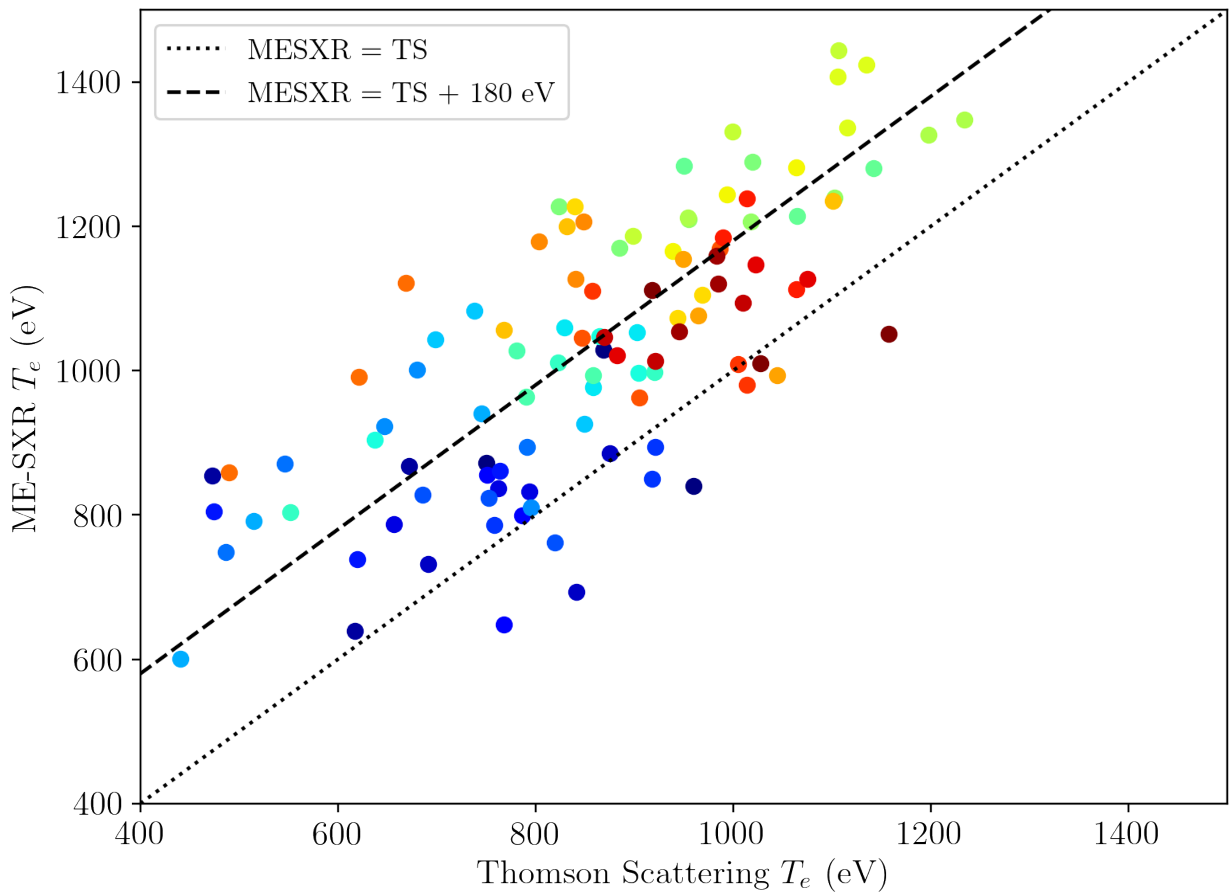

Figure 5.12: Comparison of core T_e for 300 kA PPCD plasmas as measured by the ME-SXR diagnostic in the lowE (high gain) mode vs corresponding Thomson Scattering measurements. TS measurements were derived by averaging measurements from the inner-most channels and interpolating to match the ME-SXR time points. Points of the same color correspond to different time points from the same plasma discharge. The systematic discrepancy of $\Delta T_e \approx 180$ eV is illustrated by the dashed line. Note that the 100 μm Mylar filter was installed.

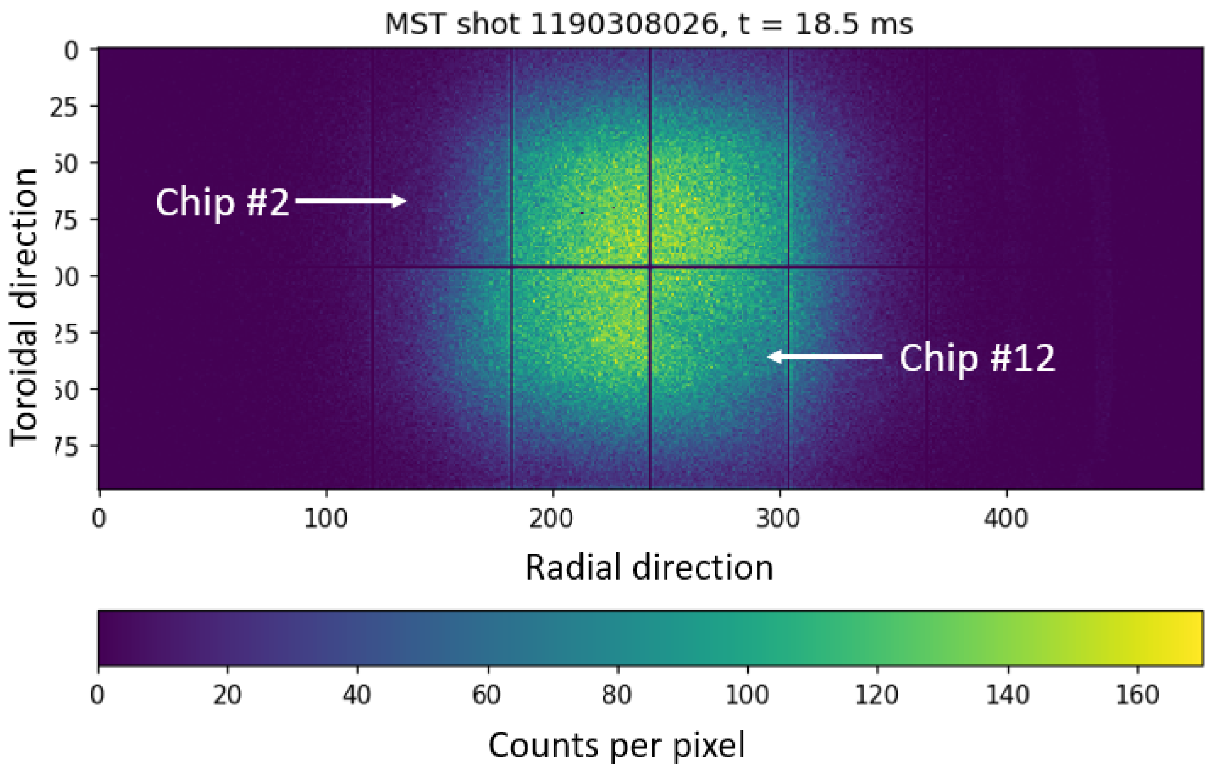

Figure 5.13: Output image displaying the characteristics of saturation behavior on our PILATUS3 module. Clear artifacts are seen on chips #2 and #12, with discontinuities visible across chip boundaries. Though subtle, these features were seen to be persistently present on all data taken with the detector until the additional Mylar filter was installed. Data was taken from a 500 kA PPCD plasma with uniform 2 keV lower threshold.

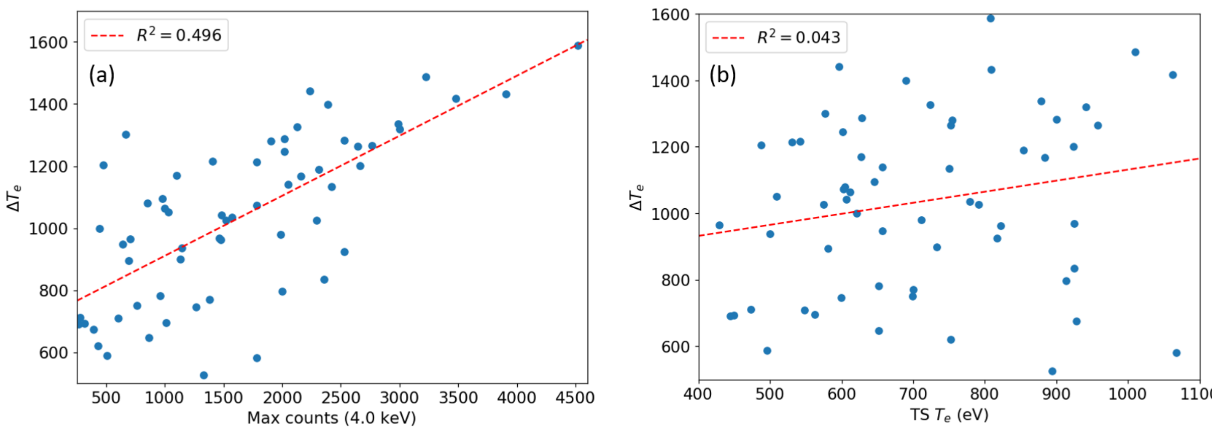

Figure 5.14: Prior to installation of a Mylar filter, the discrepancy ΔT_e between ME-SXR and TS T_e estimates scaled more strongly strongly with the photon flux (a) than the measured temperature (b).

which contaminates the desired signal. Since the ASICs are fairly independent from one another, it is unsurprising that they might behave in different and unexpected ways when saturated, leading to the observed artifacts in the data.

In order to reduce the low-energy photon flux, a 50 μm Mylar filter was installed over top of the Be window. This reduced the diagnostic sensitivity further from $\sim 5\%$ at 2 keV to $< 1\%$. This new configuration was tested by separately measuring the core T_e with Thomson scattering and the ME-SXR for a dataset of 300 kA PPCD data. The discrepancy between the two measurements was quantified as $\Delta T_e = T_e^{(ME)} - T_e^{(TS)}$. As shown in Figure 5.14, it was observed that ΔT_e was much more strongly correlated with the overall photon flux than with the measured temperature. In fact, it was found that the relationship between ΔT_e and the photon flux was nearly linear. This is strongly suggestive of pulse tail pileup, subsequent pulses form before the previous pulse has fully subsided. This results in a higher-energy apparent spectrum, and therefore a larger inferred temperature. The two types of pileup (peak and tail) are illustrated in Figure 5.15.

As a result of these observations, the the Mylar filter thickness was increased to

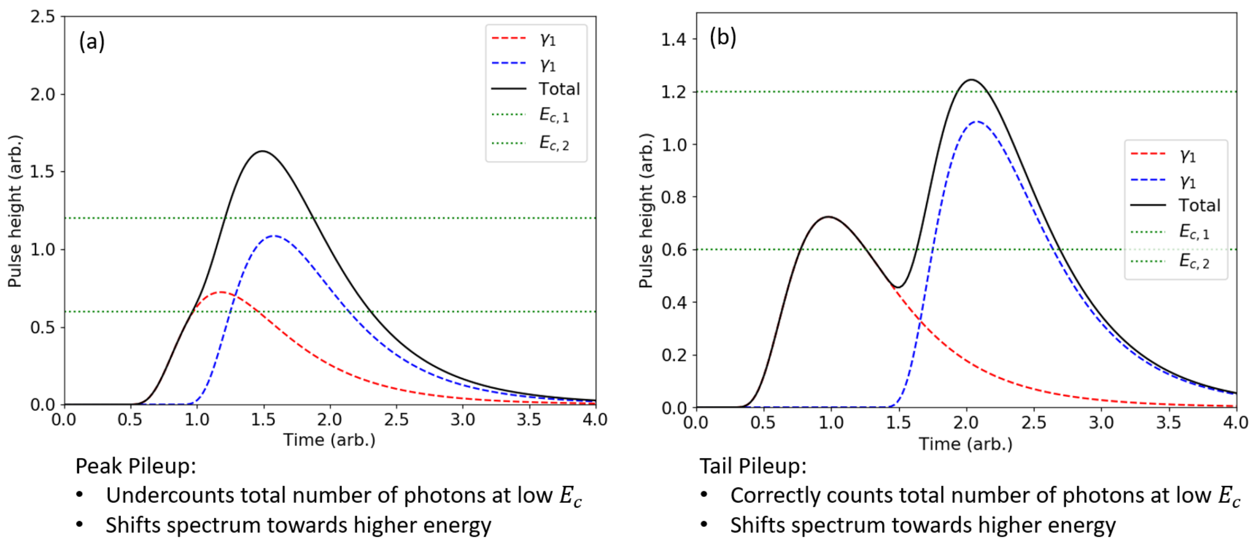

Figure 5.15: Conceptual example of the two types of pileup which can affect pulse height analysis detectors. Peak pileup (a) results in fewer detected photons at higher energy, while tail pileup (b) results in the same number of detected photons, but a higher-energy apparent spectrum.

100μm. This reduces the transmission of photons $\sim$ 2 keV from about 50% with no Mylar filter to $<$ 1%. A comparison between the composite filter transmission function for 50 μm, 100 μm, and no Mylar (just Be) is shown in Figure 5.16. Thomson scattering measurements were not available at the time of this second test, so instead we looked at how the core temperature $T_{e,0}$ measured by the ME-SXR varied with photon flux. Figure 5.17 shows that $T_{e,0}$ increased much less sharply with photon flux when the 100 μm filter was installed than when just the 50 μm filter was installed. Measured temperatures were also more closely aligned with expectations for plasmas of these settings. Based on these results, it was decided to move forward with the 100 μm filter installed. Later comparisons with Thomson scattering measurements (such as those previously shown in Figure 5.12) provided increased confidence that the pileup issue was resolved.

All data shown in remainder of this thesis was taken with the 100μm Mylar filter installed. Scientists designing future implementations of the ME-SXR system should

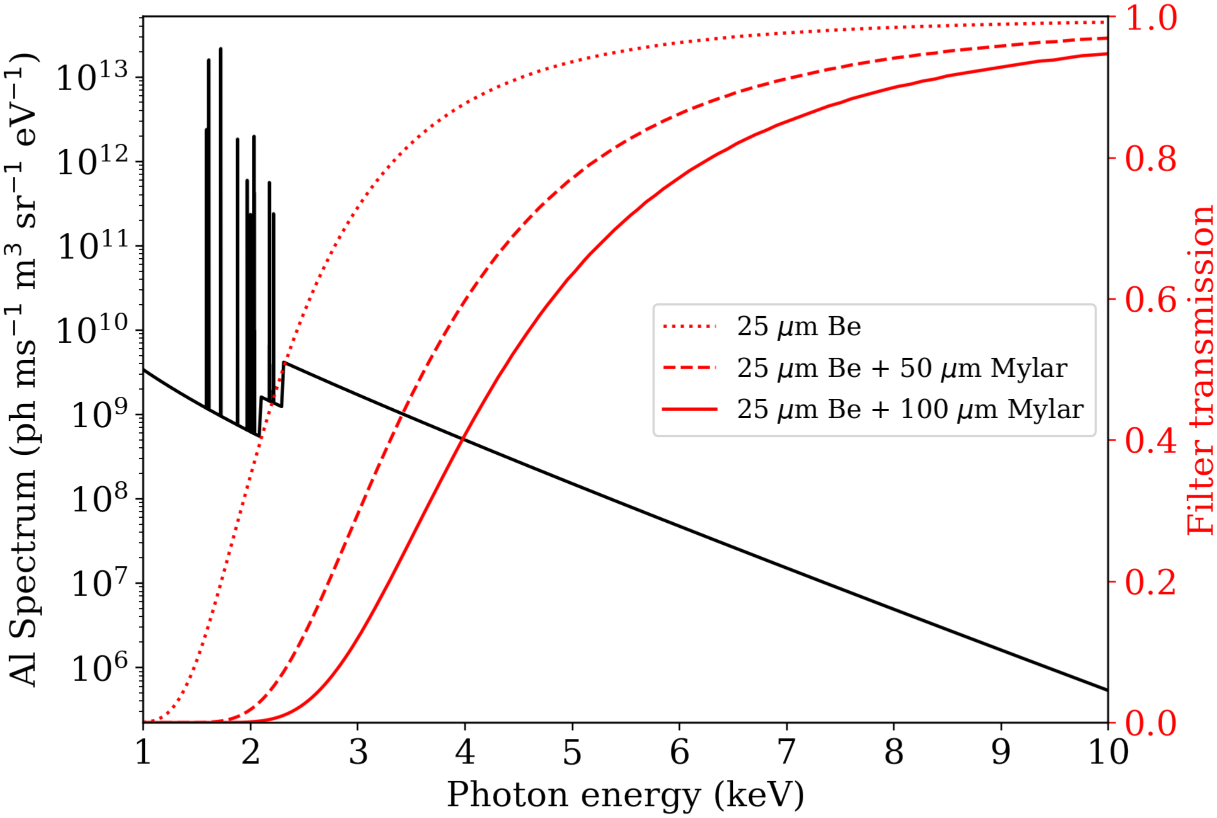

Figure 5.16: Filter transmission curves for three tested configurations: original (no Mylar), 50 μm Mylar, and 100 μm Mylar. All three curves include the 25 μm beryllium filter. Simulated aluminum spectrum at $T_e = 1$ keV is shown for reference.

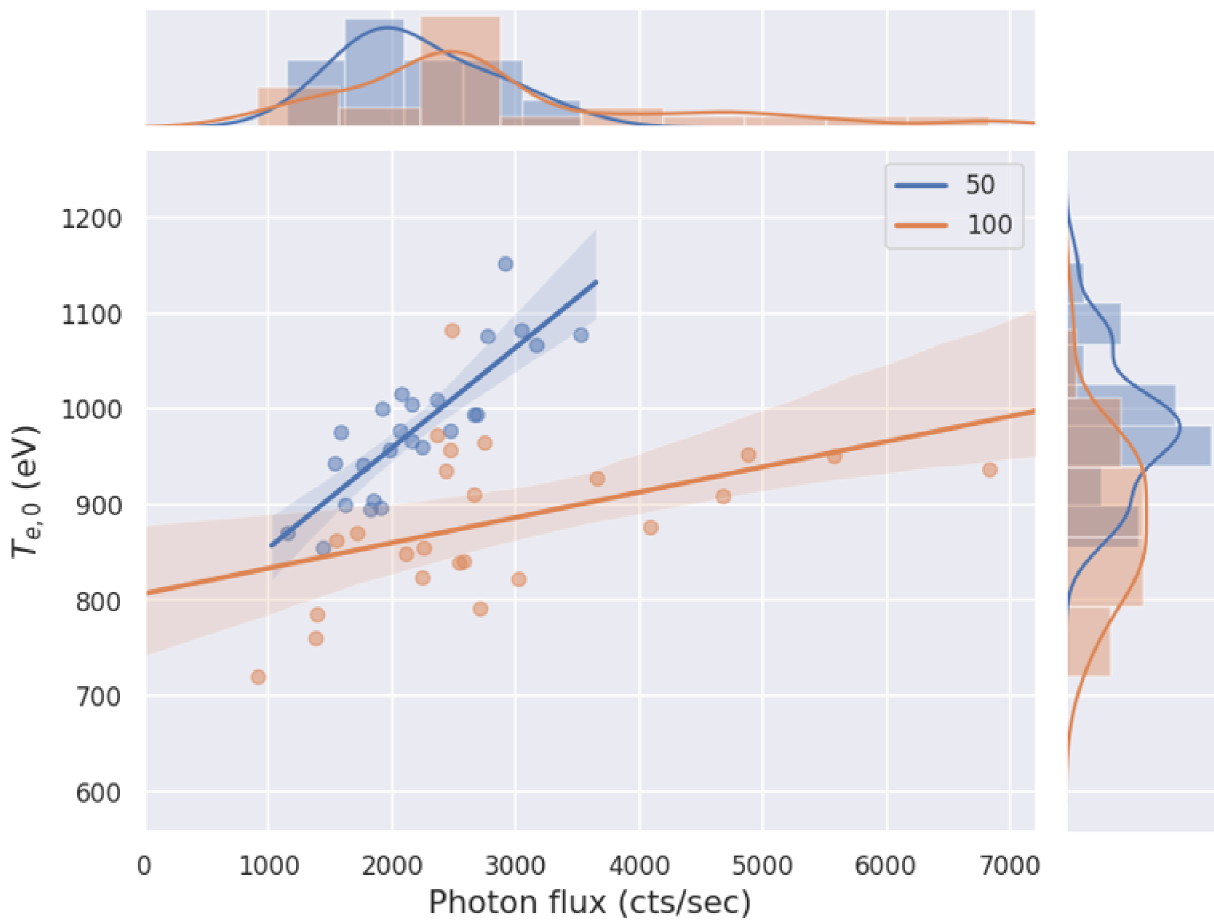

Figure 5.17: Correlation between core T_e and photon flux is much stronger for the 50 μm filter than the 100 μm, implying pileup is significant in the former.

be careful to account for sources of photons outside of the "target range," but which nonetheless reach the detector screen, when designing the pinhole and filter configurations.

5.5 Integrated data analysis

Section 5.3 provided a straightforward technique to extract T_e measurements with minimal *a-priori* assumptions. However, oftentimes additional information is available, such as concurrent measurements from other diagnostics, previous results, and knowledge of the underlying physics. This section presents a consistent way of incorporating all of this additional information into a single framework, called integrated data analysis (IDA) [3]. With this framework we will be able to simultaneously produce measurements of the electron temperature and impurity ion densities, demonstrating the versatility of the ME-SXR diagnostic.

The IDA methodology demonstrated here is based on Bayesian inference [6]. This is an approach to probability which views probability as the quantification of the degree of certainty based on the available information, rather than than the long-term frequency over many repetitions. The heart of Bayesian inference is Bayes' Rule,

$$p(\boldsymbol{\theta}|\boldsymbol{d}, I) \propto p(\boldsymbol{d}|\boldsymbol{\theta}, I)\, p(\boldsymbol{\theta}|I), \tag{5.8}$$

which states that the *posterior* $p(\boldsymbol{\theta}|\boldsymbol{d}, I)$ of a system being best-described by the parameter vector $\boldsymbol{\theta}$ given some measured data $\boldsymbol{d}$ is proportional to the product of the *likelihood* $p(\boldsymbol{d}|\boldsymbol{\theta}, I)$ of having measured that data given the parameter vector and the *prior* information $p(\boldsymbol{\theta}|I)$. The I in Bayes' rules is meant to represent the additional information that has been incorporated into the analysis, such as the choice of a particular model.

Bayes' Rule is intended to be applied iteratively whenever new information is made available. The approach to IDA employed here takes advantage of this feature to iteratively apply Bayes' Rule with simultaneous measurements from independent diagnostics, taking advantage of all available information to make the best possible estimate of the model parameters [3]. This approach allows the data analyst to exploit the ways in which different diagnostic measurements correlate the plasma properties, sometimes enabling estimates of properties (like Z_{eff}) that no individual diagnostic can effectively measure [7]. This approach will also provide a natural and consistent framework for propagating measurement uncertainty, including when the underlying distributions are not Gaussian.

The aim of this section is to provide a full, detailed example of how to implement the ME-SXR diagnostic into an IDA framework. Two versions of the analysis will be performed. The first will incorporate just the ME-SXR and Thomson scattering diagnostics, while the second will also include SXR tomography and the NICKAL2 Ross spectrometer. The first analysis will allow us to demonstrate the ME-SXR diagnostic's sensitivity to mid-Z impurities (like Al), while the second will provide the best estimates of the plasma profiles given the available data. A PPCD plasma was chosen to be the test case since PPCD plasmas have been thoroughly studied, meaning we can make use of well-informed priors. Section 5.5.1 begins by describing the PPCD plasma model. Section 5.5.2 then goes through the process of correctly selecting likelihood and prior distributions based on the best available information. Section 5.5.3 presents the results of the ME-SXR + TS analysis, demonstrating good sensitivity to the aluminum density. Finally, Section 5.5.4 presents the results of full IDA analysis, an draws comparisons. Some additional background and theory are provided in two companion Appendices, B and C, which discuss the theory of Bayesian inference and computational sampling techniques, respectively.

5.5.1 The PPCD plasma model

Previous experience on MST has found that the temperature and density profiles in high-current PPCD plasmas are well-described by so-called "α-β" models which assume radially symmetric profiles (as a function of the model magnetic coordinate ρ) controlled by two shape parameters [7]. Additional features such as islands can be added on as additional terms if needed, as may be necessary when modeling standard RFP plasma conditions [8]. However, these can typically be omitted for PPCD conditions.

The temperature profile is modeled by a simple α-β model, given by

$$T_e(\rho) = T_{e,0}\left(1 - \rho^{\alpha_T}\right)^{\beta_T}, \tag{5.9}$$

where ρ is the normalized MSTFit radial flux surface label (Equation 2.13, normalized to the minor radius). We will typically fix the value of β_T as it tends to be semi-redundant with the α_T parameter. An MSTFit reconstruction will be used to provide the mapping between ρ and standard Cartesian spatial coordinates (x, y), as well as a reconstructed $n_e(\rho)$ profile.

The density profile could have alternatively be included in the model using an α-β profile and fit using the FIR data directly. However, it was found that the resulting variance in the $n_e(\rho)$ profile is very small and has a minimal impact on the uncertainty estimates for the other profiles. Given that the impact of adding three additional parameters ($n_{e,0}$, α_n, β_n) has on the computational time is significant, it was decided to just consider the existing MSTFit profile as a fixed input.

MST plasmas typically feature several impurity species in concentration high enough to measurably impact the SXR spectrum: N and O from air; C from graphite in the limiter; B from probes, and Al from the vacuum vessel wall. Ar can also be doped into MST, but is not otherwise present in significant concentrations. It has been established that during PPCD discharges these ions are subject to a classical transport effect known as temperature screening [9]. Essentially, the presence of an ion temperature gradient leads to an expulsion of impurity ions from the core of the plasma, resulting in a hollow profile shape which peaks at the outer mid-radius. These predictions have been experimentally confirmed with charge-exchange recombination spectroscopy measurements. The ME-SXR model accounts for this phenomenon by adding an additional "hollow bump" term to the standard "α-β" shape, given by

$$n_Z(\rho) = n_{Z,0}\big(1-\rho^{\alpha}\big)^{\beta} + \delta n_Z \exp\left(-\frac{(\rho - \delta r/a)^2}{2(w_r/a)^2}\right), \tag{5.10}$$

where n_Z refers to the total impurity species density (including all ionization states), δn_Z refers to the increased accumulation at the hollow bump for the impurity species Z, $\delta r/a$ is the normalized location of the bump and w_r/a is the normalized width of the bump. We typically treat α_Z, β_Z, δr, and w_r as being the same for all impurity species.

It should be noted that we are not directly modeling the physics of ion transport. Instead previous analysis has informed us that classical transport tends to result in the formation of hollow profiles, and we have therefore modified our model to include this as a phenomenological feature. This is the kind of "additional information" that is referred to by the symbol I in the notation of the likelihood function $p(\boldsymbol{d}|\boldsymbol{\theta}, I)$.

Rather than treat all impurity ion densities as free parameters, we will typically treat n_C as an input to the model and constrain the other low-Z impurities based on empirically-established ratios [7]: $n_O/n_C = 0.9$, $n_N/n_C = 0.3$, and $n_B/n_C = 0.3$. As these species all tend to be fully-ionized and do not feature recombination steps or lines in the SXR spectrum, they essentially just serve to contribute to the overall Z_{eff}. This means that any analysis routine using this model would not be able to distinguish between, for instance, an increase in n_C vs an increase in n_O. So we circumvent this by fixing all but one of these densities, and leave n_C as a free parameter to ensure that we correctly capture the behavior of Z_{eff}. Al in the MST is not fully ionized and contains multiple constraining features in the SXR spectrum, so n_{Al} is left as a free parameter.

Finally, the neutral density was modeled by a parametric estimation of A. Xing's results discussed in Section 4.2.2. This is given by

$$n_0(\rho) = n_{0,0} + n_{0,1}\rho^{\alpha_0} \tag{5.11}$$

where the three parameters $(n_{0,0}, n_{0,1}, \alpha_0)$ were chosen so that $n_0(0) = 2 \times 10^{14}$ m^{-3}, $n_0(0.6) = 10 \times 10^{14}$ m^{-3}, and $n_0(1) = 450 \times 10^{14}$ m^{-3}. This results in a neutral density profile which is flat throughout the core but grows exponentially by two orders of magnitude near the edge. These parameters are fixed and are not included in the fitting procedure.

Examples for each of the input profiles discussed here using parameters typical of 400 kA PPCD were shown in Figure 4.13. The values used to generate this figure are given in Table 5.1. These are the default input values for the ME-SXR model and should be assumed whenever no other values are specified. As implemented, the model has six free parameters represented in the parameter vector,

$$\boldsymbol{\theta} = (T_{e,0}, \alpha_T, n_{Al,0}, n_{C,0}, \delta n_{Al}, \delta n_C). \tag{5.12}$$

The remaining parameters are held fixed at the default values. This set of parameters was settled upon as a compromise between model flexibility (including more parameters) and computational requirements (including more parameters requires more steps to converge).

5.5.2 Likelihood and priors

A summary of all the parameters which describe the plasma in the ME-SXR model, and their default values, is given in Table 5.1. Most of the priors taken to be uniform (see Equation B.11), with bounds given by the [Min,Max] values listed in the table. This

PPCD model parameters

Parameter	Units	Default	Min	Max
$T_{e,0}$	eV	1100	0	2000
α_T	N/A	7	2	14
β_T	N/A	7		Fixed
$n_{Al,0}$	10^{19} m^{-3}	2e-3	1e-4	1e-2
$n_{C,0}$	10^{19} m^{-3}	6.8e-3	$\mathcal{N}(2.8e-3, 1.4e-3)$	
α_Z	N/A	12		Fixed
β_Z	N/A	4		Fixed
δn_{Al}	10^{19} m^{-3}	2.5e-3	1e-4	1e-2
δn_C	10^{19} m^{-3}	1.8e-2	1e-4	5e-2
$\delta r/a$	N/A	0.635		Fixed
w_r/a	N/A	0.09		Fixed

Table 5.1: List of PPCD model parameters, including units, default values, and [Min, Max] range for uniform priors. Parameters which are fixed in the model are noted.

equates to an assertion that the probability of that parameter taking on any value within this range is equal, while any value outside of this range has a probability of zero. These bounds were mostly chosen by physical considerations ($T_{e,0}$ must be positive) or constrained by previous expectations ($T_{e,0} > 2$ keV is implausible for these bank settings).

One exception to this is the core carbon density, which will be assigned a normal distribution based upon an ensemble of prior charge exchange recombination spectroscopy measurements [7]. This is given by

$$\ln p(n_{C,0}|\mu_C, \sigma_C, I) = -\frac{1}{2}\left(\frac{n_{C,0} - \mu_C}{\sigma_C}\right)^2 - \frac{1}{2}\ln 2\pi\sigma_C^2, \tag{5.13}$$

where $\mu_C = 2.8$e-3 and $\sigma_C = 1.4$e-3, measured in the units shown in Table 5.1.

A likelihood function was chosen for each diagnostic (ME-SXR, SXR tomo., and NICKAL2). Each likelihood model, which quantifies the probability that one might have measured the data d given some known parameters $\boldsymbol{\theta}$, was chosen by combining the relevant diagnostic forward model with a statistical model of the measurement noise. For

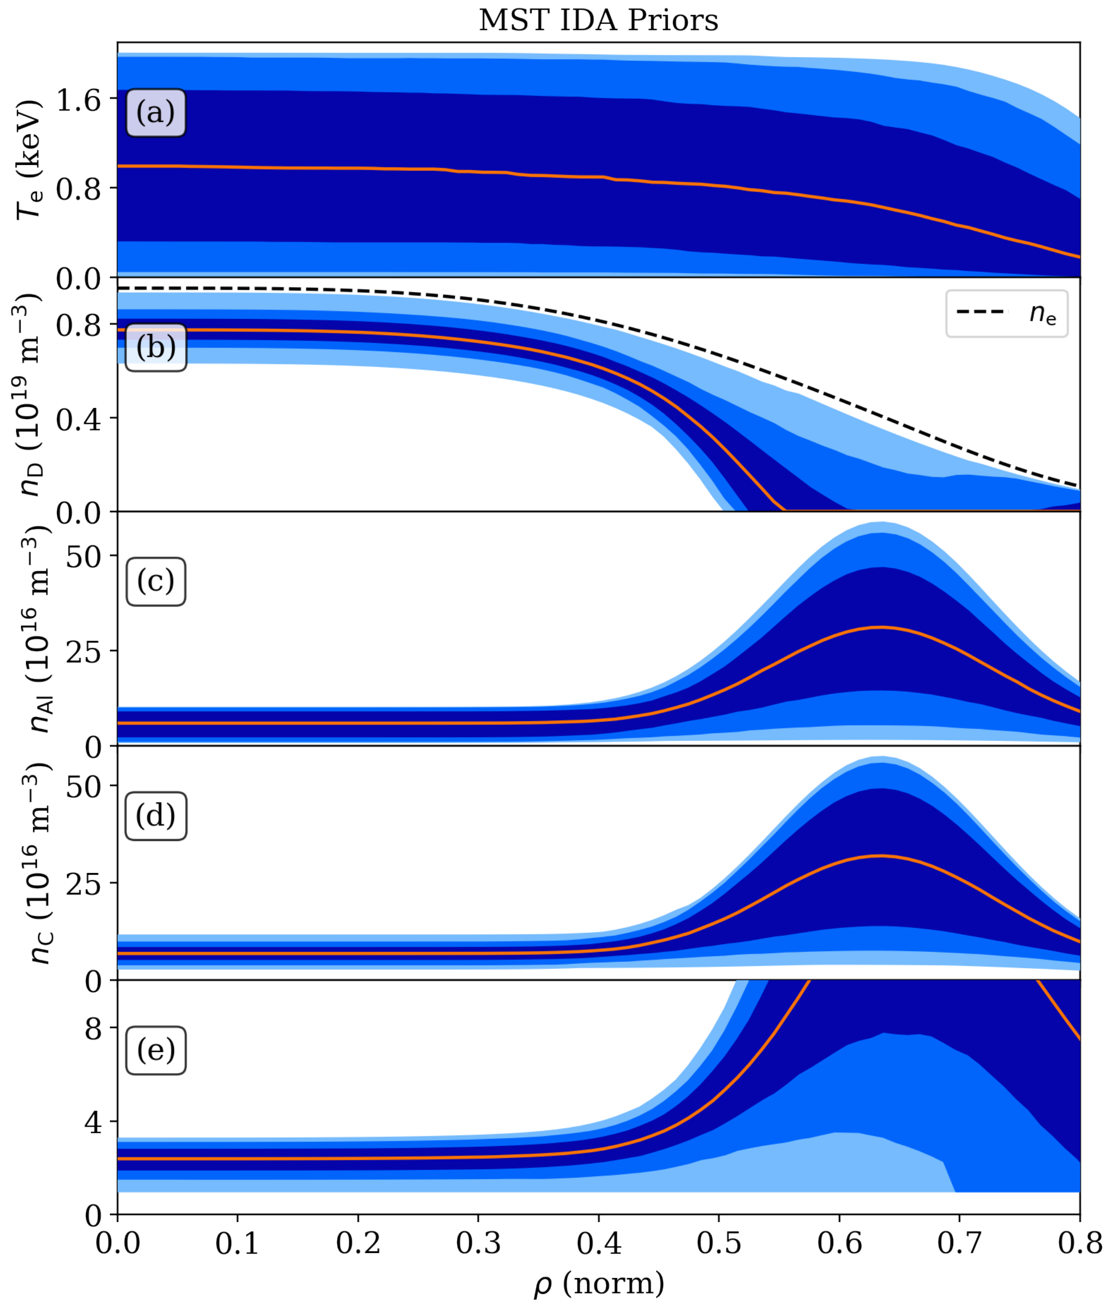

Figure 5.18: Prior distributions over the (a) electron temperature, (b) deuterium density, and (c) aluminum density, (d) carbon density, and (e) ion-effective charge (Z_{eff}) profiles. The orange line represents the median profile and the shaded regions encompass the 1-, 2-, and 3-σ credibility regions.

each case the measurement noise was assumed to be Gaussian, giving a general form of

$$\ln p(\boldsymbol{d}|\boldsymbol{\theta}, I) = -\frac{1}{2}\sum_{i}^{N}\left(\frac{d_i - f(p_i, \boldsymbol{\theta})}{\sigma_i}\right)^2, \tag{5.14}$$

where $f(p_i, \boldsymbol{\theta})$ is the diagnostic forward model for the i$^{\text{th}}$ pixel given the chord radius p_i, and $\sigma_i^2 = \sigma_{d,i}^2 + \sigma_{m,i}^2$ where $\sigma_{d,i}$ is the measurement noise and $\sigma_{d,i}$ is the systematic uncertainty in the model. The diagnostic forward model and model uncertainty σ_m for the ME-SXR was described in Chapter 4. A similar model had previously been developed for the SXR tomography and NICKAL2 diagnostics [10], and the associated σ_m was taken to be 2% [11].

Although the noise model for the ME-SXR diagnostic is more properly described by a Poisson distribution, given the relatively high count rates during PPCD the Gaussian approximation $\sigma_{d,i} \approx \sqrt{N}$ was considered to be adequate. Prior to the analysis the the ME-SXR 1D profiles (like those shown in Figure 3.26) were smoothed and interpolated using Gaussian process regression interpolation (so-called "kriging") [12, 13], implemented using Sci-kit Learn [14]. This is a nonparametric regression scheme which estimates the intermediate values from the existing values based on a smoothness criteria (defined by a kernel). This method allows for an accurate estimate on the uncertainty of the interpolated data points, assuming Gaussian statistics. Data was interpolated to exactly 60 shared lines of sight, significantly reducing the computational overhead required for each evaluation of the model.

The likelihood were combined via an iterative application of Bayes' Rule, yielding the posterior distribution. For the ME-SXR + TS analysis this is given by

$$p_1(\boldsymbol{\theta}|\boldsymbol{D}_1, I) \propto p_{\text{MESXR}}(\boldsymbol{d}_{\text{MESXR}}|\boldsymbol{\theta}, I)\, p_{\text{TS}}(\boldsymbol{d}_{\text{TS}}|\boldsymbol{\theta}, I)\, p(\boldsymbol{\theta}|I), \tag{5.15}$$

where $\boldsymbol{D}_1 = (\boldsymbol{d}_{\text{MESXR}}, \boldsymbol{d}_{\text{TS}})$. For the full IDA framework, the likelihood is given by

$$p_2(\boldsymbol{\theta}|\boldsymbol{D}_2, I) \propto p_{\text{MESXR}}(\boldsymbol{d}_{\text{MESXR}}|\boldsymbol{\theta}, I)\, p_{\text{SXT}}(\boldsymbol{d}_{\text{SXT}}|\boldsymbol{\theta}, I) \\ \times p_{\text{N2}}(\boldsymbol{d}_{\text{N2}}|\boldsymbol{\theta}, I)\, p_{\text{TS}}(\boldsymbol{d}_{\text{TS}}|\boldsymbol{\theta}, I)\, p(\boldsymbol{\theta}|I), \tag{5.16}$$

where $\boldsymbol{D}_2 = (\boldsymbol{d}_{\text{MESXR}}, \boldsymbol{d}_{\text{SXT}}, \boldsymbol{d}_{\text{N2}}, \boldsymbol{d}_{\text{TS}})$. The posterior distribution describes how well a given set of parameters $\boldsymbol{\theta}$ describes the data $\boldsymbol{D}$, so the process of model-fitting is replaced with drawing many samples from this distribution and analyzing their statistical properties.

5.5.3 Results: ME-SXR + TS

This analysis was performed for a single time point (averaged over 1 millisecond) near the end of the enhanced confinement period of a 300 kA PPCD plasma. The posterior distribution $p_1(\boldsymbol{\theta}|\boldsymbol{D}_1, I)$ was sampled using the `emcee` MCMC sampling software [15] until good converge was achieved. The results of this analysis are summarized in the corner plot shown in Figure 5.19. This type of plot, common in Bayesian sampling problems, shows the marginal distributions on the diagonals and contours of the joint two-parameter semi-marginal distributions on the off-diagonals (all but the two specified parameters have been marginalized out). This allows the user to see both how well-constrained the various parameters are, all-things-considered, as well as see how strongly they correlate with other parameters.

We see that some of the model parameters are strongly correlated, especially $n_{Al,0}$ and $n_{C,0}$. This makes sense, as both parameters affect ME-SXR count rate. We also see a correlation between core impurity densities and the amplitudes of the hollow profile bumps. These correlations provide a significant opportunity for these results to be improved with more measurements, as will be done in the next section.

N samples were drawn from the posterior distribution, $\{\hat{\boldsymbol{\theta}}_1, \hat{\boldsymbol{\theta}}_2, \ldots, \hat{\boldsymbol{\theta}}_N\}$, which were used to produce an ensemble of N plasma profiles, i.e. $\{T_e(\rho|\hat{\boldsymbol{\theta}}_1)\}, \{T_e(\rho|\hat{\boldsymbol{\theta}}_2)\}, \ldots, \{T_e(\rho|\hat{\boldsymbol{\theta}}_N)\}$ for the temperature profile, each of which were calculated over a grid of 100 ρ points.

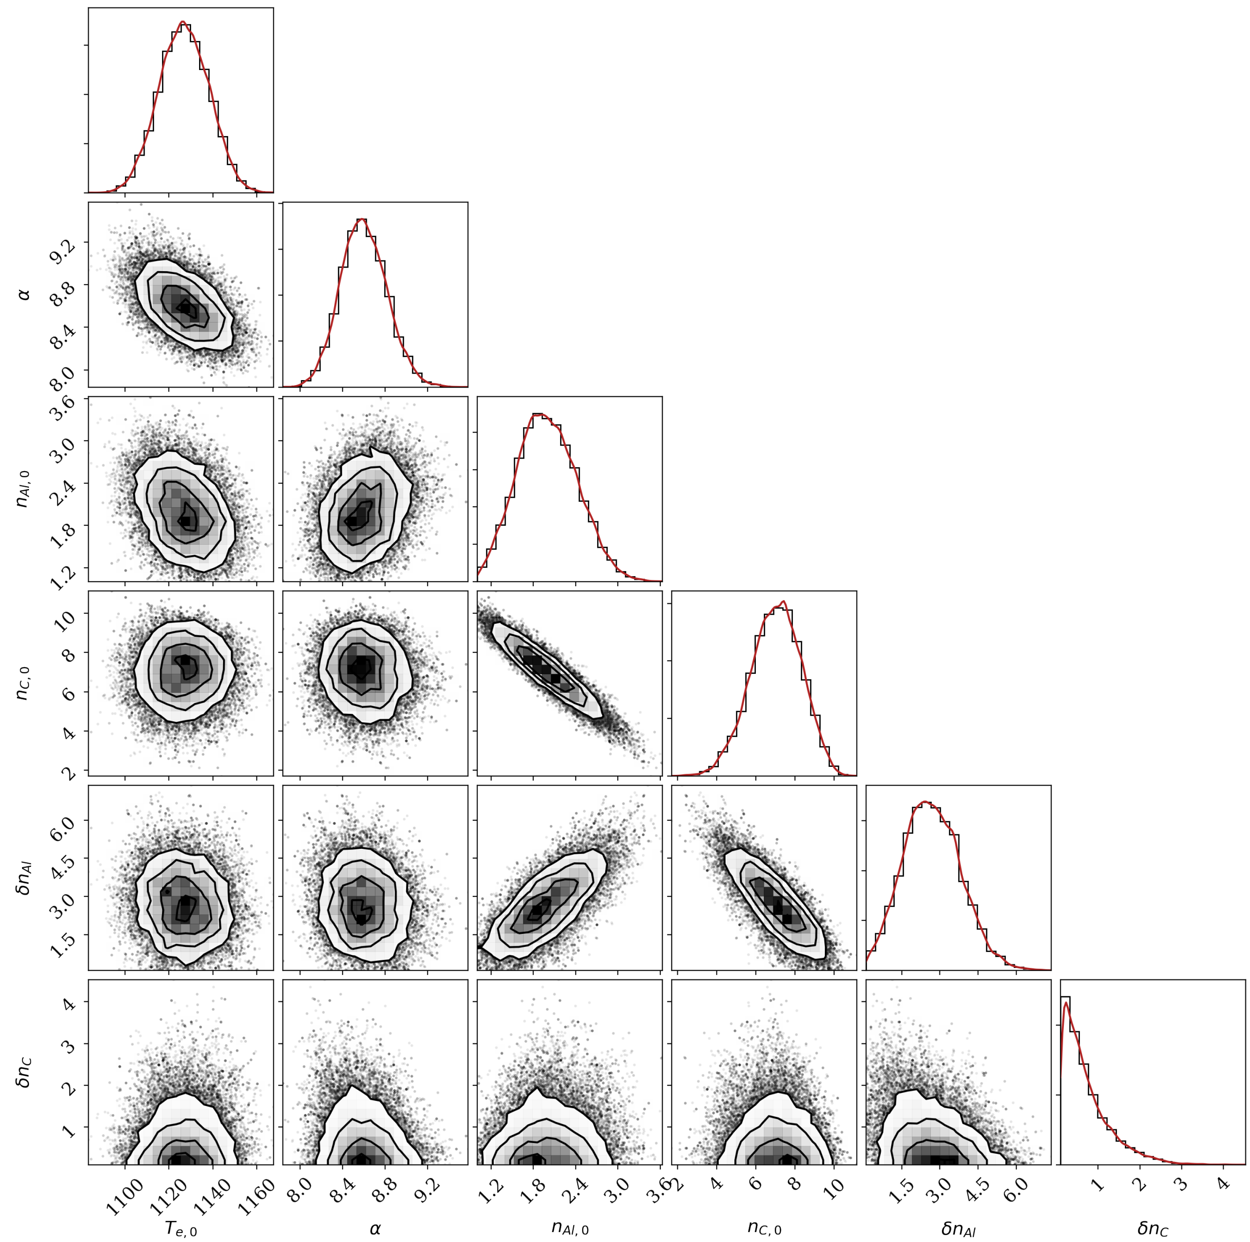

Figure 5.19: Corner plot showing marginal (diagonals) and first-order joint (off-diagonals) distributions for parameters in the joint posterior distribution.

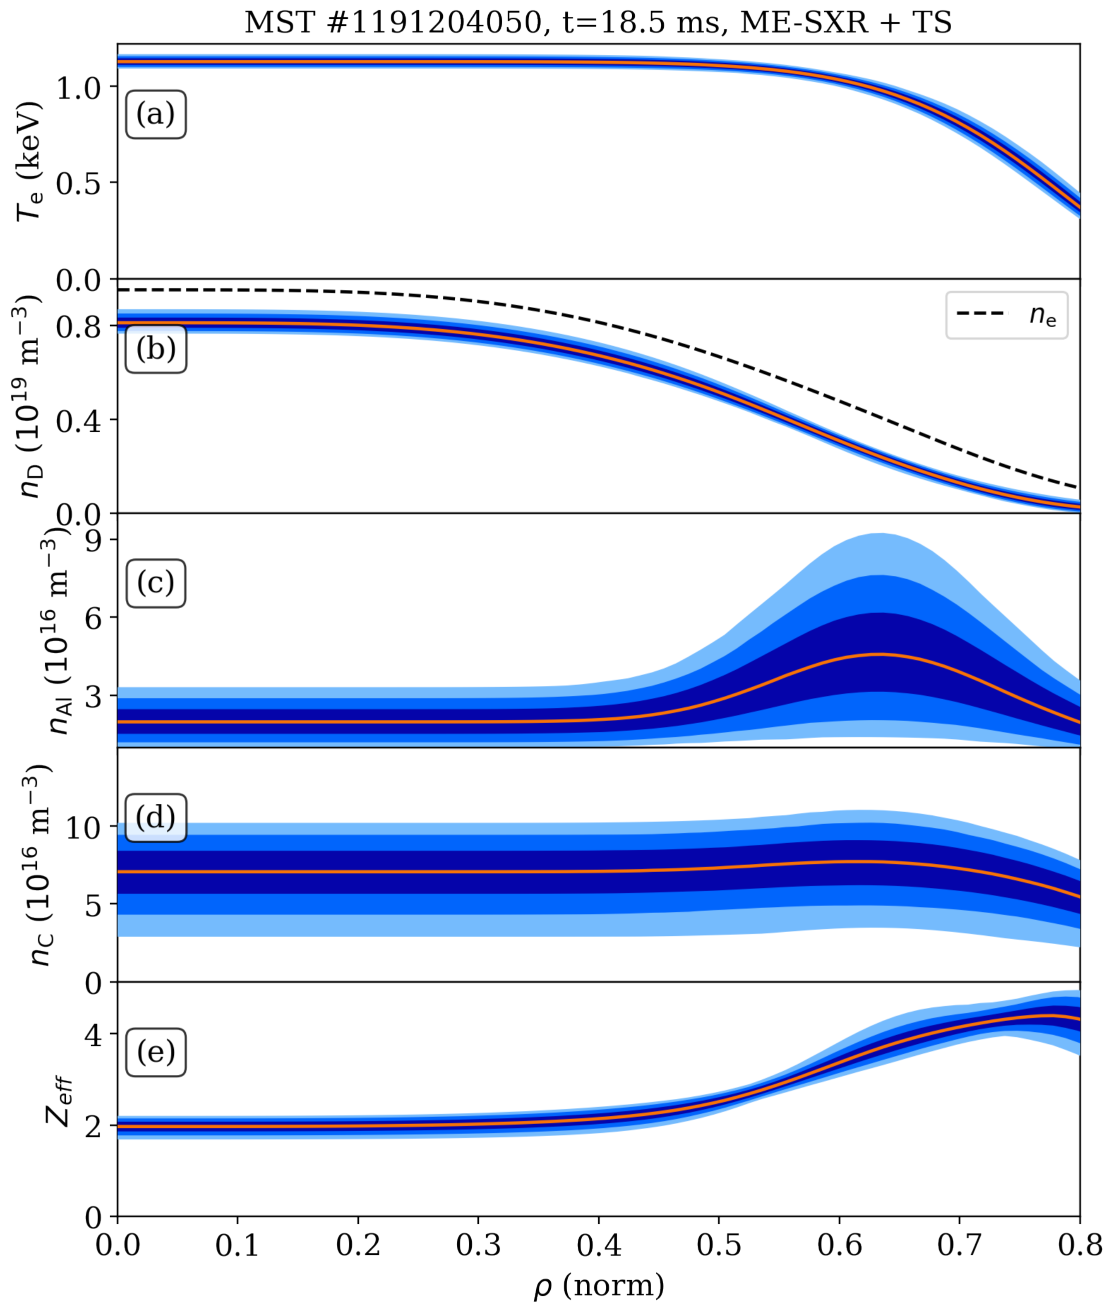

Figure 5.20: Posterior distributions for the ME-SXR + TS analysis, showing (a) electron temperature, (b) deuterium density, and (c) aluminum density, (d) carbon density, and (e) ion-effective charge (Z_{eff}) profiles. The orange line represents the median profile and the shaded regions encompass the 1-, 2-, and 3-σ credibility regions of all possible profiles.

This ensemble was used to estimate the median profile, $\langle T_e(\rho) \rangle$, and the 65%, 95%, and 99.7% credibility regions (the Bayesian equivalent to 1-, 2-, and 3-σ confidence intervals [16]) at each point. These profile samples are shown in Figure 5.20 (a), (c), and (d). The ensemble of profiles can also be used to compute the profiles for derived quantities, like n_D and Z_{eff}, shown in (b) and (d) respectively. These profiles are the primary result of this analysis.

The integration of ME-SXR and Thomson scattering data allows for a highly accurate estimate of the core electron temperature, $T_e = 1125 \pm 12$ eV. The ion densities have also been well-constrained compared to the priors, with core values of $n_{Al} = 1.99 \pm 0.44 \times 10^{16}$ m^{-3} and $n_C = 7.01 \pm 1.3 \times 10^{16}$ m^{-3}. However, because the n_C profile is actually a stand-in for all of the low-Z impurities (C, B, O, and N), its individual value is not necessarily physically meaningful. Instead, the important result is the Z_{eff} profile, which is well-constrained to $Z_{\text{eff}} = 1.97 \pm 0.08$ in the core. This is consistent with previous estimates [7].

This analysis does not, however, do a good job constraining the ion density profiles outside of the core. This is because the ME-SXR model, with the 100 μm filter, has difficulty discerning between n_{Al} and n_C in the lower-signal regions. This can be seen in the strong correlations between n_{Al}, n_C, and δn_{Al} in the corner plot. The overall estimate for Z_{eff}, however, is relatively well-constrained. The high uncertainty in the individual measurements presents a significant opportunity for additional diagnostics to improve these results. Even so, the obtained profiles represent a significant improvement over the priors $p(\boldsymbol{\theta}|I)$.

The best way to assess the quality of a fit is to compare the the output of the models with the original measured data. This too is accomplished through sampling. For each diagnostic, the forward model is calculated for each profile in the ensemble. This produces an ensemble of synthetic measurements for each channel, which are then analyzed statistically. Because these ensembles tend to be very nearly Gaussian, they are

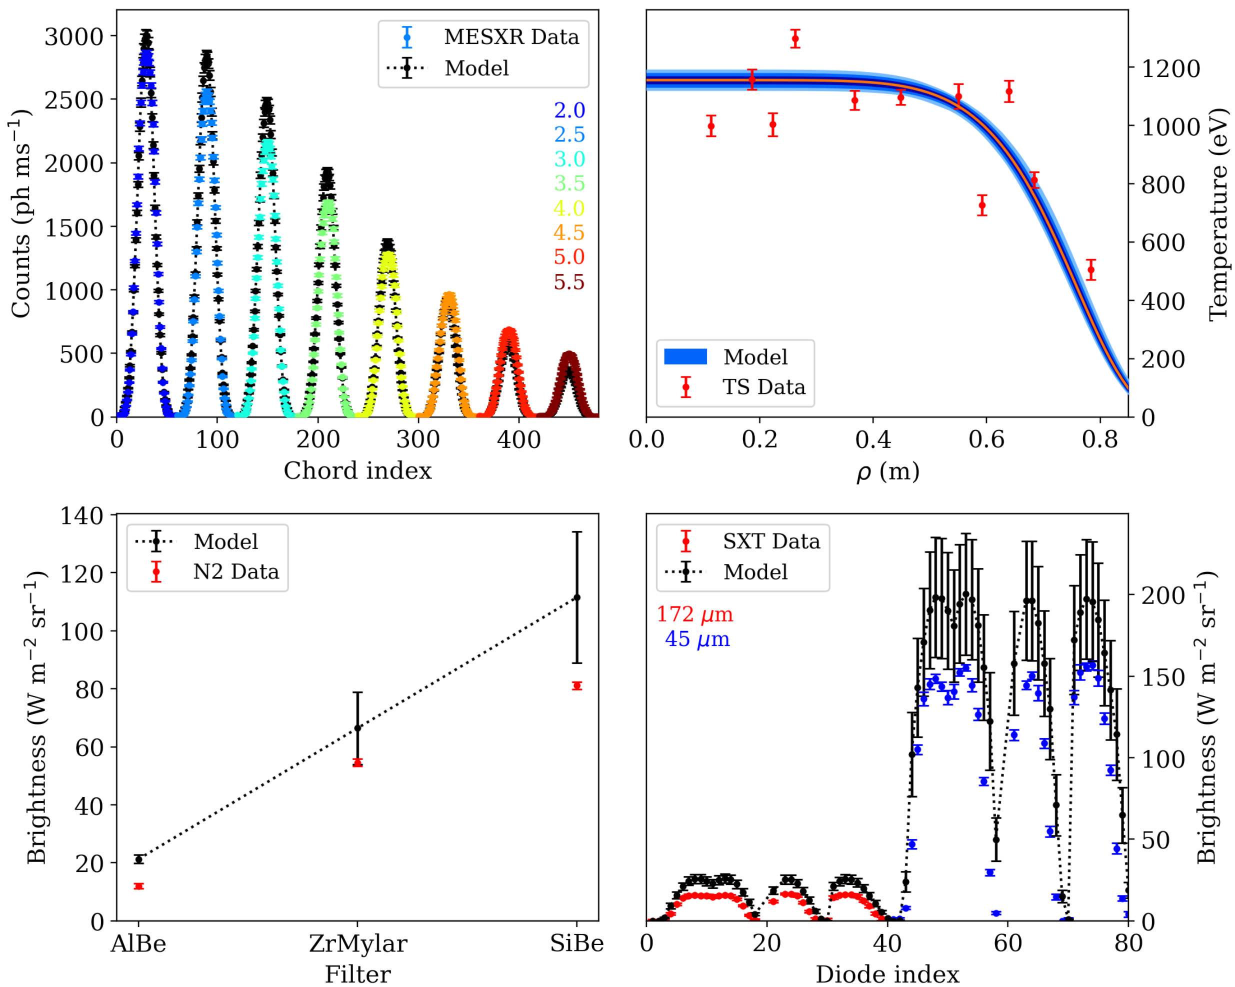

Figure 5.21: Comparison between diagnostics data and model results for (from top left, clock-wise): ME-SXR, Thomson scattering, SXR tomography, and NICKAL2. Results were constrained using only ME-SXR and TS measurements.

well-characterized by the mean and standard deviation. Figure 5.21 shows this data/-model comparison not only for the ME-SXR and Thomson scattering diagnostics, which were included in the analysis, but also for the SXR tomography and NICKAL2, which were not. Agreement for the included diagnostics is good, while the uncertainty is high for the remaining diagnostics. This is because both the SXR tomography and NICKAL2 diagnostics are highly sensitive to the aluminum density, so the large uncertainty on δn_{Al} results in a significant variation in their signals. This implies that incorporating these measurements into the analysis will significantly constrain the ion profiles. This also clearly shows that the ME-SXR sensitivity to aluminum lines, with the 100 μm Mylar filter installed, is relatively weak. However the overall sensitivity to impurities (i.e., Z_{eff}) is strong.

Overall, these results serve as a demonstration that the ME-SXR measurements, when combined with Thomson scattering, can be used to reconstruct T_e, n_{Al}, and Z_{eff} profiles which are well-constrained in the core. These results are both more accurate and more informative than the direct T_e inversion method discussed in Section 5.3, although they require more information and assumptions about the underlying plasma equilibrium. Uncertainty in the ion densities is still large in the edge, though, providing an opportunity for additional diagnostics to further constrain the results. This will be explored in the next section.

5.5.4 Results: Full IDA

The analysis was repeated with the full likelihood function $p_2(\boldsymbol{\theta}|\boldsymbol{D}_2, I)$, which incorporates the SXR tomography and NICKAL2 data and forward models. The posterior was again sampled, with the resulting profiles and credible intervals shown in Figure 5.22. In comparison to the ME-SXR + TS analysis, the addition of new measurements has done little to affect the T_e profile. However, the ion density profiles have been significantly refined with n_C now clearly demonstrating a hollow profile while n_{Al} nearly flat. The core values have also been further refined, especially n_{Al}. However, the change to the resulting Z_{eff} profile is modest. The new profiles are all within the uncertainty bands of the previous analysis (Figure 5.20).

Figure 5.23 repeats the comparison between the forward model samples and the original data. The agreement is not perfect, but agreement is much better than in the previous less-constrained analysis. In particular, uncertainty in the SXR tomography and NICKAL2 models has been dramatically reduced, corresponding to the reduction of uncertainty in n_{Al}. Some level of disagreement is to be expected when simultaneously fitting multiple models, each with their own sources of uncorrelated systematic uncertainty. This is in-fact a feature of IDA, and is the reason we included the σ_m terms in the likelihood functions. The fact that all four diagnostics are in good agreement with the synthetic measurements, and that the resulting profiles are in good agreement with prior results, lends a significant amount to confidence in the ME-SXR forward model.

Finally, it is of interest to make a direct comparison between these results and the ME-SXR + TS analysis. This is shown in Figure 5.24, which compares the marginal distributions for the core values of T_e, n_{Al}, n_C, and Z_{eff} for both analyses, as well as their priors. Both analyses show a similar level of uncertainty in T_e, although the central value shifts somewhat. The most striking difference is n_{Al}, where the addition of SXR tomography measurements has reduced the uncertainty by an order of magnitude. Also notable is that for the ME-SXR + TS analysis the marginal $p(n_C|\boldsymbol{D}_1, I)$ is similar to the prior, implying that the value was not strongly constrained by the data. However by imposing an informed prior based on previous measurements, we were able to infer n_{Al} to reasonable accuracy.

First and foremost, the results of these two analyses demonstrate that the ME-SXR diagnostic, when constrained by additional information (such as Thomson scattering measurements) can be used to infer accurate impurity density profiles for the plasma core to mid-radius. They also demonstrate that by integrating additional information

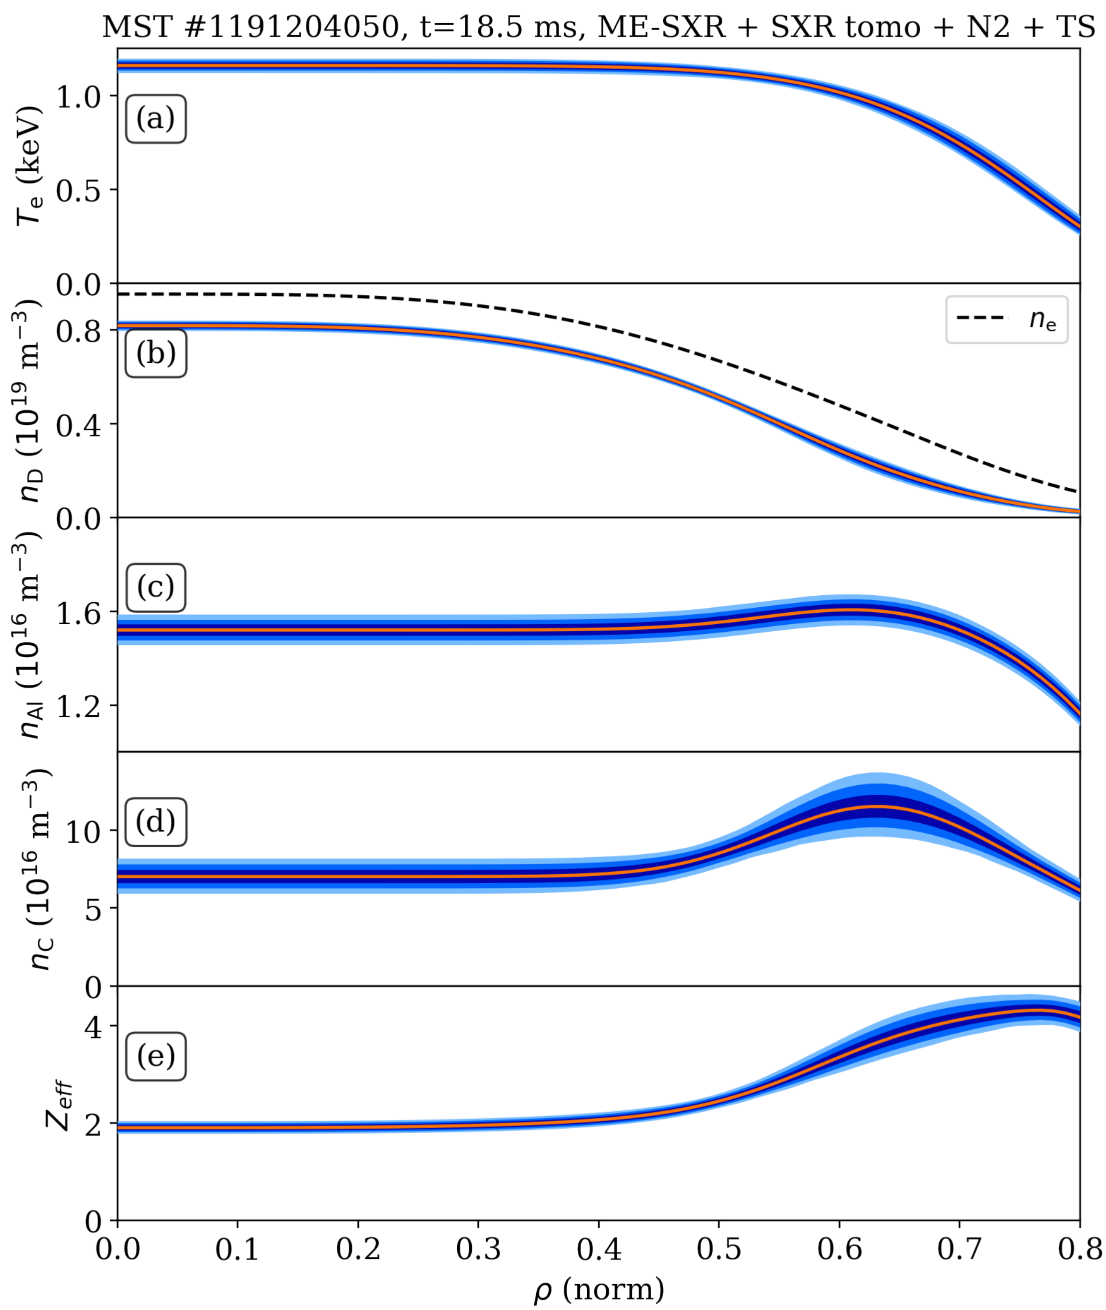

Figure 5.22: Posterior distributions for the full IDA framework, showing (a) electron temperature, (b) deuterium density, and (c) aluminum density, (d) carbon density, and (e) ion-effective charge (Z_{eff}) profiles. The orange line represents the median profile and the shaded regions encompass the 1-, 2-, and 3-σ credibility regions of all possible profiles.

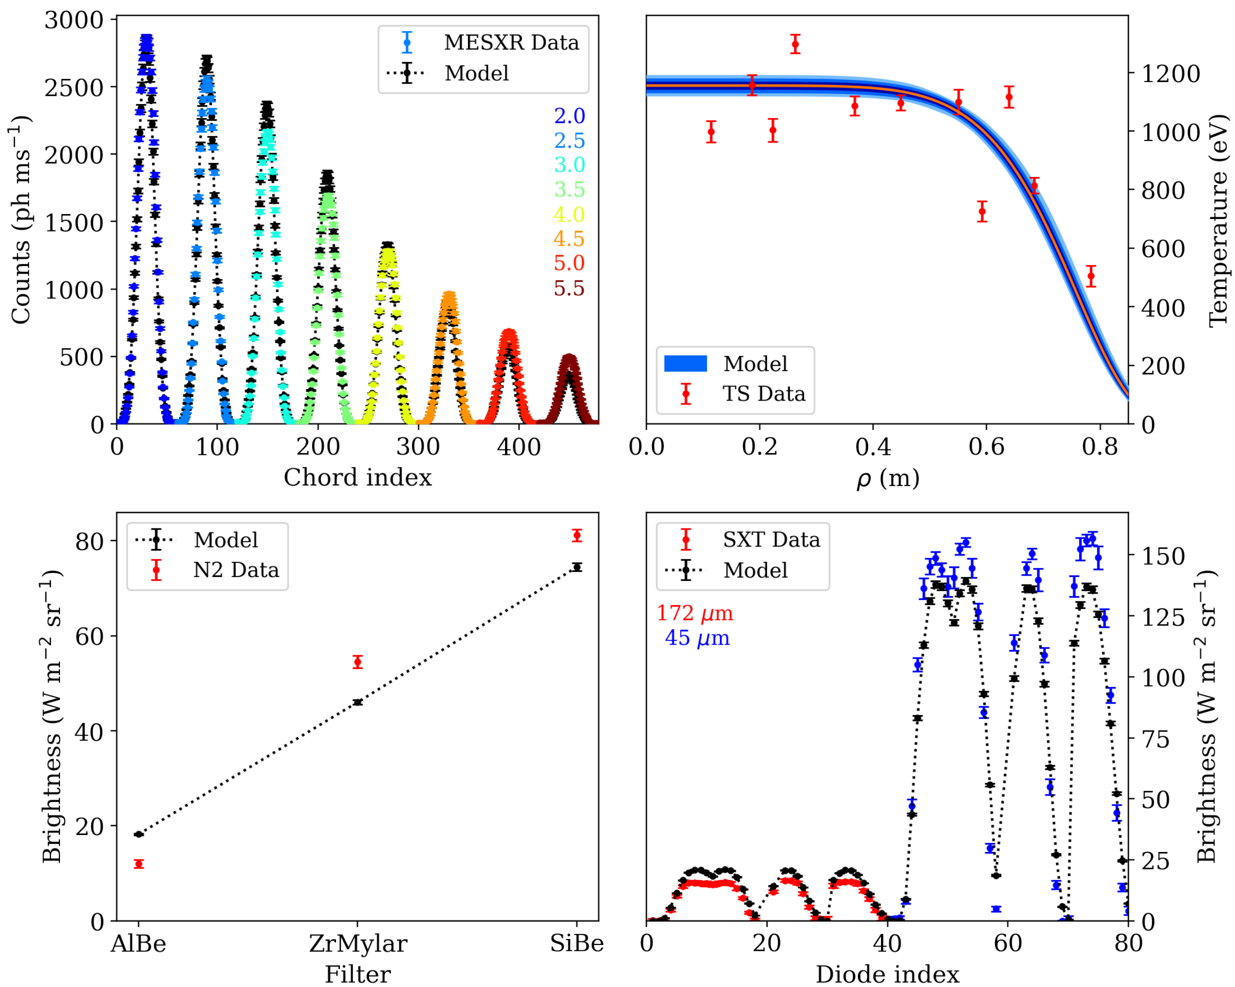

Figure 5.23: Comparison between diagnostics data and model results for (from top left, clock-wise): ME-SXR, Thomson scattering, SXR tomography, and NICKAL2. Results were constrained using all four diagnostics.

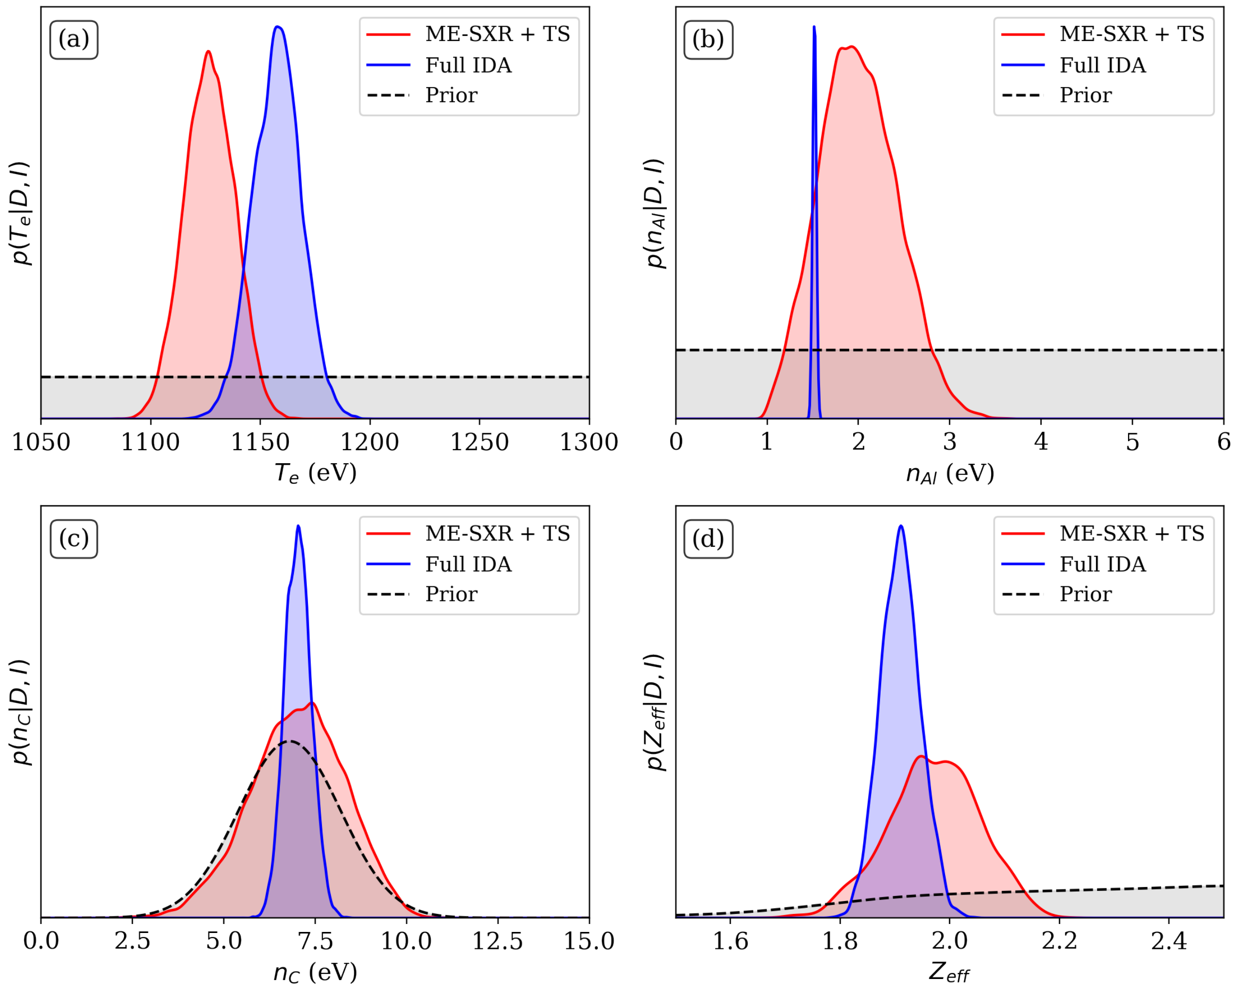

Figure 5.24: Comparison between the marginal posterior distributions for the core (a) electron temperature, (b) aluminum density, (c) carbon density, and (d) Z_{eff} for both analyses. Prior distributions are also shown.

into the IDA framework, these profiles can be further refined and extended out towards the plasma edge. Profiles are inferred in a way which is self-consistent, and the uncertainty analysis is an automatic part of the methodology. This is a powerful methodology, which will form the basis of the QSH analysis performed in Section 6.4. Another important conclusion is that Figure 5.23 shows that it is possible to simultaneously produce synthetic ME-SXR and SXR tomography measurements using physically-realistic profiles which are quantitatively consistent with the measured data. Since the SXR tomography forward model has previously been well-vetted, this provides a significant amount of confidence in the ME-SXR forward model developed in Chapter 4.

Bibliography

[1] R. E. Bell, “Inversion technique to obtain an emissivity profile from tangential line-integrated hard x-ray measurements,” *Review of Scientific Instruments*, vol. 66, no. 1, pp. 558–560, 1995. [Online]. Available: https://doi.org/10.1063/1.1146350

[2] J. K. Anderson, C. B. Forest, T. M. Biewer, J. S. Sarff, and J. C. Wright, “Equilibrium reconstruction in the Madison Symmetric Torus reversed field pinch,” *Nuclear Fusion*, vol. 44, no. 1, pp. 162–171, 2004. [Online]. Available: https://doi.org/10.1088/0029-5515/44/1/018

[3] L. M. Reusch, M. D. Nornberg, J. A. Goetz, and D. J. Den Hartog, “Using integrated data analysis to extend measurement capability (invited),” *Review of Scientific Instruments*, vol. 89, no. 10, 2018. [Online]. Available: https://doi.org/10.1063/1.5039349

[4] D. Li, J. Svensson, H. Thomsen, F. Medina, A. Werner, and R. Wolf, “Bayesian soft X-ray tomography using non-stationary Gaussian Processes,” *Review of Scientific Instruments*, vol. 84, no. 8, pp. 0–10, 2013. [Online]. Available: https://doi.org/10.1063/1.4817591

[5] R. O’Connell, D. J. Den Hartog, C. B. Forest, J. K. Anderson, T. M. Biewer, B. E. Chapman, D. Craig, G. Fiksel, S. C. Prager, J. S. Sarff, S. D. Terry, and R. W. Harvey, “Observation of Velocity-Independent Electron Transport in the Reversed Field Pinch,” *Physical Review Letters*, vol. 91, no. 4, pp. 8–11, 2003. [Online]. Available: https://doi.org/10.1103/PhysRevLett.91.045002

[6] U. Von Toussaint, “Bayesian inference in physics,” *Reviews of Modern Physics*, vol. 83, no. 3, pp. 943–999, 2011. [Online]. Available: https://doi.org/10.1103/RevModPhys.83.943

[7] M. Galante, L. Reusch, D. Den Hartog, P. Franz, J. Johnson, M. McGarry, M. Nornberg, and H. Stephens, “Determination of Z_eff by integrating measurements from x-ray tomography and charge exchange recombination spectroscopy,” *Nuclear Fusion*, vol. 55, no. 12, p. 123016, 2015. [Online]. Available: https://doi.org/10.1088/0029-5515/55/12/123016

[8] M. B. McGarry, “Probing the relationship between magnetic and temperature structures with soft x-rays on the Madison Symmetric Torus,” Ph.D. dissertation, University of Wisconsin-Madison, 2013.

[9] S. T. Kumar, D. J. Denhartog, K. J. Caspary, R. M. Magee, V. V. Mirnov, B. E. Chapman, D. Craig, G. Fiksel, and J. S. Sarff, “Classical impurity ion confinement in a toroidal magnetized fusion plasma,” *Physical Review Letters*, vol. 108, no. 12, pp. 3–6, 2012.

[10] L. Reusch, P. Franz, D. Den Hartog, J. Goetz, M. Nornberg, and P. VanMeter, “Model validation for quantitative X-ray measurements,” *Fusion Science and*

Technology, vol. 74, no. 1-2, 2018. [Online]. Available: https://doi.org/10.1080/15361055.2017.1404340

[11] J. Johnson, "Implementing Bayesian Statistics and a Full Systematic Uncertainty Propagation with the Soft X-Ray Tomography Diagnostic on the Madison Symmetric Torus (undergraduate thesis)," University of Wisconsin-Madison, Madison, Tech. Rep., 2013.

[12] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: The MIT Press, 2006.

[13] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. Boca Raton, FL: CRC Press, 2014.

[14] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, oct 2011. [Online]. Available: http://scikit-learn.sourceforge.net

[15] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, "emcee: The MCMC Hammer," *Publications of the Royal Astronomical Society of the Pacific*, vol. 125, pp. 306–312, 2013. [Online]. Available: http://dan.iel.fm/emcee.

[16] M. D. Nornberg, D. J. Den Hartog, and L. M. Reusch, "Incorporating Beam Attenuation Calculations into an Integrated Data Analysis Model for Ion Effective Charge," *Fusion Science and Technology*, vol. 74, no. 1-2, pp. 144–153, 2018. [Online]. Available: https://doi.org/10.1080/15361055.2017.1387008

Chapter 6

Evolution of the helical RFP

The MST features a robust selection of x-ray diagnostics which can be used on a routine basis: the new ME-SXR detector is highly versatile, and can be used to diagnose thermal properties, impurity content, and non-thermal populations; the SXR tomography diagnostic, now operated with thin 45/172 μm Be filters, which features a strong sensitivity to both continuum and aluminum line emissions; and the fast x-ray camera, which is specifically for non-thermal populations. In this chapter this entire suite of diagnostics will be applied to the quasi-single helicity, or QSH, regime of MST plasmas. Using an integrated analysis framework, time resolved 2D measurements of T_e structure and evolution have been produced for the first time in a QSH plasma on the MST. These measurements suggest the presence of a transport barrier around the helical core. A brief period of greatly enhanced confinement is also observed when the secondary mode activity is suppressed. A brief overview of the contents of this chapter is provided below.

Section 6.1 reviews the existing literature on QSH plasmas in the MST and RFX-mod, and discusses the distinct phases of the saturated QSH state. Section 6.2 discusses a theoretical model of the QSH state based on shear-suppression of energy transfer, and present presents recent measurements that are consistent with this model. Section 6.3

uses the ME-SXR and FXR detectors to analyze the buildup, sustainment, and decline of runaway populations in the QSH state. Section 6.4 uses SXR tomography and FIR data to analyze the time-evolving T_e, n_e, and n_{Al} profiles over the duration of a QSH flattop, revealing a period of significantly enhanced confinement. Finally, Section 6.5 presents recent observations of high-frequency turbulent fluctuations which correlate to the position of the helical structure.

6.1 Quasi-single helicity plasmas in the MST

The term *quasi-single helicity* (QSH) is used to describe the scenario in which plasma's magnetic spectrum is dominated by a single core-resonant tearing mode [1]. This is opposed to "multi-helicity" (MH), a term which has been adopted to refer to standard RFP plasmas in which the magnetic energy is more-evenly distributed across many $m = 1$ modes. The "quasi-" prefix is included to emphasize the fact that the mode spectrum is not purely single helicity — the non-dominant modes have finite amplitude and contribute significantly the physics. The "strength" of a QSH state is commonly quantified by the so-called spectral index [2],

$$N_S = \left[\sum_n \left(\frac{b_n^2}{U} \right)^2 \right]^{-1}, \tag{6.1}$$

where b_n is the amplitude of the n$^{\text{th}}$ magnetic perturbation and $U = \sum_n b_n^2$ is the total magnetic energy. This is a measure of how concentrated the magnetic energy is within a single mode, and will approach one as $b_{n'}^2 \gg b_{n \neq n'}^2$. By convention, a plasma is considered as QSH when $N_S < 2$. All plasmas with $N_S > 2$ are considered to be MH.

QSH spectra can occur under various operating conditions in the MST. QSH states may form spontaneously during the sawtooth cycle of standard reversed plasmas, typically with a large $n = 6$ mode surviving for a few milliseconds before crashing back to MH. The improved confinement state of PPCD also technically meets the requirement

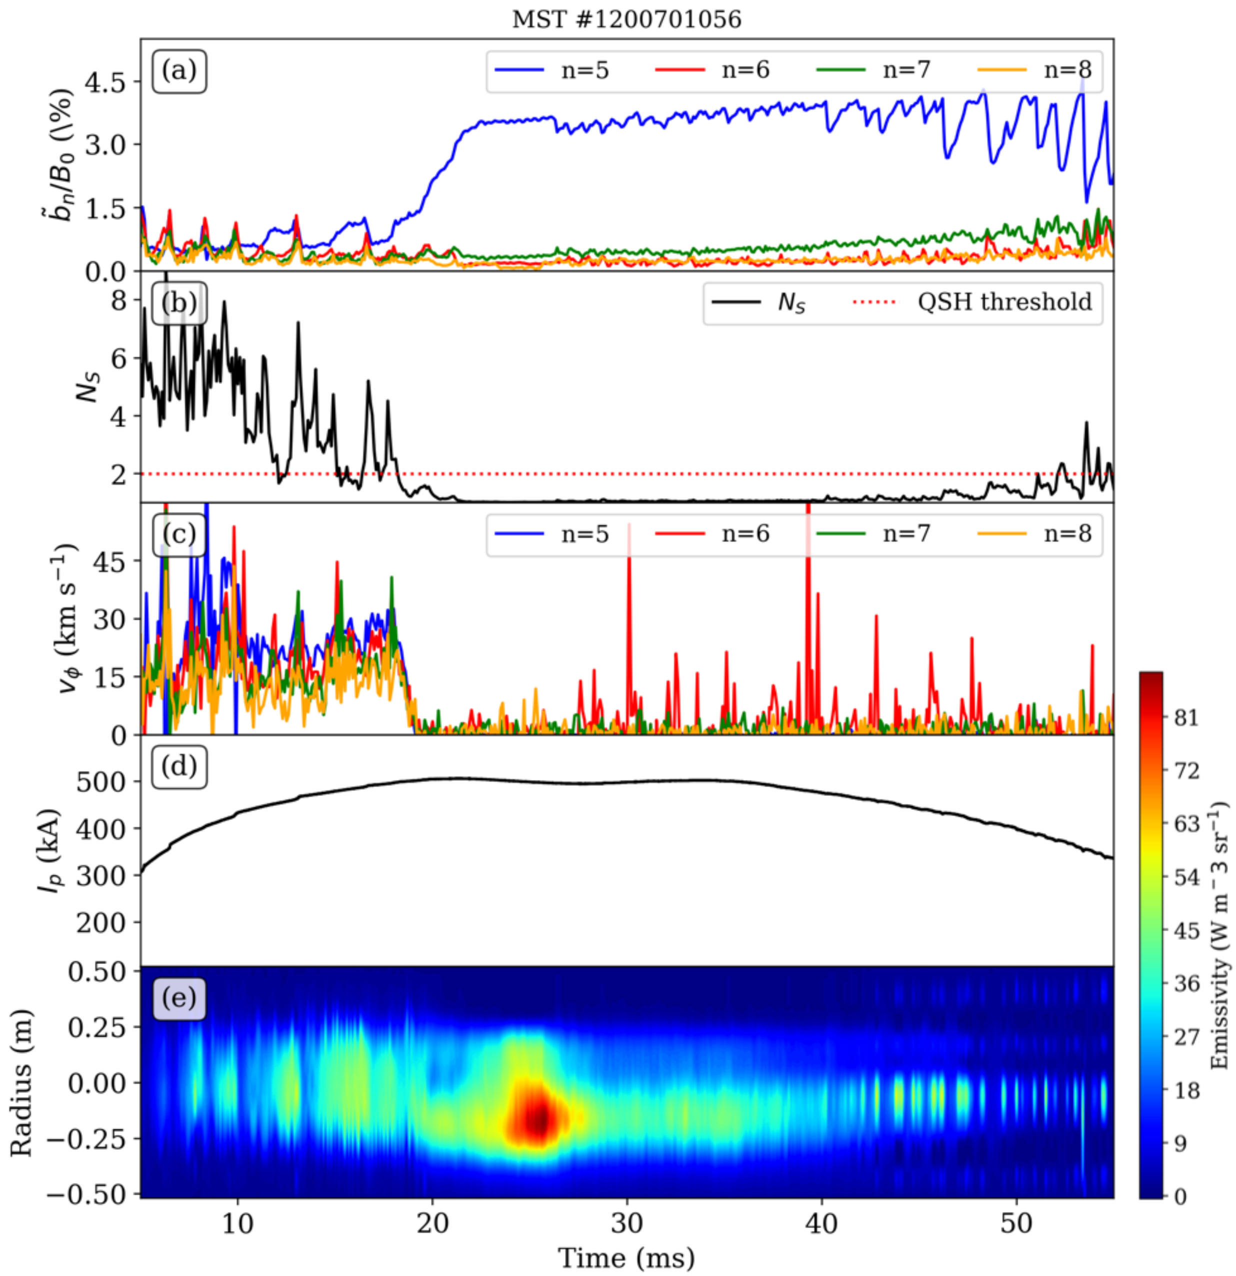

Figure 6.1: Set of plots demonstrating the evolution of a high-performance $F = 0$ QSH plasma, showing: (a) core-resonant $m = 1$ magnetic mode amplitudes normalized to the equilibrium field, (b) spectral index, (c) mode rotation velocity, exhibiting locking, (d) plasma current, and (e) a core slice of the SXR emissivity.

for QSH ($N_S < 2$), however it is rarely referred to as such because the overall mode amplitudes are very low, including the dominant $n = 6$ mode [3]. In recent years, however, the term QSH has mostly come to refer to the very large, long-lived $n = 5$ modes that tend to form when MST is operated in the "non-reversed" ($F = 0$) mode, meaning that the toroidal field has been set to exactly zero at the shell. In the standard RFP, three-wave mode coupling to an $m = 0$ mode is known to provide a significant pathway for energy transfer between the core-most unstable tearing modes and other higher-n modes in the plasma [4], so removal of the $m = 0$ resonant surface from the plasma is presumed to greatly reduce this coupling and therefore promote the accumulation of energy. Plasma rotation tends to rapidly slow down and lock as the mode amplitude grows large due to interaction with eddy currents induced in the conducting shell [5]. This typical evolution is shown in the various panels of Figure 6.1.

Plasma with a strongly QSH spectrum have been observed to exhibit a three-dimensional helical structure [6], much like that of a stellerator [7]. The plasma undergoes a change in magnetic topology where the island associated associated with the dominant mode grows to such amplitude that it encompasses the magnetic axis, shedding its separatrix in the process [8]. Such a plasma is said to be in a *single helical axis* (SHAx) state. Not all QSH plasmas are in a SHAx state, but a plasma in a SHAx state must necessarily exhibit a QSH mode spectrum. High current $F = 0$ plasmas in the MST tend to both exhibit QSH and SHAx properties, so the term "QSH" is often used as a shorthand for this configuration. However, properly QSH and SHAx are distinct terms.

SHAx equilibria are associated with a restoration of good flux surfaces within the helical perturbation [9]. On RFX-mod this has been demonstrated to lead to improved confinement and the introduction of an internal transport barrier around the helical structure [10]. This is illustrated in Figure 6.2, which compares the typical stochasticity of MH with the restored good flux surfaces of SHAx. Soft x-ray tomography measurements have been used to directly observe internal helical structure in both devices. This was

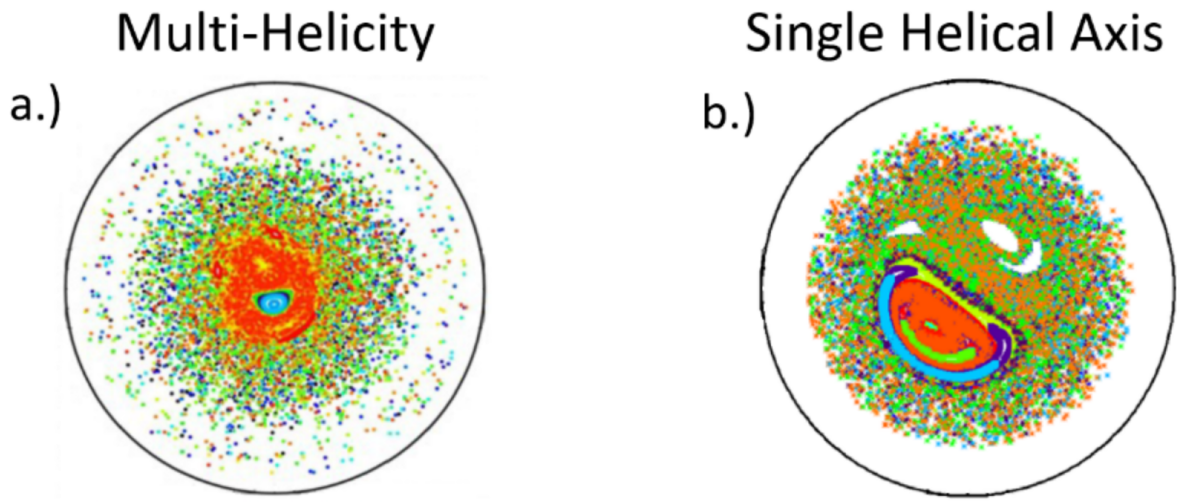

Figure 6.2: Representative Poincaré plots for a poloidal cross-section of the MST showing (a) multi-helicity and (b) quasi-single helicity SHAx euilibria. Figure reproduced from Munaretto, *et al.* [11].

shown in Figure 2.11.

RFX-mod has robustly demonstrated that electron thermal confinement is significantly enhanced during the SHAx state [12]. This can be seen in Figure 6.3, which shows a nearly two-fold increase to the core temperature. This improvement is also associated with a substantial increase in the edge temperature gradient to $\nabla T_e > 1$ keV m^{-1}. Gyrokinetic modeling suggests that such strong temperature gradients are expected to drive microtearing modes, a class of very short wavelength electromagnetic turbulence, unstable [13]. Direct measurements of the magnetic spectrum using $\dot{b}$ coils provides some support for this claim, observing that coherent fluctuations form around 200 kHz during QSH periods with $n \approx 200$ [14]. Similar observations have not previously been explored in the MST. Because of these significant improvements in electron confinement, QSH has been suggested as the possible operating regime for a hypothetical RFP fusion reactor [15].

Studies of ion confinement during QSH have had somewhat more mixed results, however. Studies at RFX-mod have confirmed a general improvement to the confinement

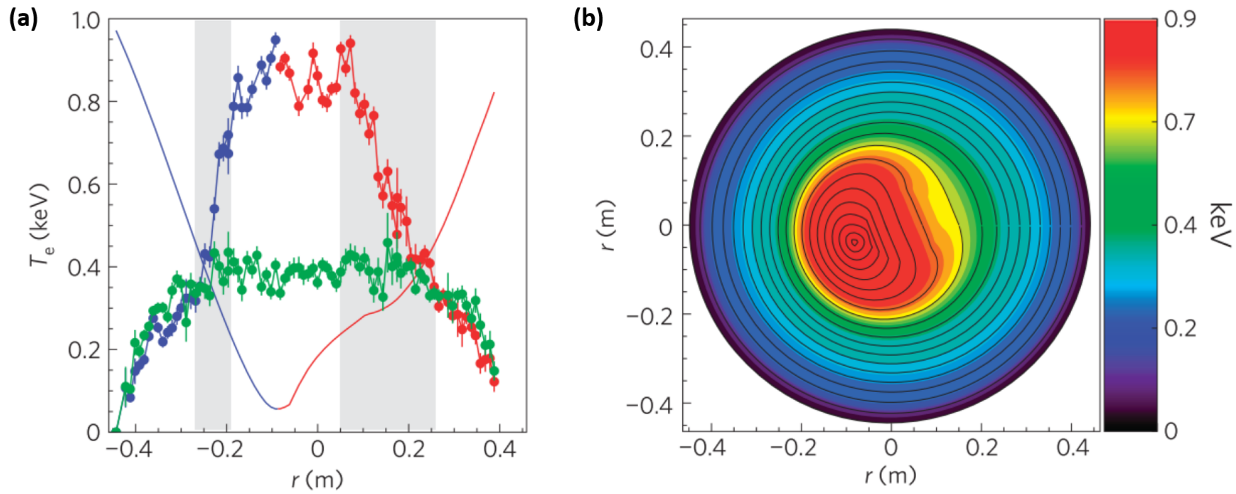

Figure 6.3: Electron temperature confinement is significantly improved on the RFX during SHAx. Panels show: (a) red and blue data points correspond to opposite sides of the magnetic axis during SHAx, and green is a typical MH plasma; (b) The reconstructed 2D T_e profile exhibits a helical symmetry with strong gradients. Figure reproduced from Lorenzini, *et al.* [15].

of thermal ions [16]. However they have also noted to tendency for an external transport barrier to form which results in the buildup of impurities into hollow profiles with peaks well outside of the core [17]. This tendency has not been observed on the MST. Multiple studies on the MST using neutral beam injection have shown that fast ions exhibit generally worse confinement in the SHAx regime than in MH [18]. This fact poses a serious obstacle for SHAx as a fusion reactor concept. However, it has since been seen that the poor confinement is related to remnant secondary mode activity [19]. Since secondary mode amplitude tends to decrease with increasing current, it is possible that sufficiently high current might suppress this instability and achieve sufficient fast ion confinement for a reactor.

The transition to, and persistence within, the SHAx state in $F = 0$ conditions has been strongly correlated to the Lundquist number,

$$S = \frac{\tau_R}{\tau_A} \propto \frac{I_p T_e^{3/2}}{n_e^{1/2}}, \tag{6.2}$$

where τ_R is the resistive time and τ_A is the Alfvén time [20]. This can be seen in Figure 6.4, which shows a nearly-linear relation between S and QSH persistence until saturating at high S. Note that this plot does not control for variation in T_e, which is likely responsible for some of the observed variance. The dependence on Lundquist number helps to explain why RFX (which has a typically operates at both higher I_p and higher n_e than the MST) and the MST observe relatively similar dynamics. QSH states can be observed with I_p as low as 300 kA in the MST, however in such cases these states tend to be intermittent and very short-lived.

Previous observations of thermal electron profile evolution during SHAx in the MST have been somewhat rare. A 2015 study by S. Munaretto, *et al.* [11] used V3Fit (see Section 2.2.2) constrained by FIR and Thomson scattering measurements. That study found the electron density to be relatively smooth and well-confined during the SHAx period, while T_e displayed a surprising intermittency. These results are reproduced in Figure 6.5. The same study also demonstrated that a radial magnetic perturbation (RMP) [20] applied by a ring of magnetic coils around the poloidal gap could be used to effectively re-introduce stochasticity in the core. An RMP has also been used to control the plasma's locking phase. This technique was not used for the data presented in this thesis.

6.1.1 Phenomenology of the QSH flattop

In previous work, the evolution of non-reversed QSH/SHAx plasmas in the MST have typically been divided into a small number of distinct phases:

1. **The early phase**: From the beginning of the plasma discharge until the start of the rise phase. Plasmas typically exhibit MH behavior, but will occasionally exhibit

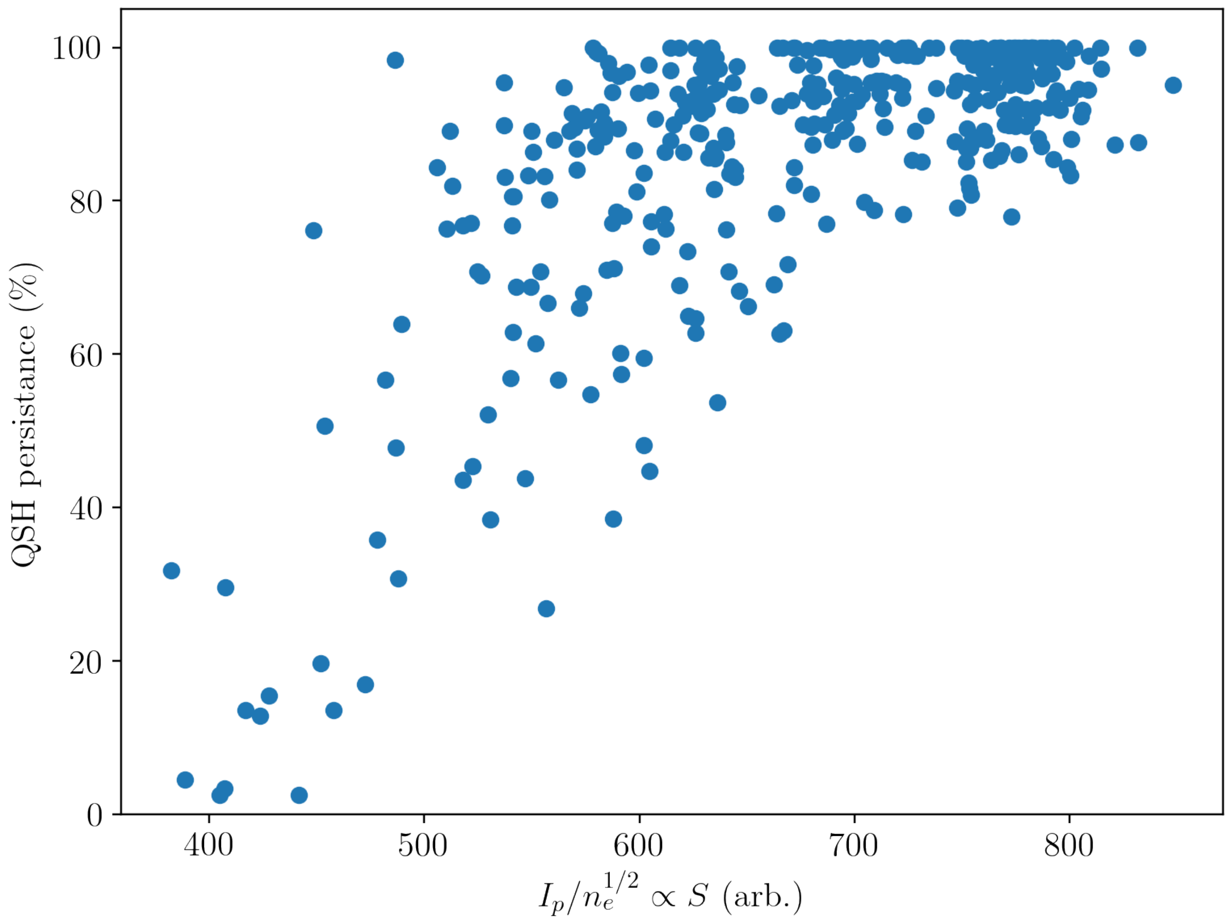

Figure 6.4: QSH persistence vs $I_p/n_e^{1/2}$, which is a proxy for Lundquist number. Here persistence is defined as the percentage of time that $N_S < 2$ during the current flattop.

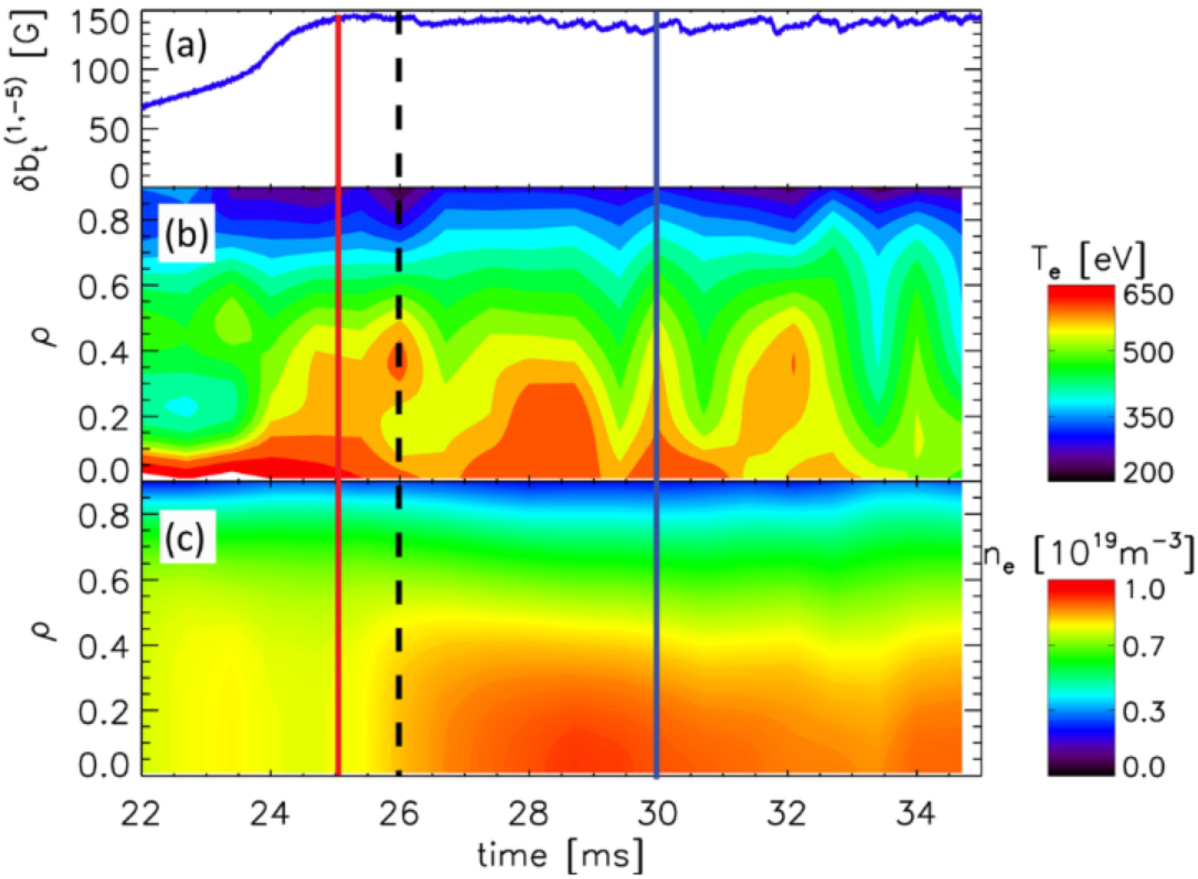

Figure 6.5: Previous observations of electron density (b) and temperature (c) evolution during the saturation of a large single magnetic mode (a) in the MST. Reproduced from S. Munaretto, *et al.* [11].

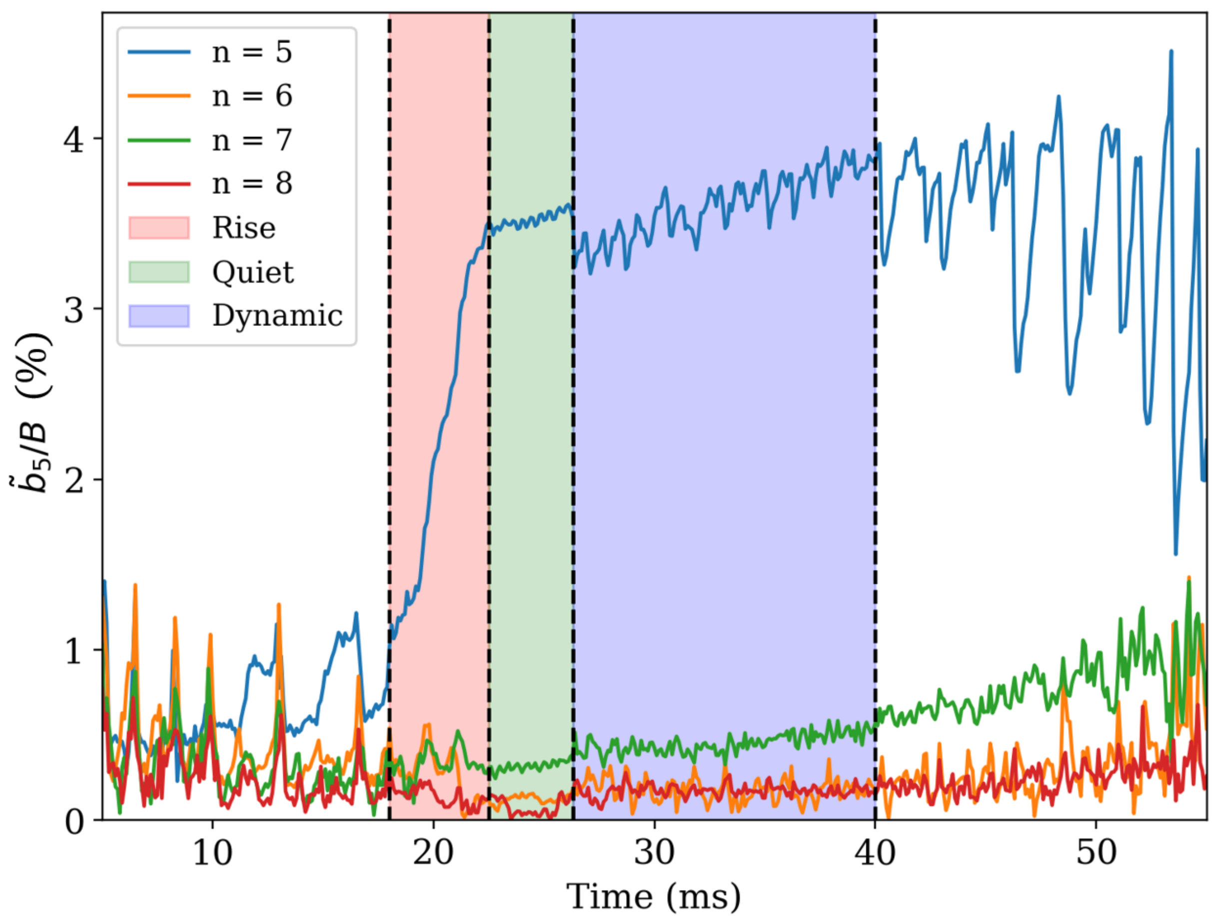

Figure 6.6: The evolution of the dominant $n = 5$ mode can be further subdivided into the rising, quiet flattop, and dynamic flattop phases.

short intervals of small amplitude, short-lived QSH.

2. **The rise phase**: The period during which $n = 5$ mode is growing rapidly to large amplitude. Plasma rotation typically locks during this phase, and the bifurcation to a SHAx equilibrium is believed to occur.

3. **The flattop phase**: The growth of the $n = 5$ mode stops and the amplitude remains large for tens of milliseconds. $N_S \approx 1$ for this entire phase, and secondary mode activity is minimal. This phase ends as the plasma current begins to decrease and the equilibrium breaks down.

This thesis claims, and will demonstrate, that the flattop phase should actually be further divided into two distinct phases:

3a. **The quiet flattop phase**: A period of minimal secondary mode activity, when the $n = 5$ mode amplitude is very large and stable. This phase always immediately follows the rise phase, and lasts for a few ms.

3b. **The dynamic flattop phase**: The rest of the flattop, marked by a resumption of low-amplitude secondary mode activity. The dominant mode typically exhibits more substantial oscillations in amplitude.

Not every MST QSH plasma features a quiet phase, but they occur in most high-quality discharges (for example, the discharge shown in Figure 6.5 does feature a quiet phase). As with many QSH features, it appears preferentially in low-density discharges. Plasmas also tend to lack quiet phases when the QSH state forms unusually late in the discharge or when the rise phase takes unusually long. The thing which makes the distinction between the quiet and dynamic phases important, and the thing which caught my attention in the first place, is that the quiet phase corresponds to a noticeable enhancement in core properties like T_e, n_e, and n_Z. Since the plasma is ohmically heated, the increase in T_e is suggestive of improved confinement.

The reduction in confinement which occurs at the transition from the quiet to dynamic flattop phases is readily seen in a tomographic inversion of the SXR emissivity. Figure 6.7 shows the same shot as featured in Figure 6.6. Two time slices only 4 milliseconds apart are singled out, one during the quiet phase at $t = 25$ ms and another in the early dynamic flattop at $t = 29$ ms. The panel in the lower left shows the full tomographic inversion during the quiet phase, which exhibits a large, bright emissivity structure. The panel in the lower right shows that only four milliseconds the emissivity has dropped by a factor of two and the width of the structure has significantly narrowed. The middle panel shows that this narrower structure remains consistent throughout the rest of the dynamic flattop. It also confirms that the loss of the hot island corresponds exactly with the end of the quiet phase.

We will also see in Section 6.3 that the end of the quiet phase tends to trigger a sudden reduction in fast electron confinement, providing further evidence for a reduction in the size of the good flux surfaces. We will also find in Section 6.4.3 that the quiet phase is associated with an increase in electron temperature, electron density, and impurity ion densities. It is clear that understanding the quiet phase, in particular how to extend it, could contribute significantly to the development of a hypothetical QSH-based RFP reactor.

We do not yet fully understand what triggers the transition between the quiet and dynamic flattop phases. The dynamic flattop is defined by the resumption of small amounts of tearing mode activity. This might be related to a change in the current profile, as some distinctions can be seen in the shape of the I_p flattop for plasmas which feature a quiet phase relative to those that do not. The transition is frequently seen occur around the time that the current from the third capacitor bank becomes large, which could be responsible for disrupting a marginally-stable equilibrium. There is also some indication (Section 6.5) that the large T_e gradient can trigger an increase in turbulence, which may be responsible for enhancing the coupling between the dominant and secondary modes.

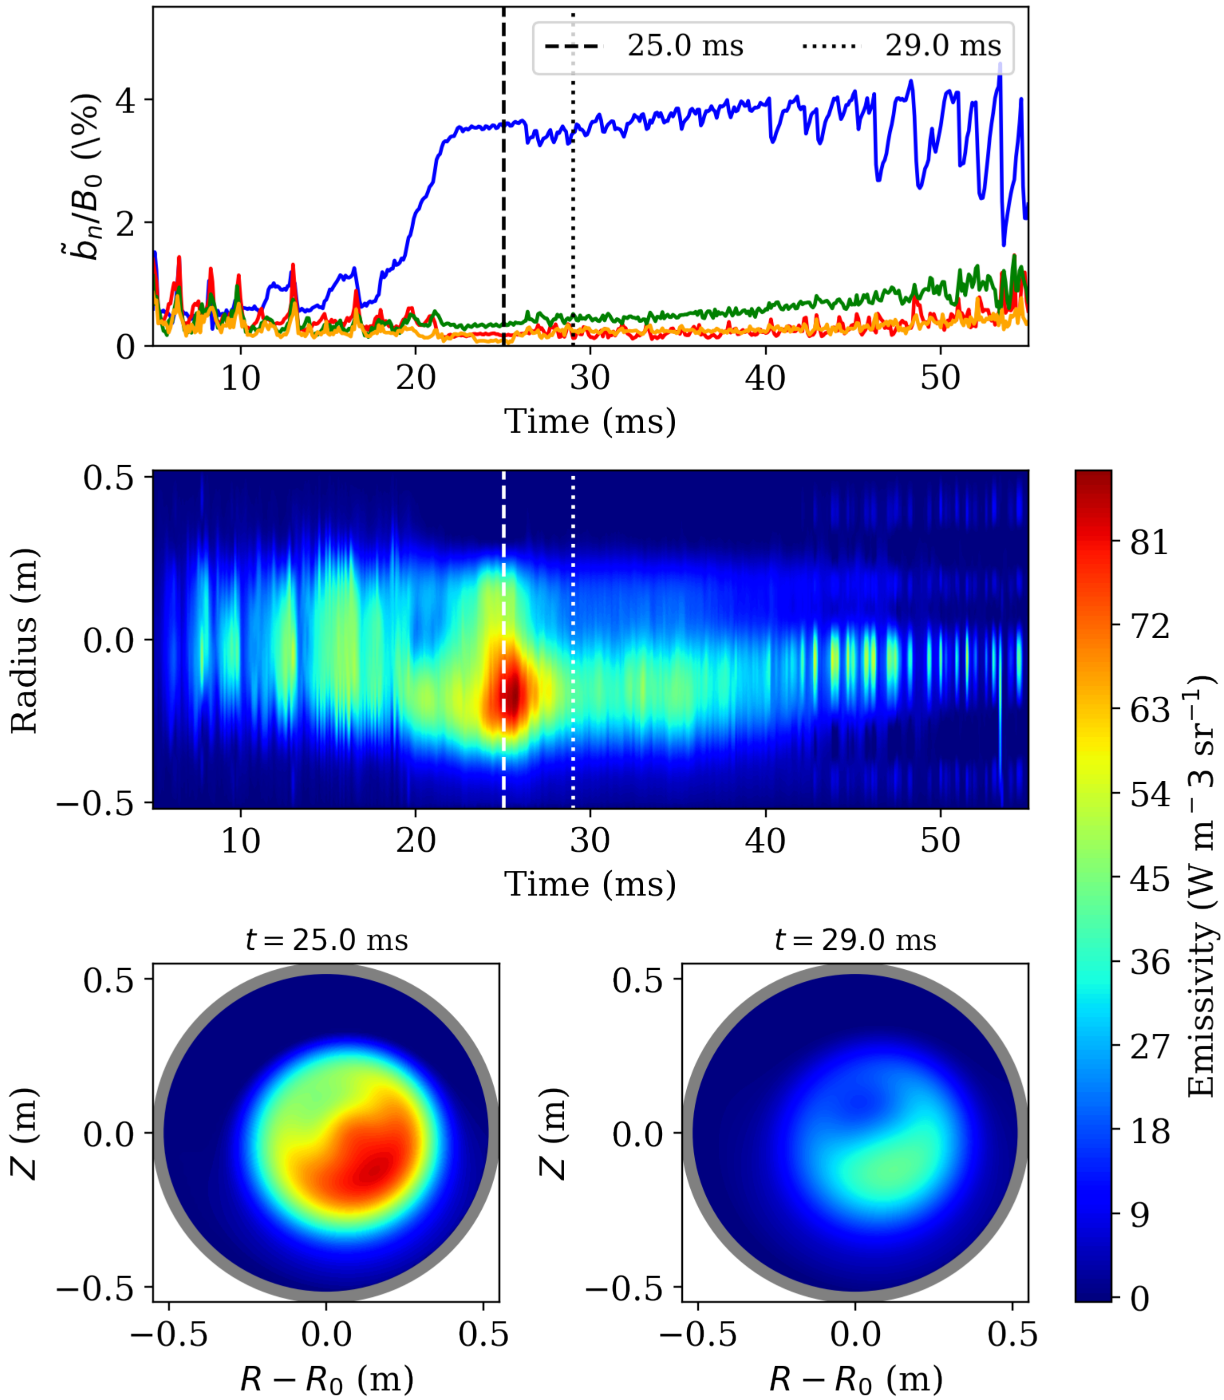

Figure 6.7: Demonstration of enhanced SXR emission during the quiet flattop phase. Panels show: (top) dominant and secondary mode amplitudes with two time slices indicated by vertical lines; (middle) Radial slice of the inverted SXR tomographic inversion over time; (bottom) 2D SXR tomographic inversions at the two indicated time slices.

A more compelling hypothesis, based on increased flow shear due to mode locking, is discussed further in Section 6.4.4. Future work will be necessary in order to test these hypotheses.

6.2 Theoretical models of QSH dynamics

A fully-consistent physics-based model of QSH formation and sustainment has proven to be an elusive goal for the RFP community. Early models found that by imposing a high-dissipation regime, cylindrical RFP-like plasmas could made to exhibit a single-helicity equilibrium in simulations [21]. Simulations using the SpeCyl code found that QSH tended to form at low Hartmann number [22],

$$Ha = \frac{S}{\sqrt{Pm}} \tag{6.3}$$

where S is the Lundquist number and Pm is the magnetic Prandtl number, which corresponds to a high-dissipation regime. One consequence of this model is the formation of an electrostatic drift velocity $\boldsymbol{v}_D = \boldsymbol{\nabla}\phi \times \boldsymbol{B}/B^2$ which shares the dominant mode helicity and sustains a steady-state MHD dynamo [23]. These simulations were not necessarily considered to imply that high dissipation is the driving mechanism of QSH formation in the real experiments, but mainly served to force a simulation to "emulate" the SHAx state. Nonetheless, features of these experiments have been used to interpret experimental data [24, 25].

An immediate problem with this framework can be seen in Equation 6.3, which shows a linear scaling between the Hartmann and Lundquist numbers. However, as shown in Figure 6.4, QSH formation and persistence has been observed to prefer large Lundquist number. A low Hartmann number regime would imply QSH should form preferentially at low plasma current, which seemingly contradicts experimental observations[1]. Furthermore, the value for Ha required to achieve QSH in these simulations is orders of magnitude lower than what can be achieved experimentally [26]. Recent measurements of velocity flow patterns all show inconsistency with the SpeCyl results, and are more consistent with NIMROD simulations[27] run at reduced periodicity ($1/5^{\text{th}}$ of the torus) and significantly higher Hartmann number [28]. It has also been found that the application of an edge radial magnetic perturbation (RMP) can induce a SHAx state both in experiment and in theory [29, 30]. However, this does not present a clear model for the spontaneous formation of SHAx states when no RMP is applied. For a more thorough description of the various models of QSH/SHAx formation, I refer the reader to J. Boguski's thesis [26].

6.2.1 Shear suppression of secondary tearing modes

An alternative framework has been developed by P. Terry, *et al.* [31, 32] which proposes that shear (either in the magnetic field or flow velocity) is responsible for decoupling energy transfer between the dominant and secondary (sub-dominant) tearing modes. This allows energy to accumulate in the dominant mode, giving rise to the QSH state. This model exhibits multiple phases, proceeding from a multiple-helicity state through a limit cycle regime and then on to fixed-point single helicity as a current-like parameter is increased. So unlike the high-dissipation models, the shear-suppression model satisfies the experimental scaling with Lundquist number. It is also a fully dynamic model which features the QSH state as an emergent feature rather than imposing it via some arbitrary constraint.

The development of the model begins within the framework of reduced MHD. The goal of this exercise was not to provide a detailed model of QSH evolution but rather to capture the fundamental physical principles at work. Hence, a reduced model is appropriate. Note that this model was derived assuming that the $n = 6$ mode is the

[1] This also depends on the exact scaling of the viscosity with I_p, which is not clearly established.

dominant mode, but the resulting behavior should be similar for $n = 5$.

The reduced MHD model for the magnetic potential ψ and the electrostatic potential ϕ is given by

$$\frac{d\omega}{dt} + \nabla_{\|} j = 0 \tag{6.4}$$
$$\frac{d\psi}{dt} + \nabla_{\|} \phi = 0, \tag{6.5}$$

where $\omega = \nabla_{\perp}^2 \phi$ is the vorticity and $j = \nabla_{\perp}^2 \psi$ is the current. Separating the magnetic fluctuations into two fields, the dominant mode $D = |\psi_{n=6}|^2$ and secondary modes $S = \sum_{n>6} |\psi_n|^2$, and introducing a strong shearing rate associated with the dominant mode, it can be shown that the reduced MHD equations yield a predator-prey model for D and S,

$$\frac{\partial D}{\partial t} = Q_D - \frac{\sigma_1 S^2 + \sigma_2 SD}{\gamma' + a\Omega'} - \alpha_D D \tag{6.6}$$
$$\frac{\partial S}{\partial t} = Q_S + \frac{\sigma_2' DS + \sigma_1' D^2}{\gamma' + a\Omega'} - \beta S^2 - \alpha_S S, \tag{6.7}$$

where Q_D is the ohmic drive, α_D is the linear forcing, β represents the transfer of energy to unresolved modes, and σ_1, σ_1', and σ_2 are nonlinear coupling coefficients. The term in the denominator, $\gamma' + a\Omega'$, is the shear of the dominant mode. Because of the symmetries in Equations 6.4 and 6.5, the shear rate is actually determined as a composite of the magnetic and flow shears, assuming that one or the other dominates:

$$\Omega' = \frac{im}{r} \max \left[\left.\frac{\partial \phi}{\partial r}\right|_{n=6}, \left.\frac{\partial \psi}{\partial r}\right|_{n=6} \right]. \tag{6.8}$$

The shearing term can also be written in a more convenient form as

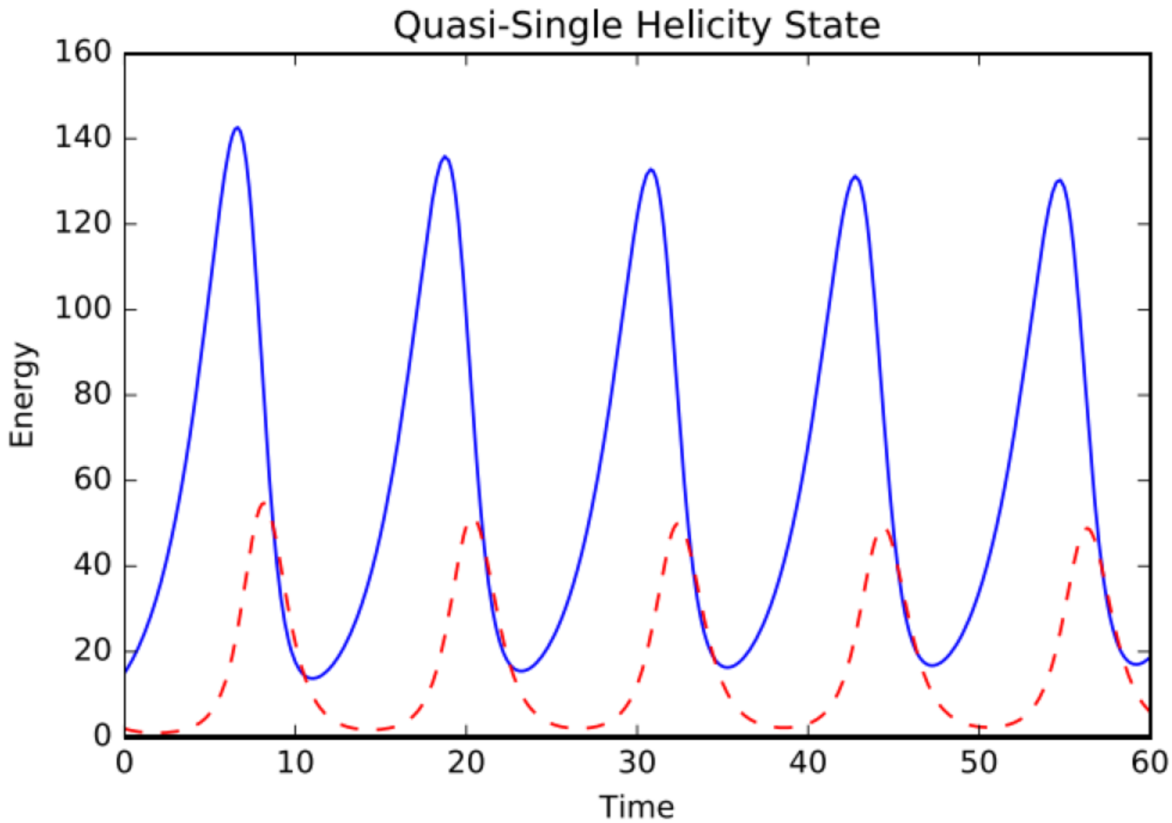

Figure 6.8: Evolution of the dominant mode D and secondary modes S during the limit cycle QSH regime. Figure reproduced from I. McKinney and P. Terry, [33].

$$\gamma' + a\Omega' = 1 + \epsilon\left(\frac{D}{D_0}\right)^{1/2}, \tag{6.9}$$

where ϵ is a normalized coefficient related to the plasma current, and D_0 is the dominant mode amplitude at the transition to QSH. It was also shown that $\varepsilon \sim B_\theta^{3/5}$, reproducing the desired scaling with plasma current. An example of the evolution of D and S which clearly demonstrates the predator-prey relationship during the limit cycle phase is shown in Figure 6.8.

Further work extended this model to account for its impact on the evolution of the electron temperature profile [33]. It was shown that the temperature profile evolves according to

$$\frac{\partial T_e}{\partial t} - \left\{ \chi_0 \left[1 + (\frac{ar}{r_0})^2 \right] + \frac{\chi_1 S}{1 + \epsilon [D f(r)]^{1/2}} \right\} \times \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial T_e}{\partial r} \right) = C(r) \tag{6.10}$$

where a is the minor radius, χ_0 is the equilibrium diffusivity, χ_1 is the flutter-induced diffusivity localized to the region outside of the dominant mode's rational surface,

$$f(r) = \frac{1}{\delta\sqrt{\pi}} \exp\left[- \left(\frac{r - r_s}{\Delta} \right)^2 \right] \tag{6.11}$$

is a Gaussian packet of width Δ which serves to localize the effect of shear suppression to a layer around the resonant surface r_s of the dominant mode, and

$$C(r) = 4\pi T_0 \chi_0 \left(1 + (\frac{ar}{r_0})^2 \right) \left(1 - (\frac{r}{r_0})^2 \right) \times \frac{1.05 r^2 - r_0^2}{(r^2 - r_0^2)^2} \tag{6.12}$$

is a steady-state heat source which was chosen phenomenologically to provide a realistic seed profile $T_e(r) = T_0(1 - r^2/r_0^2)^{1.05}$.

In this model, the quality of the thermal confinement is mostly determined by the amplitude of the secondary modes. As shown in Figure 6.9, as the mode amplitudes oscillate so does the height of the temperature profile. At its highest point the edge gradient can become quite steep, which can be interpreted as an edge transport barrier. A small phase shift is observed between ∇T_e and D, although its magnitude is somewhat dependent on the various model parameters. This type of behavior also exists in the saturated SHAx (high ϵ) state, though due to the small amplitude of the secondary modes the resulting oscillations are much less dramatic than in Figure 6.9.

As with any theoretical model, it is important to validate the shear suppression model for QSH against experimental data. Measurements of temperature profile dynamics in RFX-mod have shown evidence of large of temperature gradients which oscillate along with the dominant magnetic mode amplitude [12]. Recent measurements of flow velocity shear in the MST show that it may be large enough to suppress energy transfer with secondary modes, although the uncertainty in the calculation is large [28]. How-

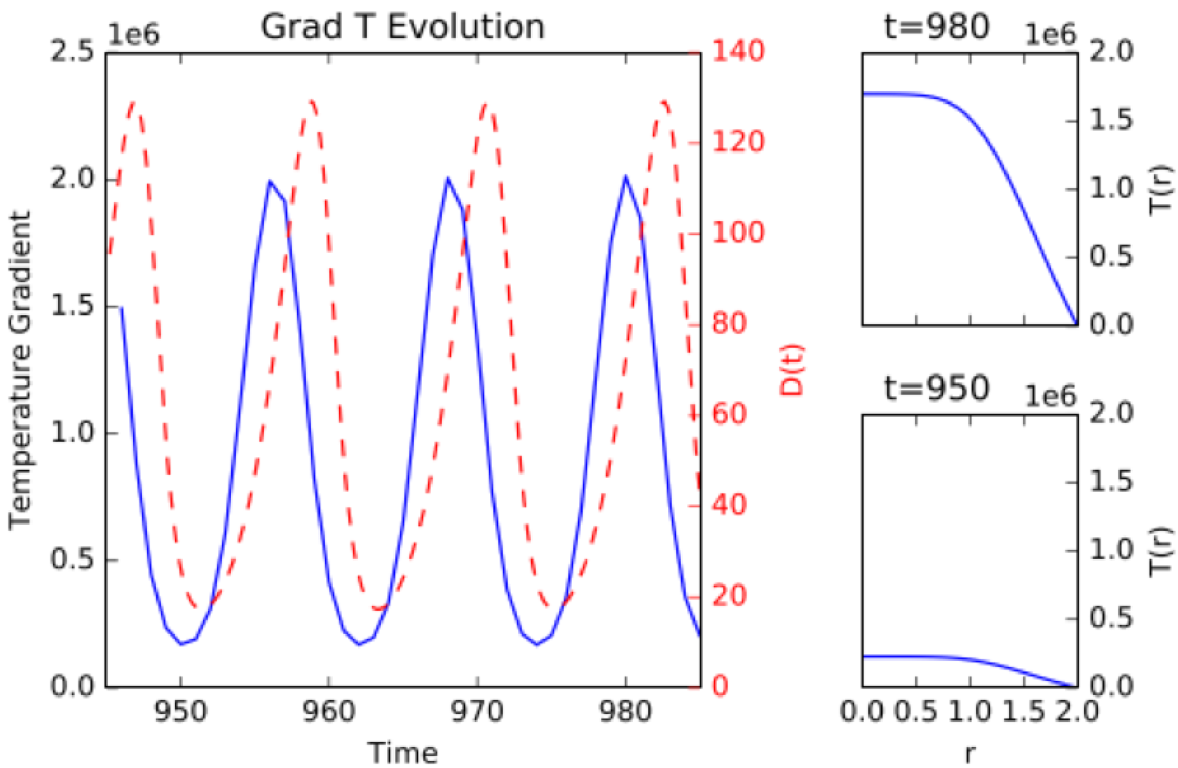

Figure 6.9: Evolution of the dominant mode D and the maximum temperature gradient. These signals are closely correlated, with ∇T_e leading D by a small phase shift due to finite heating. The small panels show the profile $T_e(r)$ for two time slices. Figure reproduced from I. McKinney and P. Terry, [33].

ever, there is still a need for additional validation. Section 6.2.2 provides two additional such observations, demonstrating a predator-prey-like relationship between the dominant and secondary modes and confirming that the link between thermal confinement and dominant mode amplitude extends to the MST.

Before moving on, it is worth taking a moment to draw out some connections between this model and other developments in plasma turbulence theory. Sheared flows have long been understood to produce transport barriers in fusion plasma devices [34]. The existence of these transport barriers is in-fact critical to the existence of modern high-performance fusion plasmas. Flow shears are known to play a key role in the transition from L- to H-mode in tokamaks [35], and predator-prey interactions between equilibrium $\boldsymbol{E} \times \boldsymbol{B}$ flow and zonal flows have been directly observed to precede the L-H transition in DIII-D [36, 37]. In many ways, the model outlined in the preceding paragraphs is a natural extension of these existing ideas to the RFP.

6.2.2 Observation of predator-prey dynamics during QSH

One of the fundamental consequences of shear suppression model for QSH is a predator-prey relationship between the dominant and secondary tearing modes. Although it can easily be seen that increases in the amplitude of the $n = 5$ mode corresponds with decreases in the secondary mode and vice-versa (as in Figure 6.6), no study had previously been conducted to test whether the relationship between fluctuations during the saturated state can be classified as a predator-prey system.

The main computational tool for this analysis is the cross-correlation, which is a measure of the similarity between two (semi-) periodic signals. For continuous signals, it is defined as

$$C(f, g) = \int f(t+\tau)\, \bar{g}(t)\, dt \tag{6.13}$$

where $\bar{g}$ is the complex conjugate of g and the offset τ measures the lag of $f(t)$ relative to $g(t)$. The phase shift between f and g is therefore given by $\Delta\phi = 2\pi f \tau$. In practice digitized signals are discrete, so the integral is replaced by a sum over the data points. This was implemented using the `numpy.correlate` Python module [38].

In order to make the various b_n signals corresponding to the measured tearing mode amplitudes easier to directly compare, the magnetics signals were normalized according to

$$y_n(t) = \frac{b_n(t) - \langle b_n \rangle}{\sigma_n} \tag{6.14}$$

where $\langle b_n \rangle$ is the average and σ_n is the standard deviation of the signal b_n. This has the added benefit of ensuring that the computed cross-correlation varies between approximately ± 0.3 no matter the size of the input signals.

We begin by considering three time windows, highlighted in Figure 6.10, and computing $C(y_6, y_5)$, $C(y_7, y_5)$, and $C(y_8, y_5)$ during each time window. The results are shown in Figures 6.11, 6.12, and 6.13, respectively. During the quiet phase we see that the secondary modes evolve together with the dominant mode, with the addition of large amounts of uncorrelated random noise. This is not the case during the two dynamic phase time windows. During the earlier window (28-32 ms), we see that all three secondary modes are strongly correlated to the dominant mode with a phase shift of approximately 180°. During the final window (34-38 ms) the situation is much the same except that the correlation between the $n = 5$ and $n = 7$ modes has been lost. This demonstrates that the transfer of energy may not always be evenly distributed between the various core-resonant tearing modes.

On the individual shot level these signals can be pretty noisy, so it is desirable to perform this analysis on an ensemble of similar time windows. This presented a challenge as QSH plasmas are highly variable. I searched through the data set and identified 15 short (3 ms) time windows which feature this ~ 1 kHz tearing activity on top of the

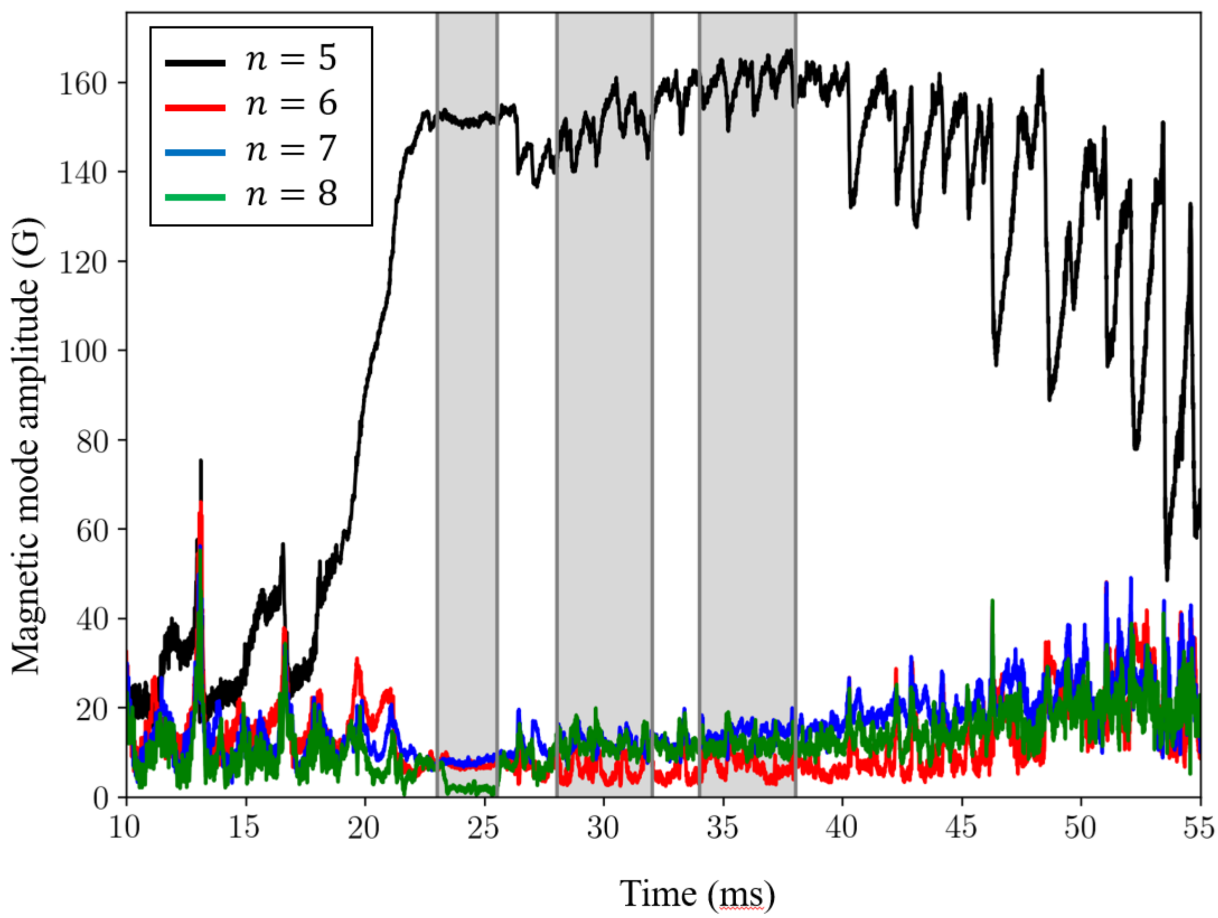

Figure 6.10: Three highlighted windows corresponding to the cross-correlation analyses shown in Figures 6.11, 6.12, and 6.13.

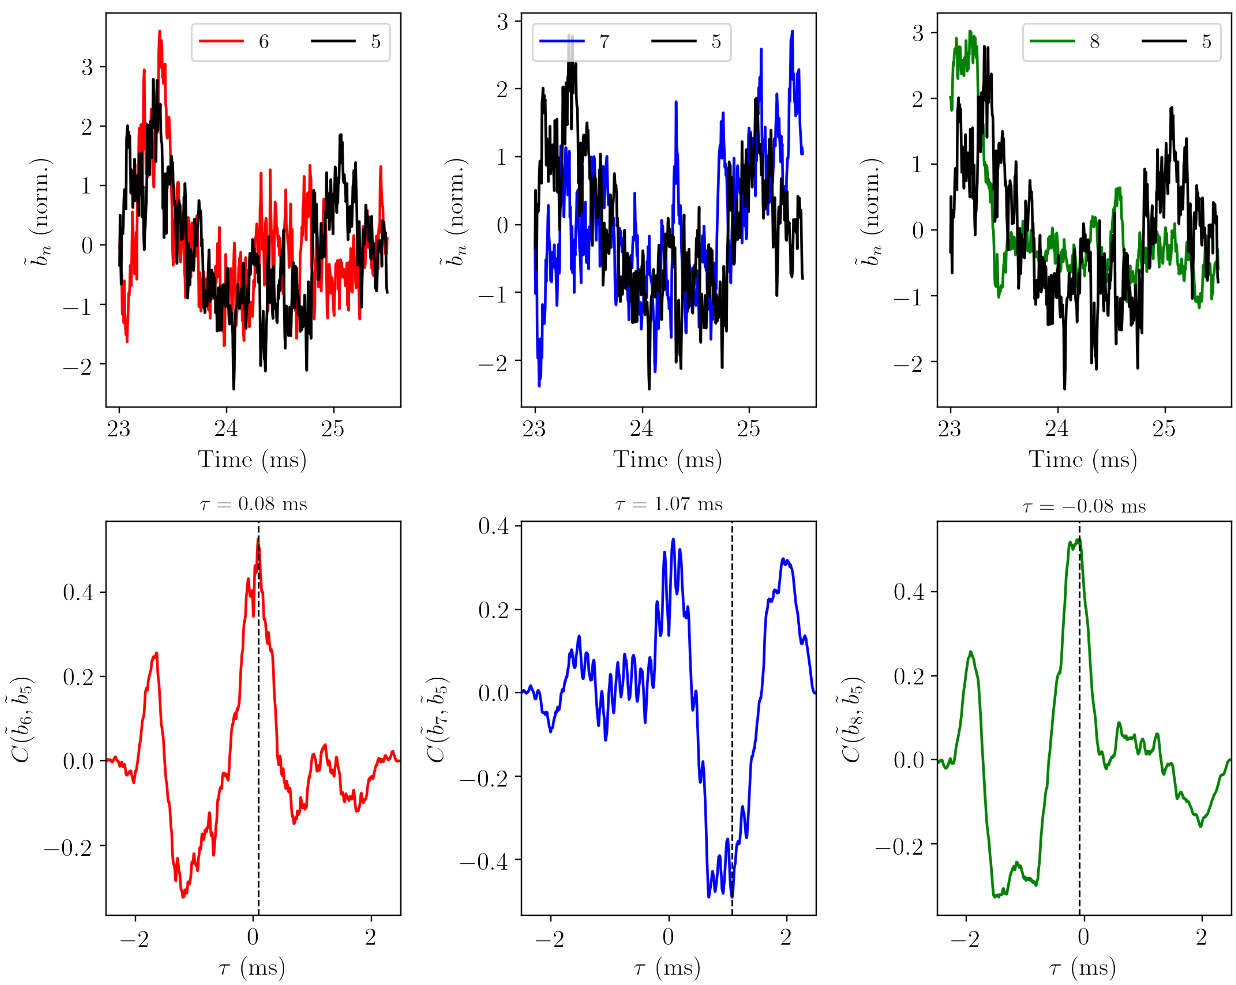

Figure 6.11: Cross correlations between the $n = 5$ and $n = 6, 7, 8$ tearing mode amplitudes during the quiet phase, 23.0 to 25.5 ms.

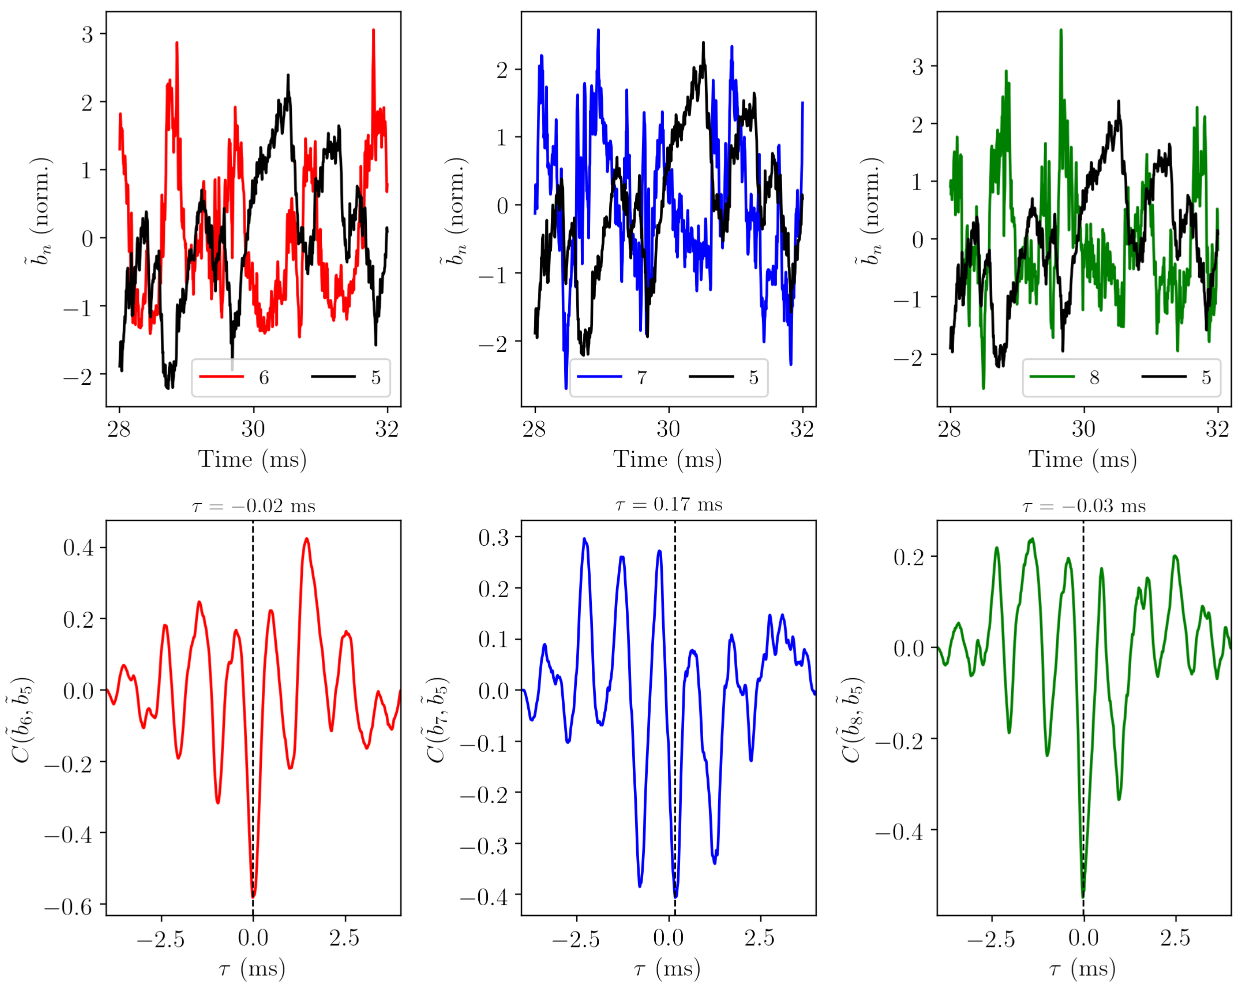

Figure 6.12: Cross correlations between the $n = 5$ and $n = 6, 7, 8$ tearing mode amplitudes during the early dynamic phase, 28.0 to 32.0 ms.

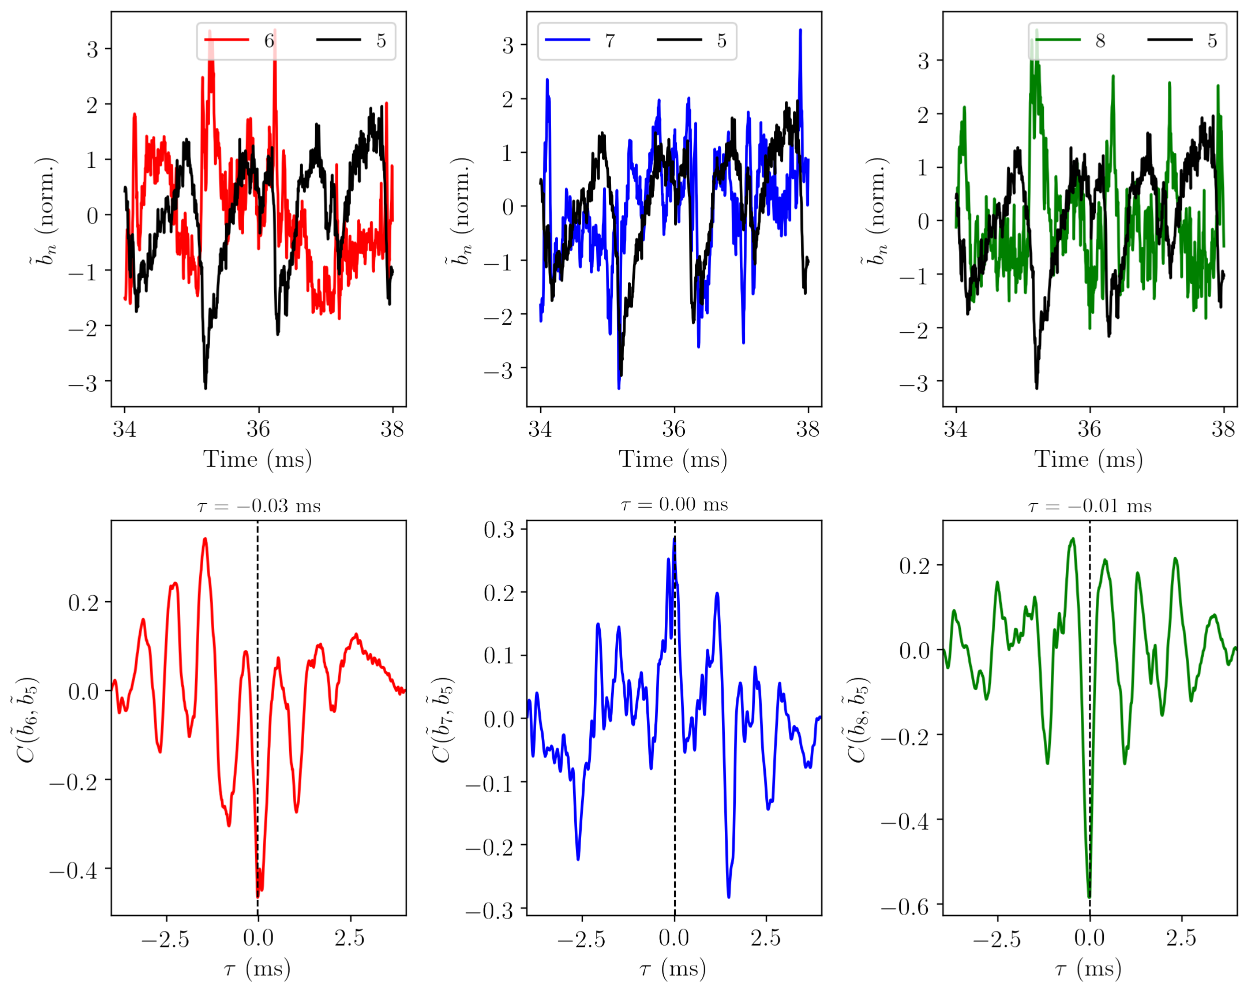

Figure 6.13: Cross correlations between the $n = 5$ and $n = 6, 7, 8$ tearing mode amplitudes during the late dynamic phase, 34.0 to 38.0 ms.

dominant mode during the dynamic flattop, like the later two windows highlighted in Figure 6.10. These perturbations tended to be most frequent in high-performance QSH plasmas. I carefully chose the starting points of these windows in order to align the $n = 5$ signals as best as possible. Using this ensemble I created $\bar{y}_5$, $\bar{y}_6$, etc., which are the average normalized tearing mode amplitudes. These signals, along with their cross correlations, are shown in Figure 6.14.

For the ensemble data set, we see that all three secondary modes are all closely correlated to the dominant mode, with a 180° phase shift. Whatever caused the discrepancy with the $n = 7$ mode in Figure 6.13 is not a consistent feature of these plasmas, suggesting that it may have been due to noise. Also notable is that there is that signals are all perfectly anti-correlated at $\tau = 0$ with no measurable lag (that is, they are exactly 180° out of phase). This means that the transfer of energy between the modes must be happening on a faster time scale than can be resolved by this analysis. The magnetics signals are digitized at 200 kHz, but the resolution of this analysis is likely much lower than that would suggest due to the variability in the ensembled signals (as represented by the semi-transparent bands in Figure 6.14). Regardless, these measurements present clear evidence of a predator-prey relationship between the dominant and secondary tearing modes during the dynamic flattop phase of QSH that is consistent with the shear-suppression model.

We also want to show a relationship between oscillations in the tearing mode amplitude and changes in thermal confinement. Although in Section 6.4.3 we will extract measurements of $T_e(\rho)$ and ∇T_e from the data, those will turn out to have enough uncertainty to make if difficult to extract a clear phase relationship with y_5. These measurements will also be difficult to ensemble, since the technique requires good flux surface reconstructions be available for every plasma in the data set. Instead we will rely directly on soft x-ray emissivity, obtained by inverting SXR tomography measurements, as a proxy.

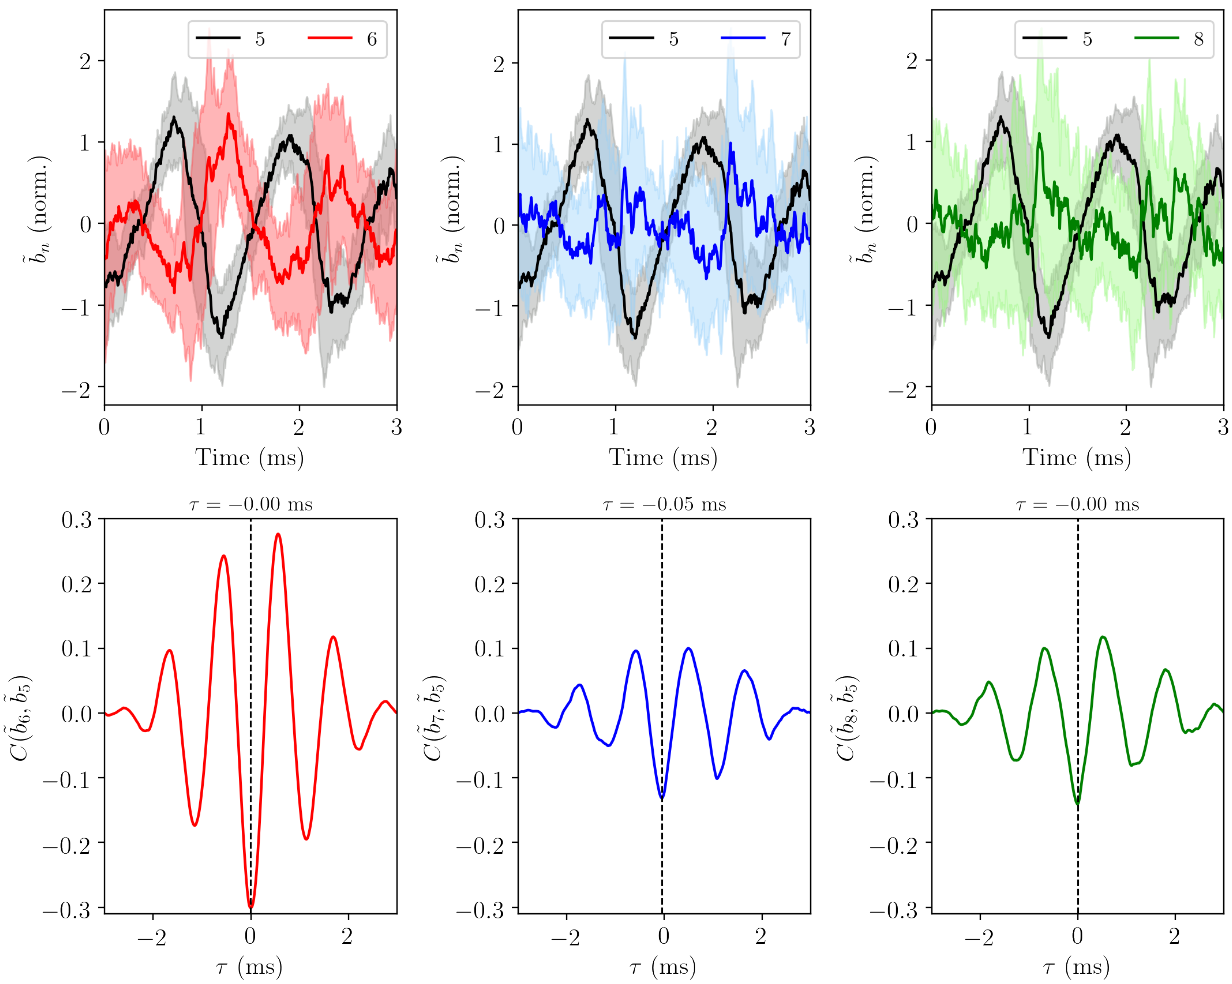

Figure 6.14: Cross correlations between the $n = 5$ and $n = 6,7,8$ tearing mode amplitudes for an ensemble of 3 ms time windows for plasmas which feature regular ~ 1 kHz oscillations during the QSH dynamic flattop phase. Semi-transparent bands represent ensemble variation.

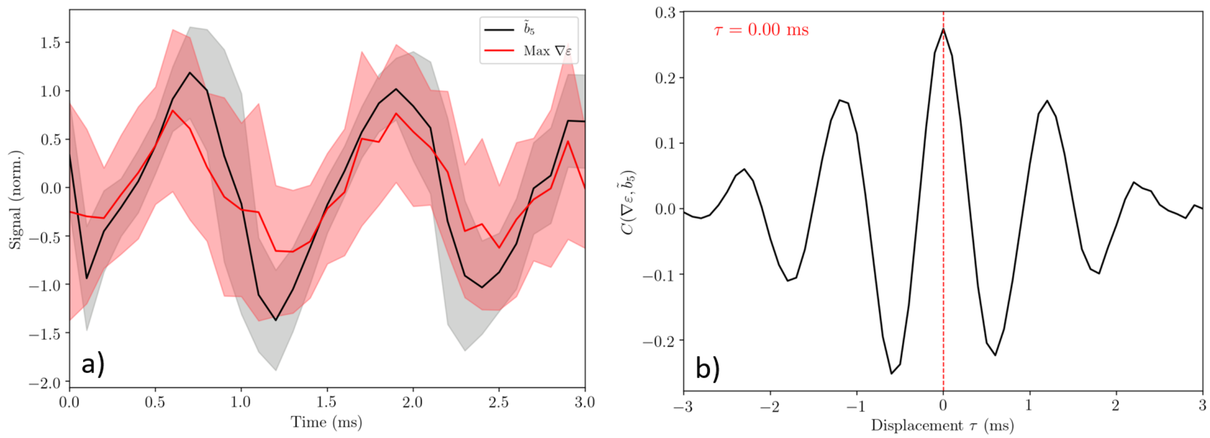

Figure 6.15: SXR emissivity is found to vary in-phase with the dominant mode. Panels show: (a) the ensembled normalized signals for the $n = 5$ tearing mode amplitude and the slope of the emissivity; (b) cross-correlation between the two signals.

As shown in Equation 2.19, variations in emissivity are likely due to variations in one of three parameters: T_e, n_e and n_i. We have seen that electron density tends to be fairly constant during QSH discharges (Figure 6.5), so this is unlikely to the the source of observed SXR variations. Although we have few measurements of ion densities during QSH, their confinement is expected to be reasonably similar to the electrons. This means that it is justifiable to interpret fast variations in emissivity (fluctuations at ~ 1 kHz) as being primarily due to changes in T_e.

For each time window in the ensemble, data from the SXR tomography 45 μm channel was re-sampled to 100 kHz. Tomographic inversions were performed for each time point to extract the emissivity $\varepsilon(\boldsymbol{x})$, like that shown in Figure 2.11. The absolute value of the gradient, $|\nabla\varepsilon|$, was computed for each time point, and the maximum value was identified. The resulting signal was normalized in the same way as the magnetic signals were,

$$y_\varepsilon(t) = \frac{|\nabla\varepsilon|_{max}(t) - \langle|\nabla\varepsilon|\rangle}{\sigma_\varepsilon} \tag{6.15}$$

and assembled into an ensemble signal $\bar{y}_\varepsilon$.

The cross correlation between emissivity and the dominant mode, $C(y_\varepsilon, y_5)$, is shown in Figure 6.15. Although the uncertainty in y_ε is high, it is clear that the b_5 and $\nabla\varepsilon$ change in phase with one another. When the dominant mode is at its highest amplitude, and correspondingly the secondary modes are at their lowest, the plasma emissivity is at its peak. Conversely, when the dominant mode is at its lowest, the emissivity is too. As with the magnetics signals, any phase shift is below the resolution of this analysis. Since we have argued that on these short time scales the behavior of the emissivity is a good proxy for the behavior of the electron temperature, we can conclude that the model for temperature evolution put forth in Equation 6.10 is consistent with observations on the MST.

6.3 Runaway electron generation in the helical RFP

The presence of runaway electrons in QSH plasmas on the MST was briefly studied in a 2010 paper by Clayton *et al.* [39]. Using the 13-chord CdZnTe HXR array, the study found that a high flux of hard x-rays was associated with the presence of a hot island in the plasma core as measured by the SXR tomography diagnostic. It was also found that the HXR flux is relatively uniform throughout the torus. This is explained as resulting from a small core region with significantly reduced transport in which the electrons accelerate before scattering into the highly stochastic surrounding plasma volume. This picture was supported by some basic ORBIT simulations, which showed that the most energetic electrons tended to be found in the $n = 5$ island.

In this section we extend upon these previous results using the ME-SXR diagnostic in conjunction with the FXR camera (Section 2.1.4). The ME-SXR detector has been set into an eight-color high SNR configuration (Config. #1 in Section 3.5) using the midE (medium gain) calibration which allows for thresholds between 4-to-14 keV. The thresholds were set to $E_c = 5, 6, 7, 8, 9, 10, 11$, and 12 keV. This configuration allows for some sensitivity to thermal electrons on the low end, but also extends into the HXR range. For non-PPCD plasmas with $T_e < 1$ keV, the signal from thermal electrons above 5-6 keV is expected to be negligible. The FXR detector was attached to a core-viewing radial channel and is calibrated for detection of photons between 5 and 30 keV. The combination of the two detectors allows for a synthesis of high spatial (ME-SXR) and spectral (FXR) resolution.

6.3.1 Observations of runaway electron generation and confinement

ME-SXR, FXR, and SXR tomography data for MST shot #1200701056 (500 kA, $F = 0$) are shown in Figure 6.16. As found in the Clayton study, HXR flux is typically stronger when a hot island is visible in the SXR tomography data. However, this observation can be expanded to fit the full picture of QSH evolution. Peak HXR flux is achieved during the rise phase, when the applied loop voltage (represented by the poloidal gap voltage, V_{pg}, in the figure) is still high. The HXR flux "calms down" but remains relatively high during the quiet phase, even though V_{pg} is now minimal. Then, at the transition to the dynamic phase, there is another burst of HXR flux. This coincides with the loss of the hot SXR island. After that point the HXR flux does not totally vanish, but remains at a much lower level than it was previously at. The highly non-thermal behavior is most evident in the ME-SXR data, where at times pixels with $E_c = 5$ keV and $E_c = 12$ keV are recording nearly identical counts. This implies the presence of very high-energy photons.

I would like to draw some extra attention to the burst in the ME-SXR data which occurs around $t = 26.5$ ms at the end of the quiet phase. This feature is a very robust predictor of a "good" quiet phase in which a very bright island can be seen in the

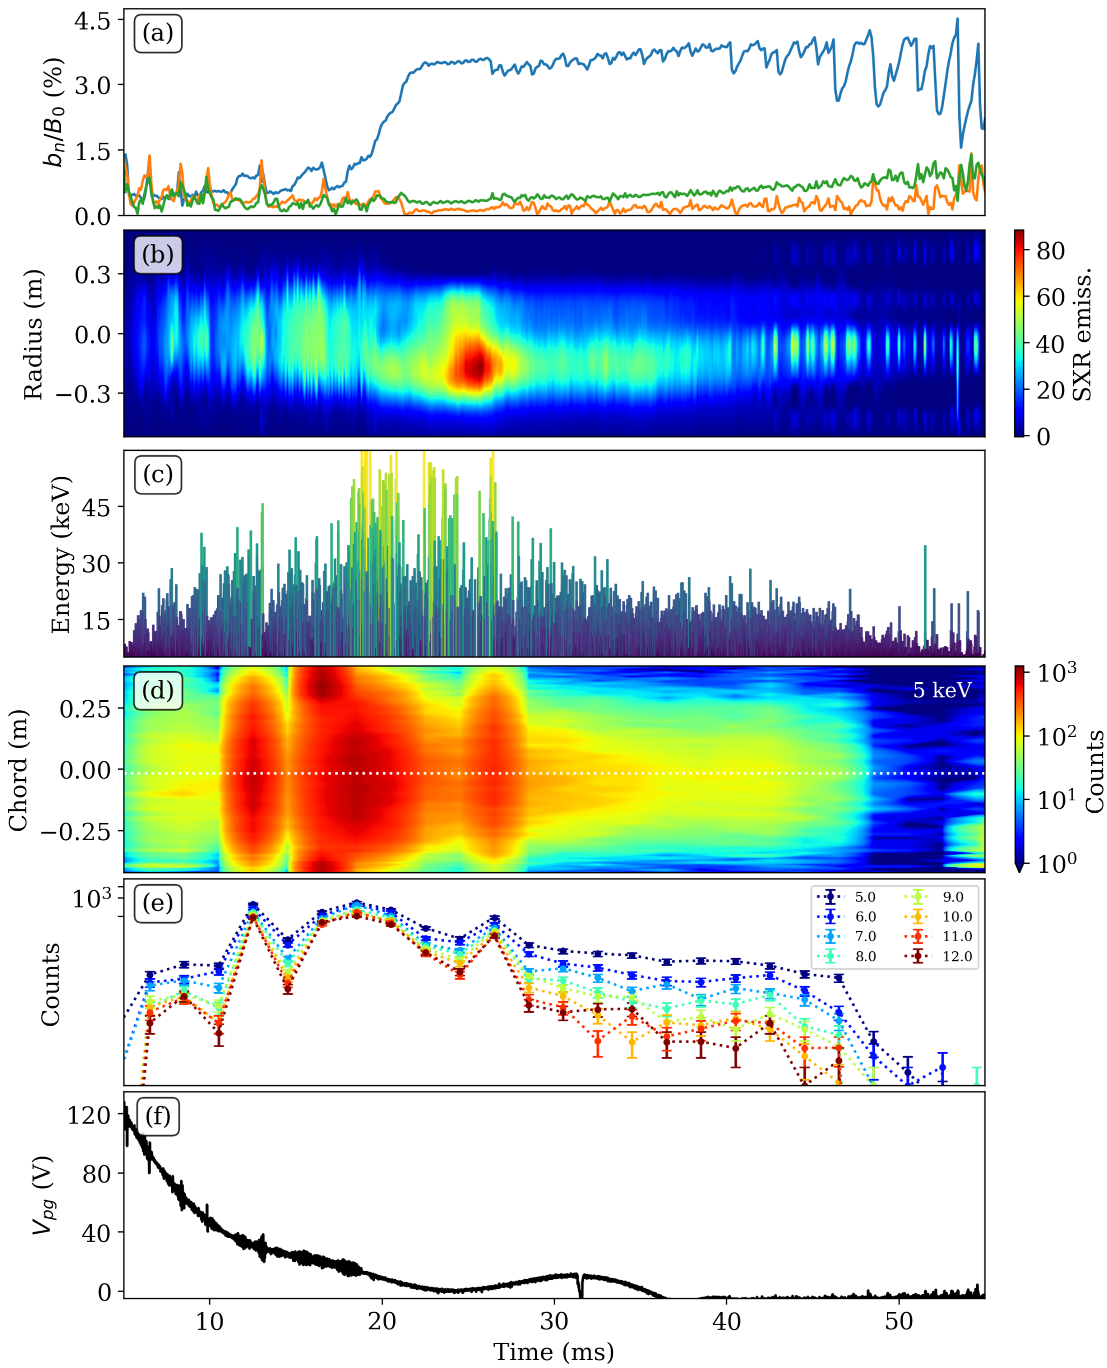

Figure 6.16: Evolution of (a) the dominant and secondary mode amplitudes, (b) the core SXR emissivity, (c) HXR flux, (d) ME-SXR profile with $E_c = 5$ keV with a logarithmic color scale, (e) single-chord ME-SXR counts along the dotted line in (d), and (f) the poloidal gap voltage.

SXR tomography data. I initially thought of it as an annoyance, likely resulting from poor confinement; it was only later that I realized that this burst was the signature of unusually good confinement. It was actually this feature which initially brought my attention to the distinction between the quiet and dynamic flattop.

As seen in Figure 6.7, there appears to be a rapid reduction in the extent of the good flux surfaces around the helical core when the quiet phase ends and secondary mode activity reactivates. The Clayton ORBIT simulation shows that fast electrons are well-confined to this region, so a rapid reduction in the volume of that region means that a substantial quantity of fast electrons suddenly find themselves in a region dominated by rapid stochastic transport. These electrons are then rapidly transported to the wall where they collide with various components, giving off target emissions which are seen by the ME-SXR detector. This provides pretty strong corroborating evidence for the general picture put forth in the Clayton paper. This burst is largely invisible to the SXR tomography diagnostic, as the 35 μm Si photodiodes are relatively transparent to photons in the low-HXR range.

Figure 6.17 (a) shows the ME-SXR signal vs threshold for several different spatial chords for $t = 28.5$ ms. The counts for all thresholds have been normalized such that $N = 1$ for $E_c = 5$ keV. This can be compared against a model spectrum, $\varepsilon \propto E^{-\gamma}$, by using the model response function developed in Section 4.3:

$$N_m(\gamma; E_c) \propto \int_0^\infty \frac{R(E; E_c)\, E^{-\gamma}}{E} dE, \tag{6.16}$$

where γ is a parameter indicating the slope of the HXR spectrum, $R(E; E_c)$ is the composite response function of Equation 4.23 with $\sigma_E = 550$ eV for midE settings, and the subscript m indicates that this is a model. Note that charge-sharing effects were included in this calculation. The resulting model output for a range of γ values is shown in Figure 6.7 (b). It should be noted that the $E^{-\gamma}$ shape of the model spectrum is heuristic and was inspired by its use in a paper by A. M. Dubois studying sawtooth energization [40].

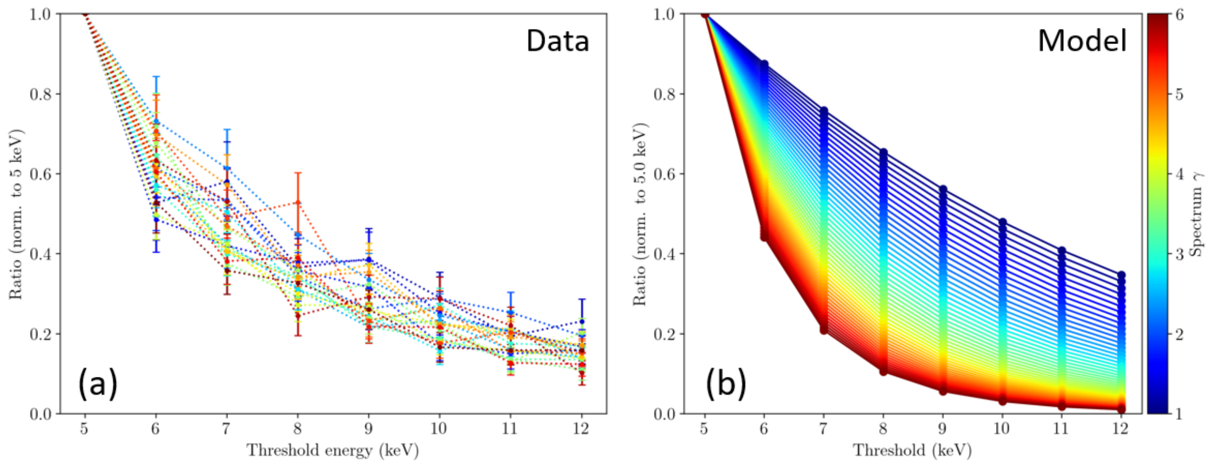

Figure 6.17: (a) ME-SXR measured spectrum for a set of pixels sharing a single line-of-sight, vs energy. The spectrum has been normalized to the counts measured with $E_c = 5$ keV. Different colors correspond to different lines-of-sight. (b) Model ME-SXR spectrum over a range of γ values.

This model can be used to infer the value of γ which best describes the data. This was accomplished using a simple Bayesian calculation (see Appendix B) with the likelihood function for each set of 8 thresholds (which approximately share a line-of-sight) given by

$$\ln \mathcal{L}(\gamma) = -\frac{1}{2} \sum_{E_c=6}^{12} \left(\frac{R_{E_c} - f(\gamma; E_c)}{\sigma_R} \right)^2 \tag{6.17}$$

where $R_{E_c} = N_{E_c}/N_5$ is the ratio of the measured counts for the threshold E_c to the counts measured for $E_c = 5$ keV, $f(\gamma; E_c) = N_m(\gamma; E_c)/N_m(\gamma; 5)$, and σ_R is uncertainty in the ratio given by

$$\frac{\sigma_R}{R} = \left[\left(\frac{\sigma_{E_c}}{N_{E_c}} \right)^2 + \left(\frac{\sigma_5}{N_5} \right)^2 \right]^{1/2} \tag{6.18}$$

where $\sigma_{E_c} = \sqrt{N_{E_c}}$ is the Poisson uncertainty on the measured data. It should be noted that Equation 6.18 is based upon a normal approximation to the Poisson distribution and can break down if the total counts for either the numerator or denominator are very low, roughly in the single digits. Priors were chosen to be uniform, $p(\gamma) \sim \mathcal{U}(-2,7)$, where some negative values were permitted in order to easily identify times when the $E^{-\gamma}$ model is inapplicable.

Posterior distributions were found to be approximately normal, and therefore well-characterized by the mean and standard deviation. These values were obtained for all chords with sufficient signal at each time point during the QSH state. Figure 6.18 shows the resulting evolution of the γ profile over time. We see that the spectrum exhibits a strong high-energy tail early on, before undergoing a rapid transition after ~ 26 ms and settling into a calmer equilibrium. This is generally consistent with the trends observed in the FXR data. Note that the plots presented throughout the rest of this section tend to show $-\gamma$ so that a higher value corresponds to a greater high-energy population.

Notably, the spectrum is almost completely spatially homogeneous, and does not feature a peak at the $n = 5$ core. This suggests that, outside of the helical core, the fast electron transport is quite rapid so that they reach the edge without having lost much kinetic energy to collisions. As a result the normalized HXR spectrum (though not necessarily the absolute intensity) is the same across the MST radius.

The spatial uniformity presents an opportunity to refine our estimate by averaging the chords at each time point. The result is shown in Figure 6.19. This makes clear the trend that is being observed. The electron distribution is highly non-Maxwellian early during the early phase sawtooth cycle and into the rise phase. Many of these data points feature $\gamma < 0$, implying that the $E^{-\gamma}$ model does not do a good job describing these data points. However it is still sufficient to infer that during this period the plasma is strongly non-Maxwellian. Once the quiet flattop phase begins, we find that γ begins to increase, implying that the energy of the fast electrons are being gradually lost. This is interrupted at a single time point by the burst at the end of the quiet phase, but then continues to drop off into the dynamic phase. The value of γ eventually stabilizes late in

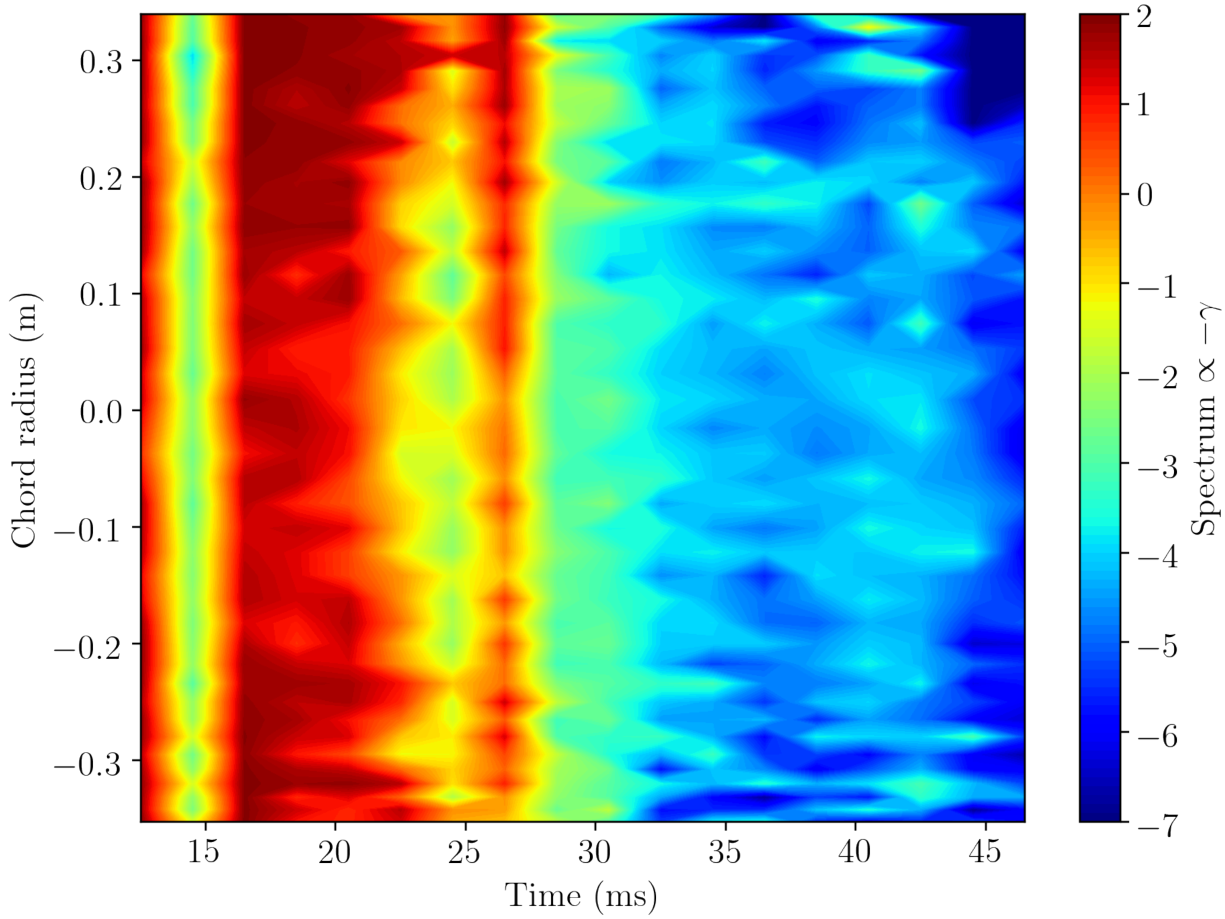

Figure 6.18: Time evolution of inferred $E^{-\gamma}$ spectrum evolution for all ME-SXR chords with sufficiently high signal. The spatial structure is almost entirely uniform.

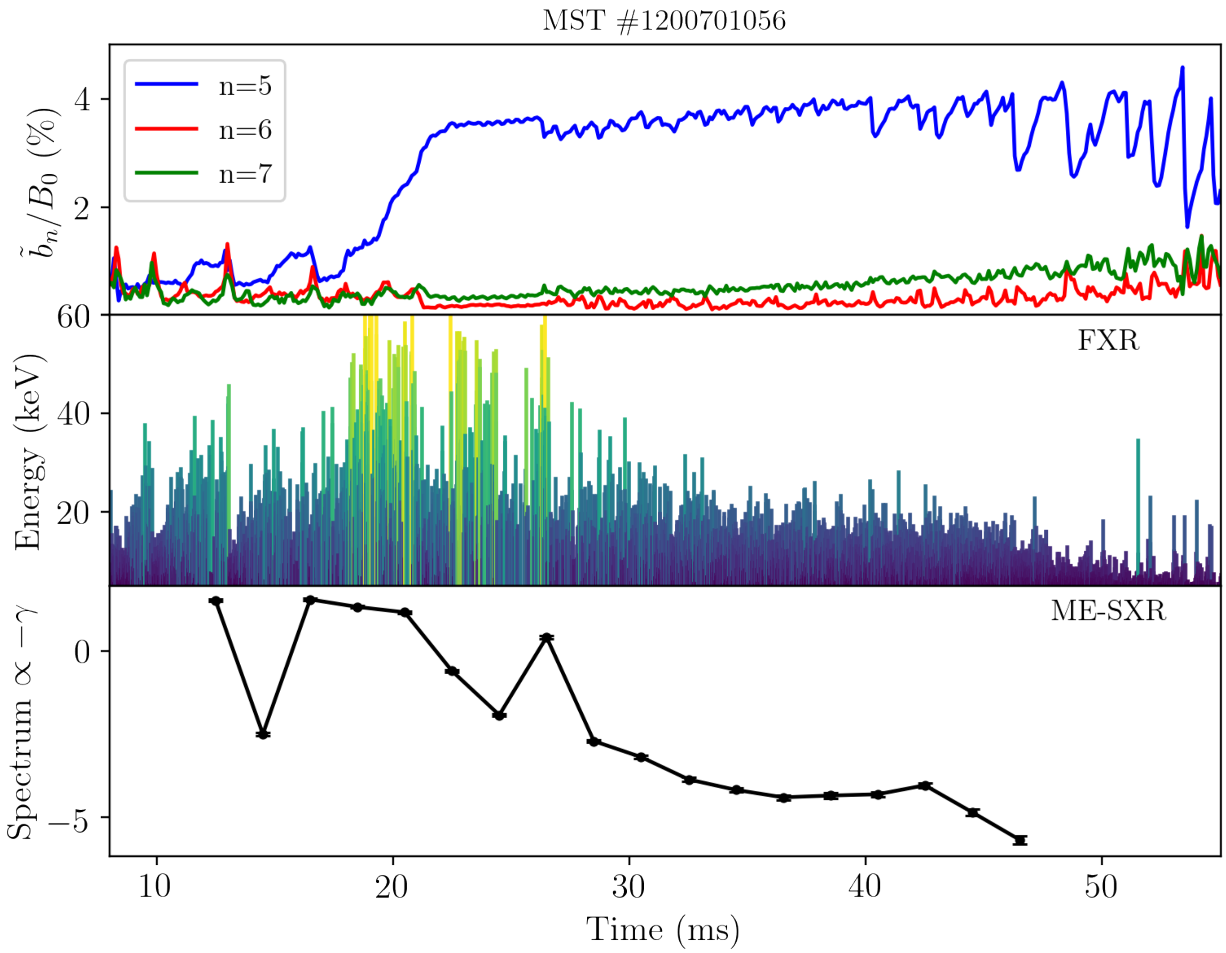

Figure 6.19: Evolution of the HXR spectrum, showing (top) the dominant and secondary mode amplitudes, (middle) FXR photon events, and (bottom) the radially-averaged spectrum coefficient, $-\langle\gamma\rangle$.

the quite phase until declining plasma current causes the QSH equilibrium to dissipate. This behavior is typical of all high-quality QSH discharges in the dataset.

6.3.2 Energy of the runaway population

In order to expand upon these observations, we want estimate the average energy of the fast electron population in each of the three phases (rise, quiet, and dynamic) of the QSH flattop. The relatively slow time resolution of the ME-SXR diagnostic makes it unsuitable for this task, so the FXR camera will instead be used. The FXR detects individual events with an accuracy of tens of nanoseconds, so these counts can easily be checked against the corresponding magnetics signals to build up and ensemble for each phase. This was done for an ensemble of 12 similar high-performance QSH discharges. The counts were then sorted into energy bins in order to estimate the spectrum. The resulting spectra for each phase are shown in Figure 6.20.

The photon energy was estimated by assuming a "Maxwellian-like" scaling for the high-energy tail, given by

$$\varepsilon(E) = C \, \exp\left(-E/E_r\right), \tag{6.19}$$

where C is treated as an arbitrary scale factor and E_r is the runaway electron population average energy. This function was fit directly to the logarithm of the HXR data using a linear least-squares algorithm, and the results also shown in Figure 6.20. At lower energies the thermal population contributes the to x-ray flux, so the fits were only performed for data to the right side of the dashed line in each plot.

We found that during the rise phase, the fast population has an average energy of $E_r = 18.2 \pm 0.7$ keV. During the quiet phase, this reduces slightly to $E_r = 15.3 \pm 1.2$ keV before dropping of substantially to $E_r = 5.0 \pm 0.1$ keV during the flattop. These results support our ongoing narrative that the most energetic fast electrons are accelerated dur-

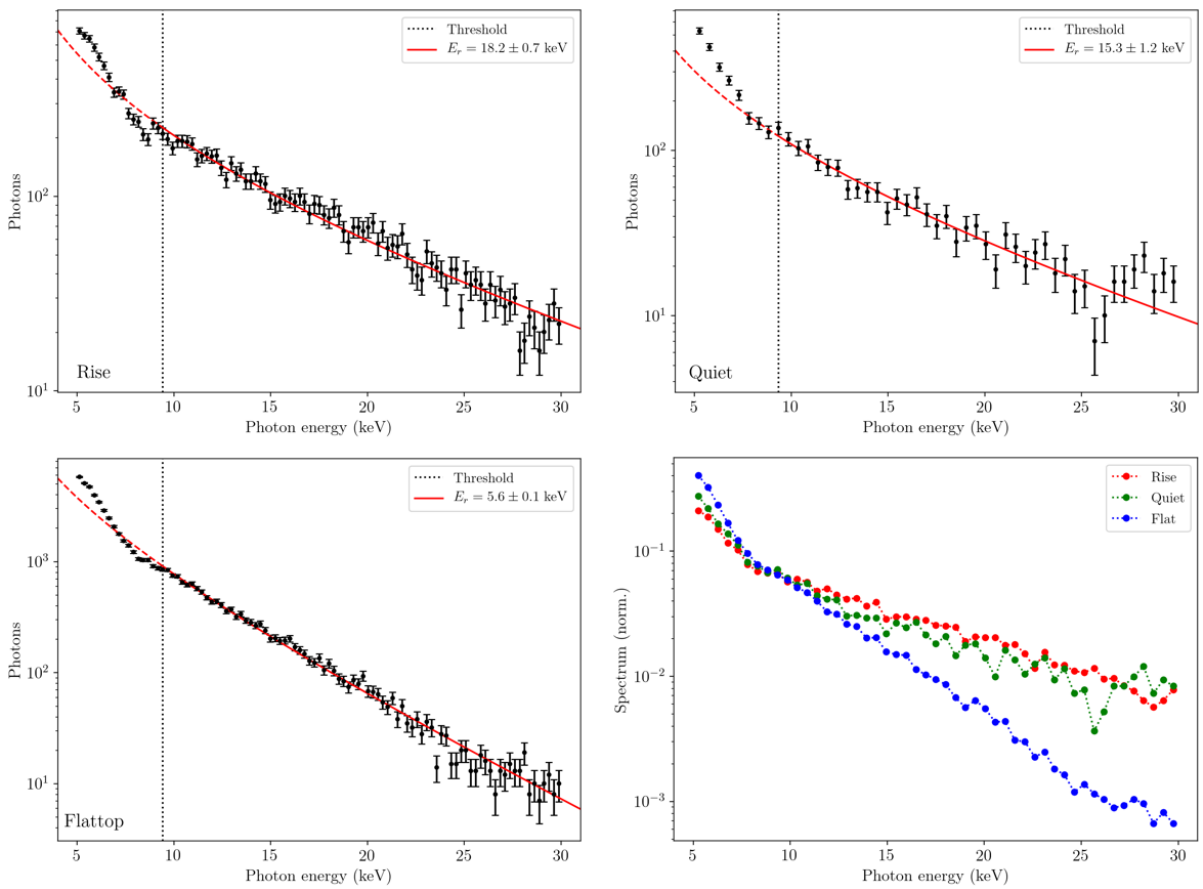

Figure 6.20: The hard x-ray spectrum for an ensemble of plasmas during the rise, quiet, and dynamic phases. The bottom right panel shows all three (normalized) spectra together on the same plot. It can be seen that the spectra for the rise and quiet phases are very similar, while the dynamic flattop spectrum is significantly reduced.

ing the rise phase and largely maintained during the quiet phase until confinement is severally degraded by the increased secondary mode activity that characterizes the dynamic flattop phase. Note that the $\pm$ ranges cited here are uncertainties in the mean, not the range of variation across the dataset.

We can use the average fast electron energy observed during the rise phase to estimate a rough approximation for the magnitude of the net electric field that would be required to generate this runaway population. We will assume that electrons are well confined within the core so that collisional momentum loss (rather than diffusion) sets the upper limit on the speed of the runaway population. The maximum speed is then given by the balance between the electromagnetic acceleration and the collisional momentum loss.

A simple model of the slowing-down force on a single particle in a plasma due to collisions [41] is given by

$$\begin{aligned} F &= \frac{d}{dt}(mv) \\ &= -\nu_e(v) m v \end{aligned} \tag{6.20}$$

where m is the particle mass, v is the particle velocity, and $\nu_e(v)$ is the velocity-dependent collisional frequency given by

$$\nu_e(v) = \frac{4\pi n_i Z_i^2 e^4}{(4\pi\epsilon_0)^2 m_e^2 v^3} \ln\Lambda, \tag{6.21}$$

where Z_i is that charge number for the background species and $\ln\Lambda$ is the Coulomb logarithm, which is approximately 15.6 for typical $F = 0$ parameters. When the background contains multiple species, the collisional effects are summed together, meaning that the $n_i Z_i^2$ terms can be replaced with $n_e Z_{\text{eff}}$ via quasi-neutrality.

Assuming that other loss terms are not important (a big assumption), the equilibrium speed is determined by the balance between the force from the net parallel electric field and the collisional losses,

$$eE_{\parallel} - \nu_e m_e v = 0, \tag{6.22}$$

which can be used to get an estimate of the magnitude of $E_{\parallel}$:

$$E_{\parallel} = \frac{1}{(4\pi\epsilon_0)^2} \frac{2\pi n_e (2 + Z_{\text{eff}}) e^3}{\left(\frac{1}{2} m_e v^2\right)} \ln \Lambda \tag{6.23}$$

For the case of the rising phase where $m_e v^2/2 \approx 18$ keV, and we will assume $Z_{\text{eff}} \approx 2$ and $n_e \approx 0.8 \times 10^{19}$ m^{-3}. The extra factor of 2 in $(2 + Z_{\text{eff}})$ accounts for electron-electron collisions. Plugging these in Equation 6.23, we get an estimate of $E_{\parallel} \approx 0.8$ V/m.

To provide some context for this estimate, we turn to a 2004 study by Piovesan, *et al.* [24]. In the study, careful correlation between magnetic field fluctuations and measurements in the IDS spectrometer all an estimation of the MHD dynamo, $E_{\text{dyn}} = \langle \tilde{v}^{(1,n)} \times \tilde{b}^{(1,n)} \rangle \approx 1 - 2$ V/m in QSH conditions. This dynamo field is relatively small but constant, unlike for multi-helicity discharges which tend to briefly feature very large dynamo fields during a sawtooth crash [42]. The dynamo field was observed to oppose the applied field of approximately $E_{\text{loop}} \approx 2$ V/m. The difference between E_{dyn} and E_{loop} is sufficient to explain the observed runaway population. This is consistent with the picture that the buildup of a runaway population is the result of reduced radial transport in the core in the presence of an applied loop voltage opposed by a constant MHD dynamo.

6.4 QSH profile evolution

This section will talk about building an analysis model to tackle the questions posed in the previous section. Although the ME-SXR diagnostic will be essential, TS, FIR, and SXT will also play a role. Although well-resolved density reconstructions of SHAx states in the MST have been produced in the past using V3Fit [11], temperature reconstructions have always been limited by the use of Thomson scattering data which tends to perform poorly at low densities (see Figure 6.5). Instead, we will rely on measurements of soft x-ray emission to simultaneously constrain the electron temperature and the aluminum density by employing a novel IDA methodology. This will result in a much higher-fidelity reconstruction of the thermal evolution of the plasma than has been achieved previously. This result will allow us to observe how the structure of the temperature profile evolves as secondary modes are suppressed during the quiet phase.

Section 6.4.1 begins by presenting the analysis framework which integrates FIR and SXR tomography data to infer the best-fit for parameterized profiles. Section 6.4.2 examines the quality of fits and develops the methodology for calculating the gradient. Finally, Section 6.4.3 presents the time-resolved profile evolution and examines how changes in the magnetic mode dynamics affect thermal confinement. Comparisons will be made to observations from RFX-mod, and the implications of these results on the shear-suppression model of QSH will be considered.

6.4.1 Building the analysis framework

This analysis was conducted primarily using data from the soft x-ray tomography (Section 2.3.2) and FIR interferometry (2.1.3) diagnostics. The magnetic equilibrium is expected to exhibit an $n = 5$ helical distortion, meaning that the geometry of the magnetic equilibrium is dependent on the toroidal angle. The SHEq code (Section 2.2.2) was used to to provide a consistent mapping between diagnostics. Figure 6.21 illustrates this geometry and shows that the helical O-point should be well-resolved for both diagnostics. Data was re-sampled to a resolution of $\Delta t = 0.5$ ms in order to focus on equilibrium evolution.

Parameterized profiles were chosen in order to obtain good agreement between the model and measured data while still providing enough constraints to avoid over-fitting

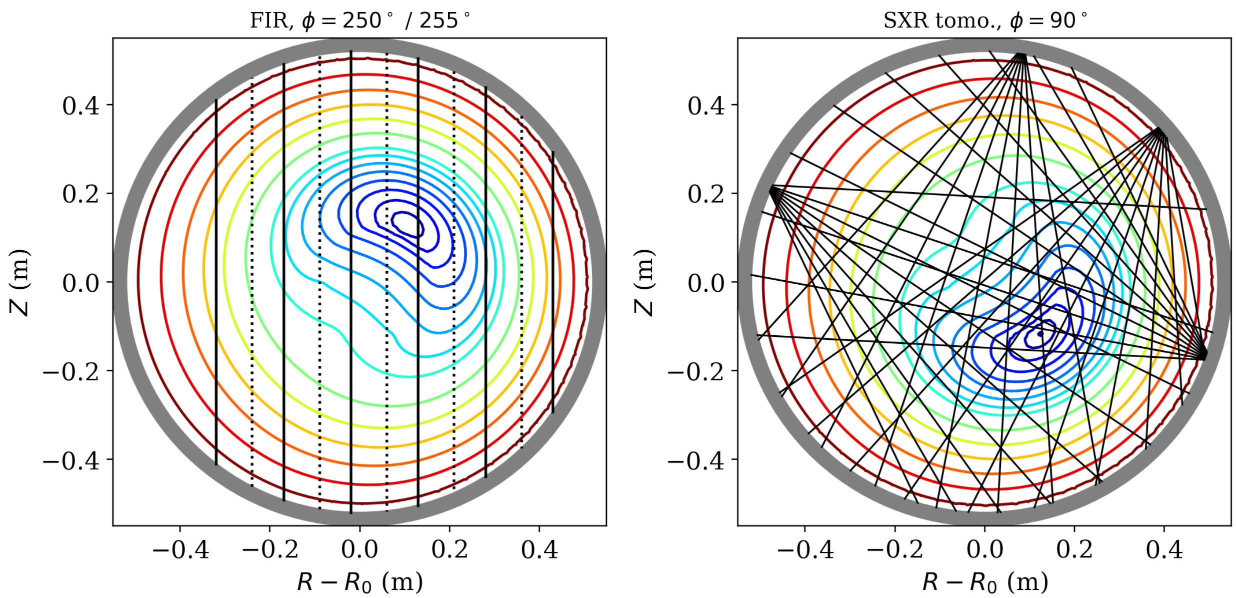

Figure 6.21: Viewing chords for the FIR interferometry (left) and SXR tomography (right) diagnostics. Flux surface contours taken from the SHEq reconstruction are also shown.

and ensure that results are physically plausible. Profile selection can be the most difficult part of the IDA process, and may require many stages of trial and error. The model electron density profile, which was inspired by a similar model used in Auriemma, *et al.* [43], is given by

$$n_e(\rho) = (n_{e,0} - n_{e,1}) \cdot (1 - \rho^{\alpha_n})^{\beta_n} + n_1(1 - \rho^{\gamma}) \tag{6.24}$$

where $n_{e,0}$ is the density at the helical axis, $n_{e,1}$ is the density at the edge, α_n and β_n are shape parameters, and γ controls the shape of the background density at the edge. In practice FIR measurements are minimally sensitive to γ, so a constant value of $\gamma = 100$ is assumed.

Selection of the temperature profile was more difficult. The typical $\alpha - \beta$ style models used in Section 5.5 tended to produce reconstructions which deviated significantly from SXR tomography data. Through trial and error a different parameterization was found which produced better agreement. This model is given by

$$T_e(\rho) = T_{e,0} \exp\left[-\left(\frac{\rho}{w_\rho}\right)^{\beta_T} \right] \tag{6.25}$$

where $T_{e,0}$ is the temperature at the helical axis, w_ρ controls the width of the core structure, and β_T is a shape parameter related to the steepness of the gradient. This parameterization was found to be flexible enough to represent both broad and narrowly-peaked thermal structures.

Very little prior information was available about impurity ion density profiles in the helical RFP. In order to avoid imposing unphysical structure, it was decided to use flat profiles,

$$n_Z(\rho) = n_{Z,0}(1 - \rho^{\alpha_Z})^{\beta_Z} \tag{6.26}$$

where $n_{Z,0}$ is the core impurity density and $\alpha_Z = 12$ and $\beta_Z = 4$ are shape parameters chosen so that the density remains flat until near the edge. As in Section 5.5, the densities of the other low-Z impurities were scaled based on the carbon density. It was found that leaving both n_C and n_{Al} as free parameters tended to yield unphysical results, so a representative value of $n_{C,0} = 6.8 \times 10^{16}$ m^{-3} was chosen based on prior RFP measurements [44]. Because the effect of low-Z impurities on SXR emission is less significant than aluminum, this is an acceptable compromise. The neutral density profile, $n_0(\rho)$, was implemented in the same way as described in Section 5.5.

The SXR tomography diagnostic has some sensitivity to background emissions from the non-thermal electron population, although it is far less sensitive to this than the ME-SXR. In order to account for this offset, two additional background emissivity terms have been included in the analysis, ε_1 and ε_2, which accounts for the offset in the 45 μm and 172 μm data respectively. Altogether, this model four free parameters for the FIR model, $\boldsymbol{\theta}_{FIR} = (n_{e,0}, n_{e,1}, \alpha_n, \beta_n)$, and six free parameters for the SXR model, $\boldsymbol{\theta}_{SXR} = (T_{e,0}, \beta_T, w_\rho, n_{Al,0}, \varepsilon_1, \varepsilon_2)$.

The soft x-ray emissivity $\varepsilon_Z^{(t)}(T_e, n_e, n_0)$ was modeled using a well-validated software model developed primarily by my colleagues P. Franz and L. Reusch [45], and the geometry of the diagnostic is was modeled using the tools developed for the ME-SXR model in Chapter 4. For each line of sight, the model solves the integral

$$f_{\mathrm{SXT}}(p_i, \theta) = \int n_e(\ell) \left[\sum_Z n_Z(\ell)\, \varepsilon_Z^{(t_i)} \big(T_e(\ell), n_e(\ell), n_0(\ell)\big) \right] d\ell \tag{6.27}$$

where t is the beryllium filter thickness associated with the i$^{\text{th}}$ chord, p_i is the chord's impact parameter, and ℓ is a measure of the distance along the chord. The model transforms $\ell \to (x, y) \to \rho$ which is used to evaluate the profiles as defined above. This will then be compared against the data $\boldsymbol{d}_{\mathrm{SXR}} = (d_1, d_2, \ldots, d_N)$ which has associated uncertainties $\boldsymbol{\sigma}_{\mathrm{SXR}} = (\sigma_1, \sigma_2, \ldots, \sigma_N)$.

The FIR diagnostic was implemented by modeling the chord geometry depicted in Figure 6.21 and directly solving for the line-averaged density

$$f_{\mathrm{FIR}}(p_i, \theta) = \frac{1}{L_i} \int n_e(\ell) d\ell, \tag{6.28}$$

where L_i is the length of the i$^{\text{th}}$ chord. A more thorough model could be constructed which models the measured phase shift $\Delta\phi$ directly via Equation 2.9. However, given that the uncertainty in the beam wavelength λ is low this was not considered to be necessary. Note that some FIR chords are at $\phi = 250^\circ$ and others are at $\phi = 255^\circ$, which is correctly accounted for in the geometry model.

All priors were chosen to be uniform, implemented according to Equation B.11. Bounds were chosen to restrict parameters to to be physically meaningful (i.e., $T_e > 0$) and within the bounds of plausibility based on previous experience (i.e. $T_e < 1$ keV for these conditions). The upper and lower bounds for each parameter is given in Table 6.1. These priors are represented as profiles in Figure 6.22, illustrating the allowable range.

All likelihoods were assumed to be Gaussian,

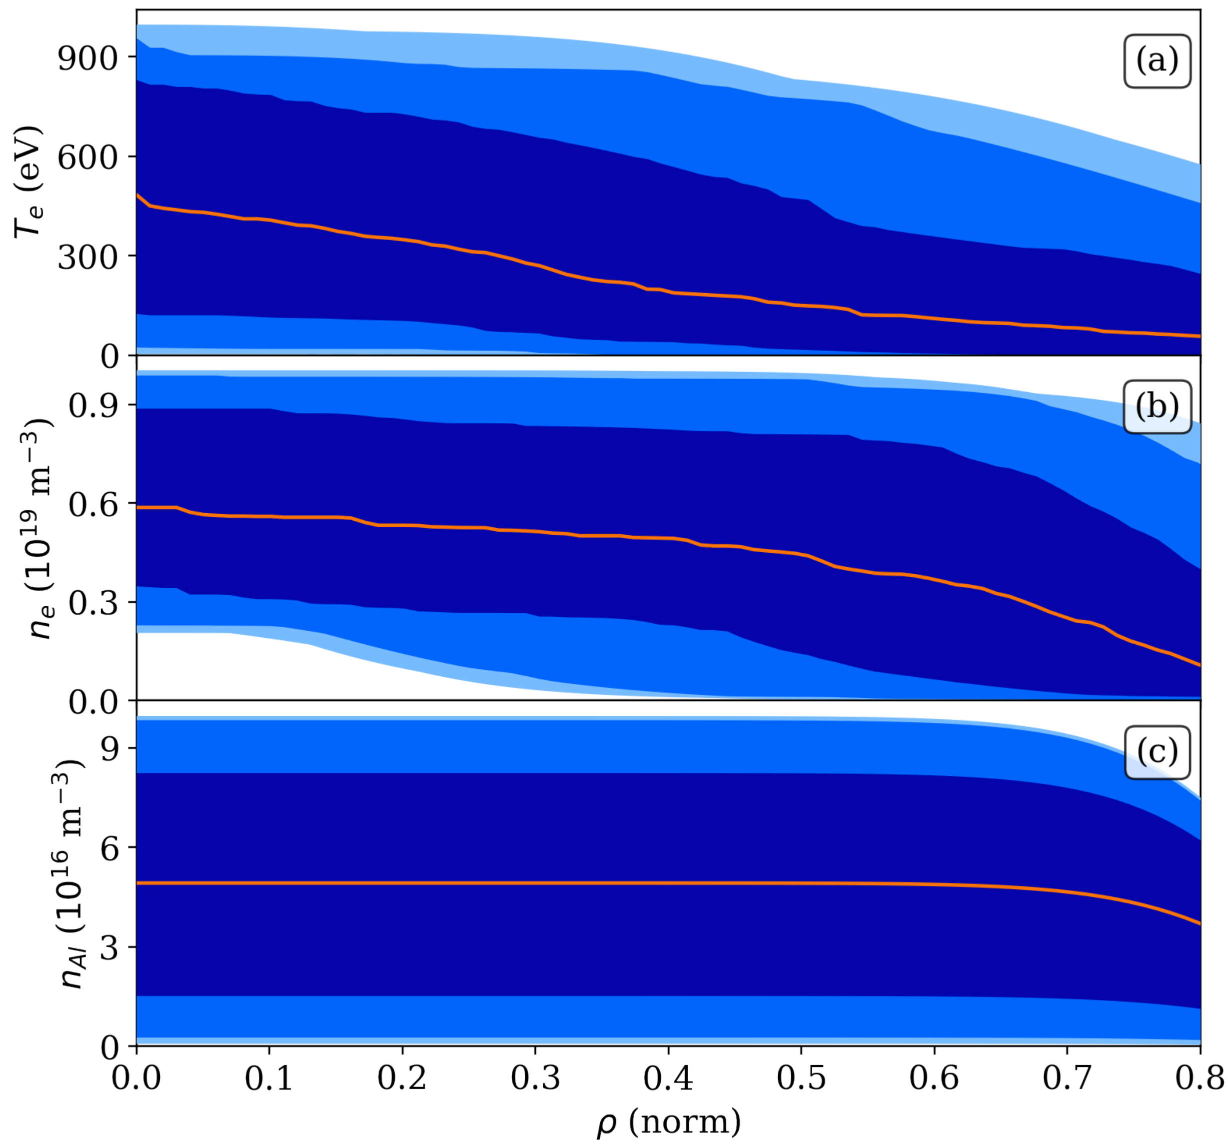

Figure 6.22: Prior distributions over the (a) electron temperature, (b) electron density, and (c) aluminum density profiles. The orange line represents the median and the shaded regions encompass 1-, 2-, and 3-σ ranges of possible profiles.

FIR Model				SXT Model			
Parameter	Units	Min	Max	Parameter	Units	Min	Max
$n_{e,0}$	10^{19} m^{-3}	0.2	1.0	$T_{e,0}$	eV	0.0	1e3
$n_{e,1}$	10^{19} m^{-3}	0.0	0.2	β_T	N/A	0.1	5.0
α_n	N/A	1.0	14.0	w_ρ	N/A	0.2	1.0
β_n	N/A	1.0	14.0	$n_{Al,0}$	10^{19} m^{-3}	1e-4	1e-2
				ε_1	W m^{-3} sr^{-1}	0	6
				ε_1	W m^{-3} sr^{-1}	0	1

Table 6.1: The parameters for the SXR tomography and FIR models with associated units and [Min, Max] prior bounds.

$$\ln p(\boldsymbol{d}|\boldsymbol{\theta}, I) = -\frac{1}{2}\sum_{i}^{N}\left(\frac{d_i - f(p_i, \boldsymbol{\theta})}{\sigma_i}\right)^2 \tag{6.29}$$

where I represents the additional knowledge that has been included in assembling the model. An additional 2% error has also been included in the SXR tomography likelihood function to account for known systematic uncertainties [46].

The likelihood and priors were combined according to Bayes rule (see Appendix B) to obtain the posterior distribution

$$p(\boldsymbol{\theta}|\boldsymbol{d}, I) \propto p(\boldsymbol{d}|\boldsymbol{\theta}, I)\, p(\boldsymbol{\theta}). \tag{6.30}$$

The posterior was sampled using the `emcee` MCMC software (see Appendix C). In this analysis the FIR and SXR tomography models were not solved simultaneously. Rather, the density profile was solved first using the FIR model and the result was taken as an input to the emissivity model. This may lead to a slight underestimation of the uncertainty on the T_e profile due to ignoring potential correlations with n_e. However, in practice the uncertainty on the density profile was seen to be very small, meaning that this effect is minimal.

6.4.2 Analyzing the fits

Samples were obtained from FIR and SXR tomography posterior distributions, which we will write as $p(\boldsymbol{\theta}_{FIR}|\boldsymbol{d}, I)$ and $p(\boldsymbol{\theta}_{SXT}|\boldsymbol{d}, I)$. Corner plots of the distributions are shown in Figures 6.23 and 6.24 respectively. The diagonal entries show the marginal distributions (Equation B.20) while the off-diagonals show the first-order correlations (all but two variables have been marginalized). The corresponding profiles are shown in Figure 6.25, along with 1-, 2-, and 3-σ uncertainty bands. Comparison to Figure 6.22 shows a significant reduction in uncertainty compared to the priors.

Most marginal distributions are seen to be nearly Gaussian, with the exception of two shape parameters, β_n and w_ρ (written as $\Delta\rho$ in the plot). The posteriors for these parameters are maximal near the boundaries of their respective priors, leading to a non-normal shape. This is not considered to be a problem, though, since at large values both parameters exhibit "diminishing returns" (that is, a increasingly large change in the parameter results in an increasingly small shift in the resulting profile). Several parameters are strongly correlated (such as $T_{e,0}$ and $n_{Al,0}$), but these correlations are approximately linear.

Next, we want to assess how well the model is able to reproduce the measured data. This is an important check which will provide us with confidence in the quality of the model. The model measurements in this section will be based on the means of the marginal posterior distributions for each parameter,

$$\langle \theta_i \rangle = \int p(\theta_i|\boldsymbol{d}, I)\theta_i d\theta_i, \tag{6.31}$$

which will then be plotted against the measured data. For convenience I will refer to these as being the "best-fit" parameters, although that is not a proper Bayesian terminology.

Figure 6.26 shows the comparison between the best-fit FIR model and the FIR data.

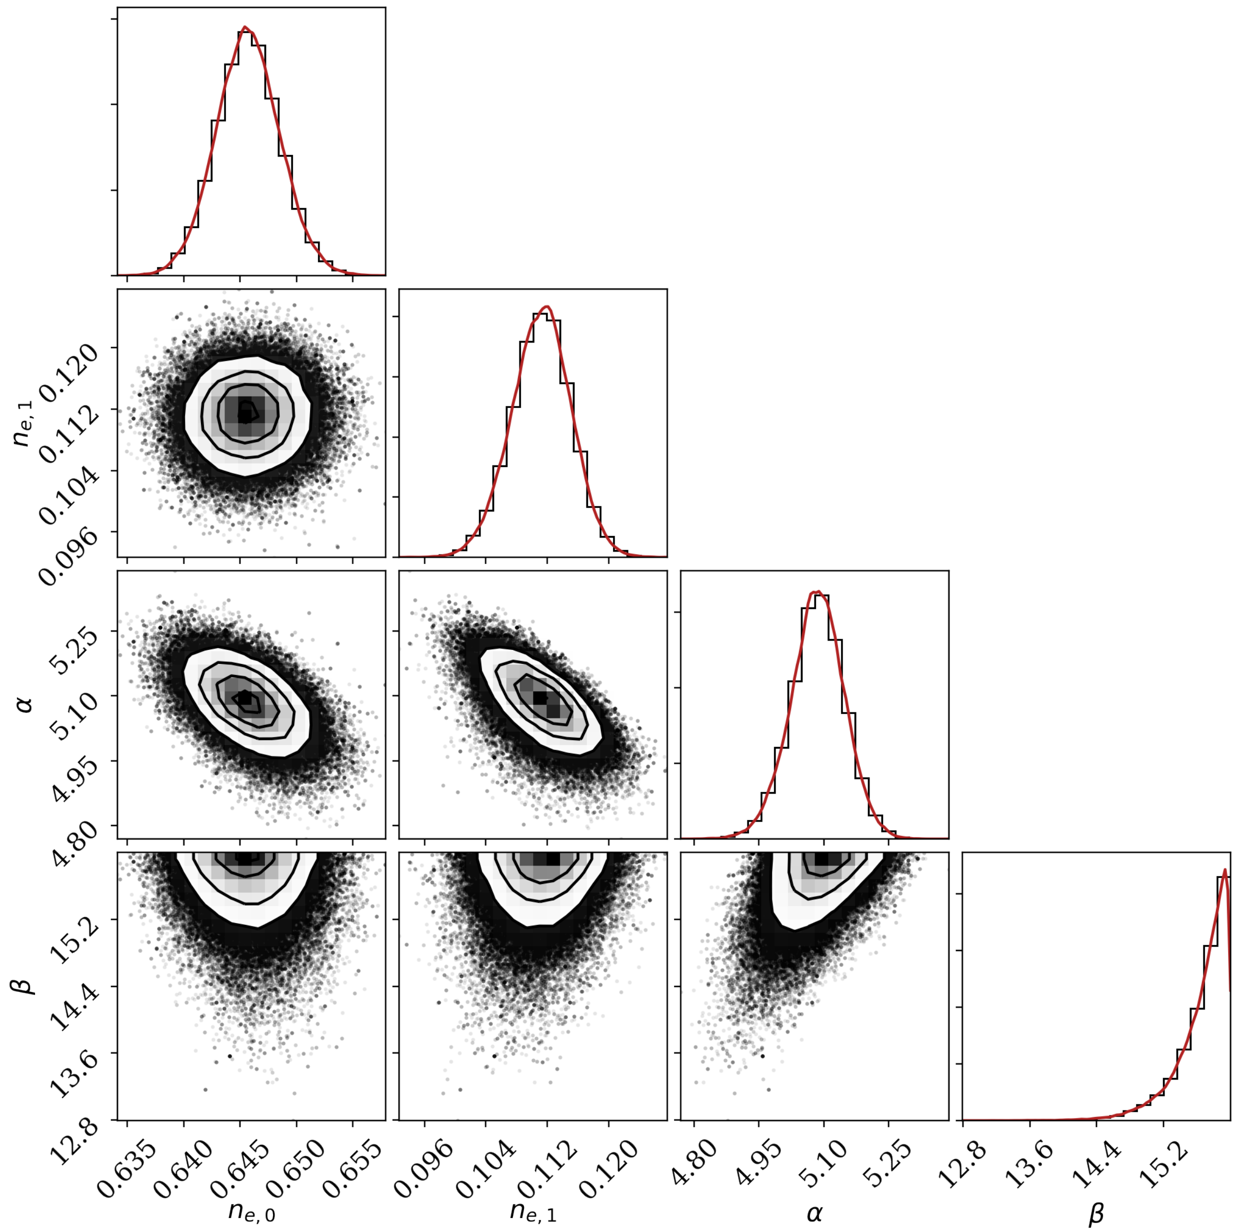

Figure 6.23: Corner plot showing marginal (diagonals) and first-order joint (off-diagonals) distributions for parameters in the FIR model posterior distribution for $t = 25$ ms.

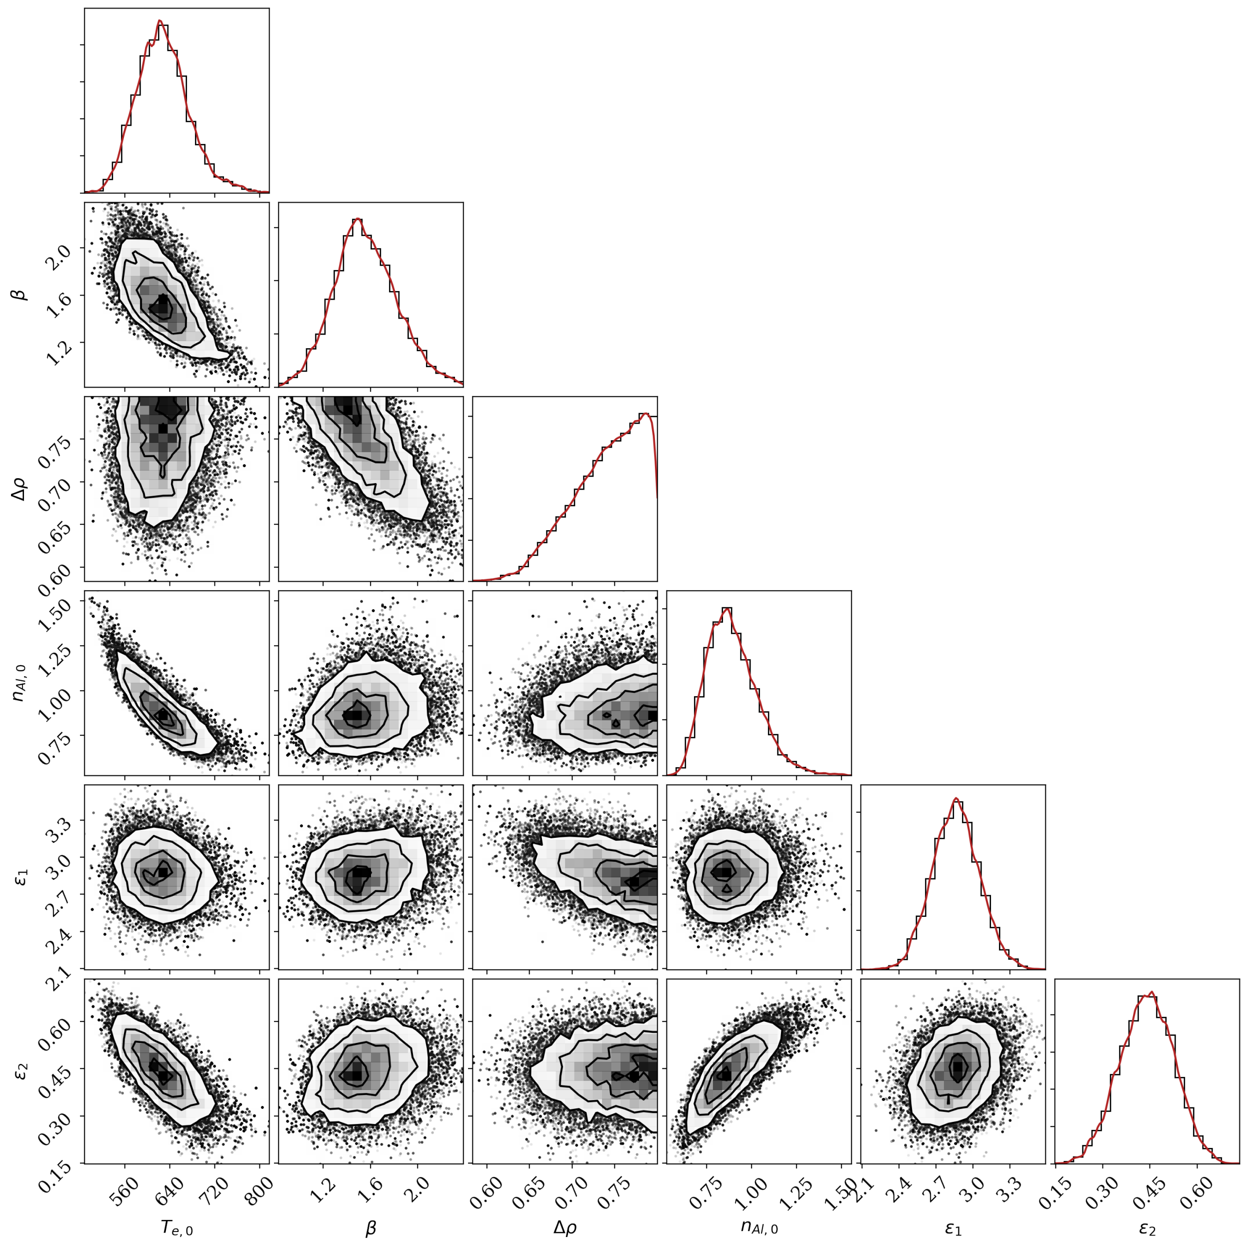

Figure 6.24: Corner plot showing marginal (diagonals) and first-order joint (off-diagonals) distributions for parameters in the SXR tomography model posterior distribution for $t = 25$ ms.

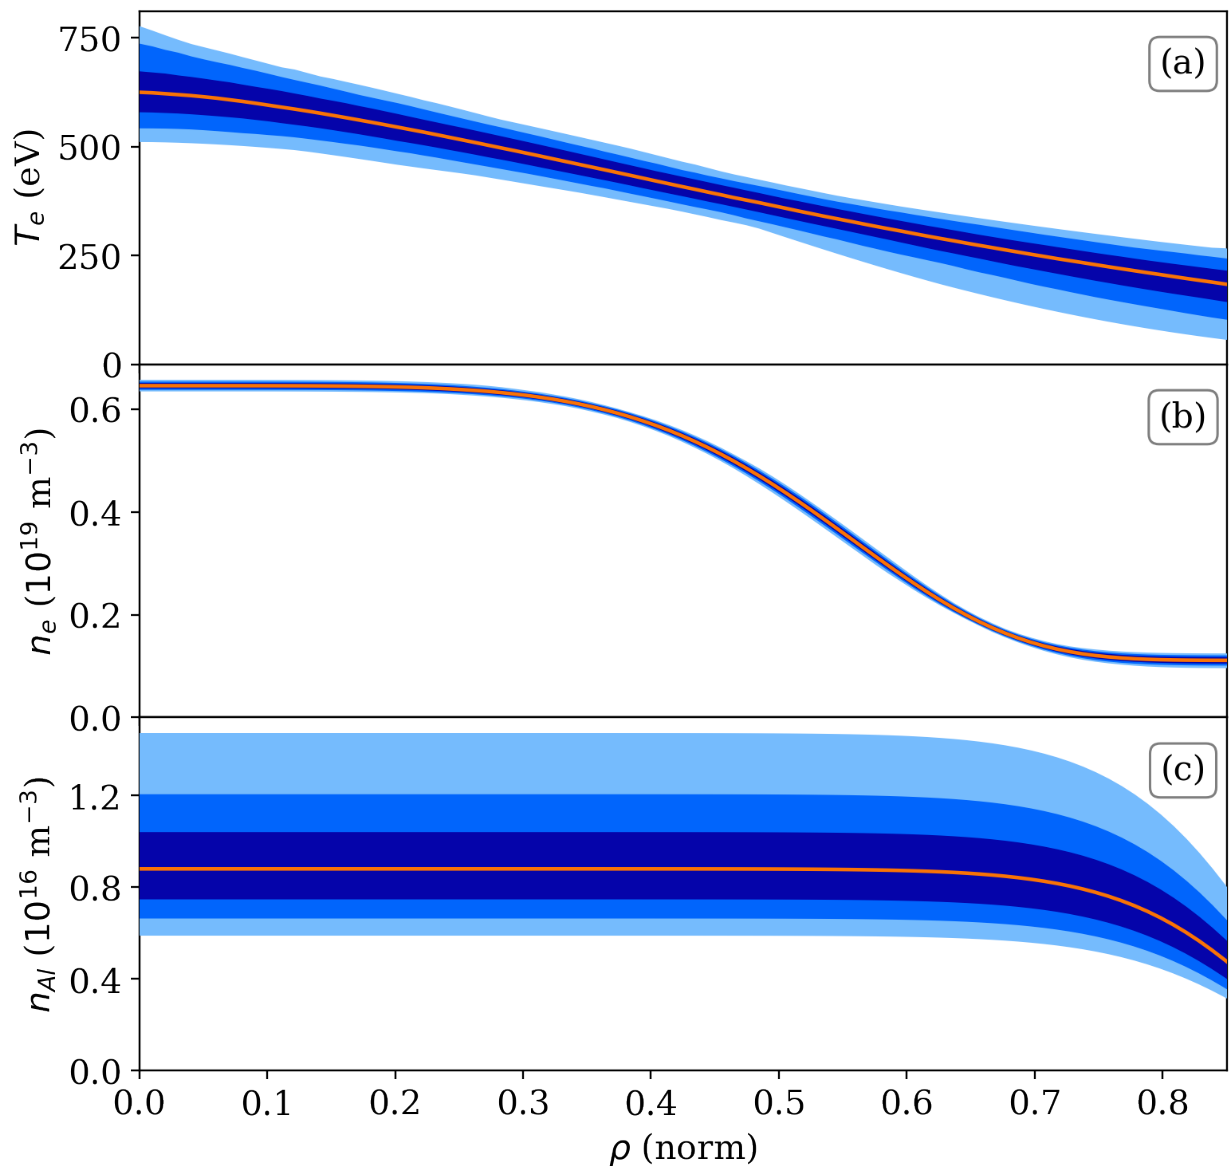

Figure 6.25: Posterior distributions over the (a) electron temperature, (b) electron density, and (c) aluminum density profiles. The orange line represents the median and the shaded regions encompass 1-, 2-, and 3-σ regions of all possible profiles.

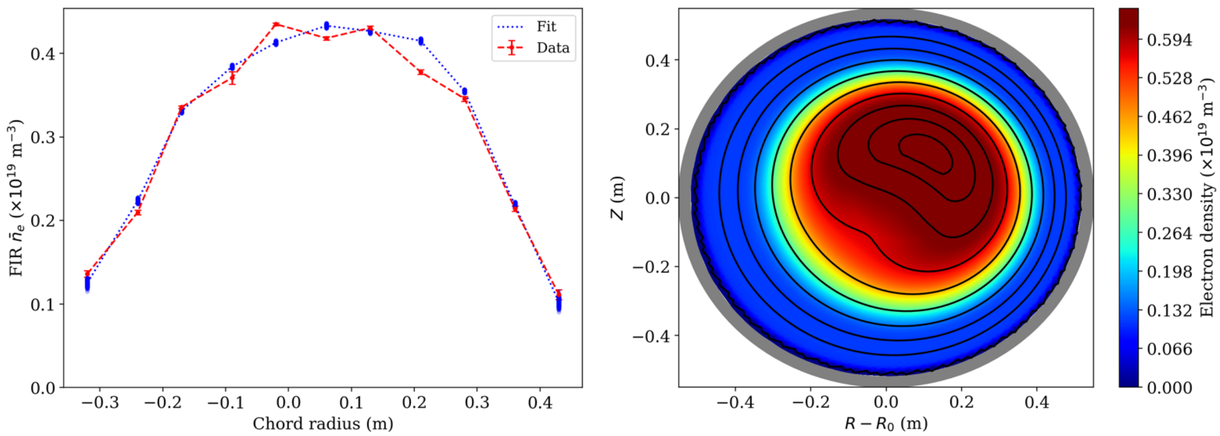

Figure 6.26: Results for the FIR fit showing (left) the measured data compared with the best-fit model and (right) the reconstructed 2D best-fit density profile.

Agreement is generally good, although the FIR data exhibits some small variation which is above the reported noise level but not accounted for in the model. Figure 6.27 shows the same comparison for the SXR tomography data. As with the FIR, the model does a very good job replicating the SXR data. Finally, Figure 6.28 shows a comparison between the reconstructed SXR emissivity and a direct inversion using the Cormack-Bessel technique (Section 2.3.3). The best-fit model closely matches the tomographic inversion while providing significantly more information about the underlying plasma. Taken together, these assessments allow us to have a high degree of confidence in the underlying model and the resulting reconstructed profiles.

To study the possibility of transport barrier formation, we want to evaluate the gradient of the best-fit profiles. As shown in Figure 6.29, the spacing (or packing) between adjacent flux surfaces is not even throughout the poloidal cross-section. The helical perturbation compresses flux surfaces between itself and the wall, and so physical profiles which are constant along flux surfaces will be steepest in this region. In order to evaluate this geometry, the profiles were evaluated along a chord which passes through the original (Shafranov-shifted axisymmetric) magnetic axis and the helical o-point. The distance along this chord is represented by the coordinate X (with a capital letter). It is assumed

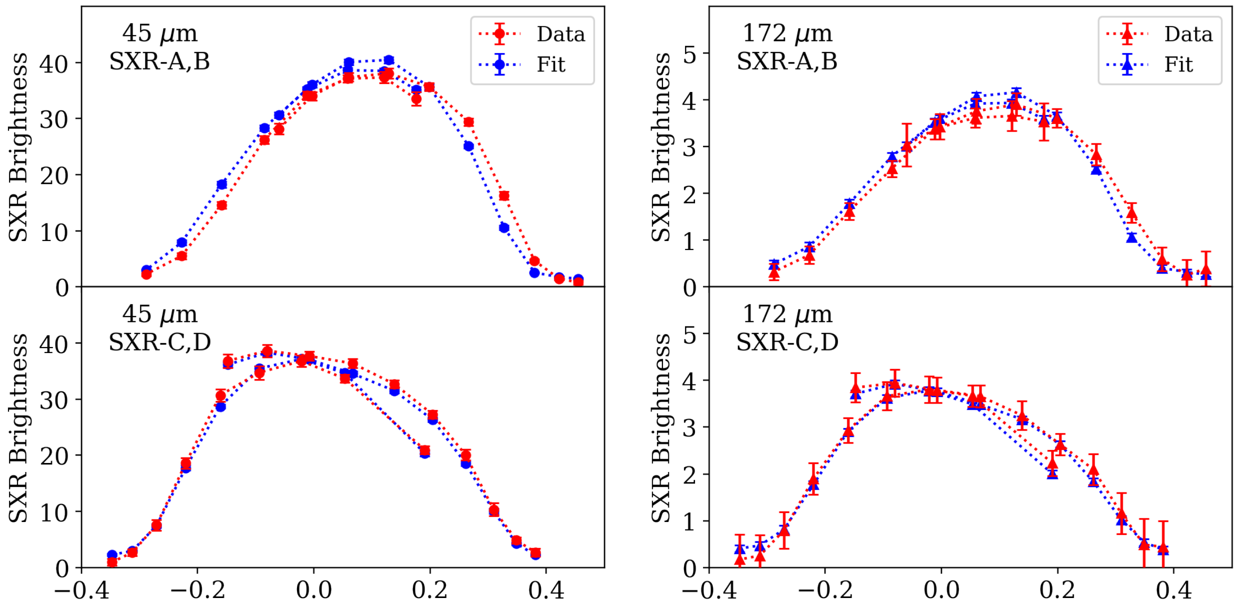

Figure 6.27: SXR tomography best-fit model compared with the input data for each diagnostic channel.

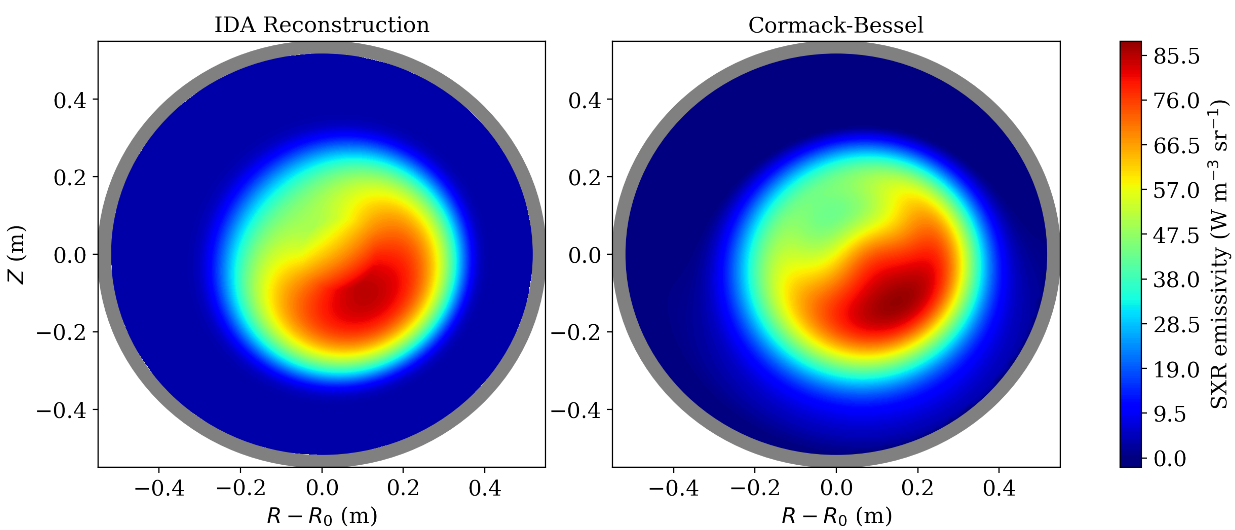

Figure 6.28: Comparison between the reconstructed SXR emissivity profile using (left) the best-fit model and (right) the Cormack-Bessel tomographic inversion.

that the maximum gradient lies along this chord such that $\max|\nabla T_e| = \max|dT_e/dX|$. The evolution of ∇T_e will be considered further in Section 6.4.3.

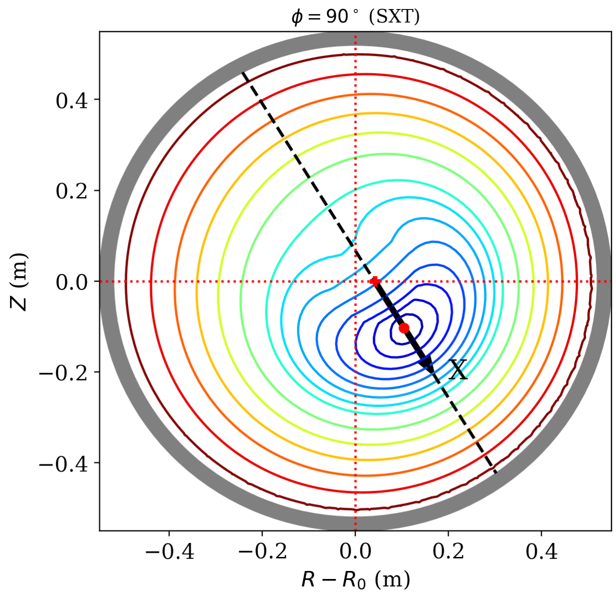

Figure 6.29: Profiles and associated gradients were evaluated along a chord passing through the location of the axisymmetric magnetic axis and the helical o-point. Distance along this line is represented by the X coordinate.

Drawing samples from a multidimensional posterior distribution defined by a complex model using an MCMC (or any other) algorithm is not a quick process. Although ideal for thoroughly analyzing a single time point, the computational requirements are prohibitive when attempting to analyze profile evolution over dozens of time points. In order to compensate for this, the MCMC sampling procedure was replaced with a Gaussian approximation using a nonlinear least-squares algorithm (based on `scipy.optimize.curve_fit` [47]) which uses bounds to enforce the uniform priors. The algorithm produces a mean parameter vector $\boldsymbol{\mu}_\theta$ which describes the best-fit and a covariance matrix $\boldsymbol{\Sigma}_\theta$ which describes the shapes of the distributions. Samples can be drawn using `scipy.stats.multivariate_normal.rvs`, which is based on Cholesky decomposition. An comparison between the resulting posterior T_e profile using the Bayesian

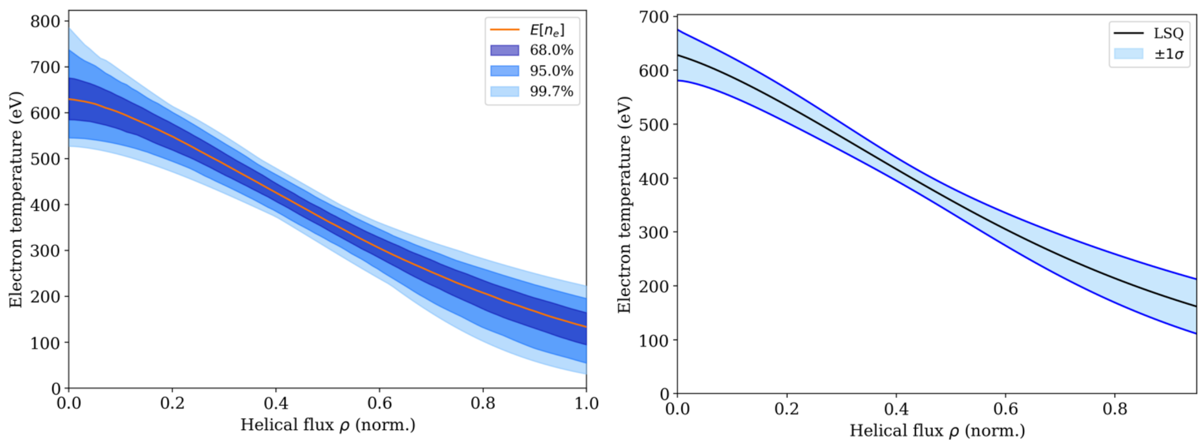

Figure 6.30: Comparison between the inferred T_e profile using Bayesian inference (left) and non-linear least squares (right).

MCMC method and the Gaussian least-squares approximation is shown in Figure 6.30. As anticipated, the posterior profiles are very nearly identical.

The next step is to extend this analysis to multiple time points. The algorithm will proceed as follows. For each time point:

1. Determine the best-fit $\boldsymbol{\theta}_{\text{FIR}}$ using the least-squares method, using the previous fit as a starting point.

2. Determine the best-fit $\boldsymbol{\theta}_{\text{SXT}}$ using the least-squares fitting method, taking $\boldsymbol{\theta}_{\text{FIR}}$ as given and using the previous fit as a starting point.

3. Evaluate the profiles across the chord X and calculate the gradients.

4. Assess uncertainties by ensembling over samples drawn from $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

5. Store the results and precede to the next time point.

It is worth pointing out that this procedure does have some weaknesses. Since the FIR and SXT models are solved individually rather than together it is not a proper

IDA procedure, and may not properly interpret the uncertainty. This analysis is also prone to over-estimating the temperature in the edge. SXR tomography signals fall off with temperature and are therefore very low in the edge region, providing a poor constraint. Therefore the edge T_e should be viewed as an extrapolation outside of the fitted region. The chosen model also does not explicitly enforce $T_e \to 0$ at the wall. This could be corrected with the addition of edge measurements (such as from Thomson scattering) and/or a modified profile model, but since this analysis is focused on core confinement (where the signal is strong) this was not considered to be a major flaw. Finally, this analysis strictly enforces a helical symmetry based on the underlying SHEq reconstructions, so any behavior which deviates from this symmetry will not be detected.

Finally, before moving on to the profile evolution results, I would like to discuss why the ME-SXR diagnostic was not used in this analysis as it was originally intended to be. As discussed throughout Section 6.3, the thick 450 μm Si absorber gives the ME-SXR a strong sensitivity to x-rays in the low-HXR part of the spectrum. Although this is not a problem when the population is mostly thermal, high-performance QSH plasmas in the MST feature a significant population of runaway electrons which constantly emit HXR radiation and complicate interpretation of the measurements. This could addressed by modeling the emissivity from a non-Maxwellian distribution, but that is a complicated task beyond the scope of this analysis. We could have also have added additional offset terms (like ε_1 and ε_2 in the SXR tomography model) to account for the HXR contamination, but that would have added eight additional free parameters to the model which would greatly increase the required computational time and the probability of over-fitting. So, the present approach was adopted where SXT data was used to extract T_e and the ME-SXR was used separately to diagnose the runaway population.

6.4.3 Temperature and density profile evolution

The results of the time-resolved analysis are summarized in Figures 6.31 and 6.32, which show the evolution of the equilibrium profiles and profile gradients, respectively. Profiles are plotted as a function of the spatial X coordinate in order to clearly demonstrate the spatial structure. Consistent with previous observations, we see in Figure 6.31 (b) that the electron density is well-confined throughout the discharge and changes very little in response to the resumption of tearing activity at $\sim$ 26 ms. At most there is a slight increase in the core electron density that builds up during the latter half of the quiet phase, but this effect is small. As was seen in Figure 6.26, the density profile is very broad and only weakly helical in character. The observed electron density gradient, $|\nabla n_e| \approx 4 \times 10^{19}$ m^{-4}, is steep, but less-so than observed in high-performance PPCD plasmas [48].

The temperature profile, however, displays a strongly helical character and evolves significantly during the course of the plasma discharge. A seen in Figure 6.31 (c), a broad, hot temperature structure develops during the quiet phase, when secondary mode activity is minimal. This structure is seen to collapse rapidly at the transition to the dynamic phase, leaving behind a hot peaked core surrounded by a cooler stochastic region. This hot core survives until the plasma begins to fall off at the end of the discharge. This is illustrated in Figure 6.33, which shows the reconstructed T_e cross-section at two time points only 4 milliseconds apart, one before the resumption of tearing activity and one after. The electron pressure, $p = n_e T_e$, is significantly increased in the quiet phase relative to the rest the plasma lifetime.

The maximum gradient reaches its zenith during the quiet period ($|\nabla T_e| \approx 3.5$ keV m^{-1}), but remains high ($|\nabla T_e| \approx 1.5$ keV m^{-1}) around the remnant hot core throughout the duration of the discharge. This is similar to the gradients observed in RFX [12], and strongly suggest the formation of a transport barrier. Note that Figure 6.32 (c) shows a very strong positive temperature gradient forming around 26 ms. This seems to be due

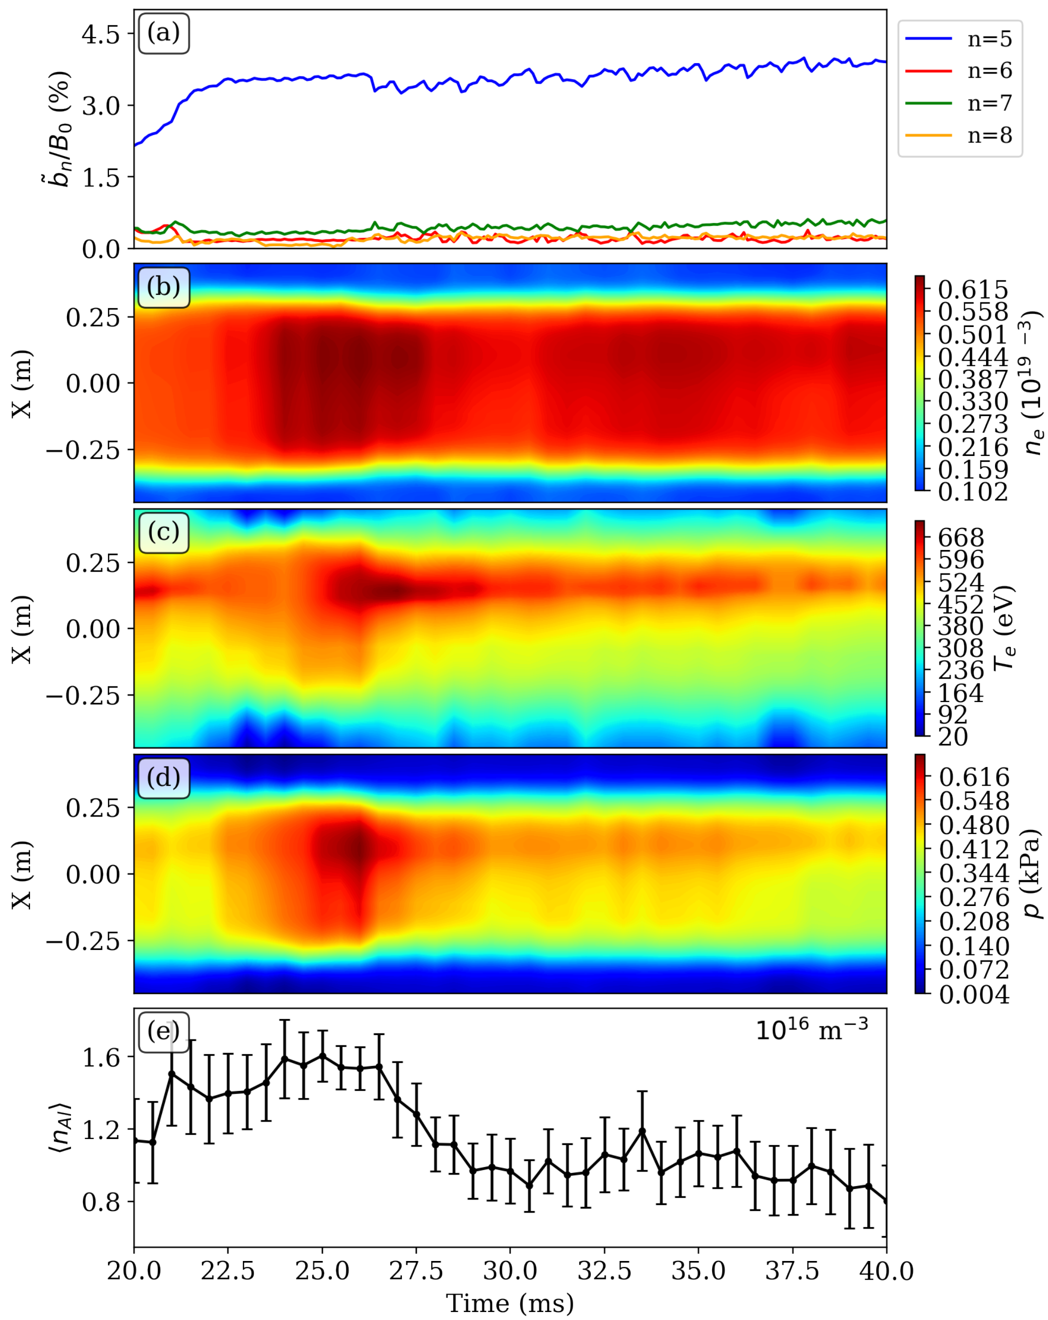

Figure 6.31: Plasma profile evolution for a single high-performance QSH plasma. Panels show: (a) dominant and secondary mode amplitudes; the reconstructed (b) electron density, (c) electron temperature and (d) electron thermal pressure profiles as a function of the X coordinate; and (e) core aluminum density.

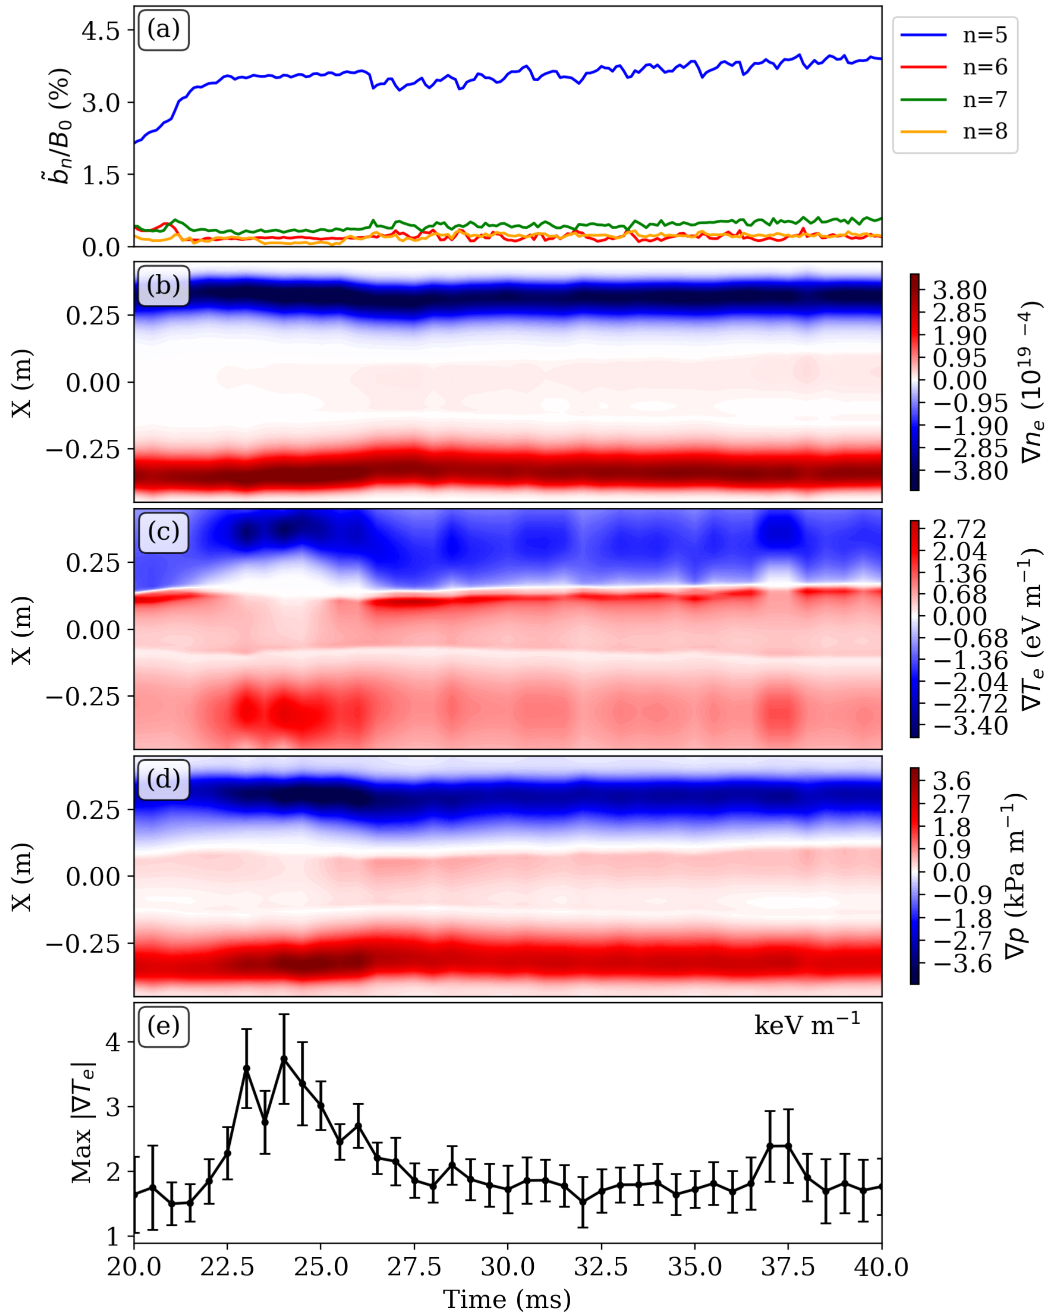

Figure 6.32: Plasma profile gradients for a the reconstructed QSH plasma shown in Figure 6.31. Panels show: (a) dominant and secondary mode amplitudes; the reconstructed (b) electron density, (c) electron temperature, and (d) electron pressure gradients as a function of the X coordinate; and (e) magnitude of the maximum T_e gradient, ignoring artifacts.

to a small, highly-localized kink in the reconstructed profile shape and is likely a quirk of Equation 6.25. This artifact was ignored in Figure 6.32 (d) and in later analyses.

Figure 6.31 also shows an $\sim$ 50% increase in the aluminum ion density during the quiet phase, suggesting a significant improvement to thermal ion confinement coinciding with the improved electron temperature confinement. This observation is consistent with the hypothesis that a broad internal transport barrier has formed during the quiet phase. Unfortunately this analysis does not resolve the n_{Al} ion spatial structure, so it is unknown to whether the improved confinement is global (like n_e) or restricted to the helical core (like T_e).

This analysis was performed on several additional MST plasmas other than the one presented in this Section. These reconstructions confirmed the results that were presented here, although it was observed that reconstructions for higher-density QSH plasmas showed overall reduced confinement. Figure 6.34 presents one such reconstruction, which is of particular interest because the electron density is significantly higher than for the plasma discussed above, and its flattop does not exhibit a quiet phase at all. Consequently, a region of broad thermal confinement does not develop, although a small hot spot does form in the plasma core. The temperature gradient is also less-steep, owing in part to the overall lower T_e, and the aluminum density decreases throughout the flattop, suggesting that ions are not well-confined.

6.4.4 Discussion: enhanced confinement during QSH

The time-evolving profiles presented in the previous section makes it clear that confinement is significantly improved during the MST's QSH/SHAx state. For plasmas which exhibit a quiet phase, this improvement is substantial: a broad T_e structure forms with steep gradients suggestive of a transport barrier, and the electron density increases, and impurity ions accumulate. It was also seen (Section 6.3) that runaway electrons are well-confined in this phase, suggesting the existence of good flux surfaces. The situation

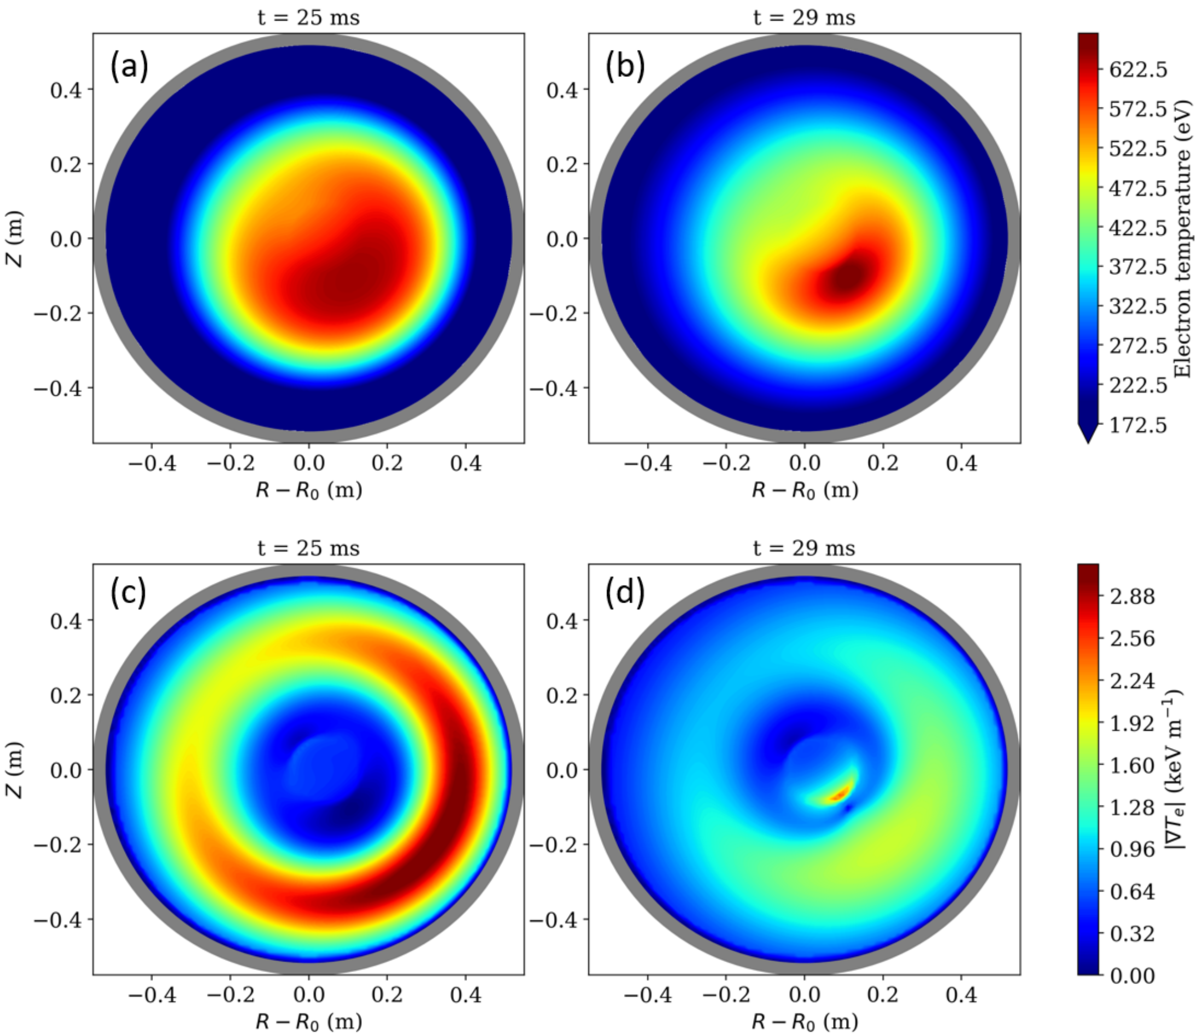

Figure 6.33: Comparison between the reconstructed 2D temperature profiles (a & b) for the best-fit parameters at 25 ms (quiet phase) and 29 ms (dynamic phase) and the associated gradients (c & d). A significant reduction in the size of the hot spot can be seen, corresponding to a relaxation in the gradient. Note that the small peak near the core in (d) is likely a fitting artifact.

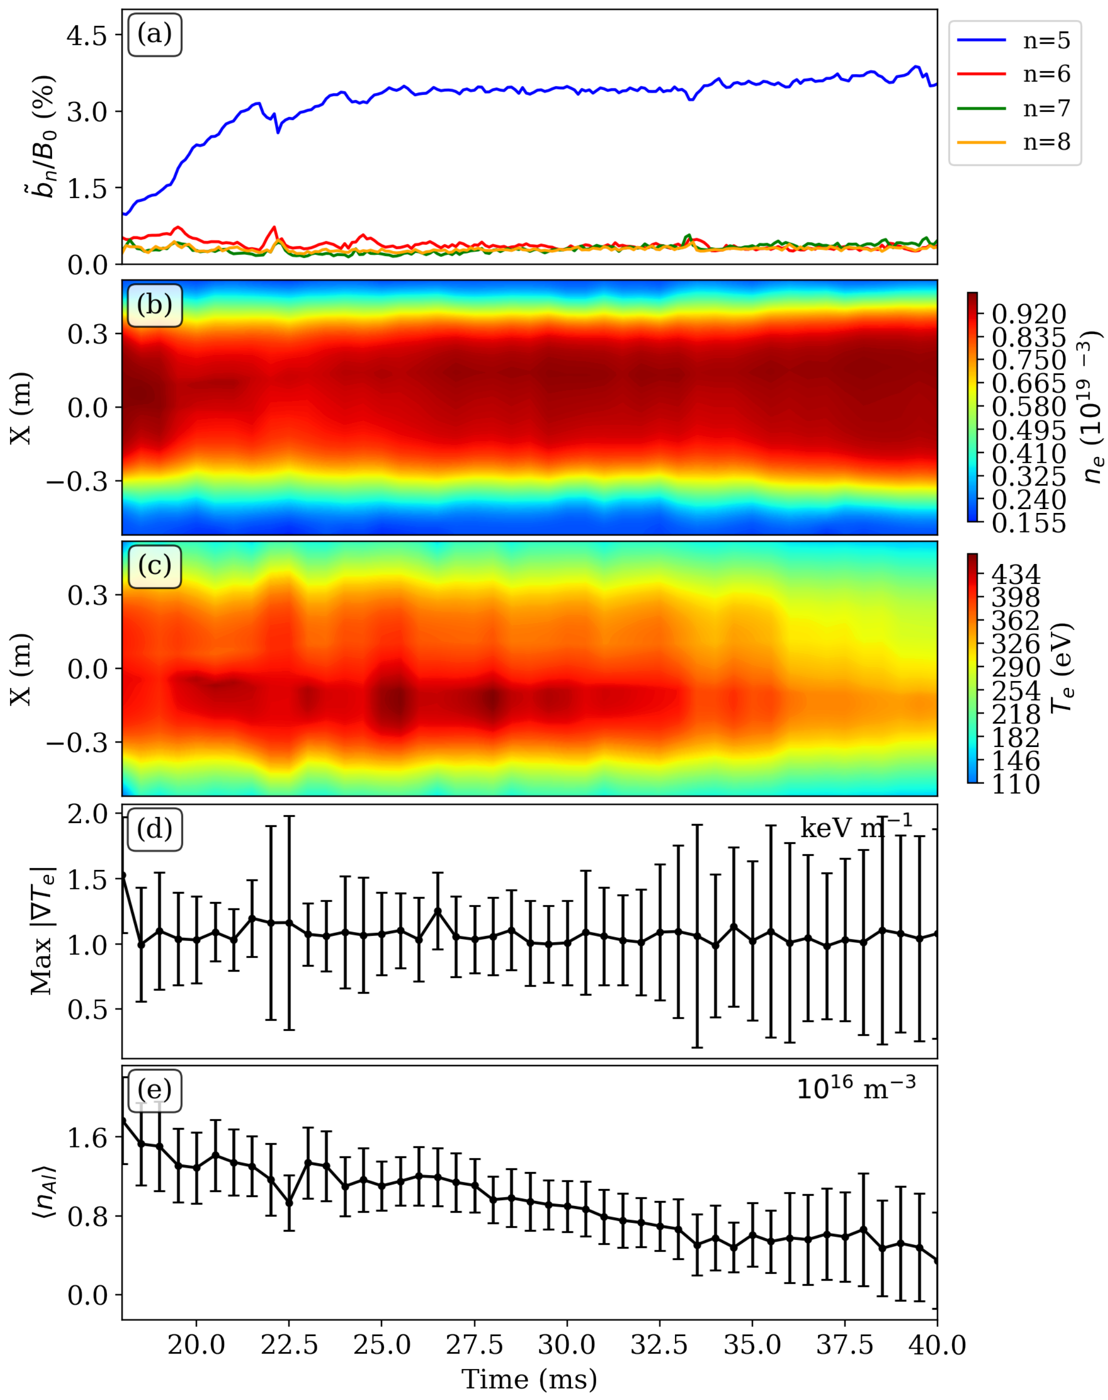

Figure 6.34: Plasma profile evolution for a single high-density QSH plasma with no quiet period. Panels show: (a) dominant and secondary mode amplitudes, the reconstructed (b) electron density, (c) electron temperature, (d) maximum temperature gradient, and (e) core aluminum density.

is somewhat akin to PPCD, but spontaneously self-organized rather than induced by applied current drive. After a few milliseconds, secondary mode activity suddenly resumes. The broad thermal structure is reduced, but a hot helical core survives. The impurity ion concentration drops, and the runaway electron population is rapidly reduced as stochasticity is restored. However, if the plasma lacks a quiet phase (instead going straight from the rising into the dynamic phase), the enhancements are reduced. It is clear that the existence of a quiet phase early in the lifetime of the QSH state is predictive not only to improved confinement during that phase, but also during the dynamic phase.

The change from a broad to a narrow structure is reminiscent of the distinction made at RFX between SHAx and DAx (double axis) equilibria [12]. DAx refers to an intermediate equilibrium in which a large magnetic island has formed but has not grown large enough to overtake the magnetic axis [49], leaving both a magnetic and helical axis in the plasma separated by a separatrix. In the SHAx state the separatrix vanishes and only the helical axis remains. The region of healed flux surfaces tends to be smaller for DAx states than SHAx states, and correspondingly the region of enhanced thermal confinement is smaller. A comparison between SHAx and DAx at RFX and MST temperature profiles from this reconstruction are shown in Figure 6.35. The similarities in the structure of the temperature profile are suggestive of similar underlying physics. Unfortunately, the SHEq reconstructions used in this analysis are not able to conclusively diagnose the existence or lack of a separatrix, so further analysis is needed.

These measurements, for the first time, provide strong evidence supporting the formation of a transport barrier during the QSH state in the MST. This is consistent with both previous studies at RFX-mod and the theoretical model described in Section 6.2.1, in which secondary tearing mode amplitudes are suppressed due to magnetic or flow shear, leading to enhanced confinement and formation of a transport barrier. Unfortunately, the uncertainty on ∇T_e is too large to establish a direct phase relation with

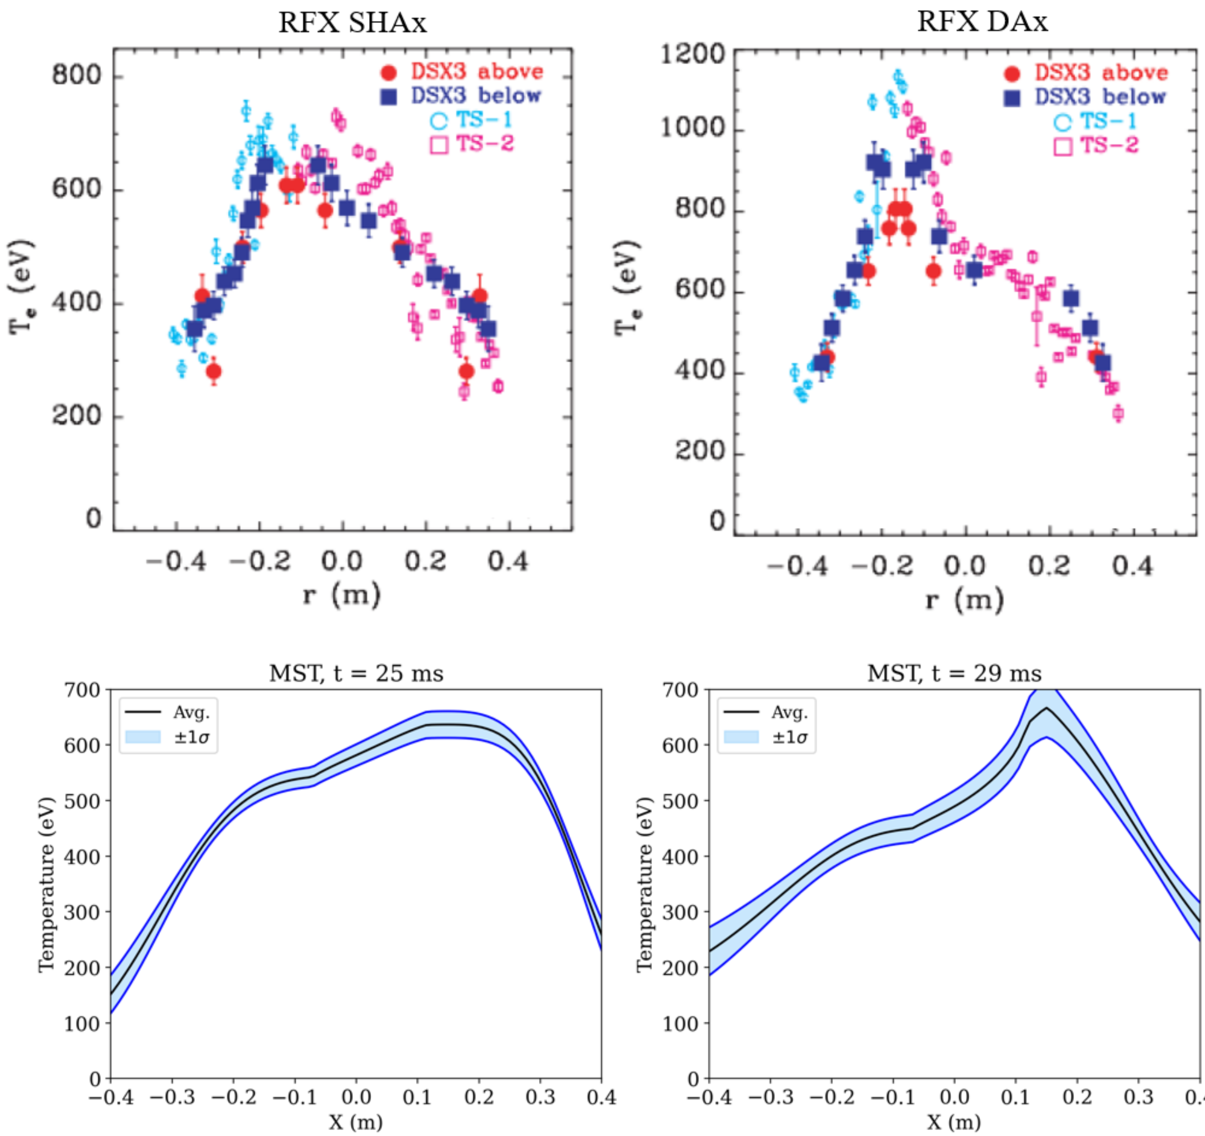

Figure 6.35: (Top) RFX T_e measurements for plasmas exhibiting SHAx and DAx equilibria, reproduced from Franz, *et al.* [12]. (Bottom) Temperature profiles during the quiet (left) and dynamic (right) flattop phases, separated by only four milliseconds.

oscillations in the dominant mode amplitude. This could be addressed in future work by ensembling over multiple reconstructions, or by incorporating additional information (such as Thomson scattering data) to better constrain the T_e profile. However, the observation that n_e changes very little over the course of the QSH state bolsters the claim made in Section 6.2.2 that oscillations in emissivity, which were observed to vary with the dominant mode amplitude, are due primarily to variation in T_e.

As previously mentioned, there have been proposals that the QSH state could play a role in the design of a hypothetical RFP fusion reactor [15]. In this context, the observation of enhanced confinement across a broad section of the plasma volume during the quiet phase is very promising. However, before this behavior can be exploited we must first understand the onset of the quiet phase, as well as what triggers its end. This is not presently understood, but the presence of predator-prey oscillations in the magnetic mode amplitudes and the SXR emissivity suggest that flow shears may be involved. This has led me to develop a hypothesis for the quiet phase based on temporarily-increased flow shear due to the sudden onset of magnetic mode locking.

In multi-helicity plasmas, it has been observed that when mode rotations temporarily lock due to a sawtooth crash, the plasma bulk velocity tends to slow down very rapidly along with it (on a $\tau < 1$ ms timescale) [50]. This is partly attributed to the fact that multiple magnetic islands are present in the plasma, meaning that a torque (arising from the error field) is applied to multiple locations in the plasma volume. This rapidly slows the fluid down through viscous coupling, before coming to rest at a lower, but non-zero, velocity. In QSH, however, there is only the one core mode with a large amplitude, meaning that the torque is probably much less uniform, concentrated near the location of the $n = 5$ island. Not only might this increase the slowing-down time, but it is possible that the non-uniformity of the force could create a strongly-sheared flow velocity profile. According to the Terry, *et al.* model (Section 6.2.1), this increased shear might serve to further suppress the transfer of energy to secondary modes. This could potentially explain both the $n = 5$ mode's sudden rise to very large amplitude and the brief enhanced confinement regime that sets in once the mode amplitude has saturated. Under this hypothesis, the resumption of low-amplitude tearing mode activity at the end of the quiet phase is because the flow profile has had time to slow down and even out such that the shear drops below some critical threshold. This also suggests that the way to extend the good confinement period is to somehow inject some momentum into the plasma rotation. For now this hypothesis remains unproven, but it provides a possible avenue for future work (Section 7.3).

6.5 Turbulent fluctuations in QSH plasmas

One of the major results of Section 6.4.3 is that low-density QSH plasmas can exhibit steep gradients, on the order $\nabla T_e \sim 2 - 4$ keV m^{-1}. Large gradients like these tend to drive turbulence in plasmas. For example, the large densities gradients that accumulate near the edge in PPCD plasmas have been found to drive trapped electron mode (TEM) turbulence during PPCD [48]. The steep ∇T_e is raises the potential of microtearing turbulence, a very high wavenumber temperature-driven cousin of the longer-wavelength current-driven instability which dominates the standard RFP spectrum. Gyrokinetic modeling of RFX-mod QSH states have suggested microtearing modes might be unstable in the transport barrier region [13], and limited experimental evidence seems to support this [14]. In light of the observations of improved confinement during the QSH period presented in the previous chapter, it is prudent to check if any new turbulent fluctuations can be observed.

Section 6.5.1 will discuss how digitized $\dot{b}_\theta \equiv \partial b_\theta / \partial t$ signals in the MST can be used to detect high-frequency fluctuations. Section 6.5.2 will apply this methodology to QSH plasmas, revealing the presence of a peak in the fluctuation spectrum around 600-1000 kHz which exhibits $n = 5$ symmetry and is only detectable when the helical structure is aligned with the magnetic coil. These results are preliminary, so further investigation will be needed. Still, they suggest that the improved confinement of the QSH/SHAx state may introduce new types of turbulent fluctuations.

6.5.1 Methodology

In order to retain sensitivity to fluctuations with frequencies in the hundreds of kHz, this analysis was conducted on the $\dot{b}_\theta$ signals measured by the toroidal array prior to passing through the integrator (which limits the bandwidth). These signals are digitized at 3 MHz, meaning that FFT analysis can be used to sense fluctuations up to 1.5 MHz. Since we are searching for fluctuations which potentially have mode numbers far above the resolution of the toroidal array ($n \gg 15$), we will focus our analysis on single $\dot{b}$ coil measurements. The signals were divided into windows of $\Delta t = 0.1$ ms, and FFTs were performed using the `numpy.fft` package [38]. For a more sophisticated analysis wavelets could be used, but given the high sampling rate of the input data this technique was found to be sufficient.

We are ultimately interested in the the frequency spectrum for the integrated magnetic field fluctuations, so we need to "integrate" the $\dot{b}$ signal by dividing each element in the FFT by its corresponding frequency. We will write the Fourier transform of the input signal as $\mathcal{F}[\dot{b}(t)] \equiv \dot{b}(\omega)$ where $\omega = 2\pi f$ is the angular frequency. Then, the measured signal is related to $b(\omega)$ by

$$\dot{b}(\omega) = i\omega b(\omega) \tag{6.32}$$

where i is the imaginary unit. The input signal $\dot{b}$(t) is real, the power spectral density is defined as

$$S(\omega) = 2\, b(\omega)\, b^*(\omega) \tag{6.33}$$

$$= 2\, \frac{\dot{b}(\omega)\, \dot{b}^*(\omega)}{|\omega|^2} \tag{6.34}$$

where b^* is the complex conjugate of b and $\omega > 0$ is assumed. Note that the $\omega = 0$ component has been dropped to avoid division by zero.

We can use two adjacent coils to estimate the toroidal mode n of an observed fluctuation. This is accomplished by calculating the wavenumber spectrum $S(k, \omega)$, where $k = n/R$ is the toroidal wavenumber and R is the major radius at the position of the coil ($R = 1.25$ m for the toroidal array). This technique has a resolution of of $k = \pm\pi/d$, so closer coils are able to resolve higher mode numbers. First, we calculate the cross-spectrum between the two signals b_1 and b_2,

$$S_{12}(\omega) = b_1(\omega)\, b_2^*(\omega) \tag{6.35}$$

which has an associated phase $\Theta_{12} = \tan^{-1} \frac{\Im S_{12}}{\Re S_{12}}$. The local wavenumber is then given by

$$K(\omega) = \frac{\Theta_{12}(\omega)}{d} \tag{6.36}$$

where $d = 2\pi R|\phi_1 - \phi_2| = 0.25$ m is the distance between adjacent coils in the toroidal array. The wavenumber spectrum is then given by

$$S(k, \omega) = I(k - K(\omega); \Delta k)\, \frac{S_1(\omega) + S_2(\omega)}{2} \tag{6.37}$$

where Δk sets the resolution on the mode number and

$$I(x, \Delta x) = \begin{cases} 1, & \text{if } 0 \le x \le \Delta x \\ 0, & \text{otherwise .} \end{cases} \tag{6.38}$$

is a "boxcar" function used to isolate k to a wavenumber window. For a more detailed (but still concise) discussion of this methodology, I refer the reader to J. Duff's Ph.D. thesis [51]. More thorough discussions can be found in standard signal processing textbooks.

6.5.2 Observation of high-frequency fluctuations during QSH

The power spectral density was computed for two coils in the toroidal array, one at $\phi = 76^\circ$ where the O-point of the helical structure is locked very close to the magnetic coil (in-phase) and the other at $\phi = 42^\circ$ where the O-point is locked on the opposite side of the vacuum vessel. This is shown in Figure 6.36, which shows the SHEq reconstruction at each toroidal angle. Although a reconstruction is more accurate (see Appendix D), we can approximate the poloidal position of the O-point as

$$\theta_{\text{op}}(\phi) = 241^\circ + 5 \cdot \phi - \delta_5, \tag{6.39}$$

where 241° is the poloidal angle of the toroidal array and δ_5 is the phase of the dominant magnetic mode, as decomposed according to Equation 2.3. Coil locations were chosen by selecting locations ϕ which minimize $\theta_{\text{op}} - 241^\circ$ (that is, as close to $\phi = \delta_5/5$ as possible).

A substantial difference was observed in the fluctuation spectra for the in-phase coil vs the out-of-phase coil. The in-phase coil, shown in Figure 6.37, shows a significant suppression of fluctuations during most of the quiet phase. Increased activity appears slightly before the end of the quiet phase and persists into the dynamic period. By the $t \approx 30$ ms, substantial activity can be seen in the 400-100 kHz range. This activity grows even stronger after $t = 35$ ms, when the plasma current begins to ramp down (Figure 6.1 (d)). This evolution can be seen clearly in the top plot of Figure 6.39, which shows the power spectral density averaged over the four windows indicated in Figure 6.37.

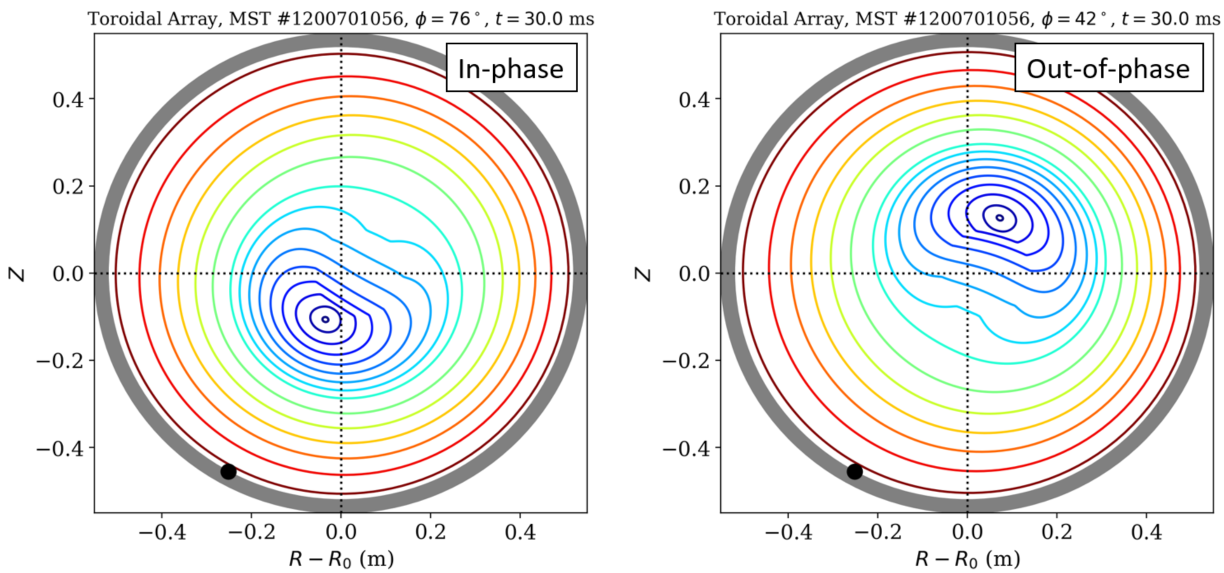

Figure 6.36: Sheq reconstructions for $\phi = 76^\circ$ and $\phi = 42^\circ$, where the black dot represents the location of the toroidal array. The O-point is locked in-phase and out-of-phase with the toroidal array, respectively.

Figure 6.38 shows the power spectral density for the out-of-phase coil. We still see the suppression of activity during the quiet phase, but the development of high-frequency activity in the dynamic phase is almost entirely absent. The bottom plot of Figure 6.39 shows that the spectrum changes only minimally as the plasma evolves. Even as the plasma current begins to ramp down, we do not see the emergence of strong activity for $f > 400kHz$. These observations have been repeated for multiple other QSH plasmas with different locking phases, using different $\dot{b}_\theta$ coils in the toroidal array. The robustness of these observations strongly suggest a connection between these observed fluctuations and the orientation QSH state. This may have to due with the coil's proximity to the steep gradient region.

If this increased activity is connected to the orientation of the helical structure, it should exhibit a clear $n = 5$ symmetry. This was tested by computing $S(\omega)$, averaged over a 30-40 ms, for all 32 coils in the helical array. This is plotted versus ϕ in Figure 6.40,

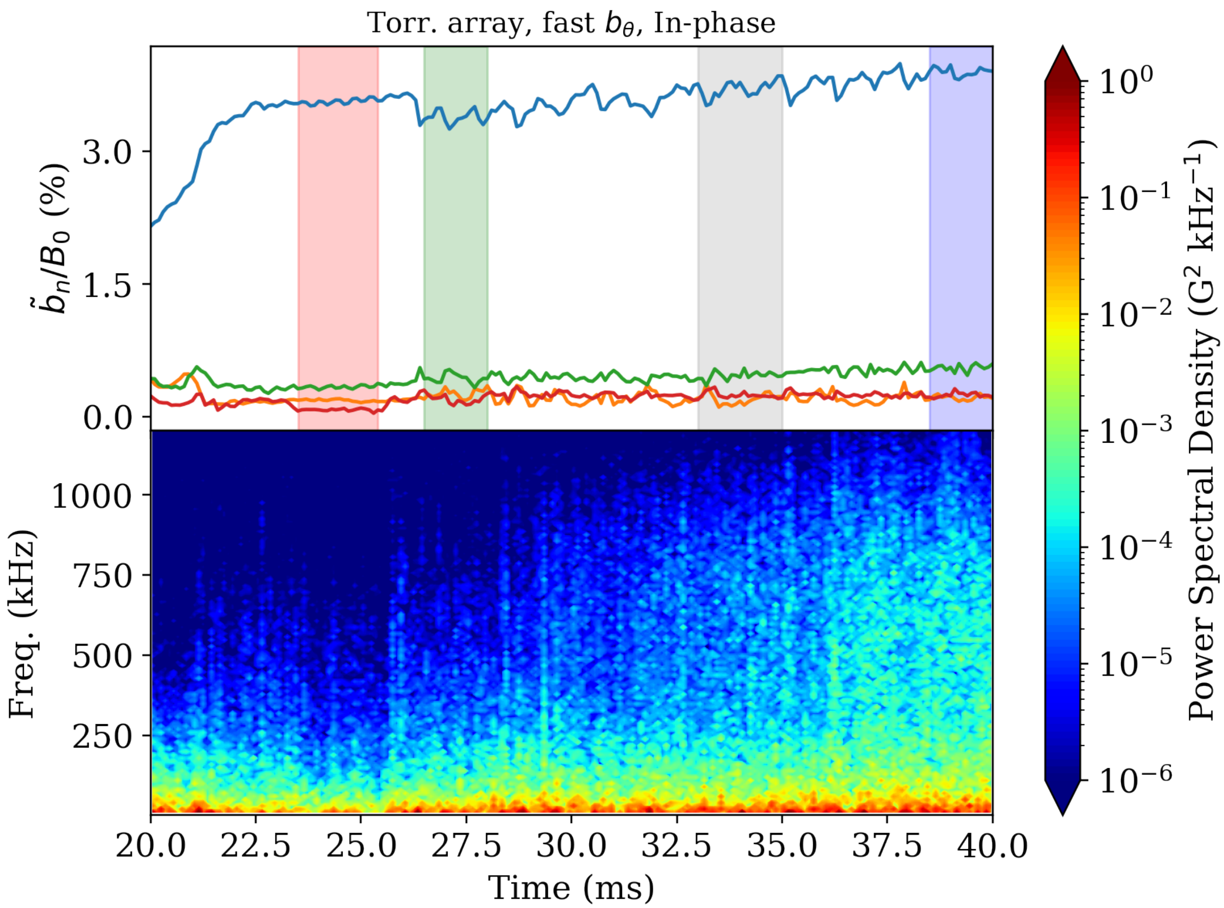

Figure 6.37: Spectrogram showing the evolution of the power spectral density of b_p over time when the helical O-point is locked near the toroidal array, showing significant high-frequency activity. The windows highlighted in the mode amplitude plot are averaged in Figure 6.39.

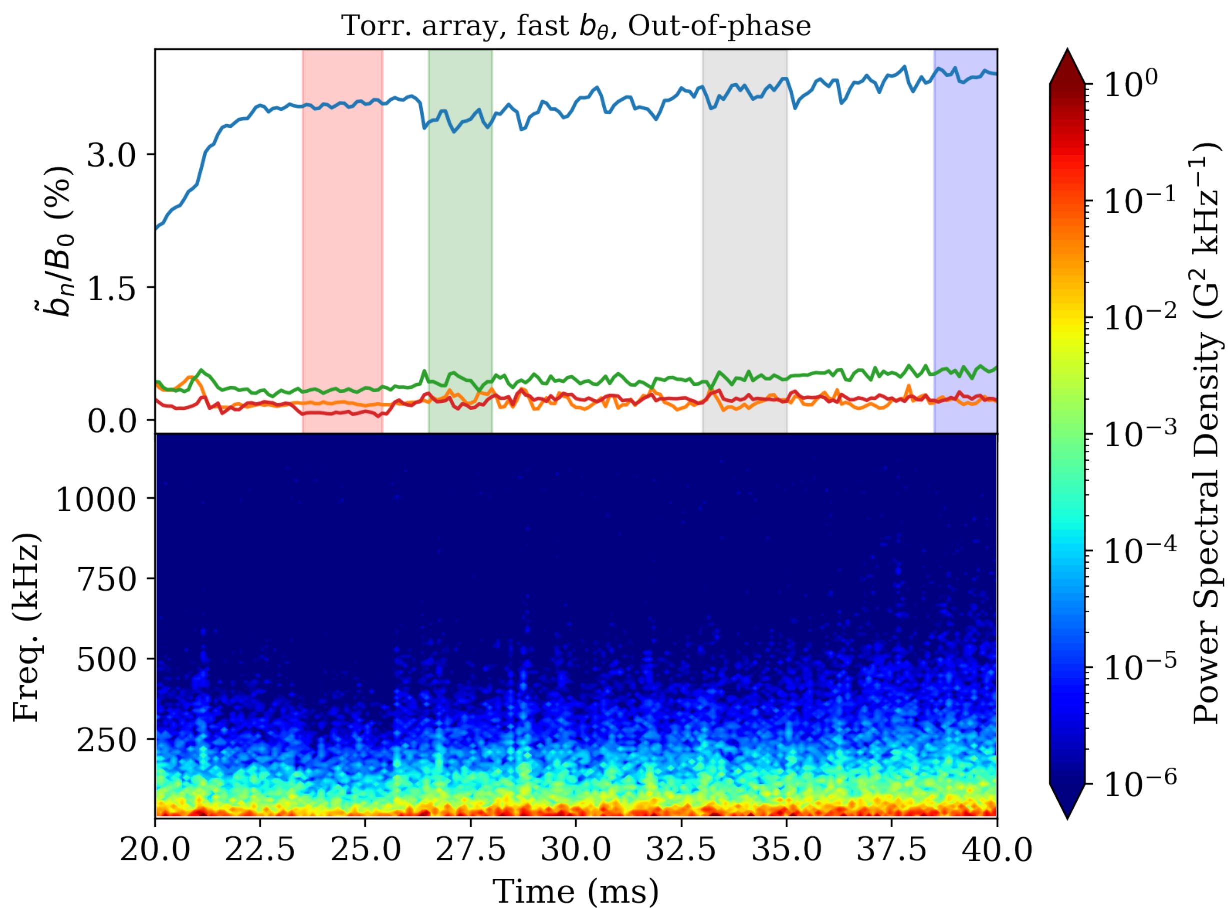

Figure 6.38: Spectrogram showing the evolution of the power spectral density of b_p over time when the helical O-point is locked away from the toroidal array. Very little high-frequency activity is observed.

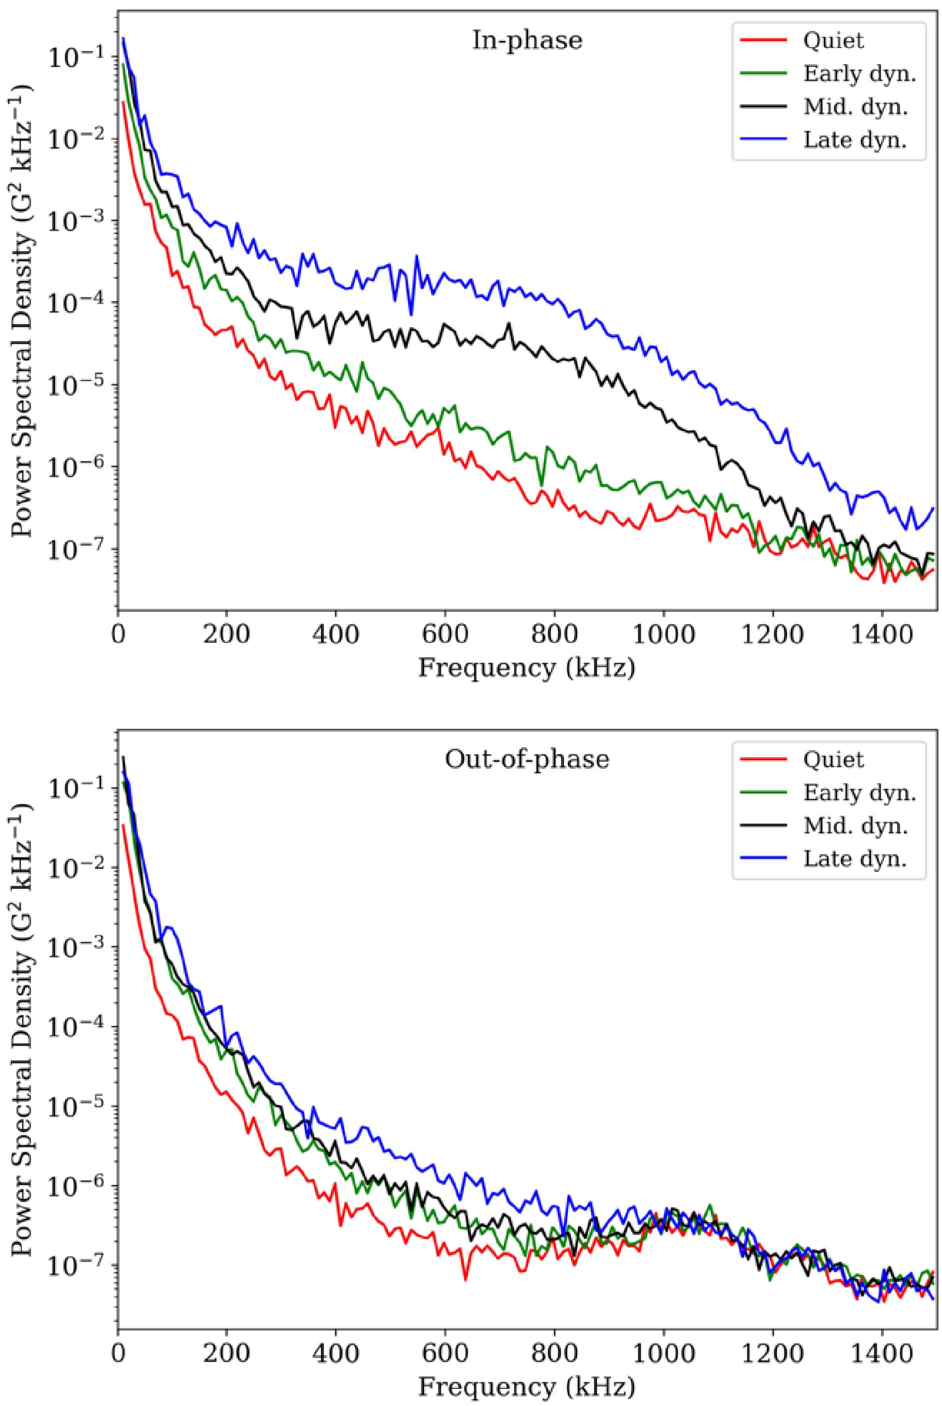

Figure 6.39: Time averaged spectra for the windows indicated in Figures 6.37 and 6.38, respectively. Note that significant high-frequency activity occurs during the dynamic phase only when the helical structure is aligned with the magnetic coil.

showing a clear periodicity. $\theta_{\text{op}} - 241^\circ$ is also overplotted, showing that the periodic increases in high-frequency activity do indeed correspond to angles where the helical structure is near the toroidal array. However, the intensity of the peaks show a toroidal variation. This may be in-part due to the finite spacing of the toroidal array coils[2] leading to variation in the distance of closest approach. However, it may also suggest the presence of additional toroidal structure (such as a superposition of multiple n modes), or a coupling to the error correction field.

The plot also shows a the same analysis but averaged from 15-20 ms, when the plasma was still in a multi-helicity (MH) state. As expected, the MH case entirely lacks the $n = 5$ periodicity seen in the QSH case, confirming that the observation of periodic increases in the spectrum is real and not due to systematic hardware effects.

Next, we want to estimate the wavenumber spectrum $S(k, \omega)$ in order to get an idea of the n numbers this new "mode" might be operating at. A small dataset of 6 low-density $F = 0$ plasmas exhibiting similar QSH properties was assembled, and Equation 6.39 was used to select the optimal coil locations. Equation 6.37 was averaged over a 5 ms window for each plasma, and then averaged together over the ensemble. The result is shown in Figure 6.41. The same analysis was performed using early time points in order to deduce $S(k, \omega)$ for equivalent multi-helicity plasmas, shown in Figure 6.42. Note that in both cases $+k$ is the ion drift direction and $-k$ is the electron drift direction.

A broad "streak-like" feature can be seen in the 400-1000 kHz region of the QSH spectrum which is absent in the MH spectrum. There is also substantial background noise in this region for both spectra. The group velocity of can be estimated fom two points within the streak:

[2] 360° / 32 coils $= 11.25^\circ$ between each coil, corresponding to a shift in θ_{op} of $5 \cdot 11.25^\circ = 56.25^\circ$.

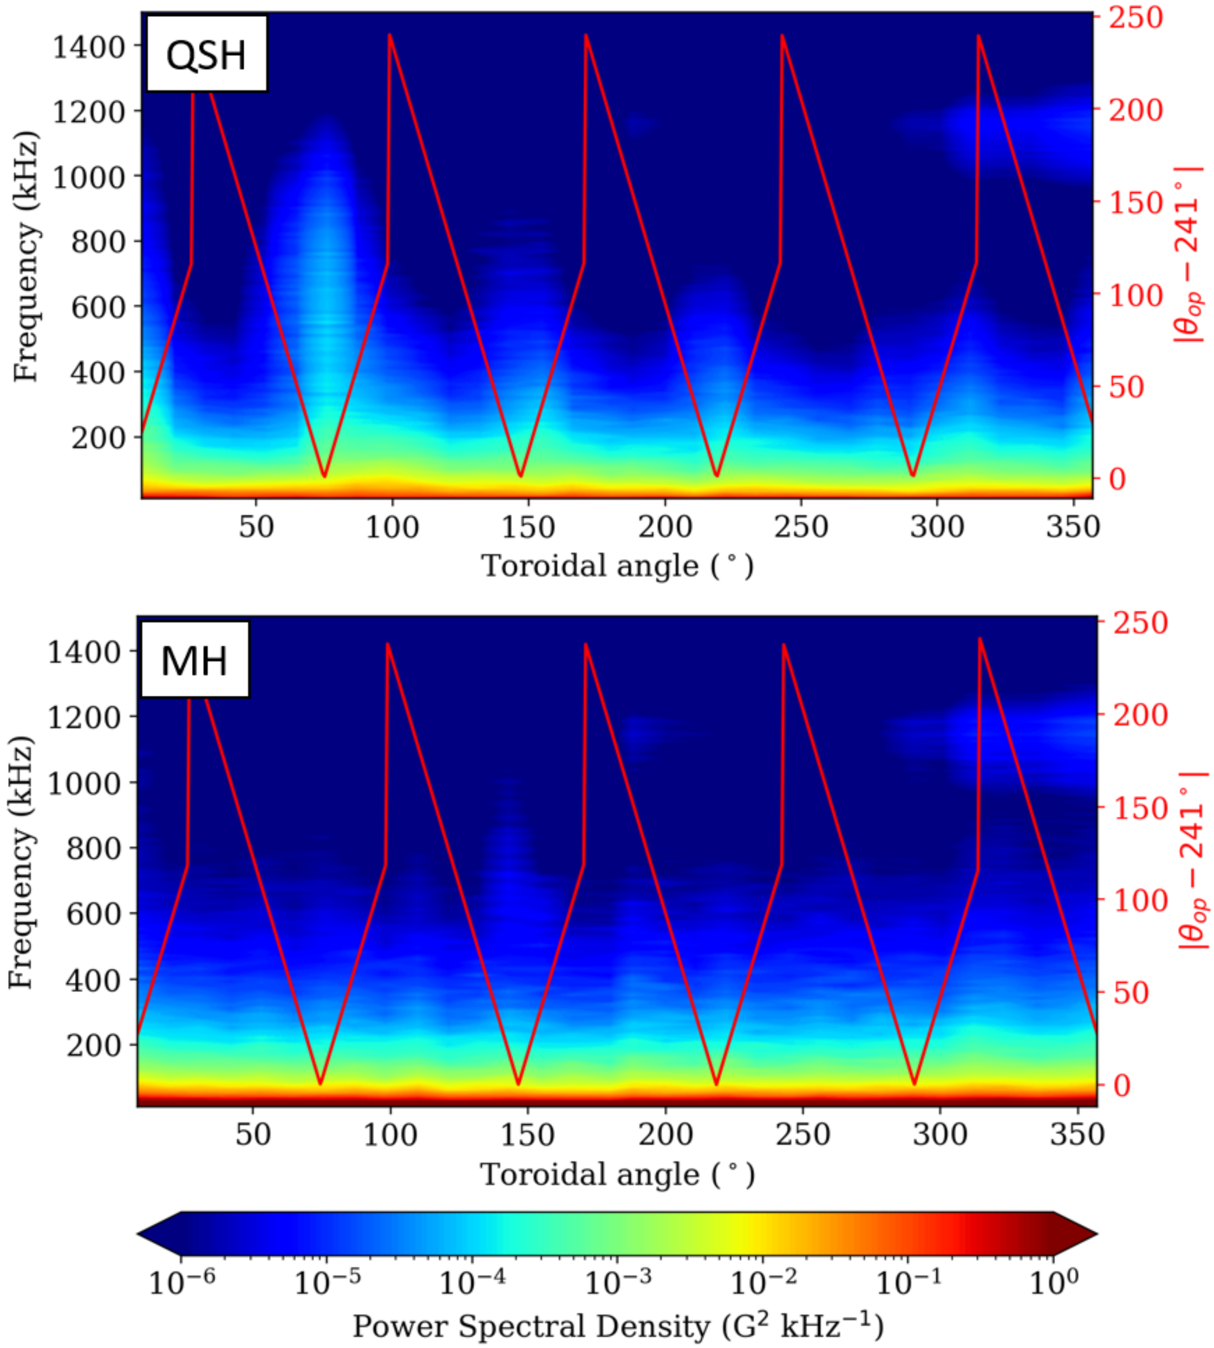

Figure 6.40: Increases in the power spectral density during QSH (top) show a clear $n = 5$ periodicity which aligns with the position of the helical O-point; during multi-helicity (bottom), no periodicity is observed. Note that the activity around $f \approx 1.1$ MHz for $\phi > 300^\circ$ is likely due switching noise from from the poloidal gap feedback system.

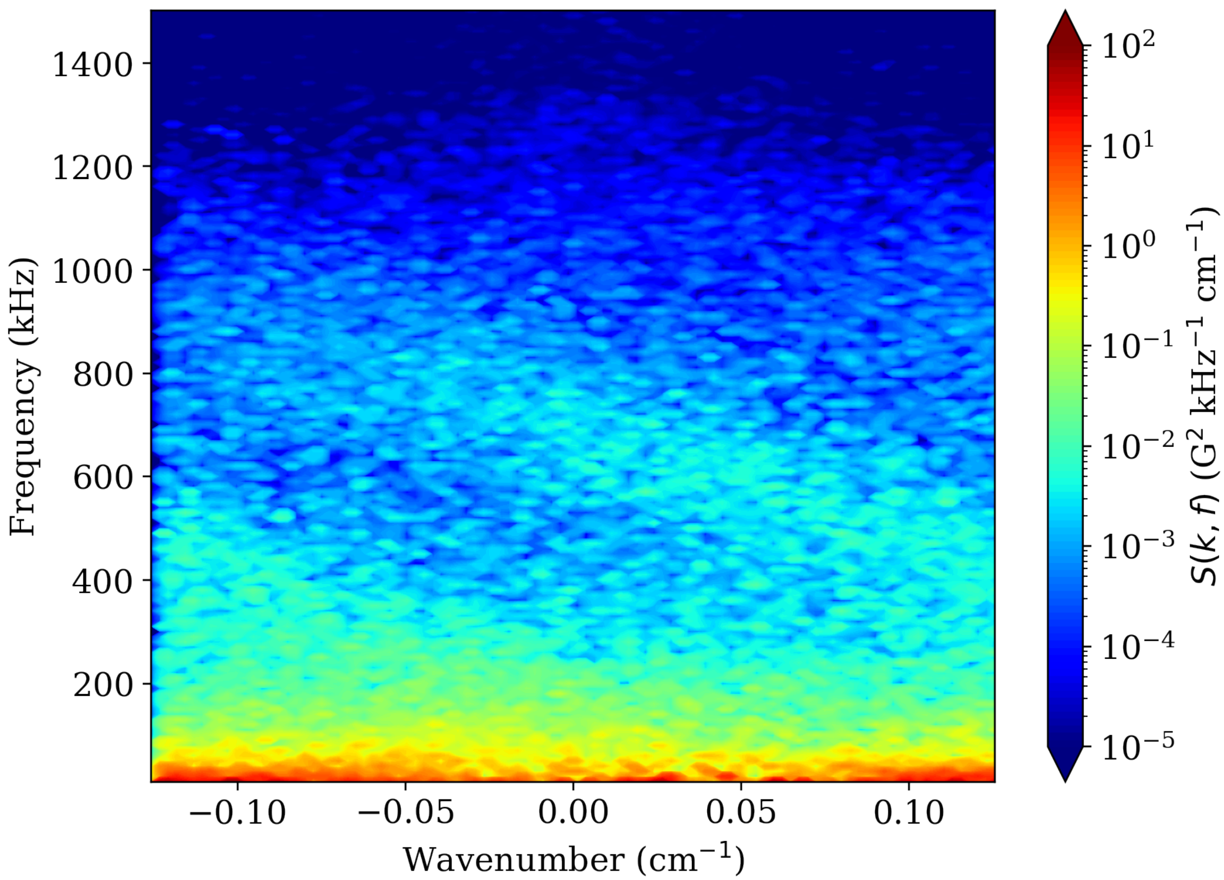

Figure 6.41: Ensemble-averaged wavenumber spectrum $S(k,\omega)$ for low-n_e QSH plasmas exhibiting high-frequency fluctuations.

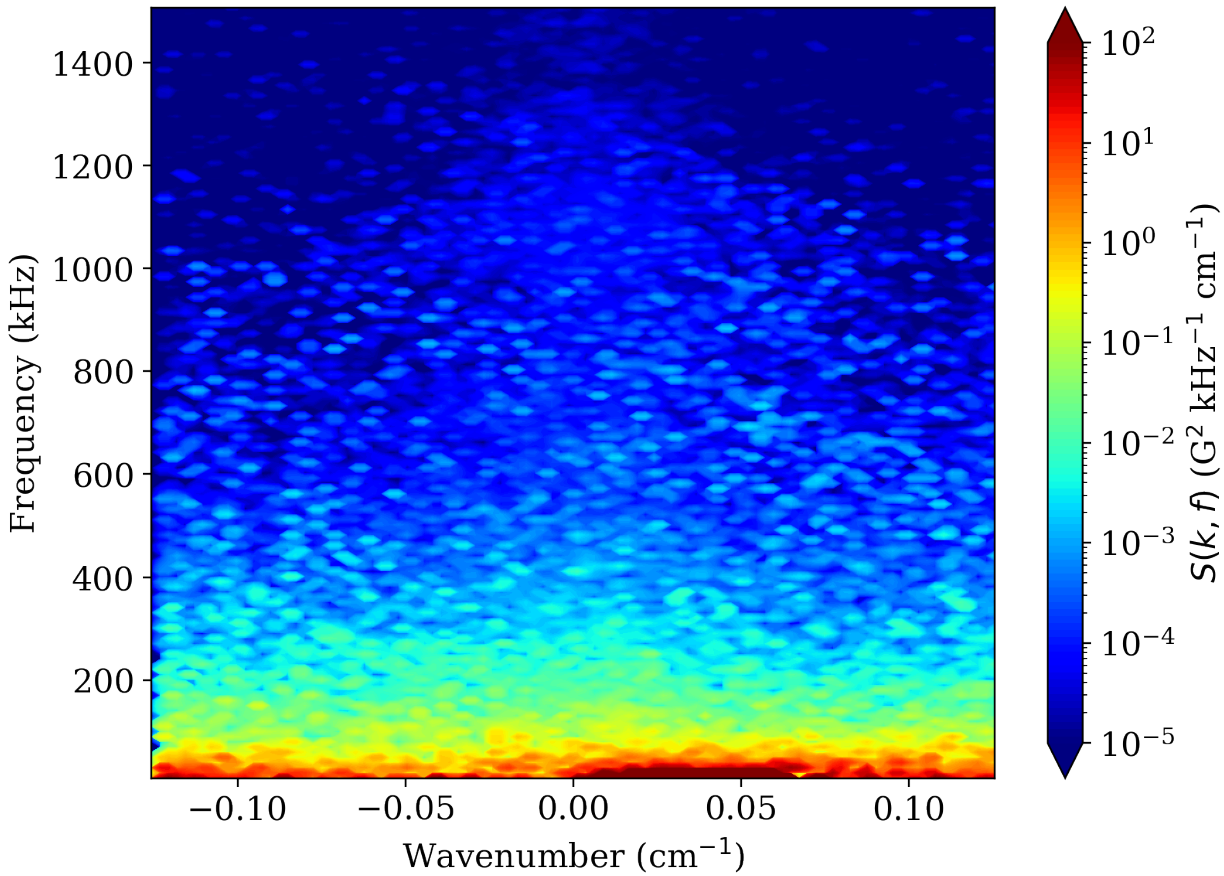

Figure 6.42: Ensemble-averaged wavenumber spectrum $S(k, \omega)$ over the same dataset as Figure 6.41 but during the early multi-helicity phase. Standard tearing mode activity can be seen at the bottom of the plot, at $f \sim$ 10-30 kHz and $k \approx +0.05$ cm^{-1}.

$$
\begin{aligned}
v_g &= \frac{\partial \omega}{\partial k} \\
&\approx \frac{2\pi(680 - 500\,\text{kHz})}{(0.08 - 0.0\,\text{cm}^{-1})} \\
&\approx 141\,\text{km s}^{-1}
\end{aligned}
\tag{6.40}
$$

The value of v_g estimated above is quite high. To put it into context, we can compare it against the diamagnetic electron drift velocity for these plasma conditions. The diamagnetic electron drift is driven by the thermal pressure gradient and has a velocity given by

$$
\boldsymbol{v}_{*e} = -\frac{\boldsymbol{\nabla} p \times \boldsymbol{B}}{e n_e B^2} \tag{6.41}
$$

where $p = n_e T_e + n_i T_i$ is the thermal pressure and e is the electron charge. Although the helical geometry somewhat complicates the the direct calculation of the cross product, by focusing solely on the region where the gradient is steepest some simplifications can be assumed. By assuming that the vectors are nearly perpendicular with $\boldsymbol{\nabla} p \approx \frac{\partial p}{\partial r}\hat{\boldsymbol{r}}$ and $\boldsymbol{B} \approx B_\theta \hat{\boldsymbol{\theta}}$, the (scalar) diamagnetic electron drift speed reduces to

$$
v_{*e} \approx \frac{\partial p}{\partial r} / \left(e n_e B\right). \tag{6.42}
$$

In the steep-gradient region we have $n_e \approx 0.3 \times 10^{19}$ m^{-3}, $B \approx 0.35$ T, and $(\partial p / \partial r) \approx 4$ kPa m^{-1} during the quiet phase and 3 kPa m^{-1} during the rest of the flattop (see Figure 6.32). This results in $v_{*e} \approx 24$ km s^{-1} in the quiet phase and ≈ 18 km s^{-1} later, which for simplicity we will average out to $v_{*e} \approx 20$ km s^{-1}. This is significantly lower than the group velocity estimate from $S(k, \omega)$ above.

One possible explanation is that the v_g estimated above was artificially high due to aliasing from wavenumbers beyond the resolution of the two-point technique, and actually corresponds to a mode with $k < 0$ (that is, propagation in the electron drift direction). If we assume that the observed instability actually travels at the diamagnetic electron drift speed, then we can estimate the range of wavenumbers we would expect the fluctuations to fall within. By setting $\omega/k = v_{*e}$, where $k = n/R$, we find $n \approx 2\pi R f / v_{*e}$, which gives $n \sim 200 - 300$ for an frequency range of $f \sim 600 - 800$ kHz. This corresponds to $k \sim 5 - 6\ \mathrm{cm}^{-1}$, which is well outside of the resolvable range for two adjacent coils in the toroidal array.

Attempts to examine the fluctuations using coils with closer spacing have not yet proven fruitful. No signatures of this new mode could be found in the dense array $\dot{b}_\phi$ measurements, though it is possible that these fluctuations are drowned out by the ambient background fluctuations that tend to dominate toroidal component measurements. Still, that means that the turbulence does not lead to an enhancement of b_ϕ. More curiously, no substantial activity in the 400-1000 kHz range was observed in the $\dot{b}_\theta$ coils on the dense array either, despite examining multiple plasmas with different locking phases. It is unclear if this is due to the ϕ dependence that was observed previously, or if something about the mode favors inboard rather than outboard measurements. This would be somewhat unusual, given a toroidal plasmas tendency to shift towards the outboard side of the vessel. Additional measurements will be needed to clarify the mode's properties and identity.

Bibliography

[1] P. Martin, L. Marrelli, G. Spizzo, P. Franz, P. Piovesan, I. Predebon, T. Bolzonella, S. Cappello, A. Cravotta, D. F. Escande, L. Frassinetti, S. Ortolani, R. Paccagnella, D. Terranova, R. Team, B. E. Chapman, D. Craig, S. C. Prager, J. S. Sarff, M. Team, P. Brunsell, J.-A. Malmberg, J. Drake, Y. Yagi, H. Koguchi, Y. Hirano, T.-R. Team, R. B. White, C. Sovinec, C. Xiao, R. A. Nebel, and D. D. Schnack, "Overview of quasi-single helicity experiments in reversed field pinches," *Nuclear Fusion*, vol. 43, pp. 1855–1862, 2003. [Online]. Available: https://doi.org/10.1088/0029-5515/43/12/028

[2] L. Marrelli, P. Martin, G. Spizzo, P. Franz, B. E. Chapman, D. Craig, J. S. Sarff, T. M. Biewer, S. C. Prager, J. C. Reardon, M. Symmetric, T. R. N. Dexter, D. W. Kerst, T. W. Lovell, and J. C. Sprott, "Quasi-single helicity spectra in the Madison Symmetric Torus," *Physics of Plasmas*, vol. 9, no. 7, pp. 2868–2871, aug 2002. [Online]. Available: https://doi.org/10.1063/1.1482766

[3] J. Sarff, A. Almagri, J. Anderson, T. Biewer, A. Blair, D. Brower, M. Cengher, B. Chapman, P. Chattopadhyay, V. Davydenko, D. Den Hartog, W. Ding, F. Ebrahimi, G. Fiksel, C. Forest, J. Goetz, R. Harvey, D. Holly, B. Hudson, A. Ivanov, T. Lovell, K. McCollam, P. Nonn, S. Oliva, R. Pinsker, S. Prager, J. Reardon, S. Terry, M. Thomas, C. Xiao, and T. MST Team, "Overview of Improved Confinement and Plasma Control in the MST Reversed Field Pinch," in *Proceedings of the Nineteenth IAEA Fusion Energy Conference*, Lyon, France, 2002.

[4] J. A. Holmes, B. A. Carreras, P. H. Diamond, and V. E. Lynch, "Nonlinear dynamics of tearing modes in the reversed field pinch," *Citation: The Physics of Fluids*, vol. 31, p. 1166, 1988. [Online]. Available: https://doi.org/10.1063/1.866746

[5] B. E. Chapman, R. Fitzpatrick, D. Craig, P. Martin, and G. Spizzo, "Observation of tearing mode deceleration and locking due to eddy currents induced in a conducting shell," *Physics of Plasmas*, vol. 11, no. 5, may 2004. [Online]. Available: https://doi.org/10.1063/1.1689353

[6] D. F. Escande, P. Martin, S. Ortolani, A. Buffa, P. Franz, L. Marrelli, E. Martines, G. Spizzo, S. Cappello, A. Murari, R. Pasqualotto, and P. Zanca, "Quasi-single-helicity reversed-field-pinch plasmas," *Physical Review Letters*, vol. 85, no. 8, pp. 1662–1665, 2000. [Online]. Available: https://doi.org/10.1103/PhysRevLett.85.1662

[7] M. Gobbin, D. Bonfiglio, A. H. Boozer, A. W. Cooper, D. F. Escande, S. P. Hirshman, J. Lore, R. Lorenzini, L. Marrelli, P. Martin, E. Martines, B. Momo, N. Pomphrey, I. Predebon, M. E. Puiatti, R. Sanchez, G. Spizzo, D. A. Spong, and D. Terranova, "Three-dimensional equilibria and transport in RFX-mod: A description using stellarator tools," *Physics of Plasmas*, vol. 18, no. 6, 2011. [Online]. Available: https://doi.org/10.1063/1.3602083

[8] W. F. Bergerson, F. Auriemma, B. E. Chapman, W. X. Ding, P. Zanca, D. L. Brower, P. Innocente, L. Lin, R. Lorenzini, E. Martines, B. Momo, J. S. Sarff, and D. Terranova, "Bifurcation to 3D Helical Magnetic Equilibrium in an Axisymmetric Toroidal Device," *Physical Review Letters*, vol. 107, no. 255001, pp. 1–5, dec 2011. [Online]. Available: https://doi.org/10.1103/PhysRevLett.107.255001

[9] D. Bonfiglio, M. Veranda, S. Cappello, L. Chacón, and G. Spizzo, "Magnetic chaos healing in the helical reversed-field pinch: Indications from the volume-preserving field line tracing code NEMATO," *Journal of Physics: Conference Series*, vol. 260, no. 1, 2010. [Online]. Available: https://doi.org/10.1088/1742-6596/260/1/012003

[10] S. Cappello, D. Bonfiglio, D. F. Escande, M. Veranda, L. Chacón, and G. Spizzo, "Equilibrium and transport for quasi-helical reversed field pinches," *Journal of Physics: Conference Series*, vol. 260, no. 012003, p. 12003, 2010. [Online]. Available: http://dx.doi.org/10.1088/0029-5515/51/10/103012

[11] S. Munaretto, B. E. Chapman, M. D. Nornberg, J. Boguski, A. M. Dubois, A. F. Almagri, J. S. Sarff, S. Munaretto, B. E. Chapman, M. D. Nornberg, J. Boguski, A. M. Dubois, and A. F. Almagri, "Effect of resonant magnetic perturbations on three dimensional equilibria in the Madison Symmetric Torus reversed-field pinch Effect of resonant magnetic perturbations on three dimensional equilibria in the Madison Symmetric Torus reversed-field pinch," *Physics of Plasmas*, vol. 056104, no. 23, 2016. [Online]. Available: http://dx.doi.org/10.1063/1.4943524

[12] P. Franz, M. Gobbin, L. Marrelli, A. Ruzzon, A. Fassina, E. Martines, and G. Spizzo, "Experimental investigation of electron temperature dynamics of helical states in the RFX-Mod reversed field pinch," *Nuclear Fusion*, vol. 53, no. 5, p. 053011, 2013. [Online]. Available: https://doi.org/10.1088/0029-5515/53/5/053011

[13] I. Predebon, F. Sattin, M. Veranda, D. Bonfiglio, and S. Cappello, "Microtearing modes in reversed field pinch plasmas," *Physical Review Letters*, vol. 105, no. 19, pp. 1–4, 2010. [Online]. Available: https://doi.org/10.1103/PhysRevLett.105.195001

[14] M. Zuin, S. Spagnolo, I. Predebon, F. Sattin, F. Auriemma, R. Cavazzana, A. Fassina, E. Martines, R. Paccagnella, M. Spolaore, and N. Vianello, "Experimental observation of microtearing modes in a toroidal fusion plasma," *Physical Review Letters*, vol. 110, no. 5, pp. 1–5, 2013. [Online]. Available: https://doi.org/10.1103/PhysRevLett.110.055002

[15] R. Lorenzini, E. Martines, P. Piovesan, D. Terranova, P. Zanca, M. Zuin, A. Alfier, D. Bonfiglio, F. Bonomo, A. Canton, S. Cappello, L. Carraro, R. Cavazzana, D. F. Escande, A. Fassina, P. Franz, M. Gobbin, P. Innocente, L. Marrelli, R. Pasqualotto, M. E. Puiatti, M. Spolaore, M. Valisa, N. Vianello, and P. Martin, "Self-organized helical equilibria as a new paradigm for ohmically heated fusion plasmas," *Nature Physics*, vol. 5, pp. 570–574, aug 2009. [Online]. Available: https://doi.org/10.1038/nphys1308

[16] I. Predebon, L. Marrelli, R. B. White, and P. Martin, "Particle-Transport Analysis in Reversed Field Pinch Helical States," *Physical Review Letters*, vol. 93, no. 14, 2004. [Online]. Available: https://doi.org/10.1103/PhysRevLett.93.145001

[17] S. Menmuir, L. Carraro, A. Alfier, F. Bonomo, A. Fassina, G. Spizzo, and N. Vianello, "Impurity transport studies in RFX-mod multiple helicity and enhanced confinement QSH regimes," *Plasma Physics and Controlled Fusion*, vol. 52, no. 095001, 2010. [Online]. Available: http://dx.doi.org/10.1088/0741-3335/52/9/095001

[18] J. K. Anderson, W. Capecchi, S. Eilerman, J. J. Koliner, M. D. Nornberg, J. A. Reusch, J. S. Sarff, and L. Lin, "Fast ion confinement in the three-dimensional helical reversed-field pinch," *Plasma Physics and Controlled Fusion*, vol. 56, no. 9, 2014. [Online]. Available: https://doi.org/10.1088/0741-3335/56/9/094006

[19] P. Bonofiglo, M. Gobbin, D. A. Spong, J. Boguski, E. Parke, J. Kim, and J. Egedal, "Fast ion transport in the quasi-single helical reversed-field pinch," *Physics of Plasmas*, vol. 022502, no. 26, 2019. [Online]. Available: http://dx.doi.org/10.1063/1.5084059

[20] S. Munaretto, B. E. Chapman, D. J. Holly, M. D. Nornberg, R. J. Norval, D. J. Den Hartog, J. A. Goetz, and K. J. McCollam, "Control of 3D equilibria with resonant magnetic perturbations in MST," *Plasma Physics and Controlled Fusion*, vol. 57, no. 10, p. 104004, 2015. [Online]. Available: http://dx.doi.org/10.1088/0741-3335/57/10/104004

[21] J. M. Finn, R. Nebel, and C. Bathke, "Single and multiple helicity Ohmic states in reversed-field pinches," *Physics of Fluids B: Plasma Physics*, vol. 4, no. 5, p. 1279, 1992. [Online]. Available: https://doi.org/10.1063/1.860082

[22] S. Cappello and D. F. Escande, "Bifurcation in Viscoresistive MHD: The Hartmann Number and the Reversed Field Pinch," *Physical Review Letters*, vol. 85, no. 18, pp. 3838–3841, 2000. [Online]. Available: https://doi.org/10.1103/physrevlett.85.3838

[23] S. Cappello, D. Bonfiglio, and D. F. Escande, "Magnetohydrodynamic dynamo in reversed field pinch plasmas: Electrostatic drift nature of the dynamo velocity field," *Physics of Plasmas*, vol. 13, no. 56102, 2006. [Online]. Available: https://doi.org/10.1063/1.2177198

[24] P. Piovesan, D. Craig, L. Marrelli, S. Cappello, and Martin P., "Measurements of the MHD Dynamo in the Quasi-Single-Helicity Reversed-Field Pinch," *Physical Review Letters*, vol. 93, no. 235001, 2004. [Online]. Available: doi.org/10.1103/PhysRevLett.93.235001

[25] P. Piovesan, D. Bonfiglio, M. Cianciosa, T.C. Luce, N.Z. Taylor, D. Terranova1, F. Turco, R.S. Wilcox, A. Wingen, S. Cappello, C. Chrystal, D.F. Escande, C.T. Holcomb, L. Marrelli, C. Paz-Soldan, L. Piron, I. Predebon, B. Zaniol, The DIII-D Team, and RFX-Mod Team, "Role of a continuous MHD dynamo in the formation of 3D equilibria in fusion plasmas Self-organized 3D equilibrium formation and its

feedback control in RFX-mod," *Nuclear Fusion*, vol. 57, no. 076014, 2017. [Online]. Available: https://doi.org/10.1088/1741-4326/aa700b

[26] J. Boguski, "Local Ion Velocity Measurements in the MST Saturated Single Helical Axis State," Ph.D. dissertation, University of Wisconsin-Madison, 2019.

[27] C. Sovinec, A. Glasser, T. Gianakon, D. Barnes, R. Nebel, S. Kruger, D. Schnack, S. Plimpton, A. Tarditi, and M. Chu, "Nonlinear magnetohydrodynamics simulation using high-order finite elements," *Journal of Computational Physics*, vol. 195, pp. 355–386, 2004. [Online]. Available: https://doi.org/10.1016/j.jcp.2003.10.004

[28] J. Boguski, M. D. Nornberg, A. F. Almagri, D. Craig, U. Gupta, K. J. McCollam, T. Nishizawa, J. S. Sarff, C. R. Sovinec, P. W. Terry, and Z. A. Xing, "Direct measurements of the 3D plasma velocity in Single-Helical-Axis RFP plasmas," *Physics of Plasmas*, vol. (in prep), 2020.

[29] D. Bonfiglio, M. Veranda, S. Cappello, D. F. Escande, and L. Chacón, "Experimental-like Helical Self-Organization in Reversed-Field Pinch Modeling," *Physical Review Letters*, vol. 111, no. 085002, aug 2013. [Online]. Available: http://dx.doi.org/10.1103/PhysRevLett.111.085002

[30] Marco Veranda, Daniele Bonfiglio, Susanna Cappello, Giovanni di Giannatale, and Dominique Frank Escande, "Helically self-organized pinches: dynamical regimes and magnetic chaos healing," *Nuclear Fusion*, vol. 60, no. 016007, 2020. [Online]. Available: https://doi.org/10.1088/1741-4326/ab4863

[31] J.-H. Kim and P. W. Terry, "Magnetic turbulence suppression by a helical mode in a cylindrical geometry," *Physics of Plasmas*, vol. 19, no. 122304, pp. 1–11, 2012. [Online]. Available: https://doi.org/10.1063/1.4769369

[32] P. W. Terry and G. G. Whelan, "Time-dependent behavior in a transport-barrier model for the quasi-single helcity state," *Plasma Physics and Controlled Fusion*, vol. 56, no. 9, 2014. [Online]. Available: https://doi.org/10.1088/0741-3335/56/9/094002

[33] I. J. McKinney and P. W. Terry, "Thermal transport dynamics in the quasi-single helicity state," *Physics of Plasmas*, vol. 24, no. 6, 2017. [Online]. Available: http://dx.doi.org/10.1063/1.4985317

[34] P. W. Terry, "Suppression of turbulence and transport by sheared flow," *Reviews of Modern Physics*, vol. 72, no. 1, pp. 109–165, jan 2000. [Online]. Available: https://doi.org/10.1103/RevModPhys.72.109

[35] E.-J. Kim and P. H. Diamond, "Zonal Flows and Transient Dynamics of the L-H Transition," *Physical Review Letters*, vol. 90, no. 185006, 2003. [Online]. Available: doi.org/10.1103/PhysRevLett.90.185006

[36] G. D. Conway, C. Angioni, F. Ryter, P. Sauter, and J. Vicente, "Mean and oscillating plasma flows and turbulence interactions across the L-H confinement

transition," *Physical Review Letters*, vol. 106, no. 6, pp. 1–4, 2011. [Online]. Available: https://doi.org/10.1103/PhysRevLett.106.065001

[37] L. Schmitz, L. Zeng, T. L. Rhodes, J. C. Hillesheim, E. J. Doyle, R. J. Groebner, W. A. Peebles, K. H. Burrell, and G. Wang, "Role of zonal flow predator-prey oscillations in triggering the transition to H-mode confinement," *Physical Review Letters*, vol. 108, no. 15, pp. 1–5, 2012. [Online]. Available: https://doi.org/10.1103/PhysRevLett.108.155002

[38] S. Van Der Walt, S. C. Colbert, and G. Varoquaux, "The NumPy array: A structure for efficient numerical computation," *Computing in Science and Engineering*, vol. 13, no. 2, pp. 22–30, 2011. [Online]. Available: https://numpy.org

[39] D. J. Clayton, B. E. Chapman, R. O'Connell, A. F. Almagri, D. R. Burke, C. B. Forest, J. A. Goetz, M. C. Kaufman, F. Bonomo, P. Franz, M. Gobbin, and P. Piovesan, "Observation of energetic electron confinement in a largely stochastic reversed-field pinch plasma," *Physics of Plasmas*, vol. 17, no. 012505, 2010. [Online]. Available: https://doi.org/10.1063/1.3292658

[40] A. M. DuBois, A. F. Almagri, J. K. Anderson, D. J. Den Hartog, J. D. Lee, and J. S. Sarff, "Anisotropic Electron Tail Generation during Tearing Mode Magnetic Reconnection," *Physical Review Letters*, vol. 118, no. 075001, pp. 1–5, 2017. [Online]. Available: https://doi.org/10.1103/PhysRevLett.118.075001

[41] J. D. Callen, *Plasma kinetic theory*, K. J. Bunkers and J. S. Kollasch, Eds. Madison: Self-published, 2015. [Online]. Available: http://homepages.cae.wisc.edu/{~}callen/plasmas.html

[42] D. J. Den Hartog, J. T. Chapman, . D. Craig, G. Fiksel, P. W. Fontana, S. C. Prager, and J. S. Sarff, "Measurement of core velocity fluctuations and the dynamo in a reversed-field pinch," *Physics of Plasmas*, vol. 6, no. 5, pp. 1813–1821, 1999. [Online]. Available: https://doi.org/10.1063/1.873439

[43] F. Auriemma, P. Zanca, W. F. Bergerson, B. E. Chapman, W. X. Ding, D. L. Brower, P. Franz, P. Innocente, R. Lorenzini, B. Momo, and D. Terranova, "Magnetic reconstruction of nonaxisymmetric quasi-single-helicity configurations in the Madison Symmetric Torus," *Plasma Physics and Controlled Fusion*, vol. 53, no. 10, 2011. [Online]. Available: https://doi.org/10.1088/0741-3335/53/10/105006

[44] M. Galante, L. Reusch, D. Den Hartog, P. Franz, J. Johnson, M. McGarry, M. Nornberg, and H. Stephens, "Determination of Z_eff by integrating measurements from x-ray tomography and charge exchange recombination spectroscopy," *Nuclear Fusion*, vol. 55, no. 12, p. 123016, 2015. [Online]. Available: https://doi.org/10.1088/0029-5515/55/12/123016

[45] L. M. Reusch, P. Franz, D. J. Den Hartog, J. A. Goetz, M. D. Nornberg, and P. VanMeter, "Model validation for quantitative X-ray measurements," *Fusion*

Science and Technology, vol. 74, no. 1-2, pp. 167–176, 2018. [Online]. Available: https://doi.org/10.1080/15361055.2017.1404340

[46] J. Johnson, "Implementing Bayesian Statistics and a Full Systematic Uncertainty Propagation with the Soft X-Ray Tomography Diagnostic on the Madison Symmetric Torus (undergraduate thesis)," University of Wisconsin-Madison, Madison, Tech. Rep., 2013.

[47] E. Jones, T. Oliphant, P. Peterson, and Others, "SciPy: Open source scientific tools for Python," 2001. [Online]. Available: http://www.scipy.org/

[48] J. R. Duff, Z. R. Williams, D. L. Brower, B. E. Chapman, W. X. Ding, M. J. Pueschel, J. S. Sarff, and P. W. Terry, "Observation of trapped-electron-mode microturbulence in reversed field pinch plasmas," *Physics of Plasmas*, vol. 25, no. 010701, 2018. [Online]. Available: https://doi.org/10.1063/1.5010198

[49] M. Gobbin, D. Bonfiglio, D. F. Escande, A. Fassina, L. Marrelli, A. Alfier, E. Martines, B. Momo, and D. Terranova, "Vanishing Magnetic Shear And Electron Transport Barriers In The RFX-Mod Reversed Field Pinch," 2011. [Online]. Available: https://doi.org/10.1103/PhysRevLett.106.025001

[50] D. J. Den Hartog, A. F. Almagri, J. T. Chapman, H. Ji, S. C. Prager, J. S. Sarff, F. J. Fonck, and C. C. Hegna, "Fast flow phenomena in a toroidal plasma," *Physics of Plasmas*, vol. 2, no. 6, pp. 2281–2285, jun 1995. [Online]. Available: https://doi.org/10.1063/1.871250

[51] J. R. Duff, "Observation of trapped-electron mode microturbulence in improved confinement reversed-field pinch plasmas," Ph.D. dissertation, University of Wisconsin-Madison, 2018.

Chapter 7

Summary and future work

A versatile new multi-energy soft x-ray (ME-SXR) diagnostic has been developed as part of a collaboration between UW-Madison and PPPL. It has been calibrated, installed, and tested on the MST for a wide range of plasma conditions. Its ability to accurately infer core T_e and n_Z has been demonstrated. It serves as a compliment to an existing suite of x-ray diagnostics, including a diode-base soft x-ray tomography array, a Ross spectrometer, and a hard x-ray camera. Together with FIR interferometry, these diagnostics were used to explore the temporal evolution of the saturated quasi-single helicity state in the MST. Time-resolved 2D T_e profiles of an MST QSH plasma were produced for the first time. Clear evidence was found of a brief period of significantly enhanced confinement.

Section 7.1 reviews the progress that has been made by the ME-SXR diagnostic program at UW-Madison. Section 7.2 reviews the physics results presented in this thesis. Finally, Section 7.3 discusses the remaining loose threads and recommends ideas for future work which naturally follow from the results presented in this thesis.

7.1 ME-SXR: A versatile new soft x-ray diagnostic

The ME-SXR diagnostic is based on a novel calibration PILATUS3 100K hybrid photon counting detector. The lower threshold for each of the detector's nearly one-hundred thousand pixels can be independently adjusted *in-situ*, permitting wide range of applications. Previously tested on C-Mod, the ME-SXR program was expanded to the MST in order to continue testing and developing the technology. MST plasmas provide ideal conditions due to their relatively high temperatures, novel magnetic configuration, and native presence of mid-Z impurities (aluminum).

An energy calibration was performed to determine the mapping between pixel trimbit settings and the corresponding energy sensitivity. This was done by exposing the detector to multiple sources of known photon energy, generated by x-ray fluorescence of a specified target, and scanning the trimbit setting. Using the calibration results, the pixel-to-pixel variation across the detector was examined. Since the trimbit setting must take an integer value between 0-63, some deviation from the requested lower threshold is unavoidable. This deviation was found to be $\Delta E < 100$ eV for the high gain (lowE) settings and $\Delta E < 200$ eV for standard gain (midE) settings. The pixel's sensitivity is well-modeled by an S-curve, and the widths were found to be $\sigma_E = 300$ eV for lowE settings and $\sigma_E = 550$ eV for midE settings. A simple model for charge sharing, the phenomenon which occurs when the energy from an absorbed photon is split between two adjacent pixels, was validated for both threshold ranges.

The ME-SXR diagnostic was installed on the MST in the Spring of 2018. A spatial calibration using an insertable probe tip with an Fe-55 source was performed in order to determine the line-of-sight of each pixel looking through the pinhole. The detector was then commissioned for routine operation with plasmas. Multiple configurations were tested, including high signal-to-noise 1D imaging, high spatial resolution 1D imaging, and energy-resolved 2D imaging. The custom calibration was tested by doping the plasma with argon gas, resulting in an increase in signal for thresholds below the 3 keV emission lines. Sensitivity to runaway electrons, which feature strong emissions in the lower hard x-ray range (~ 10 keV), was also demonstrated.

Images collected during the initial phases of ME-SXR operation showed noticeable and consistent distortions. It was eventually determined that this was due to pulse-pileup resulting from the extremely high flux of photons from Al^{+11} and Al^{+12} transition lines at $\sim$ 2 keV. Even when the chosen threshold is set high, pulses from multiple low-energy photons can pileup resulting in an apparent single high-energy photon being counted. This was resolved by installing a Mylar filter which cuts the transmission at 2 keV down by an order of magnitude. Initially a 50 μm Mylar filter was used, but it was found that sufficient numbers of low-energy photons were still transmitted as to distort the apparent spectrum when using high-gain settings. Upgrading to a 100 μm Mylar filter resolved this issue.

A sophisticated diagnostic forward model was developed to produce synthetic ME-SXR measurements given a set of plasma input profiles. This model incorporates information from the spatial and energy calibrations as well as atomic physics modeling from ADAS to simulate the underlying physics, geometry, and detector response. Charge-exchange with the neutral hydrogen population was shown to be important when simulating MST plasmas. A systematic uncertainty in the model of $\sigma_m \approx 18\%$ was estimated, due mostly to the uncertainty in the Mylar filter thickness.

The ME-SXR diagnostic was used to observe plasma temporal evolution. Photon counts were seen to evolve appropriately as the plasma heats up, and decrease as the plasma cools. A method for identifying the energy of an emission line was presented, and tested using Ar-doped plasmas. A technique for determining the electron temperature profile by directly inverting the emissivity was shown to agree with Thomson scattering measurements up to an $\sim$ 180 eV offset, which may be due to the presence of background hard x-rays from the runaway population known to develop during PPCD. The ME-SXR diagnostic was also incorporated into an integrated data analysis (IDA) framework based on Bayesian inference. Using this methodology, it was shown to be possible to simultaneously extract T_e and n_Z profiles from ME-SXR and Thomson scattering data. The resulting profiles can be further refined by adding the SXR tomography and NICKAL2 diagnostics into the analysis. The ability to simultaneously fit multiple diagnostics to their respective data using physically-reasonable profiles provides a high level of confidence for the ME-SXR forward model.

7.2 Physics results

The 2D structure and evolution of the T_e profile have been studied for the first time in a QSH plasma on the MST. This was permitted by recent advances in x-ray diagnostic modeling which allowed the SXR tomography diagnostic to be operated with thinner beryllium filters. A brief self-organized period of significantly enhanced thermal confinement was observed during a brief "quiet" phase at the start of the QSH flattop. The good confinement region was observed to contract, but not vanish, during the "dynamic" phase of the flattop, defined by the resumption of (small-amplitude) secondary tearing mode activity.

Large populations of runaway electrons were seen to form during low density QSH plasmas. The runaway population energy was seen to peak at $E_r = 18.2$ keV during the rising phase and remain high at $E_r = 15.3$ keV during the quiet flattop phase before dropping substantially to $E_r = 5.0$ keV during the dynamic flattop phase. This reinforces the view that a broad region of restored flux surfaces exist in the core during the quiet phase, but are rapidly lost as tearing mode activity resumes. These suddenly-liberated fast electrons then rapidly stream to the wall, where they create a burst of target emission. The runaway x-ray spectrum was also seen to be spatially uniform.

The physics of the transition to, and persistence of, the saturated quasi-single helicity state is not well-understood. A theoretical model put forward by Terry, *et al.* proposes that magnetic or velocity flow shear de-correlates the turbulent eddies and suppresses the transfer of energy between the core-resonant dominant and secondary tearing modes. This allows energy to accumulate in the dominant mode, driving it to large amplitudes. The proposed mechanism is similar to the L-H transition in tokamaks, which is governed by flow shear. This model also predicts a predator-prey relationship between the dominant and secondary modes, as well as the formation of a thermal transport barrier which oscillates along with the dominant mode amplitude. Direct evidence of a predator-prey relationship between $\tilde{b}_5$ and secondary mode amplitudes was provided by ensembling over small fluctuations in the QSH flattop. The fluctuations were observed to be nearly 180° out of phase. SXR tomography measurements were used to establish an in-phase relationship between the emissivity gradient $\nabla\varepsilon$ and $\tilde{b}_5$. This is strongly suggestive of the transport barrier dynamics predicted by the model. Both of these observations represents successful tests of the model and provide some confidence that the underlying mechanism (reduction of mode coupling due to shear) is an important part of QSH physics.

A novel analysis based on Bayesian inference, constrained by SXR tomography and FIR interferometry data, has allowed the observation of time-resolved T_e profiles during QSH.The T_e profile is seen to broaden during the quiet phase. Simultaneously, the n_{Al} concentration increases, and a modest increase in n_e can be seen. At the transition to the dynamic phase the hot spot rapidly collapses down to the helical core, while n_e is largely unperturbed. The helical hot spot remains through the rest of the plasma discharge. These dynamics show that T_e confinement is strongly linked to secondary mode amplitude, as predicted by the Terry, *et al.* model. It also shows strong evidence of the formation of an internal thermal transport barrier, with steep gradients developing around the hot core.

Unexpected fluctuations in the magnetic spectrum were observed to exist in the 400-1000 kHz range during high-performance QSH plasmas. These fluctuations were found to be highly localized, and could only be detected when the helical O-point was positioned near the toroidal array coils. However, some irregularities were seen in this $n = 5$ pattern, and the fluctuations could not be detected on the dense array. The wavenumber range could not be determined due to aliasing, but a rough estimate puts the mode in the $n = 200 - 300$ range. This data is reminiscent of microtearing modes in RFX-mod, which have been predicted to be unstable using gyrokinetic modeling. These are very short-wavelength fluctuations driven by steep gradients in the electron temperature profile, as have previously been observed to develop in RFX-mod and have now been observed in the MST. However, more work will be needed to confirm the identity of these fluctuations.

7.3 Suggestions for future work

The next phase of the ME-SXR project is currently underway at the WEST (Tungsten Environment Steady-state Tokamak) facility in Cadarache, France [1]. WEST is a long-pulse tokamak which aims to sustain plasma discharges up to 1000 seconds. This long duration is highly favorable for the detector's $\sim$ 500 Hz data collection rate. A hard x-ray variation of the multi-energy concept (ME-HXR), which incorporates a cadmium telluride-based PILATUS detector, is currently under development for WEST [2]. An Si-based ME-SXR diagnostic like the one described in this thesis will also be installed. This system will be used to simultaneously diagnose plasma thermal properties while also tracking core impurity accumulation, which is critical for sustaining the plasma in the tungsten-rich ITER-like environment. It is possible that two Si-based PILATUS detectors may be installed, in which case tomographic inversion techniques can also be tested. Aside from WEST, the ME-SXR is also being used in the MST to study runaway electron generation and suppression in tokamak plasmas.

There is still work to be done on ME-SXR diagnostic development. The most pressing issue is to measure the saturation behavior of the PILATUS detector using the custom calibrations. This will require a much higher-flux x-ray source than was used to perform the calibration trimbit scans, and requires a secondary measurement of the true photon count rate. Such a measurement would be invaluable, though, for determining whether the measurements are being affected by dead time or pulse-pileup. It would also be worthwhile to develop a physics-based model of the charge-sharing behavior. This could be used to determine the limits of the heuristic model presented here and to better account for the effect in a forward model.

The ME-SXR model presented in Chapter 4, while thorough, assumes that the plasma electron distribution function is Maxwellian. This is not always a good assumption, as seen with the buildup of runaway populations during some phases of QSH. It should be possible to develop a model a model based on a non-Maxwellian distribution with a high-energy tail, such as the kappa distribution. Although it would be challenging, significant progress has already been made modeling non-Maxwellian emission spectra in the context of astrophysics [3]. This could plausibly be adapted into a fusion plasma model, allowing for the simultaneous diagnosis of thermal and non-thermal properties. The greatest difficulty would likely be modeling the non-Maxwellian rate coefficients $\langle \sigma v \rangle$ for mid-to-high Z impurities, which are not commonly found in the astrophysical setting.

It would be worthwhile to continue improve the characterization of thermal transport during helical MST equilibria. The most direct task remaining is to make a direct estimate of the χ_e profile, thereby proving that the observed structure is indeed indicative of a thermal transport barrier. It would also be worthwhile to reduce the uncertainty of of the inferred T_e profiles in order to more directly correlate transport barrier strength (∇T_e) with tearing mode amplitude during flattop fluctuations. This could most readily be accomplished by incorporating Thomson scattering measurements, though this can be challenging due to the low n_e required to achieve a high-performance QSH/SHAx state. A more thorough study of impurity transport during QSH would also be interesting. Are the impurity ion density profiles broad, like n_e, or do they peak in the helical core? Finally, it would be of great interest to determine the physical mechanism behind the quiet phase, and to determine whether it can be extended by some form of active control. Measurements of flow velocity using ion Doppler spectroscopy could help determine whether the sudden increase in flow shear immediately following mode locking is a sufficient explanation.

Finally, additional work will be needed to understand the source of the high-frequency magnetic fluctuations observed when the helical structure is oriented near the $\dot{b}$ coil. Are these the result of a microtearing instability resulting from the formation of steep temperature gradients, or maybe TEM driven by ∇n_e (as in PPCD)? What is wavenumber is associated with the mode? Why are the fluctuations not not visible in dense array measurements? A good next step is use the full 64 b_θ coils in the toroidal array, decreasing the distance between adjacent coils and thereby increasing bandwidth for k. The lack of a signature in dense array measurements should also be resolved. Comparisons can be made with gyrokinetic modeling, although improved edge T_e measurements may be needed to reduce the uncertainty in the normalized gradient $a/L_T = a\nabla T_e/T_e$.

Bibliography

[1] C. Bourdelle, J. Artaud, V. Basiuk, M. Bécoulet, S. Brémond, J. Bucalossi, H. Bufferand, G. Ciraolo, L. Colas, Y. Corre, X. Courtois, J. Decker, L. Delpech, P. Devynck, G. Dif-Pradalier, R. Doerner, D. Douai, R. Dumont, A. Ekedahl, N. Fedorczak, C. Fenzi, M. Firdaouss, J. Garcia, P. Ghendrih, C. Gil, G. Giruzzi, M. Goniche, C. Grisolia, A. Grosman, D. Guilhem, R. Guirlet, J. Gunn, P. Hennequin, J. Hillairet, T. Hoang, F. Imbeaux, I. Ivanova-Stanik, E. Joffrin, A. Kallenbach, J. Linke, T. Loarer, P. Lotte, P. Maget, Y. Marandet, M. Mayoral, O. Meyer, M. Missirlian, P. Mollard, P. Monier-Garbet, P. Moreau, E. Nardon, B. Pégourí, Y. Peysson, R. Sabot, F. Saint-Laurent, M. Schneider, J. TravèreTrav, E. Tsitrone, S. Vartanian, L. Vermare, M. Yoshida, R. Zagorski, and JET Contributors, "WEST Physics Basis," *Nuclear Fusion*, vol. 55, no. 063017, 2015. [Online]. Available: http://iopscience.iop.org/0029-5515/

[2] T. Barbui, N. Pablant, C. Disch, B. Luethi, N. Pilet, B. Stratton, and P. VanMeter, "Multi-energy calibration of a PILATUS3 CdTe detector for hard x-ray measurements of magnetically confined fusion plasmas (forthcoming)," in *Proceedings of the 23rd Topical Conference on High-Temperature Plasma Diagnostics*. Santa Fe, NM: American Institute of Physics, 2021.

[3] J. Dudík, J. Kašparová, E. Dzifčáková, M. Karlický, and S. MacKovjak, "The non-Maxwellian continuum in the X-ray, UV, and radio range," *Astronomy and Astrophysics*, vol. 539, pp. 1–12, 2012. [Online]. Available: https://doi.org/10.1051/0004-6361/201118345

Appendix A

Python ME-SXR code

All of the code that I developed for ME-SXR operation, calibration, and modeling is available for future reference at https://github.com/pdvanmeter/meSXR. Some modifications may be requiered to integrate the code with local systems, MDSplus tree implementations, etc. Also note that I have not included the emissivity databases due both to their prohibitive size and because they are subject to the licensing agreement governing ADAS.

This appendix contains the Python module that I wrote to arm the detector for data collection on a runday. Although it is in the repository, I have reproduced it here because it provides a good example of how to interface with and operate the PILATUS3 detector. In particular I expect that future users will find the *camserver* class to be useful. This provides a relatively straightforward example of how to communicate with the PILATUS3 via socket connection. An example of the code's usage is also provided.

```
#!/usr/bin/env python
# -*- coding: utf-8 -*-
"""
Package: mesxr.operation
Module: camera
Author: Patrick VanMeter
Affiliation: Department of Physics, University of Wisconsin-Madison
Last Updated: January 2019

Description:
    This module is an updated version of the camera operations code. This
    version has been modified to work with the new naming conventions for
    pixel configurations and contains other feature updates.

    This code is made specifically for the PILATUS3 100K-M and might require
```

```python
    some modifications for use with other PILATUS models.
Usage:
    TBD
Acknowledgements:
    - Novimir Pablant for the original IDL code this module is loosely based upon.
    - Luis Felipe Delgado-Aparicio, for heading the PPPL/MST collaboration.
    - Daniel Den Hartog and Lisa Reusch, for advising me.
"""
import socket
import time
import os
import re
import numpy as np
import tifffile as tif
import MDSplus
import mesxr.calibration.utilities as util

# Module-level constants for the camera - these should eventually be read in
COMP_NAME = 'dec1424'
PORT = 41234
IP_ADDR = '127.0.0.1'
NUM_CHIPS = 16
M_SIZE_Y = 195
M_SIZE_X = 487
BASE_IMAGE_PATH = '/home/det/p2_det/images/MST_data'
TAKE_DATA_ADDR = 'aurora.physics.wisc.edu'

class camserver():
"""
An instance of this object is helpful in facilitating remote operation
of the camera via the camserver. This class is currently under development.
"""
    def __init__(self):
        self.sock = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
        self.online = False
        self.set_output_mode('verbose')
        self.set_timeout(None)
        self.connect(IP_ADDR, PORT)

    def __del__(self):
        if self.online:
            print('Error detected. Automatically closing socket connection.')
            self.disconnect()

    def connect(self, ip_addr, port):
        """
        Establish a connection to the PILATUS detector.
        """
        try:
            self.sock.connect((ip_addr, port))
            self.online = True
            print('Camserver connection established with ' + COMP_NAME + '.')
        except:
            print('Camserver connection cannot be established. Check settings.')
            self.online = False

    def disconnect(self):
        """
        Close the connection with the PILATUS detector.
        """
        if self.online:
            self.sock.close()
            self.online = False
            print('Disconnected from the camserver.')
        else:
            print('Alread disconnected from the camserver.')

    def execute(self, command):
        """
        Exectute a command on the camserver and wait for a return code.
        """
        if not self.quiet:
            print('CMD >> ' + command)

        if self.online:
```

```
            self.sock.send(command + '\n')
            code, message = self.recieve()
        else:
            code = -1
            message = 'OFFLINE'
            print('Must be online to execute commands.')

        if not self.quiet:
            print('p3det >> (' + str(code) + ') ' + message)

        return code, message

    def wait(self, end_code):
        """
        Wait until a specified end code is returned.
        """
        code, message = self.recieve()
        while code != end_code and code != -1:
            code, message = self.recieve()

        if not self.quiet:
            print('p3det >> (' + str(code) + ') ' + message)

        return code, message

    def recieve(self):
        """
        Read the next socket output, out to the \x18 message termination code.
        Consider implementing a timeout function if this becomes an issue.
        """
        if self.online:
            end_of_message = False
            message = ''
            while not end_of_message:
                next_char = self.sock.recv(1)
                if next_char == '\x18':
                    end_of_message = True
                else:
                    message += next_char

            # Remove the numerical code from the start of the message
            code = int(message.split()[0])
            message = ' '.join(message.split()[1:])
        else:
            code = -1
            message = 'OFFLINE'
            print('Must be online to communicate with the camserver.')

        return code, message

    def set_timeout(self, timeout):
        self.timeout = timeout
        self.sock.settimeout(self.timeout)
        print('Camserver timeout set to ' + str(timeout) + '.')

    def set_output_mode(self, mode):
        """
        Set mode to either "verbose" or "quiet". This controls whether camera
        commands and returned messages are output to the console.
        """
        if mode.lower() == 'quiet':
            self.quiet = True
        elif mode.lower() == 'verbose':
            self.quiet = False

# ---------------------------------------------------------------------------

def arm_detector(shot=0, mode='cycle', load_mds=False, write_mds=True,
                 quiet=True, timeout=None, retrigger=False, rate_corr=True,
                 acq_time=10, config_name='ppcd_8_color', calibration_name='midE',
                 n_frames=30, exp_period=0.002, exp_time=0.001, delay=0,
             image_path='/home/det/p2_det/images/MST_data'):
"""
Description
===============
```

```
    The main routine of
Parameters:
===============
- shot = (int) The shot number corresponding to the given plasma discharge. This
    is used in the resulting image filename.
- mode = (string) One of the following strings:
    'trigger' - Exposure is initiated by external trigger one time.
    'manual' - Exposure initiated by this script one time.
    'gate' - Exposure is controlled by external gate signal one time.
    'cycle' - Exposure is initiated by external trigger repeatedly. The shot number
        is also overridden by the current shot in the MDSplus tree. This is the mode
        that should be set at the beginning of a run day. Use a keyboard interrupt
        (CRTL+C) to end the loop and close the camserver connection.
- load_mds = (bool) Set to true to load default exposure settings from the MDSplus
    tree. This overrides any values set manually.
- write_mds = (bool) When True ME-SXR data is written to the MDSplus tree. When False
        data is only written to the specified output directory.
- quiet = (bool) Set to true to print all camserver outputs to the console.
- timeout = (float) Set the camera to time out if an exposure is not taken within the
        specified number of seconds. Default is None, which is no limit.
- retrigger = (bool) Set to True to enable instant retrigger technology. This
    increases the max count rate but may cause problems with the
    calibration.
- rate_corr = (bool) Enable to use the rate correction files loaded during
    calibration. This is not needed used running in the 'dectris'
    calibration mode.
- acq_time = (int) Time in seconds to wait while cycling after data is written in the
    tree before reading in the next shot number. This may need to be extended
    if a long many diagnostics are running.
- calibration_name = (string) The name used for the desired calibration.
    Use 'dectris' to load factory calibrations
- config_name = (string) The name associated with the desired trimbit configuration.
    For the 'dectris' calibration this is the command passed to the
    camserver (e.g. 'setthr 4000').
- n_frames = (int) The number of exposures to take once initialized.
- exp_period = (double) The cycle time for an exposure period, in seconds. This is
      the time from the start of one exposure until the start of the next exposure.
- exp_time = (double) The length of time for each exposure, in seconds. This should
      be shorter than exp_period.
- image_path = (str) Change this to save files in a different directory.
    This is generally useful for testing.
    Returns:
===============
    - success = (bool) True if the function executes successfully, false if
      any errors are thrown.
"""
    if load_mds:
        print('Loading default exposure settings from the MDSplus tree.')
        try:
            mesxr = MDSplus.Tree('mst_me_sxr', shot=-1)
            n_frames = mesxr.getNode(r'.CONFIG:N_IMAGES').getData().data()
            exp_time = mesxr.getNode(r'.CONFIG:EXPOSUR_TIME').getData().data()
            exp_period = mesxr.getNode(r'.CONFIG:EXPOSUR_PER').getData().data()
            delay = mesxr.getNode(r'.CONFIG:DELAY').getData().data()

            print('n_frames: ' + str(n_frames))
            print('exp_time: ' + str(exp_time))
            print('exp_period: ' + str(exp_period))
            print('delay: ' + str(delay))
        except:
            print('ERROR: Could not load settings from the model tree.')

    # Make sure the exposure period is permissible
    if exp_period <= exp_time + 0.001:
        exp_period = exp_time + 0.001

    filename = 's' + str(shot) + '.tif'
    settrims_prefix = 'autotrim_'
    setdacs_filename = 'setdacs_b01_m01.dat'
    base_shot = int(shot/1000)

    config_path = os.path.join('/home/det/meSXR/configurations/', calibration_name,
    config_name)
    setdacs_path = os.path.join(config_path, setdacs_filename)

```

```
    # Write settings to the model tree, if desired
    if write_mds:
        set_model_tree(n_frames, exp_period, exp_time, delay, config_path,
    setdacs_path)

    # Camserver commands
    command_nimages = 'nimages ' + str(n_frames)
    command_expperiod = 'expperiod ' + str(exp_period)
    command_exptime = 'exptime ' + str(exp_time)
    command_delay = 'delay ' + str(delay)

    # Select camserver exposure mode - see function description
    if mode.lower() == 'trigger':
        command_exposure = 'ExtTrigger ' + filename
    elif mode.lower() == 'manual':
        command_exposure = 'Exposure ' + filename
    elif mode.lower() == 'gate':
        command_exposure = 'ExtEnable ' + filename
    elif mode.lower() == 'cycle':
        command_exposure = 'ExtTrigger ' + filename
    else:
        print('ERROR: Mode selection not recognized. Defaulting to manual.')
        command_exposure = 'Exposure ' + filename

    # Extablish the connection to the camserver
    try:
        print('Attempting to open socket connection on ' + COMP_NAME + '.')
        cam = camserver()

        # Configure client settings
        if quiet:
            cam.set_output_mode('quiet')

        cam.set_timeout(timeout)

        if calibration_name == 'dectris':
            cam.execute('setCu')
            cam.execute(config_name)

        else:
            # Set up camera settings
            cam.execute('dacoffset off')
            cam.execute('LdCmndFile ' + setdacs_path)

            # Load the trimbit configuration
            for index in range(NUM_CHIPS):
                settrims_filename = settrims_prefix
                settrims_filename += 'b01_m01_c{0:02d}.dat'.format(index)
                cam.execute('prog b01_m01_chsel 0x0')
                cam.execute('trimfromfile ' + os.path.join(config_path,
                  settrims_filename))

            # Reset the readout chip selection
            cam.execute('prog b01_m01_chsel 0xffff')

            # Manually disable flat-field correction
            cam.execute('ldflatfield 0')

            # Set the pixel rate correction
            if rate_corr:
                print('Enabling rate correction.')
                cam.execute('ratecorrlutdir ContinuousStandard_v1.1')

        # Set up remaining exposure settings
        cam.execute(command_nimages)
        cam.execute(command_expperiod)
        cam.execute(command_exptime)

        if delay != 0:
            cam.execute(command_delay)

        # Disable instant retrigger
        if not retrigger:
            cam.execute('setretriggermode 0')

```

```
        if mode == 'cycle':
            # To get the current shot from Aurora
            conn = MDSplus.Connection(TAKE_DATA_ADDR)

            cycle = True
            while cycle:
                try:
                    # Get the next shot number from the tree
                    shot = int(conn.get('current_shot("mst")')) + 1
                    base_shot = int(shot/1000)
                    filename = 's' + str(shot) + '.tif'
                    command_exposure = 'ExtTrigger ' + filename

                    output_path = os.path.join(image_path, str(base_shot), str(shot))
                    command_imgpath = 'imgpath ' + output_path
                    cam.execute(command_imgpath)

                    # Write out the trimbit configuration
                    cam.execute('imgmode p')
                    cam.execute('imgonly trimbit.tif')

                    print('Waiting for trigger for shot ' + str(shot) + '.')
                    cam.execute(command_exposure)
                    cam.wait(7)
                    print('Data collected for shot ' + str(shot) + '.')

                    # Write data to tree if desired
                    if write_mds:
                        write_to_tree(shot, n_frames, exp_period, exp_time,
                          delay, output_path)

                    # Wait an appropriate amount of time for the shot to increment
                    time.sleep(acq_time)

                except KeyboardInterrupt:
                    cam.execute('k ')
                    print('\nME-SXR trigger manually disarmed. Cylcing stopped.')
                    cycle = False
        else:
            # Use the supplied shot number for a single exposure
            output_path = os.path.join(image_path, str(base_shot), str(shot))
            command_imgpath = 'imgpath ' + output_path
            cam.execute(command_imgpath)

            # Write out the trimbit configuration
            cam.execute('imgmode p')
            cam.execute('imgonly trimbit.tif')

            print('Beginning exposure.')
            cam.execute(command_exposure)
            cam.wait(7)

            cam.disconnect()
            return True
    except:
        print('ERROR: Camera comunication error. Check settings.')
        return False

def write_to_tree(shot, n_frames, exp_period, exp_time, delay, imgpath):
"""
This should be run after every data acquisition period. Given a shot number and
time parameters, this loads in the output tiff files, generates a time base,
and loads the data into the tree.
"""
  # Load shot data from the output tiff files
    data = np.zeros([M_SIZE_X, M_SIZE_Y, n_frames])
    for index in range(n_frames):
        fname = os.path.join(imgpath, 's{0:010d}_{1:05d}.tif'.format(shot, index))
        try:
            data[:, :, index] = tif.imread(fname).T
        except:
            print('ERROR: Data acquisition failed for shot {0:010d} frame {1:05d}.
    Setting to -1.'.format(shot, index))
            data[:, :, index] = -1

```

```python
    # Generate the time array
    time = ( np.array([n*exp_period + exp_time/2. for n in range(n_frames)]) +
      delay )*1000.

    # Write data and settings to mdsPlus
    try:
        # Connect to the tree
        #mesxr = MDSplus.Tree('me_sxr_ext', shot, 'NORMAL')
        mesxr = MDSplus.Tree('me_sxr_ext', shot, 'EDIT')

        # Write image data with time points
        imagesNode = mesxr.getNode(r'.ME_SXR_EXT:IMAGES')
        compiled_data = MDSplus.Data.compile("BUILD_SIGNAL($1,, $2)", data, time)
        imagesNode.putData(compiled_data)

        # Write camera settings and configuration data
        mesxr.write()
        print('Data for shot {0:10d} written to the tree.'.format(shot))
    except Exception as e:
        print('ERROR: Writing to MDSplus tree failed.')
        print(str(e))

def set_model_tree(n_frames, exp_period, exp_time, delay, config_path, setdacs_path):
"""
This function should be executed once at the beginning of data acquisition. It sets
various parameters in the model tree which will not be changed until the next
iteration of the acquisition loop.
"""
    # Read in the detector voltage parameters
    with open(setdacs_path) as dacs_file:
        v_cmp = np.zeros(16)

        for line in dacs_file:
            if 'VTRM' in line:
                v_trm = float(line.split()[-1])
            elif 'VCMP' in line:
                # Extract the chip number and the setting value
                line_elements = re.split('_|  |  ', line)
                vcmp_chip_index = int(line_elements[-2].split('VCMP')[-1])
                vcmp_chip_value = float(line_elements[-1])
                v_cmp[vcmp_chip_index] = vcmp_chip_value
            elif 'VCCA' in line:
                v_cca = float(line.split()[-1])
            elif 'VRF' in line and not 'VRFS' in line:
                v_rf = float(line.split()[-1])
            elif 'VRFS' in line:
                v_rfs = float(line.split()[-1])
            elif 'VCAL' in line:
                v_cal = float(line.split()[-1])
            elif 'VDEL' in line:
                v_del = float(line.split()[-1])
            elif 'VADJ' in line:
                v_adj = float(line.split()[-1])

    # Read in the trimbit and threshold maps
    trimbit_map = tif.imread(os.path.join(config_path, 'trimbits.tif')).T
    threshold_map = tif.imread(os.path.join(config_path, 'thresholds.tif')).T

    # Generate the bad pixel map
    bad_pixel_map = get_bad_pixel_map(config_path, trimbit_map.shape)

    # Write the data to the tree
    try:
        # Connect to the tree
        mesxr = MDSplus.Tree('mst_me_sxr', -1, 'EDIT')

        # Write the exposure timing settings
        n_images_node = mesxr.getNode(r'.CONFIG:N_IMAGES')
        n_images_node.putData(n_frames)
        exp_period_node = mesxr.getNode(r'.CONFIG:EXPOSUR_PER')
        exp_period_node.putData(exp_period)
        exp_time_node = mesxr.getNode(r'.CONFIG:EXPOSUR_TIME')
        exp_time_node.putData(exp_time)
        delay_node = mesxr.getNode(r'.CONFIG:DELAY')
        delay_node.putData(delay)
```

```python

        # Write the detector voltage settings
        vtrm_node = mesxr.getNode(r'.CONFIG:V_TRM')
        vtrm_node.putData(v_trm)
        vcmp_node = mesxr.getNode(r'.CONFIG:V_COMP')
        vcmp_node.putData(v_cmp)
        vcca_node = mesxr.getNode(r'.CONFIG:V_CCA')
        vcca_node.putData(v_cca)
        vrf_node = mesxr.getNode(r'.CONFIG:V_RF')
        vrf_node.putData(v_rf)
        vrfs_node = mesxr.getNode(r'.CONFIG:V_RFS')
        vrfs_node.putData(v_rfs)
        vcal_node = mesxr.getNode(r'.CONFIG:V_CAL')
        vcal_node.putData(v_cal)
        vdel_node = mesxr.getNode(r'.CONFIG:V_DEL')
        vdel_node.putData(v_del)
        vadj_node = mesxr.getNode(r'.CONFIG:V_ADJ')
        vadj_node.putData(v_adj)

        # Write the trimbit, threshold, and bad pixel maps
        threshold_node = mesxr.getNode(r'.CONFIG:E_THRESH_MAP')
        threshold_node.putData(threshold_map)
        trimbit_node = mesxr.getNode(r'.CONFIG:TRIMBIT_MAP')
        trimbit_node.putData(trimbit_map)
        bad_pix_node = mesxr.getNode(r'.CONFIG:BAD_PX_MAP')
        bad_pix_node.putData(bad_pixel_map)

        # Write camera settings and configuration data
        mesxr.write()
        print('Data for shot -1 written to the model tree.')

    except Exception as e:
        print('ERROR: Writing to MDSplus model tree failed.')
        print('Output: ' + str(e))

def initialize_model_tree(px_size=172, si_thick=450, be_thick=25):
"""
This function only needs to be executed once, or when any of the physical
parameters of the detector are changed.
"""
    try:
        # Connect to the tree
        mesxr = MDSplus.Tree('mst_me_sxr', -1, 'NORMAL')

        # Write data to the model tree
        px_size_node = mesxr.getNode(r'.CONFIG:PX_SIZE')
        px_size_node.putData(px_size)
        si_thick_node = mesxr.getNode(r'.CONFIG:SI_THICK')
        si_thick_node.putData(si_thick)
        filt_thick_node = mesxr.getNode(r'.GEOMETRY:FILT_THICK')
        filt_thick_node.putData(be_thick)

        # Write camera settings and configuration data
        mesxr.write()
        print('Model tree updated.')

    except Exception as e:
        print('ERROR: Writing to MDSplus model tree failed.')
        print('Output: ' + str(e))

def get_bad_pixel_map(config_dir, map_dims):
"""
This function loads in the bad pixel map from the calibration directory and puts
it in the expected format. It marks as bad the pixels contianed in the bad_pixels.csv
file as well as the gap pixels.
"""
    bad_coords = np.loadtxt(os.path.join(config_dir, 'bad_pixels.csv'),
      delimiter=',').astype(int).tolist()

    bad_pixels = np.full(map_dims, False, dtype=bool)
    for x in range(map_dims[0]):
        for y in range(map_dims[1]):
            if [x,y] in bad_coords:
                bad_pixels[x,y] = True
            elif util.get_chip_coords(x,y)[0] == -1:
```

```
                bad_pixels[x,y] = True

    return bad_pixels
```

The following script is a wrapper which is designed to be more user friendly. It checks the configurations directory for all available options and walks the user through the setup process.

```
#!/usr/bin/env python
"""
"""
import os
import numpy as np
import mesxr.operation.camera as cam

base_dir = '/home/det/meSXR/configurations'
test = False

# Print introductory message
print('ME-SXR runday startup script.')
print('Version 1.0')
print('Code by Patrick VanMeter, UW-Madison')

# Determine the desired energy range
calib_options = [f for f in os.listdir(base_dir) if not f.startswith('.')]
calib_options.sort()
calib_descriptions = [None for x in range(len(calib_options))]
for index, opt in enumerate(calib_options):
  try:
    with open(os.path.join(base_dir, opt, '.description.txt'), 'r') as f:
      calib_descriptions[index] = f.read().rstrip()
  except:
    calib_descriptions[index] = 'Description not available.'

print('\nAvailable energy ranges')

for index in range(len(calib_options)):
  print('{0:d}) {1:}: {2:}'.format(index, calib_options[index],
    calib_descriptions[index]))

user_input = raw_input('Select an energy range: ')
cal_index = int(user_input)
calib_choice = calib_options[cal_index]
calib_dir = os.path.join(base_dir, calib_choice)

# Determine the specific configuration
config_options = [f for f in os.listdir(calib_dir) if not f.startswith('.')]
config_options.sort()
config_descriptions = [None for x in range(len(config_options))]
for index, opt in enumerate(config_options):
  try:
    with open(os.path.join(calib_dir, opt, '.description.txt'), 'r') as f:
      config_descriptions[index] = f.read().rstrip()
  except:
    config_descriptions[index] = 'Description not available.'

print('\nAvailable configurations:')

for index in range(len(config_options)):
  print('{0:d}) {1:}: {2:}'.format(index, config_options[index],
    config_descriptions[index]))

user_input = raw_input('Select a configuration: ')
config_index = int(user_input)
config_choice = config_options[config_index]

# Determine the camera settings
exp_time = 0.001
exp_period = 0.002
delay = 0
n_frames = 30
```

```
print('\nDefault camera settings:')
print('exp_time={0:} sec., exp_period={1:} sec., delay={2:} sec., n_frames={3:}'.
    format(exp_time, exp_period, delay, n_frames))
user_input = raw_input('Are these settings OK? (y/n): ') or 'y'

if user_input[0].lower() == 'y':
  print('Using default values.')

elif user_input[0].lower() == 'n':
  while user_input[0].lower() != 'y':
    print('Please input new settings. Leave empty to keep default value.')
    exp_time = float(raw_input('exp_time (seconds): ') or exp_time)
    exp_period = float(raw_input('exp_period (seconds): ') or exp_period)
    delay = float(raw_input('delay (seconds): ') or delay)
    n_frames = int(raw_input('n_frames: ') or n_frames)

    print('New camera settings:')
    print('exp_time={0:} sec., exp_period={1:} sec., delay={2:} sec., n_frames={3:}'.
    format(exp_time, exp_period, delay, n_frames))
    user_input = raw_input('Are these settings OK? (y/n): ') or 'y'

else:
  print('Input not recognized. The default settings will be used.')

# Arm the detector
print('\nArming the detector.')
if test:
  print("cam.arm_detector(0, mode='cycle', write_mds=True, calibration_name={0:},
    config_name={1:}, n_frames={2:}, exp_period={3:}, exp_time={4:}, delay={5:})".
    format(calib_choice, config_choice, n_frames, exp_period, exp_time, delay))
else:
  cam.arm_detector(mode='cycle', write_mds=True, rate_corr=False, calibration_name=
    calib_choice, config_name=config_choice, n_frames=n_frames, exp_period=exp_period
    , exp_time=exp_time, delay=delay)
```

Appendix B

Overview of Bayesian inference and Integrated Data Analysis

Typically in experimental plasma physics data is taken with several independent diagnostics and analyzed separately. Generally, each diagnostic is designed with a specific plasma parameter (or a small number of parameters) in mind. Thomson Scattering, for instance, is considered to be a dedicated T_e diagnostic, FIR interferometry an n_e diagnostic, etc. However in reality diagnostic signals are complex and often depend on many of the same plasma parameters. For instance, though MST's SXR tomography system has traditionally been considered to be a T_e diagnostic [1], the plasma emissivity also depends on n_e and the the impurity densities $\{n_Z\}$. Traditionally, these additional dependencies are treated as a background signal and removed during analysis (e.g. by taking the ratio of measurements). However there is considerable information contained in these dependencies, if they can be correlated with the measurements from other diagnostics. This section aims to describe a consistent methodology for exploiting these correlations in order to extend measurement capabilities. In the fusion community this class of methodologies generally goes by the name of "integrated data analysis," or IDA [2].

The benefits of IDA are most apparent when two (or more) separate diagnostic signals are correlated with the plasma parameters in different ways. This concept is illustrated in Figure B.1, which is based on simplified models of the SXR tomography, Charge Exchange Recombination Spectroscopy (ChERS), and Thomson Scattering (TS) diagnostics. The SXR signal is related to T_e and n_z via the emissivity, which contains a component $\varepsilon_Z \propto Z n_Z n_e T_e^{-1/2} \exp(-E/T_e)$ for every impurity Z, such as Carbon. This means that an increase in the SXR signal can correspond to an increase in T_e, n_e, or n_Z. The measurements taken by the ChERS diagnostic, however, depend directly on n_Z with no sensitivity to T_e. By combining these measurements both T_e and n_Z can be estimated with less uncertainty than either diagnostic can achieve alone [3]. This technique has also been used to provide measurements of plasma properties for which no single dedicated diagnostic is available on MST, such as the ion-effective charge Z_{eff} [4].

The concept of IDA is agnostic to the particular details of the implementation. In this thesis I will exclusively make use of a Bayesian implementation of IDA, which is starting to gain widespread acceptance in the fusion community [5, 6]. We will see that the Bayesian approach brings with it a number of advantages, such as self-consistency and the straightforward treatment of uncertainty. The cost of these advantages is the substantial computational resources required for each analysis.

Bayesian inference is based on the *Bayesian interpretation* of probability theory. Probability is interpreted as the quantification of the degree of certainty of a proposition [7]. This can be contrasted with the *frequentist interpretation*, which views probability as the expected frequency of results over many repeated trials. This makes the Bayesian approach a natural fit applications where repeated trials are impractical or impossible, as is the case for many scientific measurements.

The foundation of Our IDA framework is Bayes' Rule [8]. This equation mathematically describes the process by which our prior knowledge is updated after new measurements are taken into consideration. Bayes' Rule is readily applied to the problem of

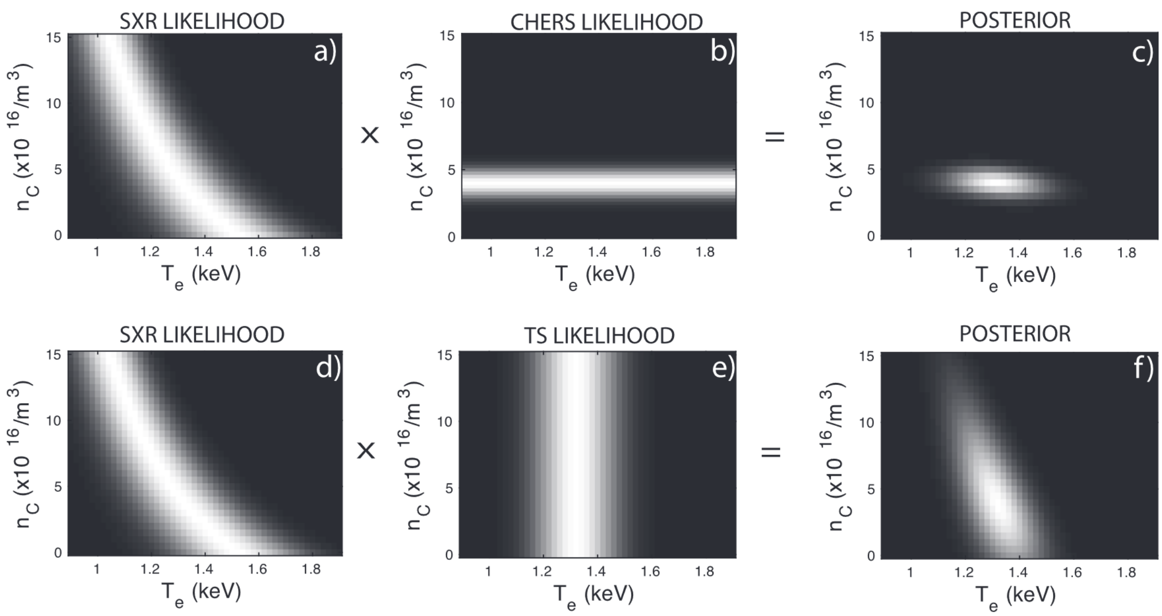

Figure B.1: Independent diagnostics work together to fully constrain the plasma properties. Figures a) and d) illustrate the correlation between carbon density and electron temperature in a model representative of a SXR tomography likelihood function. The ChERS diagnostic, b), is highly sensitive to the n_C but almost totally insensitive to T_e. By combining the information from both diagnostics together (as in Equation B.25) we get Figure c) a much more tightly constrained estimate for both T_e and n_C than either diagnostic can independently provide. This process is shown again in e) but using a Thomson Scattering likelihood which is sensitive to T_e only, resulting in the combined posterior f). Figure reproduced from L. Reusch, *et al.* [3].

parameter inference when cast in the form

$$p(\boldsymbol{\theta}|\boldsymbol{d}, I) = \frac{p(\boldsymbol{d}|\boldsymbol{\theta}, I)\, p(\boldsymbol{\theta}|I)}{p(\boldsymbol{d}|I)}, \tag{B.1}$$

where $\boldsymbol{d} = \{d_1, ..., d_N\}$ is the set of N measurements[1], $\boldsymbol{\theta} = \{\theta_1, ..., \theta_M\}$ is the set of M model parameters to be inferred, and I refers to any additional information that has been taken into consideration. The goal of Bayes' Rule is to determine the *posterior* distribution, $p(\boldsymbol{\theta}|\boldsymbol{d}, I)$, which is the conditional probability distribution over the model parameters given the measured data. This is equated to the product of the *likelihood* $p(\boldsymbol{d}|\boldsymbol{\theta}, I)$ and the *prior* $p(\boldsymbol{\theta}|I)$, normalized by the *evidence* $p(\boldsymbol{d}|I)$. We will look at each component of this equation separately before moving on.

B.1 The likelihood function

The likelihood describes the conditional probability of having obtained some particular measurements given some particular values for the model parameters $\boldsymbol{\theta}$. In the IDA context this is generally determined by a physics-based forward model relating the plasma parameters to a model of the detector output (as developed in Chapter 4) combined with a statistical model of the detector noise. Typically we will implement this by relating the probability distribution to some function $\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{d})$.

In the majority of applications measurement noise on the ith measurement will be assumed to normally distributed with variance σ_i^2 and uncorrelated with the other measurements. Recall that the probability distribution function for the normal (also sometimes called Gaussian) distribution function is given by

$$f(x|\mu, \sigma) = (2\pi)^{-1/2}\sigma^{-1} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \tag{B.2}$$

[1]Note: The subscript i here refers to a particular measurement out of a set taken simultaneously, such as different chords on a detector. It does not refer to repetitions of the same measurement over multiple experiments.

where x is some continuous scalar variable, μ characterizes the mean, and σ^2 is the variance. In such a case we write that $x \sim \mathcal{N}(\mu, \sigma)$.

Measured data is assumed to differ from the exact model my a noise term,

$$d_i = f(\boldsymbol{x}_i; \boldsymbol{\theta}) + \epsilon_i, \tag{B.3}$$

where $\epsilon_i \sim \mathcal{N}(0, \sigma_i)$ is the measurement noise and $\boldsymbol{x}_i$ is the abscissa, for instance the tangency radius of a line-of-sight. In this case the likelihood for a single measurement d_i with can be written as a normal distribution with a mean given by the forward model $\mu = f(\boldsymbol{x}_i; \boldsymbol{\theta})$ and a known variance σ_i^2,

$$p(d_i|\boldsymbol{\theta}, \sigma_i, I) = (2\pi)^{-1/2}\sigma^{-1} \exp\left[-\frac{1}{2}\left(\frac{d_i - f(\boldsymbol{x}_i; \boldsymbol{\theta})}{\sigma_i}\right)^2\right]. \tag{B.4}$$

The likelihood function we seek describes the probability of obtaining the specific set of measurements $\boldsymbol{d}$ given some specific model parameters. It is therefore given by the joint distribution over all the measurements, $p(\boldsymbol{d}|\boldsymbol{\theta}, I) = p(d_1, \ldots, d_M|\boldsymbol{\theta}, I)$. When these measurements are independent, as is typically the case[2], the distribution is separable,

$$\begin{aligned}
\mathcal{L}(\theta) &= p(d_1, \ldots, d_M|\boldsymbol{\theta}, I) \\
&= \sum_{i=1}^{N} p(d_i|\theta, I) \\
&= (2\pi)^{-N/2}\left(\prod_{i=1}^{N} \sigma_i^{-1}\right) \exp -\frac{1}{2}\sum_{i=1}^{N}\left[\frac{\left(d_i - f(x_i; \boldsymbol{\theta})\right)^2}{\sigma_i^2}\right],
\end{aligned} \tag{B.5}$$

where $f_i(\boldsymbol{\theta})$ is the forward model of the ith measurement given the model parameters $\boldsymbol{\theta}$. In practice most computational sampling algorithms require the user to specify the logarithm of the probability distribution rather than the distribution itself [9, 10]. This

[2]An example of a scenario in which measurements are not independent is when the data has undergone a smoothing procedure, such as a smoothing spline, in which measurements are adjusted by using information from other data points.

is done to avoid rounding and overflow errors due to very large or small numbers. So, taking the logarithm and simplifying:

$$\ln \mathcal{L} = -\frac{N}{2}\ln 2\pi - \ln\left(\prod_{i=1}^{N}\sigma_i\right) - \frac{1}{2}\sum_{i=1}^{N}\left[\frac{\left(d_i - f(x_i;\boldsymbol{\theta})\right)^2}{\sigma_i^2}\right]. \tag{B.6}$$

It is worth noting an assumption that has been made here. We will generally assume that the variance σ_i^2 is known exactly, generally derived from known limitations of the measurement technique (i.e. the smallest markings on a pair of calipers). However this variance is usually only an estimate, $\hat{\sigma}_i^2$, itself subject to variance from measurement noise. If this "uncertainty of the uncertainty" is small relative to σ_i then this effect is typically negligible and can be ignored. If it cannot be ignored it can be accounted for by incorporating the noise on σ_i as an additional model parameter with it's own prior. An alternative approach is to construct a likelihood function based on the Student's t-distribution rather than the normal distribution, which is equivalent to the previous approach if the priors are completely uninformative [11].

For many types of measurements the assumption of Gaussian noise is reasonable. However for detectors which count discrete events, such as photon-counting radiation detectors (like the ME-SXR) the error model is better described by Poisson statistics. The Poisson distribution characterizes the probability of registering k events (or measurements) in a given interval given a known rate (or expected value) λ. Then we say that $k \sim \mathcal{P}o(\lambda)$. It's PDF is given by

$$f(k|\lambda) = \frac{\lambda^k e^{-\lambda}}{k!}. \tag{B.7}$$

When λ becomes large the Poisson distribution converges to the normal distribution with $\sigma^2 \approx \lambda$. In many experimental contexts, when count rates are high, this is sufficient.

If we assume that a measurement subject to Poisson distributed-noise, $d_i \sim \mathcal{P}\wr(\lambda_i)$ for $\lambda_i = f(\boldsymbol{x}_i;\boldsymbol{\theta})$, then the appropriate likelihood function is a product of Poisson distributions. This is given by

$$\begin{aligned}\mathcal{L}(\boldsymbol{\theta}) &= \prod_{i=1}^{N} p(d_i|\boldsymbol{\theta}, I) \\ &= \exp\Big(-\sum_{i=1}^{N}\lambda_i\Big)\prod_{i=1}^{N}\left[\frac{\lambda_i^{d_i}}{d_i!}\right],\end{aligned} \tag{B.8}$$

where we have written $\lambda_i = \lambda_i(\boldsymbol{\theta})$ for notational simplicity. Once again, it is actually the logarithm of the likelihood which must be supplied to most computational sampling algorithms:

$$\begin{aligned}\ln\mathcal{L}(\boldsymbol{\theta}) &= -\sum_{i=1}^{N}\lambda_i + \sum_{i=1}^{N}\ln\left[\frac{\lambda_i^{d_i}}{d_i!}\right] \\ &= -\sum_{i=1}^{N}\lambda_i + \sum_{i=1}^{N}\big[d_i\ln\lambda_i - \ln(d_i!)\big] \\ &= -\sum_{i=1}^{N}\big[\lambda_i - d_i\ln\lambda_i + \ln(d_i!)\big].\end{aligned} \tag{B.9}$$

The likelihood function can also be used to determine the maximum likelihood estimate for the best fit of a function to the measured data. As the name suggests, this is the unique input $\boldsymbol{\theta}_{\mathrm{ML}}$ such that

$$\boldsymbol{\theta}_{\mathrm{ML}} = \mathrm{argmax}_{\boldsymbol{\theta}}\big[\mathcal{L}(\boldsymbol{\theta})\big]. \tag{B.10}$$

If the likelihood is taken to be a normal distribution, this is mathematically equivalent to the standard physicist technique of χ^2 minimization. This estimate can also be used as the starting point for MCMC algorithms in order to ensure quick convergence. However in the case that the likelihood is multi-modal, care must be taken to ensure that the obtained value for $\boldsymbol{\theta}_{ML}$ corresponds to a global maximum rather than a local one.

Additional terms must be incorporated to quantify the uncertainty in the model itself. In the case in which both the measurement uncertainty and model uncertainty can considered to be both independent and normally distributed with known variances, the resulting likelihood is simply another normal distribution

B.2 The prior distribution

The prior distribution $p(\boldsymbol{\theta}|I)$ quantifies the state of knowledge prior to any measurement being made. A common criticism of Bayesian inference is that the inclusion of priors introduces some level of subjectivity to the calculation. However, a substantial literature exists on the objective selection of priors based on identifying the underlying symmetries [12] or maximizing the informational entropy [7]. Without choosing a side in this philosophical controversy, it is sufficient to state that we wish to select our priors to be as objective as possible.

At the start of an analysis, it is typically the case that we wish for the prior to be as uninformative as possible. Often that means using a uniform distribution, which assigns even probability to any value within a specified range $x \in [a, b]$. The boundaries of this range should be readily justifiable, i.e. that one cannot possibly find $T_e < 0$ eV. The use of a uniform prior allows one to quantify the "eye test" of whether or not a result is plausible. We will denote a uniformly distributed random variable as $x \sim \mathcal{U}(a, b)$. The probability density function for the uniform distribution is given by

$$p(x|I) = \begin{cases} \frac{1}{b-a} & x \in [a, b] \\ 0 & x \notin [a, b] \end{cases}. \tag{B.11}$$

Normally-distributed priors are also common, when a mean μ and variance σ^2 is known (from previous measurements, for instance). In this case the prior takes the form of Equation B.2.

When the underlying model is described by a set of parameters $\boldsymbol{\theta}$, the prior function $\pi(\boldsymbol{\theta})$ is taken to be the product of the individual priors, assuming that they are uncorrelated:

$$\pi(\boldsymbol{\theta}) = p(\theta_1, \ldots, \theta_M | I) \tag{B.12}$$

$$= \prod_i^M p(\theta_i | I). \tag{B.13}$$

This simply reduces to a sum in the case that the logarithm is required:

$$\ln \pi(\boldsymbol{\theta}) = \sum_{i=1}^M \ln p(\theta_i | I). \tag{B.14}$$

The parameters a, b, μ, and σ discussed above, which parameterize the prior distribution, are referred to as hyper-parameters. In many contexts these can be assumed to be known exactly. However in some cases, such as a Gaussian prior of unknown σ, the uncertainty on the hyper-parameter must be included in the analysis via a hyper-parameter prior.

In some situations, such as when quantifying our prior knowledge of impurity content in plasma, we may not even know the proper order of magnitude to expect. In such a scenario Equation B.11 is not appropriate as the vast majority of the probability density is concentrated at the high end of the distribution. For instance, if $x \sim \mathcal{U}(10^1, 10^4)$, the vast majority of samples will have $x > 10^3$. In such a situation we may wish to employ a *log-uniform* prior. We will define $y = 10^x$ as our target random variable with the scale factor $x \sim \mathcal{U}(a, b)$. This type of scheme can be readily implemented into MCMC sampling software. We can also apply a change of variables to write the new density function directly in terms of the uniform distribution $\pi_X(x)$ and the inverse function $x(y) = \log_{10} y$:

$$\pi_Y(y) = \pi_X\big(x(y)\big)\left|\frac{d}{dy}x(y)\right| \tag{B.15}$$

$$= \frac{1}{y}\frac{\pi_X\big(x(y)\big)}{\ln 10}. \tag{B.16}$$

This result, $p(\theta_i|I) \propto 1/\theta_i$, is called a *scale-invariant prior* and is well-known and is often invoked as a hyper-parameter prior for scale-related hyper-parameters like σ [7].

B.3 The evidence

The final term in Equation B.1 is the evidence, $p(\boldsymbol{d}, I)$, sometimes also referred to as the "marginal likelihood". Unlike the other terms in the equation it is not a function, but a normalization constant which ensures that the posterior is equal to unity. It is computed by integrating the numerator over the entire parameter space. We will denote this constant as $\mathcal{Z}$, in analogy to the partition function of statistical mechanics:

$$\mathcal{Z} = \int \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{d})\,\pi(\boldsymbol{\theta}) d^N\boldsymbol{\theta}. \tag{B.17}$$

If parameter space of $\boldsymbol{\theta}$ has a large number of dimensions then direct computation of $\mathcal{Z}$ is often impractical. In practice, it is often sufficient to ignore the evidence and simply take $p(\boldsymbol{\theta}|\boldsymbol{d}, I) \propto \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{d})\pi(\boldsymbol{\theta})$ in order to obtain samples from the posterior distribution. However in some cases direct computation of the evidence is essential, such as when performing Bayesian model selection or performing hypothesis testing using Bayes factors [11].

B.3.1 The posterior distribution

Combining this all together, we can now write the posterior distribution as product of our analytical likelihood and prior functions,

$$p(\boldsymbol{\theta}) = \frac{\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{d})\, \pi(\boldsymbol{\theta})}{\mathcal{Z}}. \tag{B.18}$$

Equation B.18 is a probability distribution over a (typically) high-dimensional parameter space. In order to best characterize the distribution of a single parameter "all-else considered," we must turn to the process of marginalization. This process averages out the correlations between other variables and returns a marginal distribution for the desired parameter. For instance, in a simple two-parameter model where $\boldsymbol{\theta} = (\alpha, \beta)$, we can obtain the marginal distribution for α by marginalizing over β:

$$p(\alpha|\boldsymbol{d}, I) = \int p(\alpha, \beta|\boldsymbol{d}, I) d\beta. \tag{B.19}$$

More generally, if there are M total parameters $\boldsymbol{\theta} = \{\theta_1, \theta_2, \ldots, \theta_M\}$ then the marginal distribution over θ_1 is the multi-dimensional integral over the remaining $M - 1$ parameters:

$$p(\theta_1|\boldsymbol{d}, I) = \int \ldots \int p(\theta_1, \theta_2, \ldots, \theta_M|\boldsymbol{d}, I)\, d\theta_2 \ldots d\theta_M. \tag{B.20}$$

Marginalization also provides a way to remove so-called "nuisance" parameters from the posterior distribution. This is often the case for prior hyper-parameters, which may not contain any interesting information. In practice marginal distributions can be obtained by histogramming over samples.

Marginal distributions can be computed analytically only in the simplest possible cases, when few parameters are involved and the likelihood function has a single form (i.e., a univariate linear regression model with Gaussian noise). In the majority of cases, the problem must be approached computationally. For problems with small parameter spaces, simple grid-search methods may be sufficient. However, in many cases more sophisticated sampling methods like Markov Chain Monte Carlo [9] or nested sampling [10] are necessary. These algorithms draw samples $\hat{\boldsymbol{\theta}}_i$ from the posterior distribution, which can then be used to characterize the posterior distribution. These samples can also be used to compute expectation values. For example, the expectation value of an arbitrary function $f(\theta)$ is given by

$$\langle f \rangle \approx \frac{1}{N} \sum_{i=1}^{N} f(\hat{\theta}_i), \tag{B.21}$$

where $\hat{\theta}_i$ is the i$^{\text{th}}$ sample from the posterior distribution $p(\theta|\boldsymbol{d}, I)$. This definition works equally well when the parameter space is multi-dimensional. More information about sampling methods, including examples, can be found in the next appendix.

One of the main features of Bayesian inference is that a posterior distribution can always be updated with new information once it is available. This is accomplished by using Bayes' Rule iteratively: the posterior from the previous iteration becomes the prior when new measurements are made. An simple example of this process, inspired by a similar example in Chapter 2 of Sivia [8], is shown in Figure B.2. In this example, a coin which might be unfair (weighted) is flipped multiple times in sequence, and a result of "H" or "T" is recorded each time. After N flips, R heads have been recorded. The coin is weighted such that "H" is obtained $w \times 100\%$ of the time and "T" the $(1 - w) \times 100\%$ (a fair coin is just $w = 0.5$). The goal is to estimate the weight w of the coin as best as possible given the data available. Assuming a uniform prior $p(w|I) \propto \mathcal{U}(0, 1)$ and a binomial likelihood function $p(\boldsymbol{d}|w, i) \propto w^R(1 - w)^{N-R}$, we obtain a posterior which best quantifies our current state of knowledge. After the first flip comes up as "H", we can eliminate $w = 0$. Subsequent measurements refine our estimate. By the time 100 flips have been recorded, we are confident that the coin is not fair and is clearly weighted towards "T" ($w < 0.5$). The posterior can be refined arbitrarily further, limited only by

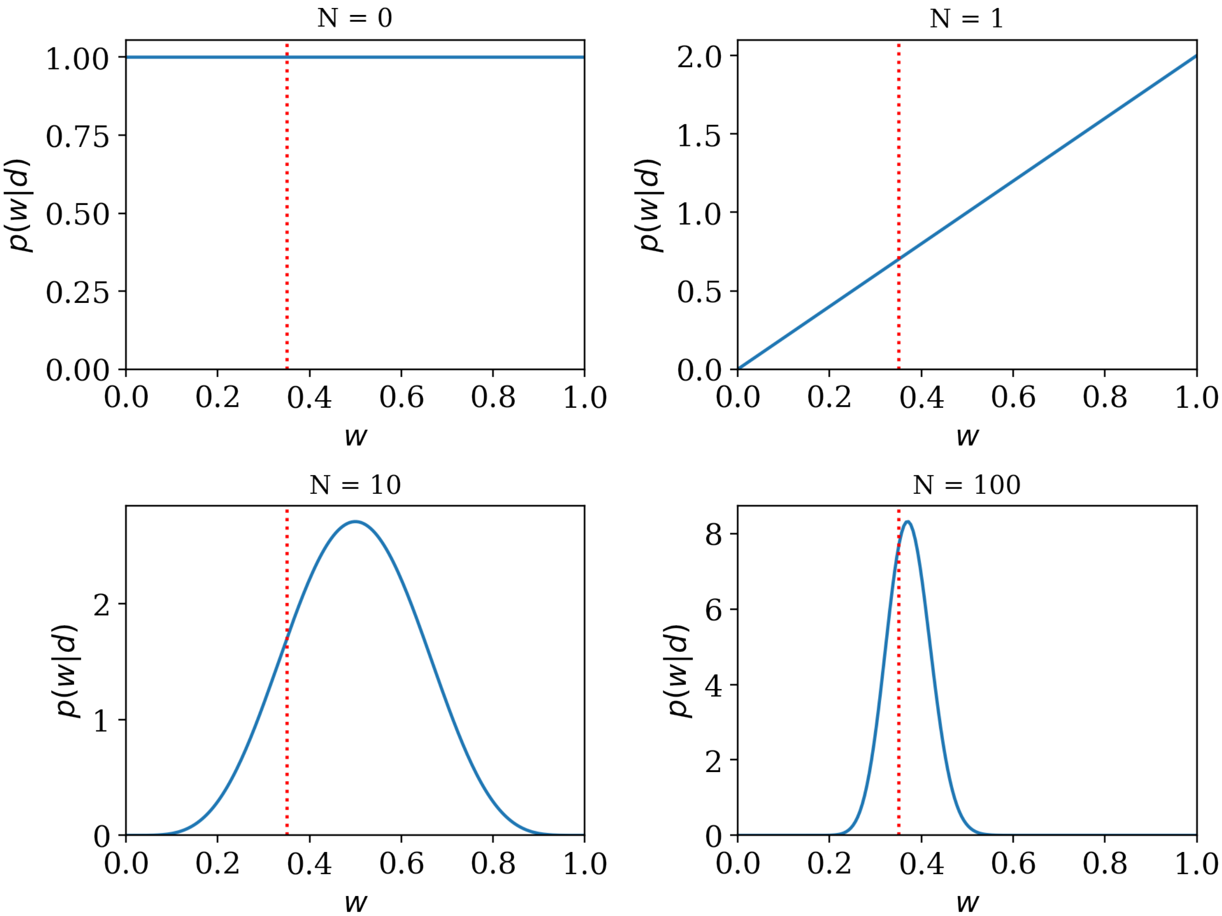

Figure B.2: The posterior distribution $p(w|\boldsymbol{d}, I)$ for the weight of an unfair coin becomes increasingly localized as it is updated with new measurements. The true weight, $w = 0.35$, is shown by the red dashed line.

how many measurements can be made.

A significant consideration of the Bayesian approach to data analysis is how to best report one's results. The full posterior distribution is the most informative way of describing the results, it is often desirable to report a single number within some kind of error bars. When the variables are normally distributed, this is readily accomplished by reporting a confidence interval $\mu \pm \sigma$, where μ is the mean and σ is the standard deviation. Under the Bayesian approach, though, it is common to work with distributions which are not normal, rendering this approach inadequate. In such cases, the most

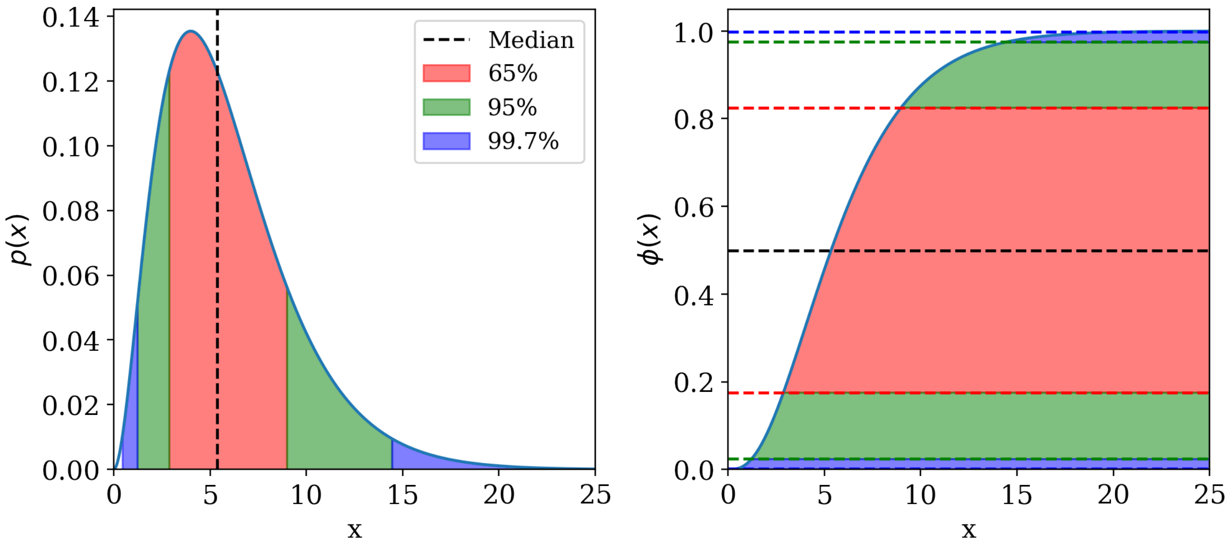

Figure B.3: 65%, 95%, and 99.7% credibility regions illustrated for an arbitrary probability distribution function $p(x)$ and its associated cumulative distribution function $\phi(x)$.

common approach is to report the median along with an asymmetric credibility region, $\bar{x}^{+\Delta x_+}_{-\Delta x_-}$ [13]. The bounds on this region are chosen to encompass some fraction $\mathcal{C}$ of the area under the distribution, typically centered about the median. Typically $\mathcal{C}$ is chosen to be one of 0.68, 0.95, or 0.997 to correspond to the 1-, 2-, and 3-σ confidence intervals commonly discussed in Gaussian statistics.

The credibility region which contains $\mathcal{C} \times 100\%$ of the distribution is computed by first finding the median $\bar{x}$ of the distribution, and then determining the upper and lower bounds x_+ and x_-. The upper bound, for instance, is chosen to ensure that the area under the distribution bounded by $\bar{x}$ and x_+ is exactly $\mathcal{C}/2$ (and likewise for x_-). This is most easily accomplished by first computing the cumulative distribution function ϕ, which for an arbitrary univariate distribution of the variable x is given by

$$\phi(x) = \int_0^x p(x')\,dx'. \tag{B.22}$$

The median is then defined as the point where $\phi(\bar{x}) = 0.5$ and bounds of the credibility region are given by $\phi(x_\pm) = 0.5 \pm \mathcal{C}/2$. The relationship between an arbitrary $p(x)$

and its associated $\phi(x)$ is shown in Figure B.3. Although phrased here in terms of analytic functions, the cumulative distribution function can be efficiently estimated given an ensemble of samples from $p(x)$ by using a numerical cumulative integration function (such as `scipy.integrate.cumtrapz` [14]).

B.4 Integrated data analysis

The concept of iteratively updating the posterior with new information provides the means through which and IDA framework can be effectively implemented. Consider the situation in which three independent diagnostics, A, B, and C, simultaneously take measurements of the same plasma. For each diagnostic we can formulate a likelihood $p_A(\boldsymbol{d}_A|\boldsymbol{\theta}, I)$, $p_B(\boldsymbol{d}_B|\boldsymbol{\theta}, I)$, *etc.*, where $\boldsymbol{\theta}$ is a shared set of parameters which describe the plasma (though each likelihood need not include every parameter in $\boldsymbol{\theta}$). Given some existing prior knowledge $p(\boldsymbol{\theta}|I)$, we can express the posterior distribution for the first measurement as

$$p(\boldsymbol{\theta}|\boldsymbol{d}_A, I) \propto p_A(\boldsymbol{d}_A|\boldsymbol{\theta}, I)\, p(\boldsymbol{\theta}|I). \tag{B.23}$$

Given that we have additional information available about the same plasma, include that information in this analysis. We do that by updating the posterior iteratively using Equation B.23 as our new prior[3]:

$$\begin{aligned} p(\boldsymbol{\theta}|\boldsymbol{d}_A, \boldsymbol{d}_B, I) &\propto p_B(\boldsymbol{d}_B|\boldsymbol{\theta}, I)\, p_A(\boldsymbol{\theta}|\boldsymbol{d}_A, I) \\ &\propto p_B(\boldsymbol{d}_B|\boldsymbol{\theta}, I)\, p_A(\boldsymbol{d}_A|\boldsymbol{\theta}, I)\, p(\boldsymbol{\theta}|I). \end{aligned} \tag{B.24}$$

This procedure can be repeated for as many diagnostics as are available. An illustration

[3] A more formal derivation of the operation can be made using the product rule of conditional probability distributions [7], but the explanation here is sufficient for developing the right intuition.

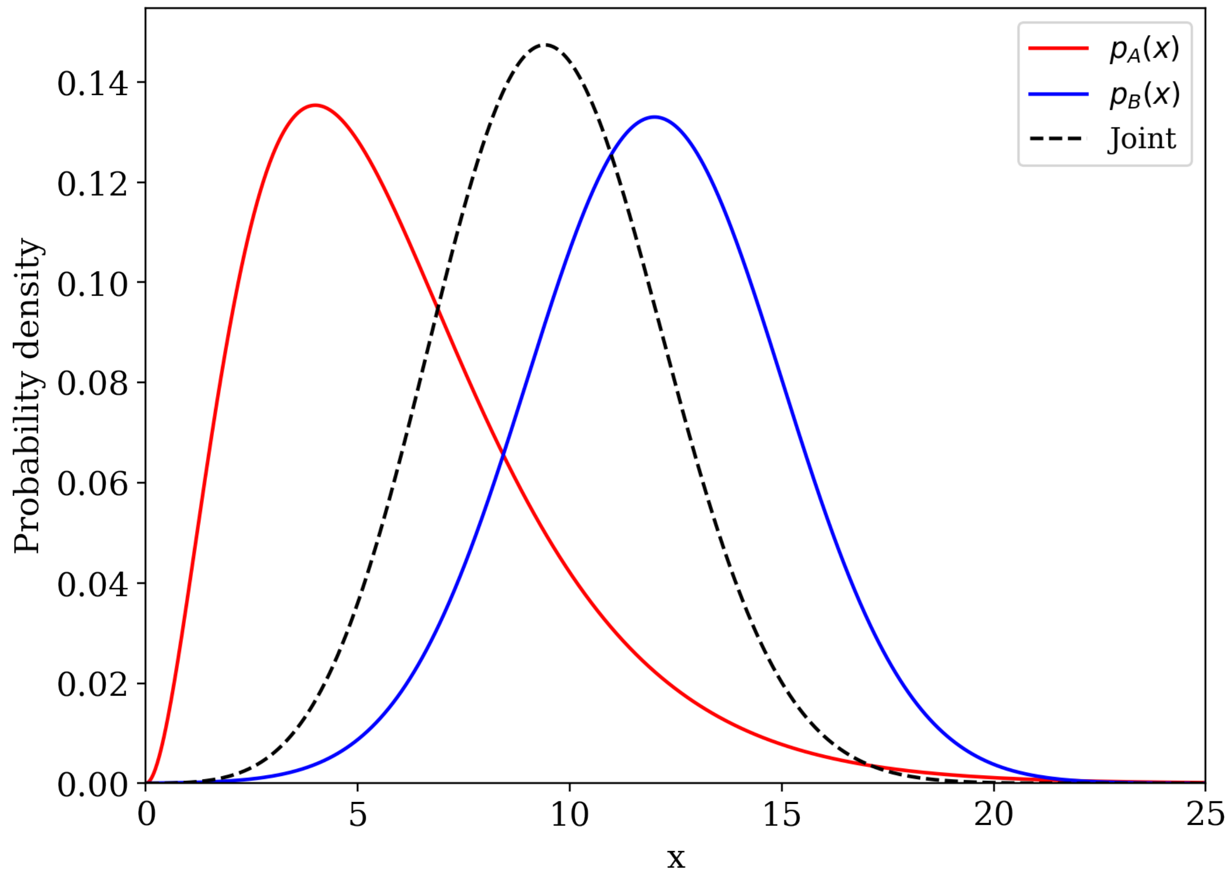

Figure B.4: Two probability distributions, $p_A(x)$ and $p_B(x)$ are combined to form a more-informed joint distribution. This illustrates what happens when multiple likelihoods are combined in the IDA process.

of the effect of this process is shown in Figure B.4, which shows the effect of multiplying two probability distributions $p_A(x)$ and $p_B(x)$ for create a new, more-informed joint distribution. It is natural to think of the resulting posterior distribution as being composed of two parts, a joint likelihood composed of the product of the individual likelihoods, and the original prior. For our three-diagnostic example, this joint likelihood is given by

$$p(\boldsymbol{D}|\boldsymbol{\theta}, I) = p_A(\boldsymbol{d}_A|\boldsymbol{\theta}, I)\, p_B(\boldsymbol{d}_B|\boldsymbol{\theta}, I)\, p_C(\boldsymbol{d}_C|\boldsymbol{\theta}, I), \tag{B.25}$$

where $\boldsymbol{D} = (\boldsymbol{d}_A, \boldsymbol{d}_B, \boldsymbol{d}_C)$ is the composite data vector. In practice, we can skip the intermediate steps and use the measured data to construct the joint likelihood immediately. Then the process of IDA is reduced to drawing samples from the resulting posterior.

In summary, the integrated data analysis framework used in this thesis proceeds as follows:

1. Take data of the same plasma using multiple diagnostics: $\boldsymbol{d}_A$, $\boldsymbol{d}_B$, ...
2. Determine a likelihood function $\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{d})$ for each diagnostic and assign appropriate priors.
3. Assemble a joint likelihood function: $\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{D}) = \prod_{i \in \{A, B, \ldots\}} \mathcal{L}_i(\boldsymbol{\theta}, \boldsymbol{d}_i)$.
4. Evaluate the posterior distribution either by direct grid search or with a numerical sampling algorithm.
5. Summarize the results using credibility regions and/or compute any desired expected values $\langle f(\boldsymbol{\theta}) \rangle$.

Bibliography

[1] M. B. McGarry, P. Franz, D. J. Den Hartog, J. A. Goetz, M. A. Thomas, M. Reyfman, and S. T. A. Kumar, "High-performance double-filter soft x-ray diagnostic for measurement of electron temperature structure and dynamics," *Review of Scientific Instruments*, vol. 83, no. 10, 2012. [Online]. Available: https://doi.org/10.1063/1.4740274

[2] R. Fischer, A. Dinklage, and E. Pasch, "Bayesian modelling of fusion diagnostics," *Plasma Physics and Controlled Fusion*, vol. 45, no. 7, pp. 1095–1111, 2003. [Online]. Available: https://doi.org/10.1088/0741-3335/45/7/304

[3] L. M. Reusch, M. D. Nornberg, J. A. Goetz, and D. J. Den Hartog, "Using integrated data analysis to extend measurement capability (invited)," *Review of Scientific Instruments*, vol. 89, no. 10, 2018. [Online]. Available: https://doi.org/10.1063/1.5039349

[4] M. Galante, L. Reusch, D. Den Hartog, P. Franz, J. Johnson, M. McGarry, M. Nornberg, and H. Stephens, "Determination of Z_eff by integrating measurements from x-ray tomography and charge exchange recombination spectroscopy," *Nuclear Fusion*, vol. 55, no. 12, p. 123016, 2015. [Online]. Available: https://doi.org/10.1088/0029-5515/55/12/123016

[5] R. Fischer, C. J. Fuchs, B. Kurzan, W. Suttrop, and E. Wolfrum, "Integrated data analysis of profile diagnostics at ASDEX upgrade," *Fusion Science and Technology*, vol. 58, no. 2, pp. 675–684, 2010. [Online]. Available: https://doi.org/10.13182/FST10-110

[6] J. Svenssson and A. Werner, "Large Scale Bayesian Data Analysis for Nuclear Fusion Experiments," in *IEEE International Symposium on Intelligent Signal Processing*. Alcala de Henares: IEEE, feb 2007, pp. 1–6. [Online]. Available: https://doi.org/10.1109/WISP.2007.4447579

[7] E. T. Jaynes, *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.

[8] D. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial*, 2nd ed. Oxford, England: Oxford University Press, 2006.

[9] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, "emcee: The MCMC Hammer," *Publications of the Royal Astronomical Society of the Pacific*, vol. 125, pp. 306–312, 2013. [Online]. Available: http://dan.iel.fm/emcee.

[10] F. Feroz, M. P. Hobson, and M. Bridges, "MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics," *Mon. Not. R. Astron. Soc*, vol. 398, pp. 1601–1614, 2009. [Online]. Available: http://www.superbayes.org.

[11] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. Boca Raton, FL: CRC Press, 2014.

[12] V. Dose, “Hyperplane Priors,” *AIP Conference Proceedings*, vol. 659, pp. 350–360, 2003. [Online]. Available: https://doi.org/10.1063/1.1570552

[13] M. D. Nornberg, D. J. Den Hartog, and L. M. Reusch, “Incorporating Beam Attenuation Calculations into an Integrated Data Analysis Model for Ion Effective Charge,” *Fusion Science and Technology*, vol. 74, no. 1-2, pp. 144–153, 2018. [Online]. Available: https://doi.org/10.1080/15361055.2017.1387008

[14] E. Jones, T. Oliphant, P. Peterson, and Others, “SciPy: Open source scientific tools for Python,” 2001. [Online]. Available: http://www.scipy.org/

Appendix C

Drawing samples from a posterior distribution

The goal of this this appendix is to provide a practical overview of some tools available for drawing samples from an arbitrary probability distribution, and to provide example code demonstrating their use. Numerous softwares and techniques exist to accomplish this goal, but this appendix will limit itself to three common choices: grid search, Markov chain Monte Carlo sampling (using `emcee`), and nested sampling (using `pyMultiNest`). The techniques are demonstrated using a simple toy model. For a thorough review of many other sampler options, I refer the reader the excellent online article by Pitkin [1], which was invaluable to the creation of this appendix.

To begin with, we must set up a scientific computing environment. The example shown in this appendix is written in Python 3.7 [2], and makes use of the Numpy [3] and Scipy [4] scientific computing packages. The examples were developed using a Jupyter notebook [5], which is a Python web interface that allows the user to separate code into executable blocks, include images inline, and embed LaTeXformatted text. We will start by importing the necessary packages:

```
import numpy as np
import scipy as sp
import scipy.stats
```

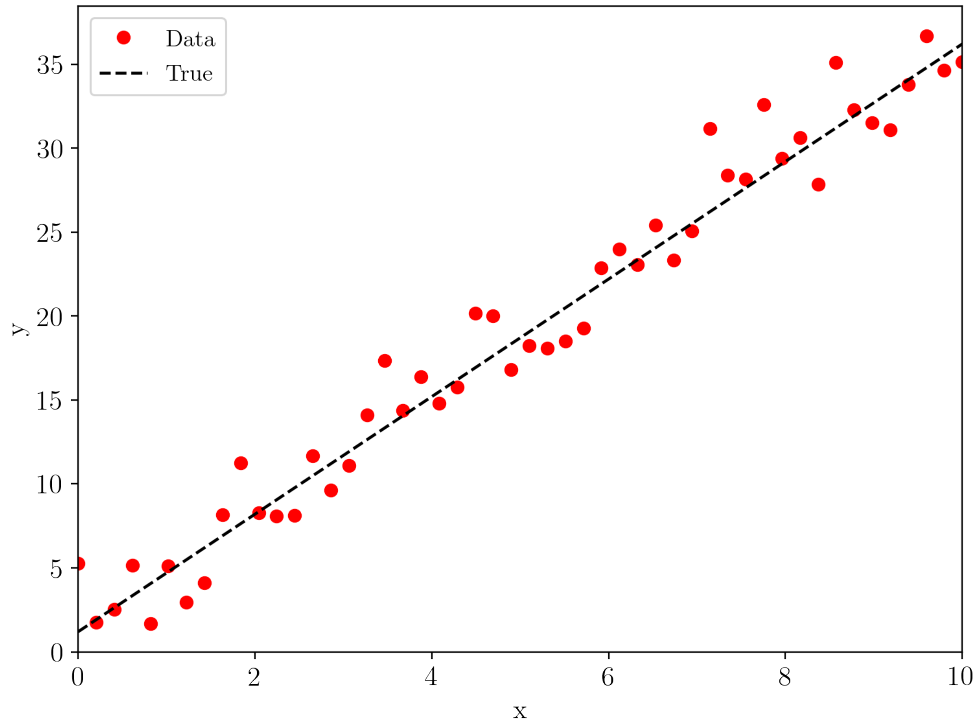

Figure C.1: Sample data used in the following examples. 50 samples were generated from Equation C.2 is normally-distributed random noise.

The examples in this appendix will be based on a simple linear toy model,

$$\boldsymbol{y}(\boldsymbol{x}; m, c) = m\boldsymbol{x} + c, \tag{C.1}$$

where $m = 3.5$ is the slope and $c = 1.2$ is the y-intercept. The simplicity of the model will allow us to focus on the details of the implementations. To simulate "real data," we also need to include measurement noise. Keeping with the simplicity of our model, we will assume that each data point is subject to normally distributed uncorrelated noise,

$$d_i = y_i(\boldsymbol{x}; m, c) + \epsilon_i, \tag{C.2}$$

where $\epsilon_i \sim \mathcal{N}(0, \sigma)$, where $\sigma = 2$ is chosen. The data is set up in the following code:

```
def y_linear(x, m, c):
    return m*x + c

# True parameters
```

```
m_true = 3.5
c_true = 1.2

# Generate the abscissa
N = 50
x = np.linspace(0, 10, num=N)

# Generate the noisy data
sigma = 2.0
data = y_linear(x, m_true, c_true) + sigma*sp.random.randn(N)
```

The resulting data, as well as the "true" line, are shown in Figure C.1. Our goal will be to use the noisy data to determine the posterior distribution for the model parameters $p(m,c|\boldsymbol{d},I)$ that best describe the true line.

Since this is a toy model, the selection of priors was somewhat arbitrary. In real problems, we want to choose priors which are "least informative" given the known constraints. In practice we will find ourselves relying mostly on uniform distributions, with bounds chosen by the constraints of physical plausibility. In this example, for the sake of variety, we will assume two different types of priors, $m \sim \mathcal{N}(\mu_m, \sigma_m)$ and $c \sim \mathcal{U}(c_{\min}, c_{\max})$, where $\mathcal{U}$ is the uniform distribution and $\mathcal{N}$ is the normal distribution. Parameters are chosen to be $\mu_m = 0$, $\sigma_m = 10$, $c_{\min} = -10$, and $c_{\max} = 10$. The parameters are assumed to be uncorrelated, so the joint prior distribution is separable:

$$\begin{aligned} \pi(\theta) &= p(m,c|I) \\ &= p(m|I)\,p(c|I). \end{aligned} \tag{C.3}$$

For the grid search and MCMC methods[1], distributions only need to be defined up to a multiplicative constant. As explained in Appendix B, we will typically work directly with the logarithm of a distribution. This can be implemented as a simple Python function using the following code:

```
m_mu = 0.0
m_sigma = 10.0
c_min = -10.0
c_max = 10.0

def ln_prior(m, c):
```

[1]Nested sampling implements likelihood and priors very differently, so all of this will need to be redefined for that example.

```
    if c_min <= c <= c_max:
        return -0.5 * ( (m - m_mu)**2 / m_sigma**2 )
    else:
        return -np.inf
```

Next, we need to define the likelihood function. This function combines the underlying data model (Equation C.1) with an model of the measurement noise. Since the measured "data" is known to be subject to uncorrelated normally distributed noise of a known variance σ^2, the likelihood function is chosen to be Gaussian. Again, we are typically interested in the logarithm of the likelihood up to a constant,

$$\ln \mathcal{L}(\boldsymbol{\theta}) = -\frac{1}{2}\sum_{i=1}^{M}\left(\frac{d_i - y(x_i; m, c)}{\sigma}\right)^2 + \text{const.} \tag{C.4}$$

This is implemented in Python code:

```
def ln_likelihood(m, c):
    return -0.5 * np.sum( (data - y_linear(x, m, c))**2 / sigma**2 )
```

Once the likelihood function and priors have been defined, they can be combined according to Bayes' Rule (Equation B.18) to form the posterior distribution. Taking the logarithm, we get

$$\ln p(\theta) = \ln \mathcal{L}(\theta) + \ln \pi(\theta) - \ln \mathcal{Z}, \tag{C.5}$$

where $\ln \mathcal{Z}$ is the logarithm of the evidence (Section B.3). If you are just interested in drawing samples from the distribution, this term can be ignored. In general, it is not calculated when sampling with MCMC-based methods. However, for a broad range of model selection and hypothesis testing problems it is an essential result of the analysis. As such, it will be calculated in the grid search and nested sampling examples. The posterior distribution is implemented in Python:

```
def ln_prob(m, c):
    lp = ln_prior(m, c)
    if np.isfinite(lp):
        return lp + ln_likelihood(m, c)
    else:
        return -np.inf
```

Now, the only remaining task is to characterize the properties of the posterior distribution we have just defined. We will consider three different approaches. Section C.1 begins with a straightforward-but-expensive grid search method, in which `ln_prob(m,c)` is evaluated directly over a large two-dimensional grid. This will allow us to evaluate the marginal distributions directly using numerical integration, and samples can be drawn by inverse transform sampling. Section C.2 will introduce MCMC sampling using the `emcee` software. This is a more sophisticated tool which enables the user to efficiently draw samples from a high-dimensional parameter space, though it does not calculate the evidence. Section C.3 introduces nested sampling and model testing using `pyMultiNest`. This technique allows the user to both draw samples from a potentially-multimodal distribution and compute the evidence, though it is somewhat less efficient and the significantly more complex than `emcee`. Finally, an example of model selection using nested sampling is presented in Section C.4.

C.1 Grid search

The most straightforward method for analyzing a probability distribution function is the grid search. Using this method, the value of the distribution is calculated over an N-dimensional grid of predefined points, where N is the number of parameters in the model. Numerical integration can then be used to directly compute the marginal distributions, moments, and credible regions. The evidence Z can also be computed directly, meaning this method can be used for model selection.

The main benefit of the grid search method is its conceptual simplicity. Because the user actually has direct access to the evaluated probability distribution, calculations (such as marginalization) closely resemble their textbook forms, and there is no need for specialized software beyond a standard scientific computing environment. However, there is a very significant drawback: the grid search method is inefficient. Substantial computational time is often spent evaluating points in regions of the parameter space with very low probability. Even worse, the method scales exponentially with the number of parameters, $\mathcal{O}(n^N)$, meaning that this method quickly becomes impractical for problems with large numbers of parameters. In my experience, I prefer to use a grid search to evaluate one- or two-dimensional distributions for which the likelihood can be evaluated quickly ($\ll 1$ second). If there are more than two parameters, it is best to use a more sophisticated method.

The following code demonstrates how to implement a grid search of the toy model (Equation C.5) using only standard `numpy` and `scipy` libraries. The grid size and bounds were chosen based on the desired level of accuracy. Multiple iterations may be required until an adequate grid has been chosen, since *a priori* we may not know where, or how localized, the peak(s) in the posterior distribution will be.

```
# The parameter grid
ms = np.linspace(0, 5, num=1000)
cs = np.linspace(-2, 6, num=1000)

# Evaluate the pdf over the grid
ln_pdf = np.zeros([len(ms), len(cs)])

for i,m in enumerate(ms):
    for j,c in enumerate(cs):
        ln_pdf[i,j] = ln_prob(m, c)

pdf = np.exp(ln_pdf)

# Caclulate the marginal distributions
pdf_c = np.trapz(pdf, x=ms, axis=0)
pdf_m = np.trapz(pdf, x=cs, axis=1)

# Scale by the evidence
Z = np.trapz(pdf_c, x=cs)
pdf /= Z
pdf_c /= Z
pdf_m /= Z
```

The posterior distribution is shown in Figure C.2 (a), in the form of a corner plot. The corner plot is a common tool for examining multidimensional probability distributions. The diagonal entries show the one-dimensional marginal distributions ($p(m|\boldsymbol{d}, I)$ and $p(c|\boldsymbol{d}, I)$, in this case) while the off-diagonals show the two-dimensional marginal distributions (just $p(m, c|\boldsymbol{d}, I)$ for this example). This allows the user to simultaneously understand how well-constrained the individual parameters are as well as how they are correlated to one another. In the case of our example, we see that m and c both seem to be well-constrained by the data. But how well, exactly?

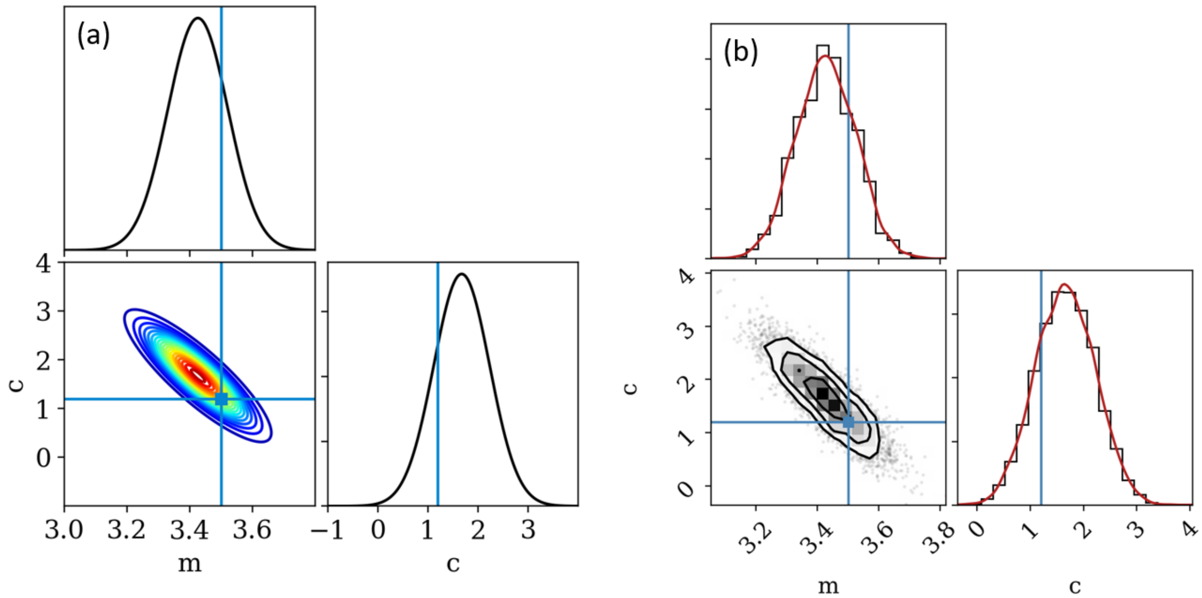

Figure C.2: Corner plots showing (a) the posterior distribution evaluated directly on the (m, c) grid, and (b) histograms generated from samples obtained using inverse transform sampling. The blue lines mark the true values.

This question is best answered by calculating the 68% credible interval for each parameter. As explained in Section B.3.1, this is the Bayesian equivalent to a 1-σ confidence interval. In fact, since both the likelihood and prior are described by normal distributions, this will turn out to be almost[2] equal to the 1σ confidence interval. However the method presented here is more general and is applicable to any arbitrary distribution. We will use the cumulative distribution function, Equation B.22, to determine the credible intervals:

```
import scipy.integrate

# Compute the CDFs
cdf_m = sp.integrate.cumtrapz(pdf_m, x=ms, initial=0)
cdf_c = sp.integrate.cumtrapz(pdf_c, x=cs, initial=0)

# Get the 0.68 credibility regions
m_mean = ms[np.argmin(np.abs(cdf_m - 0.5))]
m_upper = ms[np.argmin(np.abs(cdf_m - (0.5 - 0.68/2)))]
m_lower = ms[np.argmin(np.abs(cdf_m - (0.5 + 0.68/2)))]

```

[2]Since $p(c|I)$ is zero outside of the range $[c_{\rm min}, c_{\rm max}]$, the posterior distribution is technically not a normal distribution. However, since the posterior is localized to a region of parameter space far from the boundary, this distinction is insignificant.

```python
c_mean= cs[np.argmin(np.abs(cdf_c - 0.5))]
c_upper = cs[np.argmin(np.abs(cdf_c - (0.5 - 0.68/2)))]
c_lower = cs[np.argmin(np.abs(cdf_c - (0.5 + 0.68/2)))]
```

This calculation returns intervals of $m = 3.43^{+0.10}_{-0.10}$ and $c = 1.67^{+0.55}_{-0.55}$, which do in-fact encompass the true values. Again, the upper and lower errors are equal in this case because the distribution is Gaussian. For arbitrary distributions this is not always true.

Some computational tasks are most readily performed by drawing samples from a distribution, $\hat{\boldsymbol{\theta}}$, which allow for techniques based on ensemble averaging. Unlike the sampling methods presented later in this appendix, the grid search analysis presented above does not produce samples. In this case, since the posterior is known to be Gaussian it is possible to estimate the mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ of the distribution from the data and use a packaged multivariate normal distribution sampler (such as the one included in `scipy.stats`). However, that is not generally applicable to all distributions. Instead, we will make use of inverse transform sampling, a trick which allows us to map a uniformly distributed random variable $u \sim \mathcal{U}(0, 1)$ to an arbitrary distribution using the cumulative distribution function. For a one-dimensional distribution $p(x)$, we transform a uniform sample $\hat{u}$ into the desired parameter $\hat{x}$ by

$$\hat{x} = \phi^{-1}(\hat{u}), \tag{C.6}$$

where $\phi(x)$ is the cumulative distribution function.

For multidimensional distributions, such as our toy model, this method is a bit trickier. We must first use some other transform to map the multidimensional distribution $p(m, c|\boldsymbol{d}, I)$ to a one dimensional distribution of some other variable z which enumerates the entire space. Since $p(m, c|\boldsymbol{d}, I)$ is evaluated on a finite grid, the simplest choice is to enumerate each point in the grid and generate a new cumulative distribution using a cumulative sum function scaled by the area/volume of the grid spacing. This is demonstrated in the following code, using the built-in `numy.ravel` function to provide a consistent enumeration:

```
# Create a flattened version of the CDF
pdf_flat = pdf.ravel()
cdf_flat = np.cumsum(pdf_flat) * (ms[1] - ms[0]) * (cs[1] - cs[0])

# Draw samples from the distribution using inverse transform sampling
num_samples = 5000
us = np.random.rand(num_samples)
ns = np.array([np.argmin(np.abs(u - cdf_flat)) for u in us])

ns_m, ns_c = np.unravel_index(ns, shape=pdf.shape)
samples = np.zeros([num_samples, 2])
samples[:,0] = ms[ns_m]
samples[:,1] = cs[ns_c]
```

It is initially tempting to just implement the one-dimensional scheme for the marginal distributions $p(m|\boldsymbol{d}, I)$ and $p(c|\boldsymbol{d}, I)$, but this will eliminate the covariance between samples, artificially broadening the posterior distribution. As with other aspects of the grid search method, inverse transform sampling quickly becomes intractable as the number of dimensions increases beyond two or three.

Figure C.2 (b) shows a corner plot composed of histograms of the sampling results. This was created using the `corner` software package [6], which was also used throughout this thesis. This shows that samples provide a good estimation of the parent distribution, allowing them to be reliably used in further computations.

We will use the ensemble of posterior samples $\{(\hat{m}, \hat{c})_i\}_{i\in[1,M]}$ to illustrate the uncertainty range of our fit results. At each point x_i along the abscissa we will evaluate

$$E[y_i] \approx \frac{1}{M}\sum_{j=1}^{M} \hat{y}_{ij} \tag{C.7}$$

$$\mathrm{Var}[y_i] = \frac{1}{M}\sum_{j=1}^{M} \left(y_{ij} - E[y_i]\right)^2 \tag{C.8}$$

$\hat{y}_{ij} = y(x_i; \hat{m}_j, \hat{c}_j)$ and where M is the number of samples drawn from the posterior distribution. The result is shown in Figure C.3. The dashed “Fit” line is the expectation value at each point and the uncertainty bands are given by $\sigma = \sqrt{\mathrm{Var}[x]}$.

```
xs = np.linspace(0, 10, num=1000)
ys = np.zeros([num_samples, len(xs)])

for i in range(num_samples):
    ys[i,:] = y_linear(xs, samples[i,0], samples[i,1])

```

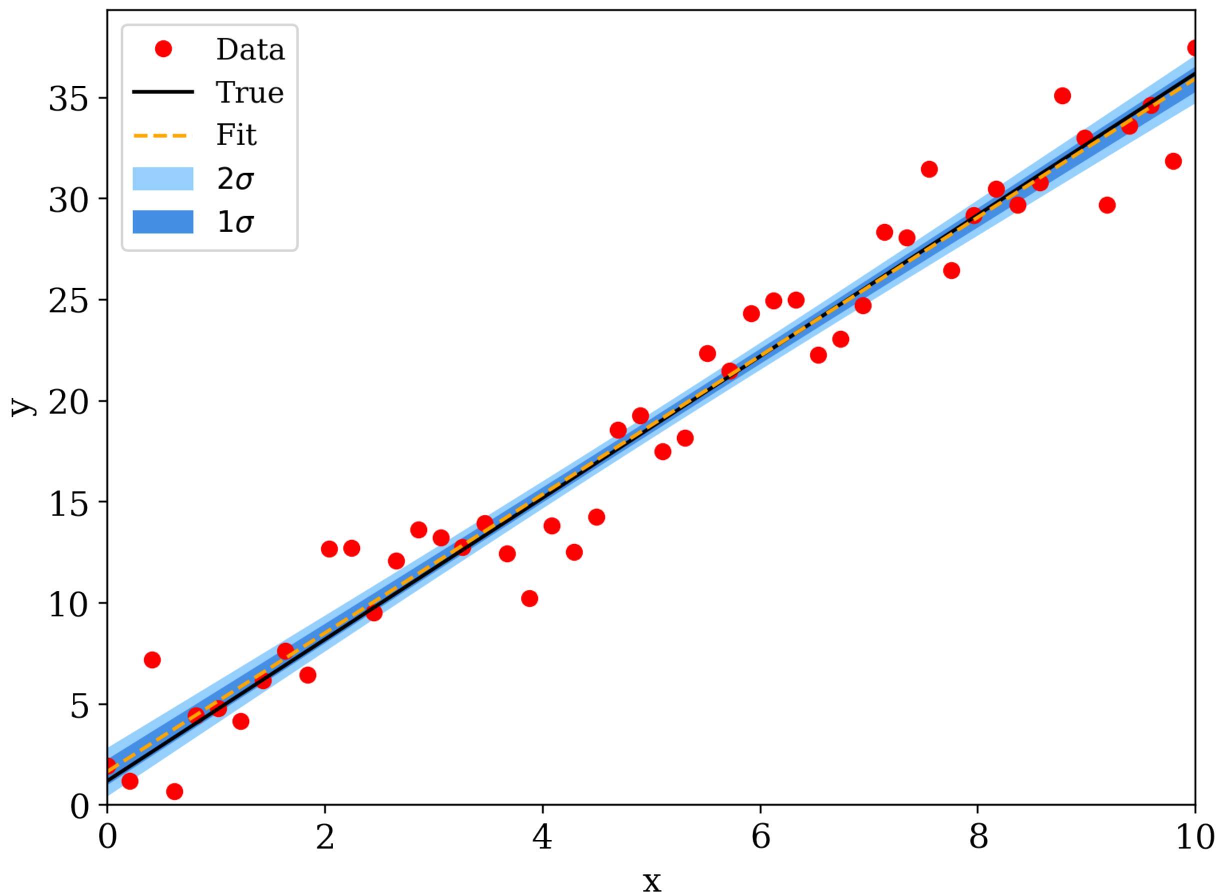

Figure C.3: Comparison of the posterior results with the original data points. The 1- and 2-σ bands were determined by averaging over the set of samples $\{(\hat{m}, \hat{c})_i\}$.

```
y_mean = np.average(ys, axis=0)
y_std = np.std(ys, axis=0)
```

C.2 Markov Chain Monte Carlo (MCMC) using `emcee`

In many cases, evaluating the posterior distribution over a predefined grid of points is impractical. This could be because the parameter space is too large ($N \gtrsim 3$), because evaluation of the likelihood is slow, or both. Regardless of the reason, what we want is an efficient method for producing samples $\hat{\boldsymbol{\theta}}$ from the distribution, which can then be used to discern its properties. The first such method we will consider is Markov chain

Monte Carlo sampling (MCMC).

Without going into too much detail, MCMC is a general class of algorithms in which a sequence of random variables $\theta^1, \theta^2, \ldots$ are produced in such a way that the distribution of samples converges to a target probability distribution. This sequence has the property that future steps in the sequence (θ^t) depend only on the present position (t^{t-1}), which is called the Markov property. Numerous algorithms fall under this category, including Gibbs sampling, the Metropolis-Hastings algorithm, and Hamiltonian Monte Carlo. A more detailed discussion on the theory and technical details of MCMC can be found in Gelman, *et al.* [7], and other textbooks on Bayesian statistics.

In this section we will implement an example analysis using `emcee`, a Python module which implements a distinctive affine-invariant ensemble sampler which has favorable convergence properties when evaluation of the evidence is computationally expensive [8]. We will initiate a number of independent Markov chain, which `emcee` calls "walkers", which will explore the parameter space and generate a set of samples $\{\hat{\boldsymbol{\theta}}_i\}$ which can be used to characterize the system. The use of multiple simultaneous independent Markov chains makes parallelization straightforward. MCMC analyses typically require a "burn-in"time, which is the number of steps required for the Markov chain to converge to the posterior distribution. This can be quantified via the autocorrelation time, τ_f, which can itself be estimated from the previous steps in the chain (see A. Sokal [9] for a more detailed discussion). For our purposes, it will suffice to visually examine the path of the Markov chain through parameter space in order to deduce convergence.

The `emcee` EnsembleSampler object requires our posterior distribution be slightly redefined so that it accepts a single parameter array (which I will call "theta"):

```
def ln_prob(theta):
    m, c = theta
    lp = ln_prior(m, c)
    if np.isfinite(lp):
        return lp + ln_likelihood(m, c)
    else:
        return -np.inf
```

In order speed up the convergence process, it is often desirable to start the walkers out in a region near the peak of the distribution. A good way to do this is to first make a maximum likelihood estimate $\boldsymbol{\theta}_{\mathrm{ML}}$ (Equation B.10). This be be done using the `scipy.optimize` module:

```
from scipy.optimize import minimize

# Inital guess and bounds
theta0 = np.array([0, 0])
bounds = [(-25, 25),
          (c_min, c_max)]

# Maximum likelihood estimate
nll = lambda theta: -ln_likelihood(*theta)
soln = minimize(nll, theta0, bounds=bounds)
theta_MLE = soln.x
```

For our toy model this results in $m_{\mathrm{MLE}} = 3.427$ and $c_{\mathrm{MLE}} = 1.673$. We will start each of our walkers at an initial position which is a small random displacement away from $\boldsymbol{\theta}_{\mathrm{MLE}}$. Note that this strategy works best for unimodal distributions. For multimodal distributions, it may be a better idea to distribute starting positions according to the prior distributions.

The python code for analyzing the toy model with `emcee` is given below. We will (arbitrarily) use 50 walkers and extend each chain for 10,000 steps, far longer than necessary for convergence. Note that although we will not need it here, the EnsembleSampler class also supports multiprocessing via a Python "Pool" object.

```
import emcee

# General emcee settings
n_steps = 10000
nwalkers = 50
ndim = len(theta0)

# Initial sampling positions
pos0 = np.zeros([nwalkers, ndim])
for ii in range(nwalkers):
    pos0[ii,:] = theta_MLE + 0.05 * theta_MLE * np.random.randn(ndim)

# Run the ensemble sampler
sampler = emcee.EnsembleSampler(nwalkers, ndim, ln_prob)
sampler.run_mcmc(pos0, n_steps, progress=True)
```

A plot of the walkers' initial movements through the parameter space is shown in Figure C.4 (a). This shows that the chain's behavior becomes stationary quite quickly, before even 100 steps. We Can expect the Markov chain to be well-converged after that point, and so will omit the earlier steps from the final analysis as burn-in. The following code

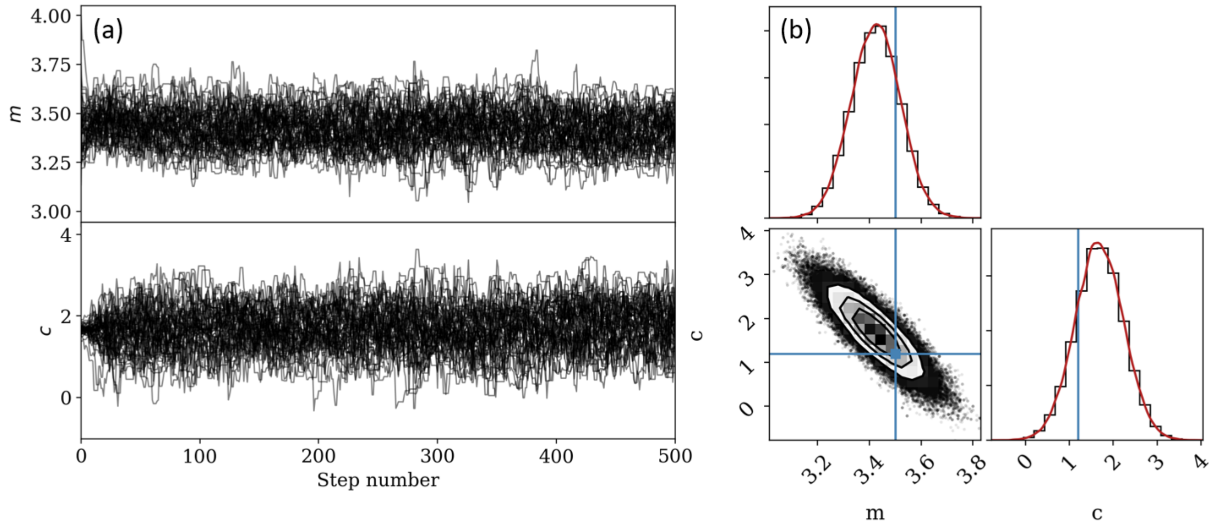

Figure C.4: Results of the `emcee` run, showing (a) walker position in the parameter space for the first 500 steps of the Markov chain, and (b) histogram corner plot of the posterior distribution.

drops the first half of the samples [3] and merges the independent walker chains to form a single set of samples:

```
samples = sampler.get_chain(discard=5000, flat=True)
```

The corner plot for these samples is shown in Figure C.4 (b). Unsurprisingly, the results are nearly identical to those seen in the grid search example (it is the same posterior distribution, after all). These samples can be used to calculate $E[y_i]$ and $\mathrm{Var}[y_i]$ according to Equations C.7 and C.8, respectively. The resulting mean and uncertainty bands are nearly identical to those shown in Figure C.3, so an additional plot has not been included here.

C.3 Nested sampling with `pyMultiNest`

Nested sampling is a technique that was originally created to efficiently evaluate the evidence Z for general Bayesian inference problems [10]. A side effect of the algorithm's

[3]Since chains have no "memory," there is no harm in dropping more early steps than necessary.

design is that it also produces samples from the distribution. The "nested" part of nested sampling refers to the way that the algorithm subdivides the sampling space into multiple sub-domains, which allows it to much more effectively sample multimodal distributions than MCMC-based methods. However, it is generally less efficient for producing large numbers of samples from unimondal or weakly multimodal distributions.

Nested sampling works by reducing the N-dimensional integral for $\mathcal{Z}$ (Equation B.17) to a single-dimensional integral over the "prior volume" X, given by

$$X(\lambda) = \int_{\mathcal{L}(\boldsymbol{\theta})<\lambda} \pi(\theta)\, d^N\boldsymbol{\theta} \tag{C.9}$$

where λ defines contours of constant likelihood $\mathcal{L}(\boldsymbol{\theta}) = \lambda$. The integral for $\mathcal{Z}$ is then reduced to

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X)\, dX. \tag{C.10}$$

The goal, then, is to evaluate $\mathcal{L}(X)$ for a finite set of X values between 0 and 1 and then approximate $\mathcal{Z}$ using numerical integration. In this example we will use the pyMultiNest code, a Python wrapper for the popular MultiNest algorithm [11].

The MultiNest algorithm requires that your priors be parameterized in terms of a unit hypercube prior, $p(u_j|I) \sim \mathcal{U}(0,1)$ for each model parameter u_j. This requirement is in part because the algorithm depends on the parameter space having a finite and easily quantified volume. Taken on its face, this seems to be a very strict requirement. However we can get around this by introducing a transformation $f(\mathbf{u})$ between the hypercube parameters $\mathbf{u} = \{u_1, u_2, \ldots, u_M\}$ and the physical parameters $\boldsymbol{\theta} = \{\theta_1, \theta_2, \ldots, \theta_M\}$, $f : \mathbf{u} \mapsto \boldsymbol{\theta}$ which properly conserves the probability density. The algorithm will perform this transformation before evaluating the likelihood.

The transformation is described in detail in Feroz, *et al.* [11], but we will go over the basics. The goal is to preserve the total probability over a given volume of the parameter space:

$$\int du_1\, du_2 \dots du_M = \int d\theta_1\, d\theta_2 \dots d\theta_M \pi(\theta_1, \theta_2, \dots, \theta_M). \tag{C.11}$$

We will assume that the the joint prior probability distribution is seperable, i.e. $\pi(\theta_1, \theta_2, \dots, \theta_M) = \Pi_{j=1}^{M} \pi_j(\theta_j)$. In such a case, we can satisfy the conservation of probability density by requiring that

$$du_j = \pi_j(\theta_j) d\theta_j. \tag{C.12}$$

Integrating this, we obtain a way to convert between parameters:

$$u_j = \int_{-\infty}^{\theta_j} \pi_j(\theta_j') d\theta_j'. \tag{C.13}$$

In our linear toy model, we have two different kinds of priors: (non-unit) uniform and normal. Let's first work through the transformation of the uniform prior over c, which is more straightforward. Following the conventions of the `GaussianSolver` code outlined later in this section, we will label this hypercube parameter u_2. We begin with the distribution over the physical parameter, $c \sim \mathcal{U}(c_{min}, c_{max})$, defined as

$$\pi(c) = \begin{cases} \frac{1}{c_{max} - c_{min}} & c \in [c_{min}, c_{max}] \\ 0 & c \notin [c_{min}, c_{max}] \end{cases}, \tag{C.14}$$

and then apply Equation C.13 to get

$$\begin{aligned} u_2 &= \int_{-\infty}^{c} \pi(c') dc' \\ &= \int_{c_{min}}^{c} \frac{dc'}{c_{max} - c_{min}} \\ &= \frac{c - c_{min}}{c_{max} - c_{min}}. \end{aligned} \tag{C.15}$$

Finally, this can be inverted to give the forward transformation $u_2 \rightarrow c$,

$$c = u_2 \cdot (c_{max} - c_{min}) + c_{min}. \tag{C.16}$$

In this derivation I have assumed $\theta \in [c_{min}, c_{max}]$. This is because $u = 0$ for all values of $c < c_{min}$ and $u = 1$ for all $c > c_{min}$, effectively collapsing each of those spaces down to a single point in u-space with zero probability weight. So for and draw of $u \sim \mathcal{U}(0, 1)$ we only need to worry about the non-zero subdomain.

Next we turn to m, which is normally distributed $m \sim \mathcal{N}(\mu_m, \sigma_m)$. As before, we begin with the definition for the probability density function,

$$\pi(m) = \frac{1}{\sigma_m \sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{m-\mu_m}{\sigma_m}\right)^2}, \tag{C.17}$$

and plug this into Equation C.13. The integral is readily solved using a change of variables, $z = \frac{m-\mu_m}{\sigma_m}$. APllting this,

$$u_1 = \int_{-\infty}^{m} \frac{1}{\sigma_m \sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{m-\mu_m}{\sigma_m}\right)^2} dm \tag{C.18}$$

$$= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{z} e^{z'^2/2} dz' \tag{C.19}$$

$$= \Phi\left(\frac{m - \mu_m}{\sigma_m}\right), \tag{C.20}$$

where $\Phi(z)$ is the cumulative distribution function for standard normal distribution. The inverse function, $\Phi^{-1}(z)$, can be calculated by a computer[4]. This gives the desired mapping,

$$m = \sigma_m \Phi^{-1}(u_1) + \mu_m. \tag{C.21}$$

[4] There seems to be some inconsistency regarding the name for $\Phi^{-1}(z)$. Some sources call this the "probit" function, while `scipy` implements it as "ndtri".

Next, we want to check that these transforms do indeed result in the correct distributions. We can confirm this via brute-force by drawing a large number of samples $\hat{u}_i$ from a uniform distribution, using Equations C.16 and C.21 to map each sample to $\hat{m}_i$ and $\hat{c}_i$, and binning the results into a histogram. The sampling step is implemented in the following Python code:

```python
from scipy.special import ndtri

us = np.random.rand(100000)
ms = m_sigma*ndtri(us) + m_mu
cs = us*(c_max - c_min) + c_min
```

The resulting histograms are shown in Figure C.5, with the target distributions (Equations C.14 and C.17) overplotted for comparison. This verifies that the transforms implemented in Equations C.16 and C.13 behave as intended. In fact, these transforms are closely related to the inverse transform sampling methodology discussed back in Section C.1.

With these transformations in-hand, we can run the `pyMultiNest` analysis. There are several ways to go about this, but I will follow the example of the Pitkin article [1] and create a new class which inherits from the `pymultinest.Solver` class. This provides a clean and consistent interface for the analysis. One more detail to note is that nested sampling requires the full normalized likelihood function (constants included), so we will need to implement Equation B.6 in full. The code for the new "GaussianSolver" class is given below.

```python
import pymultinest
from pymultinest.solve import Solver
from scipy.special import ndtri
LN2PI = np.log(2.*np.pi)

class GaussianSolver(Solver):
    """
    A simple straight line model, with a Gaussian likelihood.

    Args:
        data: an array containing the observed data
        abscissa: an array containing the points at which the data were taken
        modelfunc: a function defining the model
        sigma: the standard deviation of the noise in the data
        **kwargs: keyword arguments for the run method
    """
    # define the prior parameters
    cmin = -10.  # lower range on c
    cmax = 10.   # upper range on c
    mmu = 0.     # mean of the Gaussian prior on m
    msigma = 10. # standard deviation of the Gaussian prior on m

```

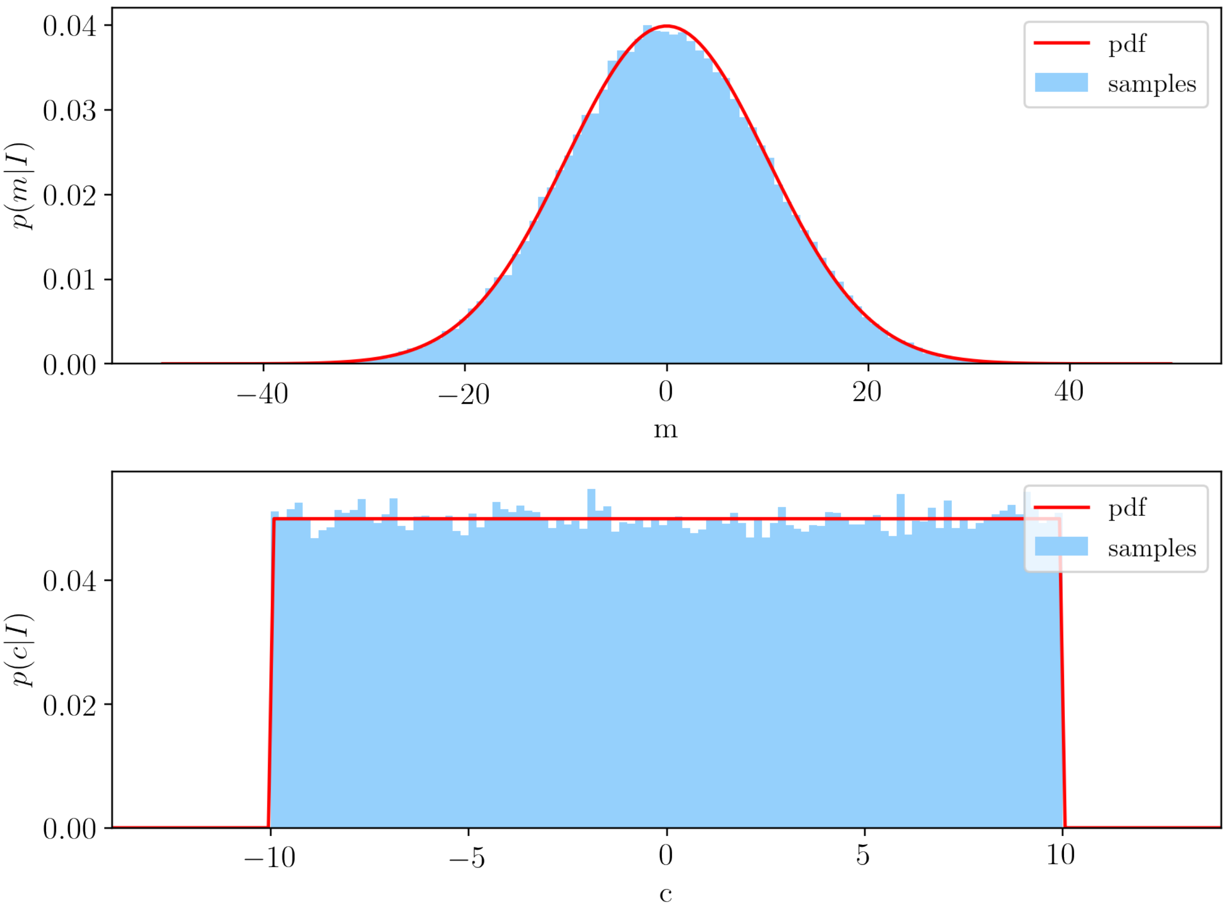

Figure C.5: Samples from the prior distribution $p(m, c|I) = p(m|I)\, p(c|I)$ are generated by transforming M samples $\{(\hat{u}_1, \hat{u}_2)_i\}_{i\in[i,M]}$ from the the unit hypercube. These samples (blue histograms) are compared against the original analytic distributions (red lines). These transformations are required for the `MultiNest` algorithm.

```python
    def __init__(self, data, abscissa, modelfunc, sigma, **kwargs):
        # set the data
        self._data = data          # oberserved data, d_i
        self._abscissa = abscissa # data points, x_i
        self._sigma = sigma        # standard deviation
        self._ndata = len(data)    # number of data points
        self._model = modelfunc    # model function

        # log sigma here to save computations in the likelihood
        self._logsigma = np.log(sigma)

        Solver.__init__(self, **kwargs)

    def Prior(self, cube):
        """
        Transfrom from the unit hypercube to true parameters
        Args:
            cube: an array of values drawn from the unit hypercube
        Returns:
            an array of the transformed parameters
        """
        mprime = cube[0]
        cprime = cube[1]

        m = self.mmu + self.msigma*ndtri(mprime)
        c = cprime*(self.cmax-self.cmin) + self.cmin

        return np.array([m, c])

    def LogLikelihood(self, cube):
        """
        Args:
            cube: an array of parameter values.
        Returns:
            the log likelihood value.
        """
        m = cube[0]
        c = cube[1]

        # calculate the model
        model = self._model(x, m, c)
        norm = -0.5*self._ndata*LN2PI - self._ndata*self._logsigma
        chisq = np.sum(((self._data - model)/(self._sigma))**2)

        return norm - 0.5*chisq
```

The analysis is performed as soon as the solver object is created. This can be done with single line of Python. We will also go ahead and get the posterior distribution samples from the solver object.

```python
# Solver parameters
nlive = 1024 # number of live points
ndim = 2     # number of parameters
tol = 0.5    # stopping criterion

# Run the solver
solution = GaussianSolver(data, x, y_linear, sigma, n_dims=ndim, n_live_points=nlive,
     evidence_tolerance=tol)

# Get the results
logZ = solution.logZ
logZ_err = solution.logZerr
samples = solution.samples
```

Once the analysis is complete, we can get a summary of the analysis by passing the solver object to the `print()` function,

```python
print(solution)
```

which ouptuts:

```
Model in "(temporary directory)" (2 dimensions)
Evidence ln Z = -108.9 +- 0.1
Parameter values:
   Parameter  1 : 3.584 +- 0.097
   Parameter  2 : 0.623 +- 0.559
```

Note that this example was run over a different synthetic data set than the previous two examples (using the same y_i points, but with different randomly generated noise). This is the primary reason that the parameter estimates are slightly different than the previous two examples. Once again, we can use the samples to make corner plots and calculate $E[y_i]$ and $\mathrm{Var}[y_i]$ according to Equations C.7 and C.8. As expected, the results are nearly indistinguishable from those in the previous two examples, and so are not repeated here.

Unlike the previous analyses, `pyMultiNest` has also provided a well-constrained estimate of the log-evidence, $\ln \mathcal{Z} = -108.916 \pm 0.083$. Although $\mathcal{Z}$ is not used in the IDA methodology presented in this thesis, it is useful quantity with many applications in Bayesian inference. The most direct application of the evidence, model selection, is demonstrated below in the final example of this appendix.

C.4 Model selection with `pyMultiNest`

This final example extends the previous section to demonstrate how `pyMultiNest` can be used to perform basic model selection. Instead of normally-distributed measurement noise, we will subject each measurement to Poisson noise. Then, we will attempt to fit the data using both a Poisson likelihood function and a Gaussian likelihood function. By calculating the evidence $\mathcal{Z}$, we will see that the Poisson model provides a better description of the data. This provides a Bayesian alternative to common frequentist technique of comparing each model's reduced χ^2.

The true model, $\boldsymbol{y}(\boldsymbol{x}; m, c)$, will again be represented by Equation C.1. Instead of adding normally-distributed noise, each "measurement" d_i are perturbed from y_i by drawing a sample from $d_i \sim \mathcal{P}\wr(y_i)$ from the Poisson distribution, $\mathcal{P}\wr(\lambda)$, given by Equation B.7.

The proper likelihood function, then, should also be a Poisson distribution. The correct form for a Poisson $\ln \mathcal{L}(\boldsymbol{d}|m,c)$ was previously derived in Equation B.9. Additionally, we will need to modify our priors slightly to be consistent with the new noise model. Since a Poisson process cannot ever return fewer than zero counts, the priors were changed to $c \sim \mathcal{U}(0,10)$ and $m \sim \mathcal{U}(0,10)$[5]. The resulting posterior distribution was implemented into a new solver class, the "PoissonSolver," using the following code:

```
class PoissonSolver(Solver):
    """
    A simple straight line model, with a Gaussian likelihood.

    Args:
        data (:class:`numpy.ndarray`): an array containing the observed data
        abscissa (:class:`numpy.ndarray`): an array containing the points at which the data were taken
        modelfunc (function): a function defining the model
        **kwargs: keyword arguments for the run method
    """

    # define the prior parameters
    cmin = 0.    # lower range on c
    cmax = 10.   # upper range on c

    mmin = 0.    # Lower range on m
    mmax = 10.   # Upper range on m

    def __init__(self, data, abscissa, modelfunc, **kwargs):
        # set the data
        self._data = data          # oberserved data
        self._abscissa = abscissa # points at which the observed data are taken
        self._ndata = len(data)    # number of data points
        self._model = modelfunc    # model function

        # d!, to speed up calculation of the likelihood
        self._ln_d_fact = np.log(sp.special.factorial(self._data))

        Solver.__init__(self, **kwargs)

    def Prior(self, cube):
        """
        The prior transform going from the unit hypercube to the true parameters. This function
        has to be called "Prior".

        Args:
            cube (:class:`numpy.ndarray`): an array of values drawn from the unit hypercube

        Returns:
            :class:`numpy.ndarray`: an array of the transformed parameters
        """
        # extract values
        mprime = cube[0]
        cprime = cube[1]

        m = mprime*(self.mmax-self.mmin) + self.mmin  # convert back to m
        c = cprime*(self.cmax-self.cmin) + self.cmin  # convert back to c
```

[5]The Gaussian prior was also changed since it always permits a finite probability that $m < 0$, causing problems with the likelihood calculation.

```

        return np.array([m, c])

    def LogLikelihood(self, cube):
        """
        The log likelihood function. This function has to be called "LogLikelihood".

        Args:
            cube (:class:`numpy.ndarray`): an array of parameter values.

        Returns:
            float: the log likelihood value.
        """
        # extract parameters
        m = cube[0]
        c = cube[1]

        # calculate the model
        model = self._model(x, m, c)

        return -np.sum(self._ln_d_fact + model - self._data*np.log(model))
```

The new solver was initiated the same way as before

```
nlive = 1024 # number of live points
ndim = 2     # number of parameters
tol = 0.5    # stopping criterion

# run the algorithm
solution1 = PoissonSolver(data, x, y_linear, n_dims=ndim, n_live_points=nlive,
    evidence_tolerance=tol)

# Get the log likelihood and samples
logZ_pois = solution1.logZ
logZerr_pois = solution1.logZerr
samples_pois = solution1.samples

# Print the summary
print(solution1)
```

```
Model in "(temporary directory)" (2 dimensions)
Evidence ln Z = -139.0 +- 0.1
Parameter values:
   Parameter  1 : 3.425 +- 0.162
   Parameter  2 : 1.583 +- 0.594
```

The "GaussianSolver" class was also updated with the modified priors and new measurement uncertainties $\sigma_i = \sqrt{d_i}$ and executed again, returning:

```
Model in "(temporary directory)" (2 dimensions)
Evidence ln Z = -156.3 +- 0.1
Parameter values:
   Parameter  1 : 3.411 +- 0.150
   Parameter  2 : 0.737 +- 0.475
```

The corner plots for the "PoissonSolver" samples are shown in Figure C.6 (a), and the "GaussianSolver" samples are shown in Figure C.6 (b). Although both models are able to infer m to a reasonable accuracy, the Gaussian model is not able to fully resolve the marginal distribution over c. This is also reflected in the less-accurate estimate for the parameter. This discrepancy is due to the fact that the measurements are quite low ($d_i < 5$) near the y-intercept, meaning that the Gaussian approximation $\sigma^2 \approx d_i$ begins

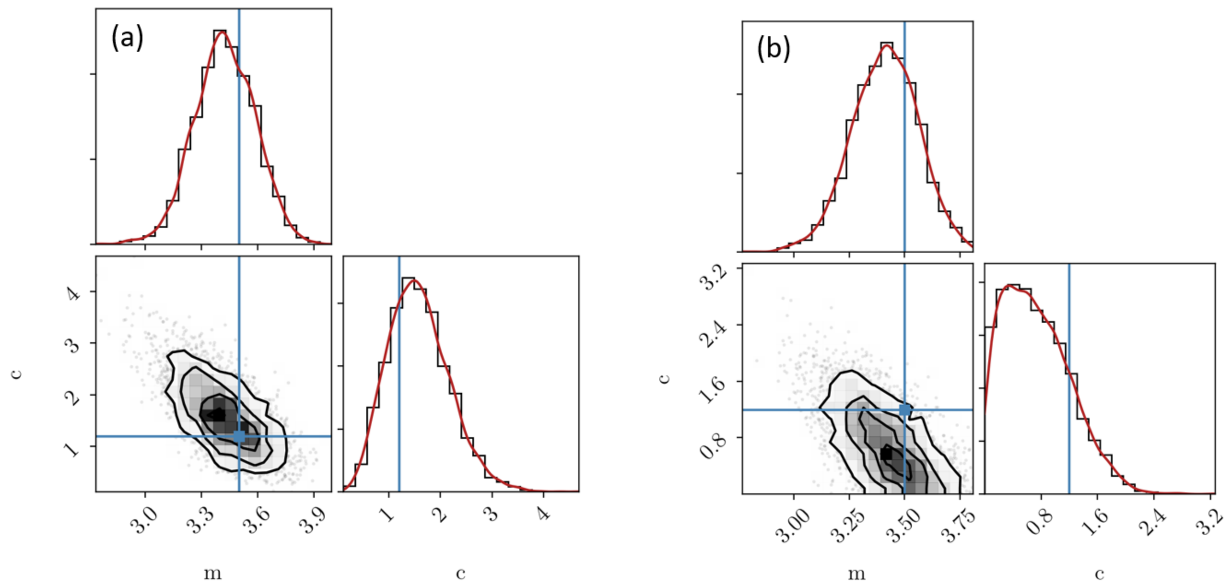

Figure C.6: Corner plots from the posterior distribution assuming (a) a Poisson likelihood function and (b) a Gaussian likelihood function. Due to the low count rates d_i for low x_i, the Gaussian model is not able to fully resolve the distribution over c.

to break down. If the overall count rate was significantly higher, the results of the two solver methods would be nearly indistinguishable. The best-fit models for each case, with associated 1σ uncertainty bands, are shown in Figure C.7 along with the data and true model. The Poisson model and the true value is nearly overlapping, which Gaussian model somewhat underestimates the true model for all x.

Although it may seem obvious from just looking at the corner plots and the fits, the evidence provides us a definitive way to determine which model is in better agreement with the data. For the Poisson model we found that $\ln \mathcal{Z}_P = -139.0 \pm 0.1$, and for the Gaussian model it was $\ln \mathcal{Z}_G = -156.3 \pm 0.1$. The better model is the one for which $\mathcal{Z}$ (not the logarithm) is greater. For our case $\mathcal{Z}_P > \mathcal{Z}_G$, meaning that the Poisson model does indeed explain the data better than a model based on Gaussian noise.

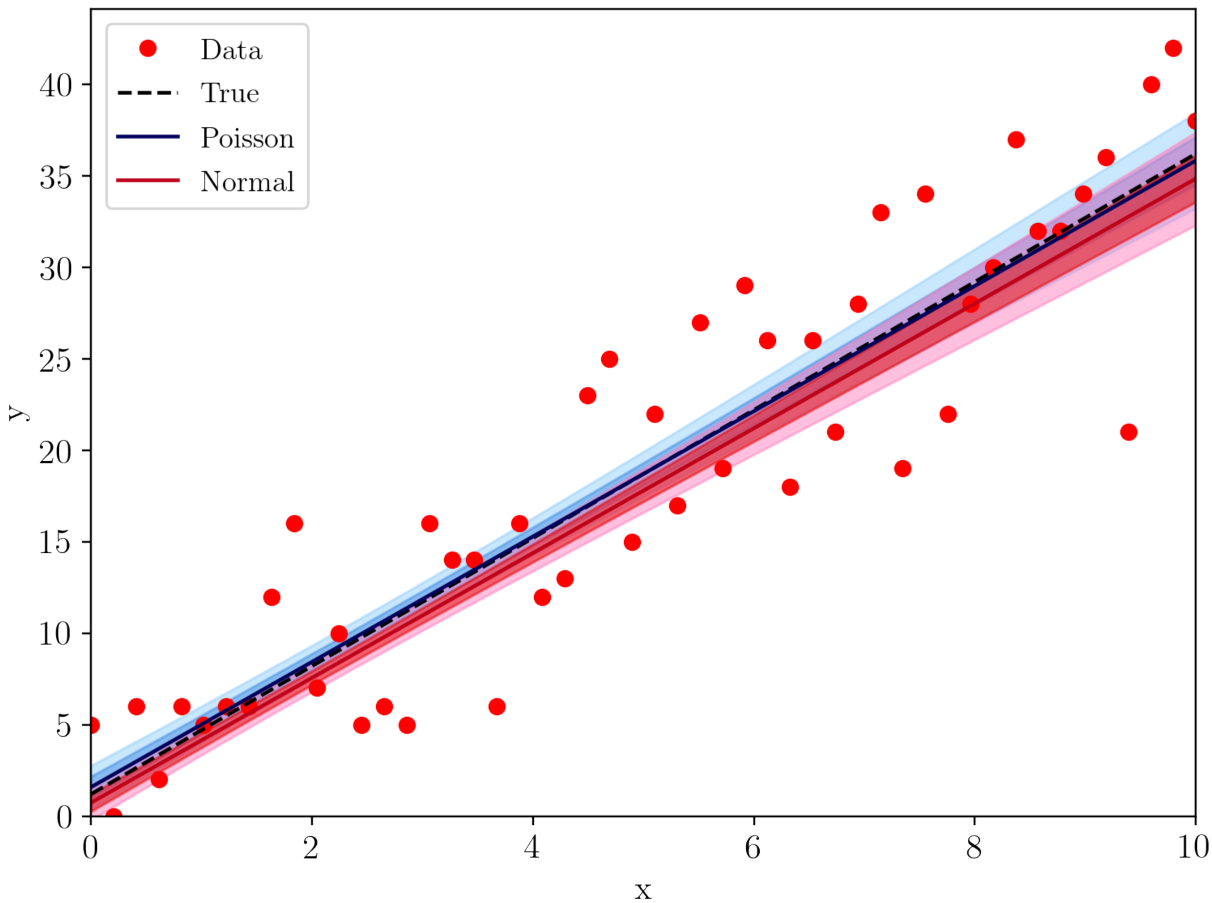

Figure C.7: The original data points, true model, and inferred models given (a) a Poisson likelihood function and (b) a Gaussian likelihood function. The shaded regions correspond to $\sigma = \sqrt{\mathrm{Var}[y_i]}$, as calculated from the posterior samples.

C.5 Conclusion: which sampler do I use?

The final question remaining is perhaps the most basic: for a given Bayesian inference problem, which sampling method should I use? Although this choice is often guided by preference, there are some basic guidelines I have developed over the course of my graduate school career. These are not meant as hard rules, nor is it intended to be comprehensive (there are dozens of types of samplers I have never even used!). Instead, it is meant as a starting point for someone who has been newly introduced to Bayesian integrated data analysis and wants to try it out for themselves.

My suggestions are summarized as a flow chart in Figure C.8. Basically, if the model only has one or two parameters, a simple grid search will typically be the most efficient method. Otherwise the choice depends on whether you need to calculate the evidence, or if the distribution is strongly multimodal (that is, the distribution contains multiple peaks which are distant from one another in parameter space). In both cases, nested sampling is the best choice. Otherwise, if all you need are samples from the distribution, MCMC will typically be the most efficient (and simplest) option This is especially true if the likelihood is expensive to calculate.

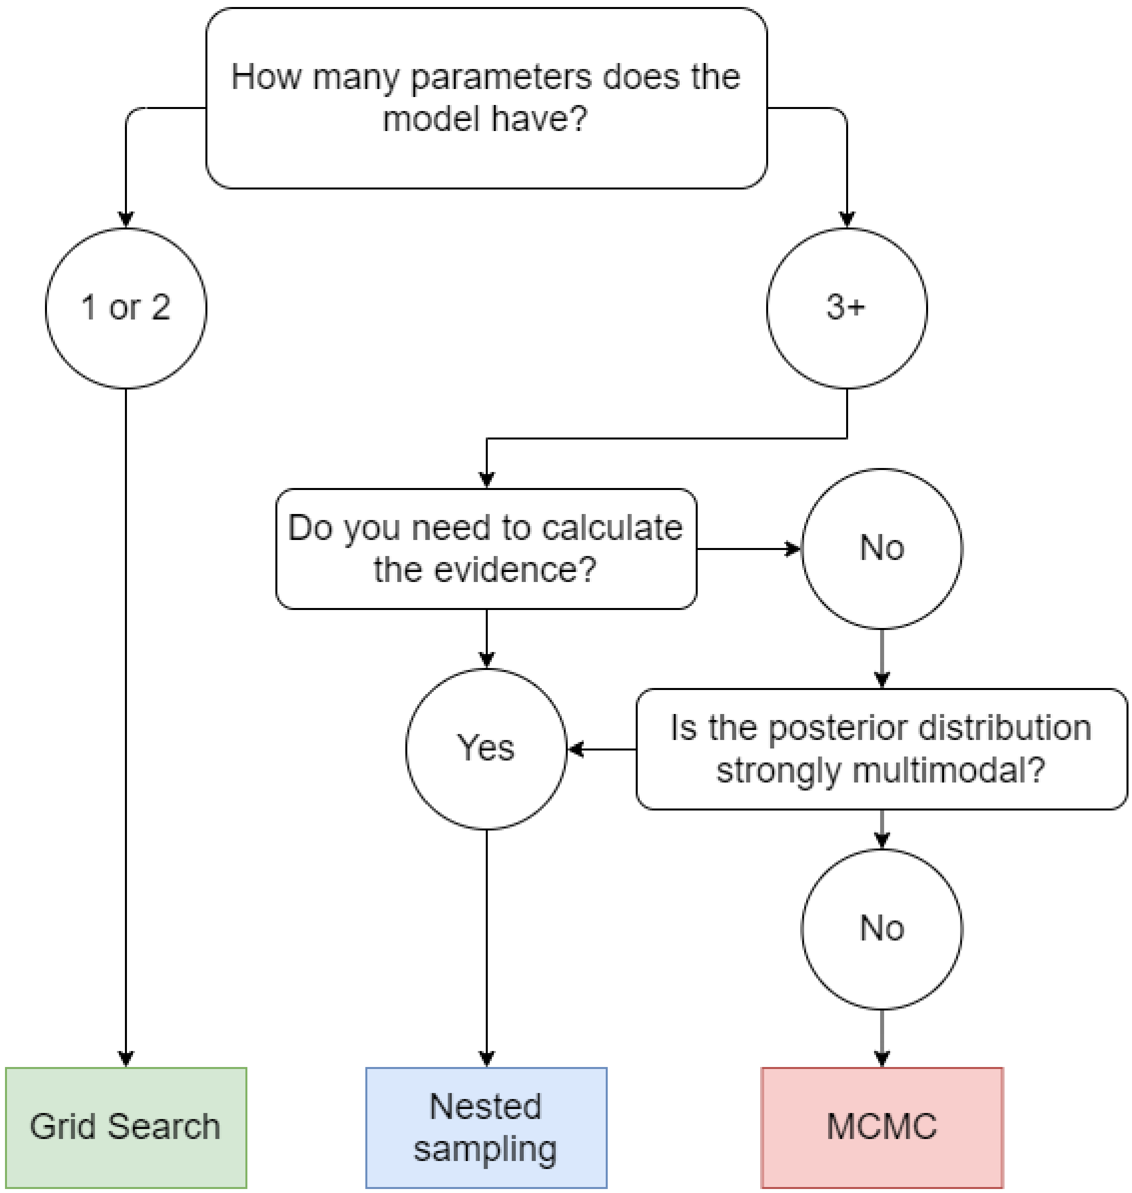

Figure C.8: Flow chart for selecting a sampling method.

Bibliography

[1] M. Pitkin, "Samplers, samplers, everywhere..." feb 2018. [Online]. Available: http://mattpitkin.github.io/samplers-demo/pages/samplers-samplers-everywhere/

[2] "Python Language Reference, version 3.7." [Online]. Available: https://www.python.org/

[3] S. Van Der Walt, S. C. Colbert, and G. Varoquaux, "The NumPy array: A structure for efficient numerical computation," *Computing in Science and Engineering*, vol. 13, no. 2, pp. 22–30, 2011. [Online]. Available: https://numpy.org

[4] E. Jones, T. Oliphant, P. Peterson, and Others, "SciPy: Open source scientific tools for Python," 2001. [Online]. Available: http://www.scipy.org/

[5] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing, "Jupyter Notebooks—a publishing format for reproducible computational workflows," *Proceedings of the 20th International Conference on Electronic Publishing*, pp. 87–90, 2016. [Online]. Available: doi.org/10.3233/978-1-61499-649-1-87

[6] D. Foreman-Mackey, "corner.py: Scatterplot matrices in Python," *Journal of Open Source Software*, vol. 1, no. 2, p. 24, jun 2016. [Online]. Available: https://doi.org/10.21105/joss.00024

[7] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. Boca Raton, FL: CRC Press, 2014.

[8] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, "emcee: The MCMC Hammer," *Publications of the Royal Astronomical Society of the Pacific*, vol. 125, pp. 306–312, 2013. [Online]. Available: http://dan.iel.fm/emcee.

[9] A. D. Sokal, "Monte Carlo Methods in Statistical Mechanics: Foundations and New Algorithms," in *Functional Integration*, C. DeWitt-Morette, P. Cartier, and A. Folacci, Eds. Boston, MA: Springer, 1997, ch. 6.

[10] J. Skilling, "Nested Sampling for General Bayesian Computation," *Bayesian Analysis*, vol. 1, no. 4, pp. 833–860, 2006.

[11] F. Feroz, M. P. Hobson, and M. Bridges, "MULTINEST: an efficient and robust Bayesian inference tool for cosmology and particle physics," *Mon. Not. R. Astron. Soc*, vol. 398, pp. 1601–1614, 2009. [Online]. Available: http://www.superbayes.org.

Appendix D

Quantifying the effect of toroidicity on tearing mode phase

This appendix concerns a project that I worked on during the earlier part of my tenure as an MST graduate student to quantify a perceived phase shift between soft X-ray emission structure observed by the SXT system and magnetic perturbations measured by the toroidal array. It is not directly related to the rest of my work described in this thesis (other than that it concerns SXT measurements), but is included here because it is potentially useful for future reference.

This investigation began as an extension of an internal report by Alberto Ruzzon [1], a graduate student at Consorzio RFX, considering both MST and RFX-mod data. This study assumed a cylindrical relation between the phase of the magnetic perturbations measured at the wall and those measured by the SXT system in the core, then quantifies the percieved discrepancy into single parameter $\Delta\delta$. RFP devices are often said to be "cylindrical-like," so this kind of cylindrical phase mapping has been commonly used. On MST $\Delta\delta$ (with reference to B_p signals in 500 kA PPCD) was found to have a value of about 25°, with a relatively large variation between shots. On RFX-mod it was found that $\Delta\delta \sim 5^\circ - 10^\circ$, with somewhat more observed variation between individual cameras.

The study does not attempt to explain the origin of this shift. I began with the goal of verifying this study, accounting for additional systematic effects, and explaining this supposed discrepancy.

The discrepancy is explained as a physical manifestation of poloidal mode coupling inherent in a system with toroidal symmetry. This effect results in a perceived shift in tearing mode phase in the core (where x-ray signal is strong) vs the outer edge (where the b-dot coils are located). This effect was explored in a series of NIMROD simulations performed by J. Sauppe, which is the subject of Section D.1. Section D.2 discusses the methodology for measuring $\Delta\delta$ using the SXR tomography diagnostic, Section D.3 considers the ideal cylindrical reference case, Section D.4 discusses the additional systematic effects which were accounted for, and Section D.5 discusses the results.

D.1 NIMROD simulations

The impact of poloidal mode coupling on phase is an inherently non-linear phenomenon, so direct comparison to analytical theory is difficult. Thankfully, early in the analysis I learned that another graduate student, J. Sauppe, had previously run a series of NIMROD simulations comparing cylindrical and toroidal geometries using MST-like conditions. To my knowledge these results remain unpublished, so their description here will be brief. I will summarize the results only to the extent necessary to understand my measurements.

NIMROD is a resistive MHD code with two-fluid and kinetic capabilities [2] which can be used to model fusion-relevant magnetically-confined plasmas. For the work discussed here, a simplified single-fluid model was used. The code solves momentum equation, Faraday's law, and the low-frequency Ampere's law,

$$m_i n\left(\frac{\partial \boldsymbol{v}}{\partial t} + \boldsymbol{v}\cdot\boldsymbol{\nabla}\boldsymbol{v}\right) = \boldsymbol{J}\times\boldsymbol{B} - \boldsymbol{\nabla}\cdot\left(m_i n \nu \underline{\underline{\boldsymbol{W}}}\right) \tag{D.1}$$

$$\frac{\partial \boldsymbol{B}}{\partial t} = -\boldsymbol{\nabla}\times\left(-\boldsymbol{v}\times\boldsymbol{B} + \eta\boldsymbol{J}\right) \tag{D.2}$$

$$\mu_0 \boldsymbol{J} = \boldsymbol{\nabla}\times\boldsymbol{B} \tag{D.3}$$

where m_i is the ion species mass, n is the number density, ν is the viscosity, $\underline{\underline{\boldsymbol{W}}}$ is the rate-of-strain tensor, and η is the resistivity. Boundary conditions were chosen consistent with frozen-in flux and no-slip constraints. Other parameters were chosen to be consistent with standard operating conditions in the MST.

Computations were performed assuming both a cylindrical and a toroidal equilibrium. In both cases general coordinates (ρ, θ, ϕ) were defined such that ρ is a normalized flux surface label ($\boldsymbol{B}\cdot\boldsymbol{\nabla}\rho = 0$), θ is the normal poloidal angle, and ϕ is the normal toroidal angle (in cylindrical system $\phi = 2\pi z/L$). In the cylindrical simulation there is no coupling between poloidal modes (that is, both m and n are good "quantum numbers"). As a consequence, the safety factor is only a function of radius, $\boldsymbol{B}\cdot\boldsymbol{\nabla}\phi/\boldsymbol{B}\cdot\boldsymbol{\nabla}\theta \equiv q(\rho)$. In the toroidal case, however, nearby poloidal modes m couple to one another meaning that in general the field line pitch is defined locally as $\boldsymbol{B}\cdot\boldsymbol{\nabla}\phi/\boldsymbol{B}\cdot\boldsymbol{\nabla}\theta \equiv \hat{q}(\rho,\theta)$. It is possible to define a transformation into so-called "straight field-line" coordinates (θ_f, ϕ_f) in which $\boldsymbol{B}\cdot\boldsymbol{\nabla}\phi_f/\boldsymbol{B}\cdot\boldsymbol{\nabla}\theta_f = q(\rho)$. This transformation is shown in Figure D.1. This provides a sense of the phase shift between the plasma core and the edge. However it should be noted poloidal mode-coupling still occurs even when modes are expressed in terms of straight field-line coordinates.

Coupling between poloidal harmonics, as observed in the toroidal case, can significantly effect the phase of a tearing perturbation. Near the plasma core, resonant modes are broadly similar to the cylindrical case. However, as the radius is increased a phase shift occurs. One result of this effect is that a $\sim 30^\circ$ phase shift develops between

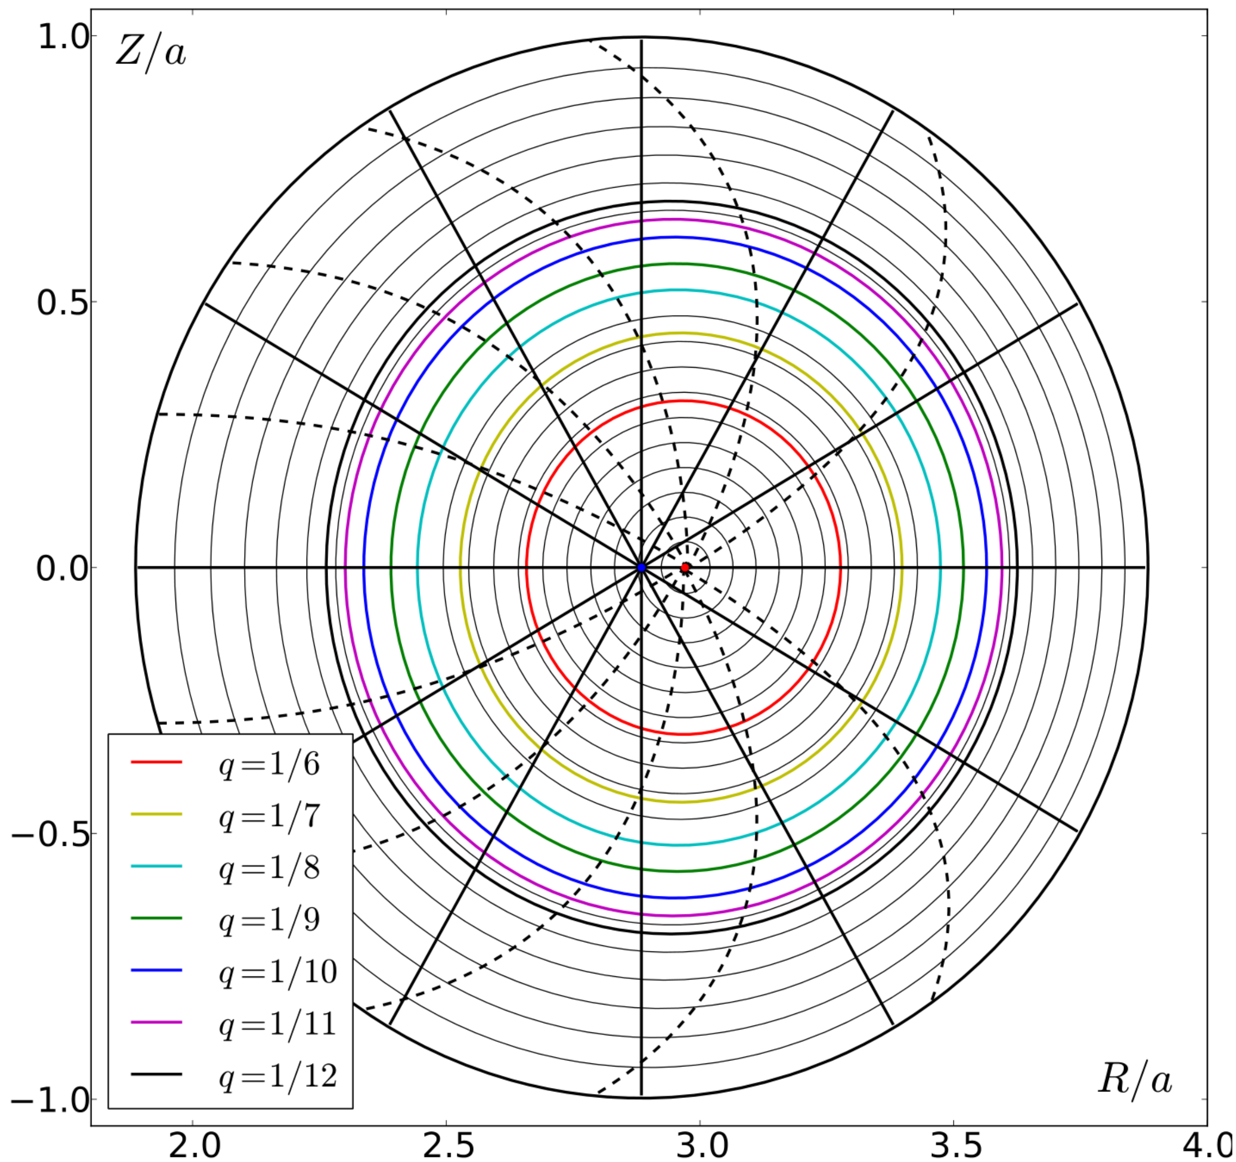

Figure D.1: Lines of constant ρ (rational surfaces in color) and lines of constant straight-field angle θ_f (dashed) and geometric angle θ_g (solid), evenly spaced by $\pi/6$. Courtesy of J. Sauppe.

toroidal and poloidal components of $\tilde{b}$. As shown in Figure D.2, this effect depends on the poloidal location of the measurement, with the current location of the magnetics toroidal array happening to be near the maximum. This is also expected to lead to a noticeable phase shift between the soft x-ray tomography diagnostic (which measures mostly the plasma core) and the edge magnetics.

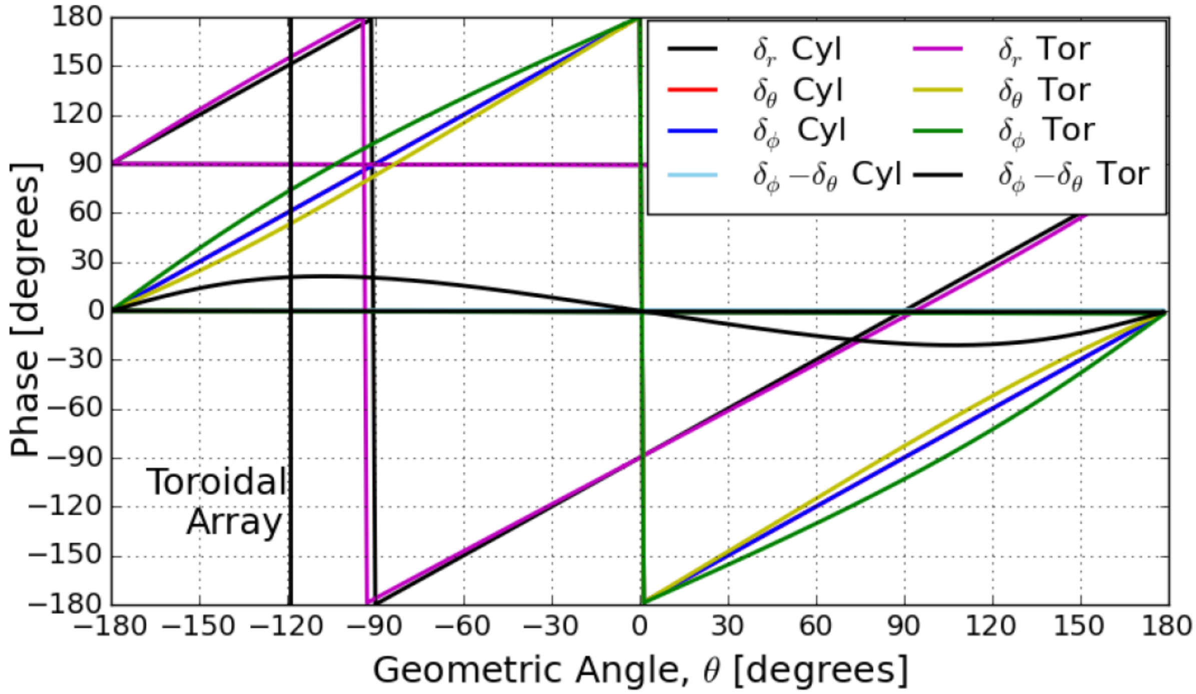

Figure D.2: Eigenmode phases of b_r , b_θ , and b_ϕ and phase difference, $\delta_\phi - \delta_\theta$, at $r = a$ for the linear computations in cylindrical and toroidal geometry. Courtesy of J. Sauppe.

D.2 Measuring SXR phase

A simple model was developed relating the phase of the O-point of a soft X-ray emissivity island, θ_{op}, to the impact parameter associated with the line of sight passing from a given SXT camera through that O-point. See Figure 1 for a schematic of how the impact parameter is defined. This model does not require performing full inversions of the emissive profile. To do so, we will make a few assumptions:

1. A chord passing from a camera through the O-point has a larger line-integrated brightness signal than any other chord originating from the same camera.

2. The "fanning out" of the camera lines of sight is ignored. That is, the model assumes that all lines of sight for a camera are parallel as shown in Figure 2. This effect is small and has a minimal effect on the overall phase.

3. The island is assumed to rotate circularly around a Shafranov-shifted axis, and the camera poloidal positions θ^{ℓ} are determined with respect to this axis.

In addition to these assumptions, we used the standard convention for the sign of the impact parameter in which chords to the outboard side of the geometric axis are assigned p>0 and chords to the inboard side are assigned $p < 0$. Putting this together, Equation A.1 was developed:

$$p^{\ell}_{max}(t) = A^{\ell}_0 \mp (A^{\ell}_1)^2 \sin\left(\theta_{op}(t) - \theta^{\ell}\right) \tag{D.4}$$

where p^{ℓ}_{max} is the impact parameter of a chord passing through both camera ℓ and the emissive O-point at any time t, $\ell \in \{A, B, C, D\}$ is the SXT camera label, θ^{ℓ} is the camera poloidal position with respect to the axis of rotation, A^{ℓ}_0 is the average impact parameter, and $(A^{\ell}_1)^2$ is the amplitude of oscillation. A_1 is squared so that when it is fit to data it will always be positive. The plus (+) sign connected the two terms is used when $\ell = D$, otherwise the minus (-) sign is used.

Next, we considered specifically the case in which the plasma tearing mode spectrum is dominated by a single large $(m = 1, n)$ mode (QSH, but not SHAx, plasmas). In this scenario, it is expected that the emissive O-point will align with the magnetic O-point Source?. Adopting this as an additional assumption, we can then write the phase of the SXT impact parameter model in terms of the magnetic phase δ_n of the dominant mode

n. The result is a model relating SXT brightness measurements (left side) to magnetics phase information (right side) via three fit parameters:

$$p^{\ell}_{max}(t) = A^{\ell}_0 \mp (A^{\ell}_1)^2 \sin(A^{\ell}_2 - \delta_n(t)) \tag{D.5}$$

In Equation D.5 we have introduced an additional parameter, A^{ℓ}_2, representing the constant phase shift between the magnetics phase at $\phi = 0°$ and the SXR phase at $\phi = 90°$. For now we will treat this as an arbitrary fit value. An expression for the value of A^{ℓ}_2 in a purely cylindrical plasma will be developed in Section D.3.

A database of time intervals from 55 PPCD shots ($n = 6$) and 66 non-reversed shots ($n = 5$) was collected, all at 500 kA. The intervals were chosen which feature rotating plasmas with $N_S < 2$ for PPCD and $N_S < 2.5$ for non-reversed. The impact parameter passing from a given camera through the O-point, p^{ℓ}_{max}, at a given time step was determined by fitting the brightness measurements to a quadratic and determining its maxima, as shown in Figure D.3. This is data forms a time series which is then fit to Equation D.5, with $A^{\ell}_{0,1,2}$ as the fit parameters and δ_n coming from the standard MST magnetic mode analysis routine. This process is illustrated in Figure D.4. Not that all data considered here was recorded during a period in 2017 when SXR-C was not operational due to a mechanical failure.

For each camera during each shot a value for the phase A^{ℓ}_2 was determined via nonlinear fit. These values were then averaged by camera, giving a set of three phases $\langle A^{\ell}_2 \rangle$. This was done for both δ_θ and δ_ϕ (later referred to by component index $i = \theta, \phi$). This process was performed separately for the PPCD and non-reversed datasets. The results of these fits are described later in Table A.1 (in Section A.4). Before that, Section A.2 describes the expected value of A^{ℓ}_2 in a cylindrical model, and Section A.3 describes the treatment of additional systematic effects related to hardware delay and geometry.

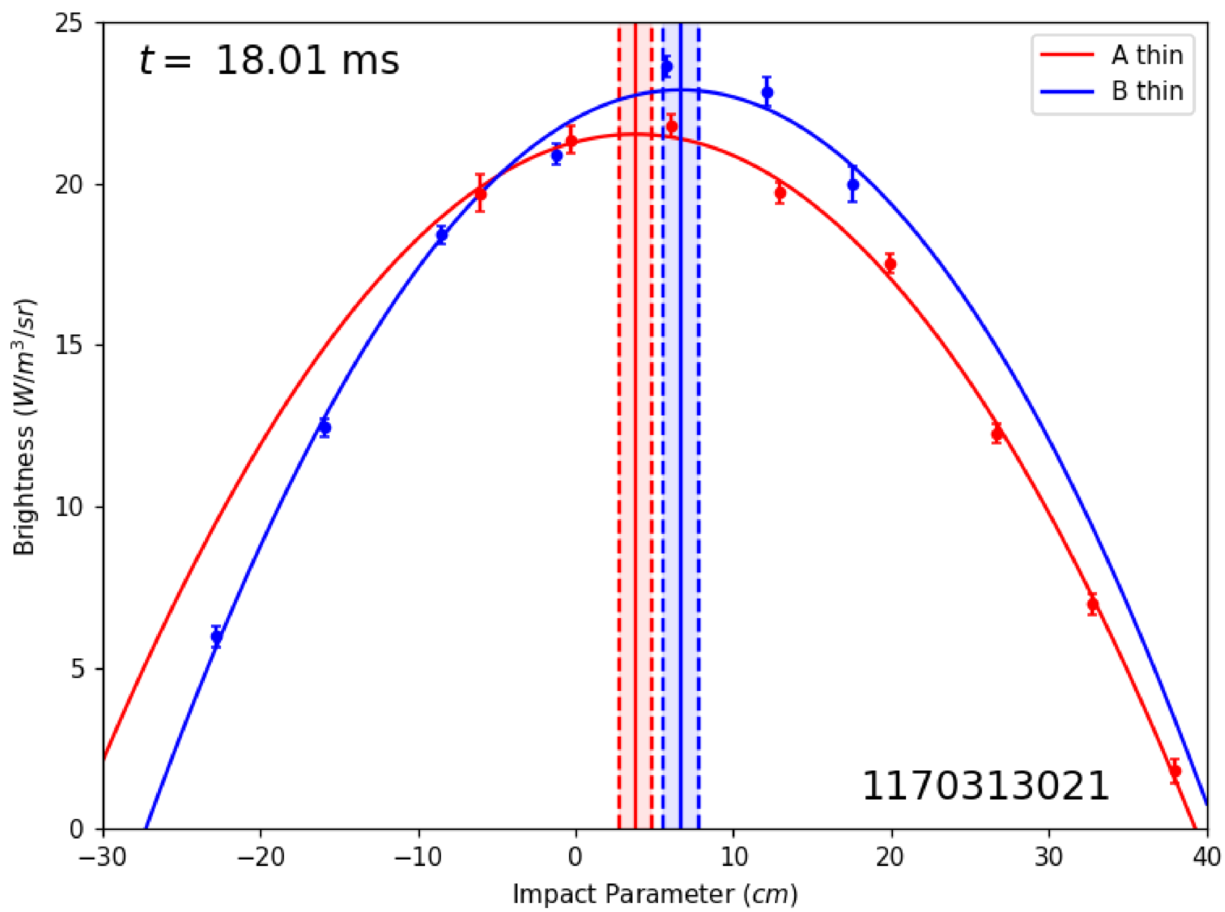

Figure D.3: The impact parameter corresponding to the brightest line of sight, p_{max}, is determined by fitting a quadratic to SXT brightness measurements and determining its maxima. This permits interpolation between the actual camera measurement chords. Dashed lines characterize uncertainty

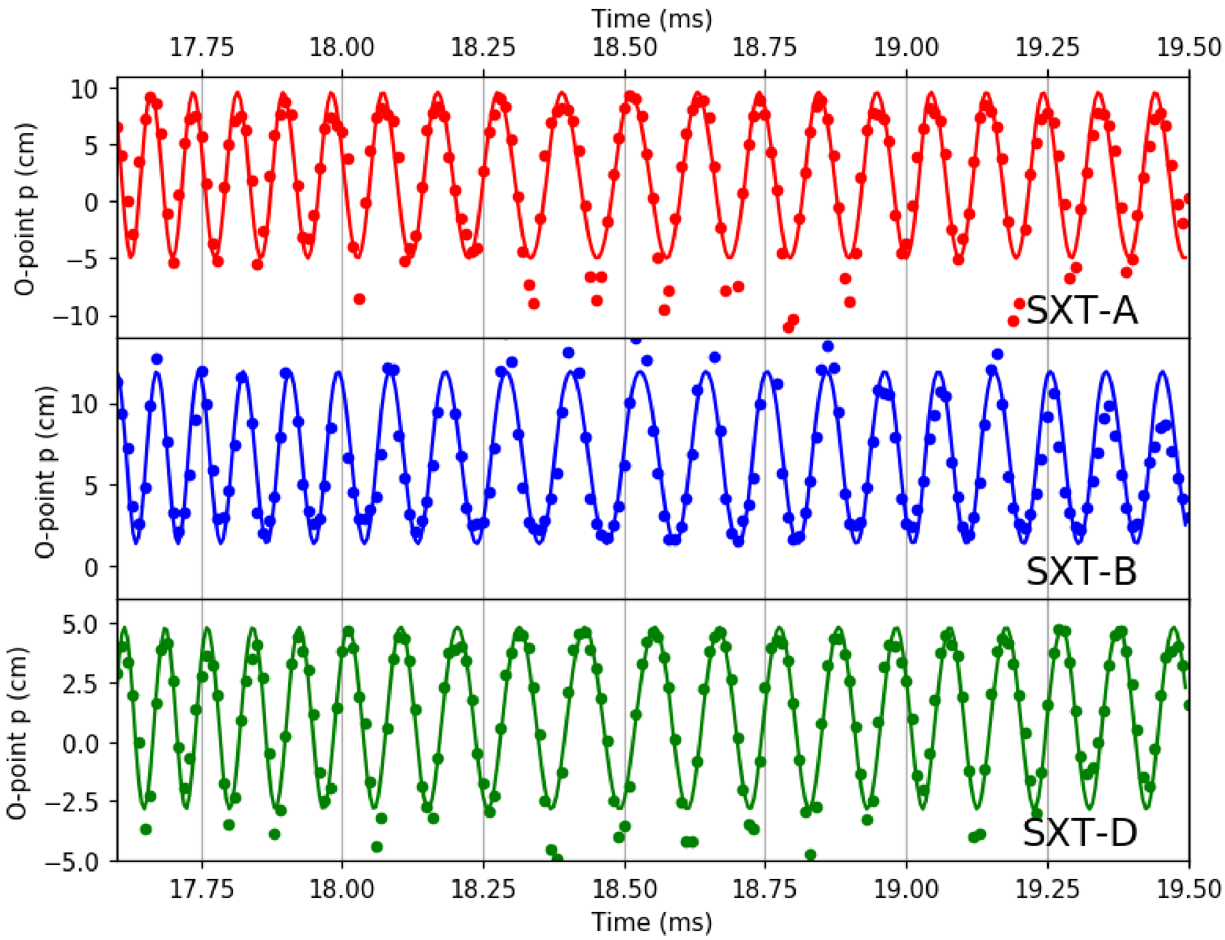

Figure D.4: Impact parameter p_{max} plotted versus time for all four cameras (this illustration is from an older dataset). The solid line is the model described by Equation D.5. The qualitative time-series agreement of the model with SXT data provides strong support in favor of the assumption relating X-ray emission to the magnetics O-point.

D.3 SXR phase in the cylindrical approximation

The δ_n referred to in the previous sections is defined by the conventions of the magnetic toroidal array mode analysis code. This code decomposes perturbations to the magnetic field at $\theta = 241^\circ$ as follows:

$$\tilde{b}_i(\phi, t) = \sum_{n>0} c_n \cos\left(n\phi - \delta_{i,n}(t)\right) \tag{D.6}$$

Since the array exists only at a single poloidal angle, it is not sensitive to poloidal harmonics. However, with the expectation that these perturbations are largely due to tearing modes we can assume that they are dominantly $m = 1$ in nature.

Now we must introduce a complication to the analysis: MST angles nominally form a left-handed coordinate system. That is, $\hat{\boldsymbol{r}} \times \hat{\boldsymbol{\theta}} = -\hat{\boldsymbol{\phi}}$ for the vector components described in Equation A.3. Since my goal was to compare these results to Nimrod simulations performed in a right-handed coordinate system, the MST data must be converted. There are two natural options to accomplish this: either swap the direction of $\hat{\boldsymbol{\theta}}$ or swap the direction of $\hat{\boldsymbol{\phi}}$. Here I will adopt the second option and define a new direction $\hat{\boldsymbol{\phi}}' = -\hat{\boldsymbol{\phi}}$. This also means that $\tilde{b}_{\phi'} = -\tilde{b}_{\phi}$ leading to a π phase shift. A consequence of converting to a right-handed coordinate system is that m and n for tearing modes now have the same sign, i.e. for $m = 1$ we have $n > 0$. The use of primed notation to designate the RH coordinate system is used consistently throughout the rest of this appendix, and the reader is cautioned to pay careful attention to this potentially confusing notation.

For a single-helicity magnetic perturbation the SXR structure can be assumed to have a toroidal symmetry determined by n, and will be aligned with the magnetics array when

$$\phi_0' = -\frac{\delta_{i,n}}{n}, \tag{D.7}$$

where $n > 0$ and $i \in \{\theta, \phi'\}$.

In a cylindrical RFP equilibrium, magnetic field lines lie on toroidal flux surfaces nested about the geometric axis. The trajectory of a field-line on a rational surface $q = -m/n$ is readily related to geometric coordinates. If we consider a field-line passing through the magnetic O-point at $(\theta = 241°, \phi'_0)$, then we can say that at some other ϕ' the O-point can be found at θ_{op}:

$$\frac{\Delta\phi'}{\Delta\theta} = \frac{\phi' - \phi'_0}{\theta_{op} - 241°} \tag{D.8}$$

$$= -\frac{m}{n}, \tag{D.9}$$

which gives

$$\theta_{op}(\phi') = 241° - n(\phi' - \phi'_0). \tag{D.10}$$

Equation D.10 makes use of the fact that unstable tearing modes in the core of MST are expected to have helicity $m = 1$. It is also assumed that $n > 0$ for the right-handed coordinate system we are now working in. This can be combined with Equation D.7 to obtain an equation relating the O-point position θ_{op} to magnetic phase $\delta_{i,n}$ at any give ϕ'. We are specifically interested in $\phi' = \phi'_{SXR} = -90°$, which gives

$$\theta_{op} = 241° + n \cdot 90° - \delta_{i,n}. \tag{D.11}$$

Subtracting the camera poloidal position θ^ℓ immediately yields an analytic expression for A_2^ℓ in a cylindrically-symmetric plasma:

$$A_2^\ell \equiv 241° + n \cdot 90° - \theta^\ell \tag{D.12}$$

MST, however, is not a cylindrical device and should not be expected to behave precisely as one. Determining $\langle A_2 \rangle$ from SXR/magnetics data as described in Section A.1

and comparing it Equation D.12 permits quantification of the degree to which MST data deviates from the cylindrical model on average. This will be expressed by the parameter $\Delta\delta$, defined as

$$\Delta\delta_i^\ell \equiv A_2^\ell - \langle A_2^\ell \rangle_i \tag{D.13}$$

where A_2^ℓ is defined by Equation D.12, ℓ is the SXT camera label, $i \in \{\theta, \phi'\}$ is the label designating the component of the magnetic field δ was taken from, and $\langle A_2^\ell \rangle_i$ is the average measured phase for the given camera label and vector component. Although $\delta_{\theta,n}$ and $\delta_{\phi',n}$ are in phase in a cylindrical system, this is not generally true in a toroidal system, so different $\Delta\delta$ values will be measured for each component. The results are given and in Table 1 in Section A.4.

To correctly interpret the meaning of $\Delta\delta$, we can consider the form of the phase of the cosine in Equation D.5 if we substitute the measured average A_2. That is, on average the brightest line of sight for a given camera will oscillate like

$$p_{max}^\ell \sim \sin\left(\langle A_2^\ell \rangle - \delta_{i,n}\right) \tag{D.14}$$

$$\sim \sin\left(A_2^\ell - \delta_{i,n} - \Delta\delta_i^\ell\right). \tag{D.15}$$

The Nimrod simulations that I will compare my data with in Section A.4 define δ with the opposite sign compared to the convention adopted by the MST mode analysis software. To make the comparison more straightforward we can switch to this convention $\hat{\delta}_{i,n} = -\delta_{i,n}$ yielding

$$p_{max}^\ell \sim \sin\left(A_2^\ell + \hat{\delta}_{i,n} - \Delta\delta_i^\ell\right) \tag{D.16}$$

Written this way it is clear that $\Delta\delta_i$ represents the **phase by which the SXR structure lags** that which would be expected from the phase of a specified component of the magnetic field perturbation in a cylindrical model.

D.4 Systematic effects

Before the SXT can be fit to the magnetics model as described in Section A.1, some known systematic effects need to be accounted for. The amplifiers through which the photodiode signals pass before being digitized were known to have a delay long enough to affect the phase of signals oscillating at typical tearing mode frequencies. This delay, as well as the delay associated with the integrators in the magnetics data collection system, needs to be properly accounted for to reliably measure a phase shift between the two signals. This section also considers the appropriate values to use for the camera poloidal positions θ^{ℓ}.

When SXR emission strikes a photodiode in an SXT camera that diode generates a signal which is sent along a transmission line to an amplifier where its amplitude is increased by a preset factor (typically 10^8) before continuing to a digitizer which converts the analog input into digital information and records that into the data tree. There is a frequency-dependent phase shift on signals passing through the amplifiers, which for the frequency range of interest (5-20 kHz) is approximately linear. This delay was measured by driving an oscillating current through an LED stored inside each SXT camera on the same board as the photodiodes. The signal is given an overall positive offset to avoid reverse-biasing the LED. Data was then collected as normal. The signal driving the LED was also collected directly via a spare channel in the digitizers using cables of the same length as used in the diode data collection. Data was collected for each over a range of driving frequencies, and FFT analysis was used to determine the shift in phase between the driving signal and the amplifier output. Results are shown in Figure D.5.

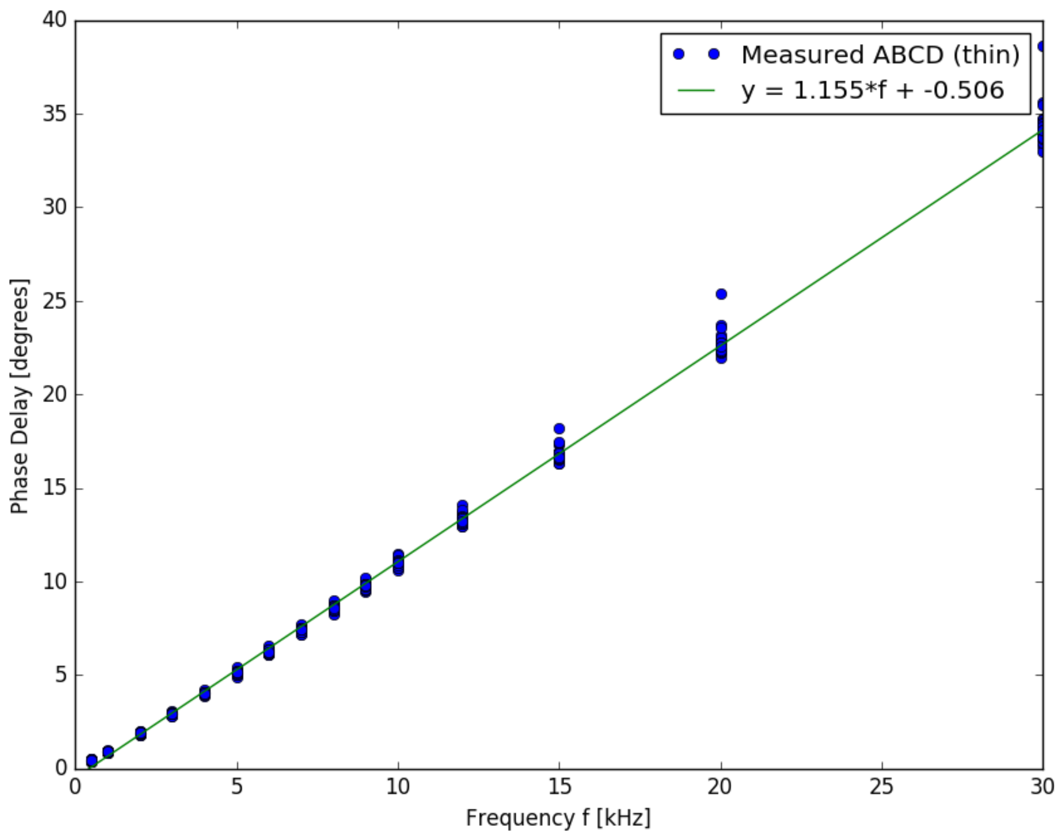

Figure D.5: Data was collected for all thin-filter channels over a range of LED frequencies spanning from 0.5 to 30 kHz. Over this range the amplifier phases shift is highly linear. The y-intercept in the resulting fit equation was permitted to be non-zero to account for nonlinear behavior at low frequency.

Signal delay of a signal through the magnetics array is also important. Although the digitizers used in this system are the same as those used in the SXT system (so any delay should affect both systems equally), there is a non-negligible delay through the integrator circuit which converts voltage measurements (proportional to dB/dt) to signals proportional to B. To a reasonable degree of accuracy this circuit can be modeled as a simple RC integrator circuit with a pole at $f_0 = \omega_0/2\pi \approx 250$ kHz. This circuit can be shown to have a phase lag between the input and output signals given approximately by $\delta_{mag} \approx f/f_0 = (0.229^\circ/\text{kHz})f$ for input signals $f \gg 250$ kHz.

It was also noticed that the results of the fitting procedure were sensitive to the choice of the magnetic axis location (the size of the Shafranov shift), which is used to calculate the poloidal position of the SXR cameras (θ^ℓ). Equilibrium reconstructions were not generally available for the entire dataset, so the magnitude of the shift was chosen to minimize the variance in the resulting $\Delta\delta$. This is shown in Figure D.6.

D.5 Results

Results for both the non-reversed ($n = 5$ core-resonant) and PPCD ($n = 6$ core-resonant) datasets are summarized in Figure D.7. Somewhat surprisingly, there was a substantial difference between the tow scenarios. For the non-reversed ($n = 5$ core-resonant) data it was found that on average $\Delta\delta_\theta = 9.7^\circ \pm 0.5^\circ$ and $\Delta\delta_\phi = -19.4^\circ \pm 0.5^\circ$. This is in good agreement with the NIMROD calculations, as illustrated in Fig. D.8, and demonstrates that the phase of the toroidal component of a core tearing mode is more significantly affected by toroidicity than the poloidal component. Although the simulations did not feature unstable n=5 eigenmodes, since mode number does not strongly affect the phase of core-resonant eigenmodes this is a useful comparison.

This same analysis was performed on on the PPCD dataset, which feature relatively large n=6 tearing modes (though the amplitude of all modes are suppressed). From this

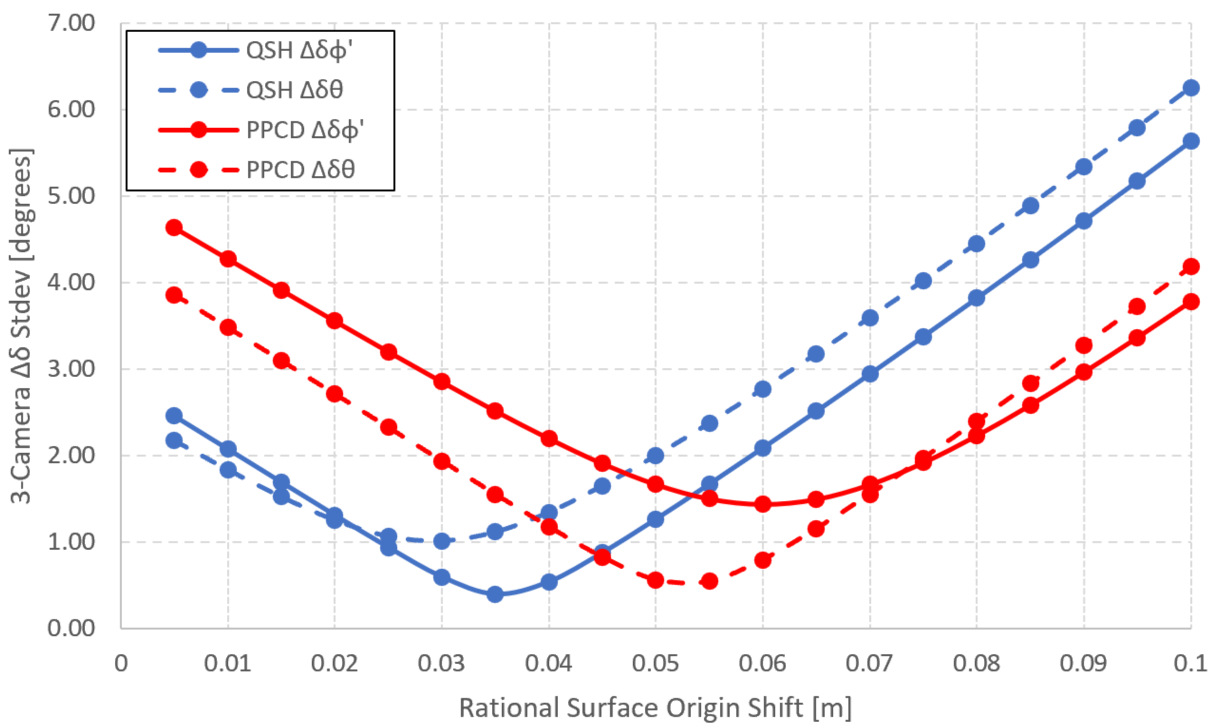

Figure D.6: The optimal magnitude of the Shafranov shift (used to determine θ^{ℓ}) was chosen to minimize the standard deviation of the resulting $\Delta\delta$ dataset.

	ℓ =	A	B	D	Overall
	Expected	14.60	334.00	83.50	-
	$\langle A_2^{\ell}\rangle_{\theta}$	354.03	314.64	62.62	-
PPCD (n=6)	**$\Delta\delta^{\ell}_{\theta}$**	20.57	19.37	20.88	**20.27**
	Error	0.53	0.40	0.46	0.27
	$\langle A_2^{\ell}\rangle_{\varphi'}$	18.98	341.24	89.69	-
	$\Delta\delta^{\ell}_{\varphi'}$	-4.39	-7.24	-6.19	**-5.94**
Shift (cm)	**Error**	0.58	0.38	0.48	0.28
6	**$\langle A2^{\ell}\rangle_{\varphi'} - \langle A2^{\ell}\rangle_{\theta}$**	24.96	26.61	27.07	26.21
	Expected	287.10	247.33	352.16	-
	$\langle A_2^{\ell}\rangle_{\theta}$	278.54	236.81	342.17	-
Nonrev. (n=5)	**$\Delta\delta^{\ell}_{\theta}$**	8.56	10.52	9.99	**9.69**
	Error	0.74	1.04	0.96	0.53
	$\langle A_2^{\ell}\rangle_{\varphi'}$	306.76	266.05	371.99	-
	$\Delta\delta^{\ell}_{\varphi'}$	-19.66	-18.71	-19.82	**-19.40**
Shift (cm)	**Error**	0.69	0.94	0.93	0.50
3	**$\langle A2^{\ell}\rangle_{\varphi'} - \langle A2^{\ell}\rangle_{\theta}$**	28.22	29.24	29.81	29.09

Figure D.7: Average results of the fits for SXR-A,B, and D for both PPCD and $F = 0$ RFP scenarios. SXR-C was not included in the analysis.

data it was found that on average $\Delta\delta_\theta = 20.3^\circ \pm 0.3^\circ$ and $\Delta\delta_\phi = -5.9^\circ \pm 0.3^\circ$, which displays an additional $\sim 10^\circ$ shift from the NIMROD simulations and the non-reversed data. It is possible that this discrepancy emerges from the relatively small dominant mode amplitude, resulting in a weaker correlation between the dominant mode phase and the SXR structure. It is also possible that the discrepancy is due to some additional physics particular to PPCD (i.e., the current drive) which the NIMROD simulations do not capture.

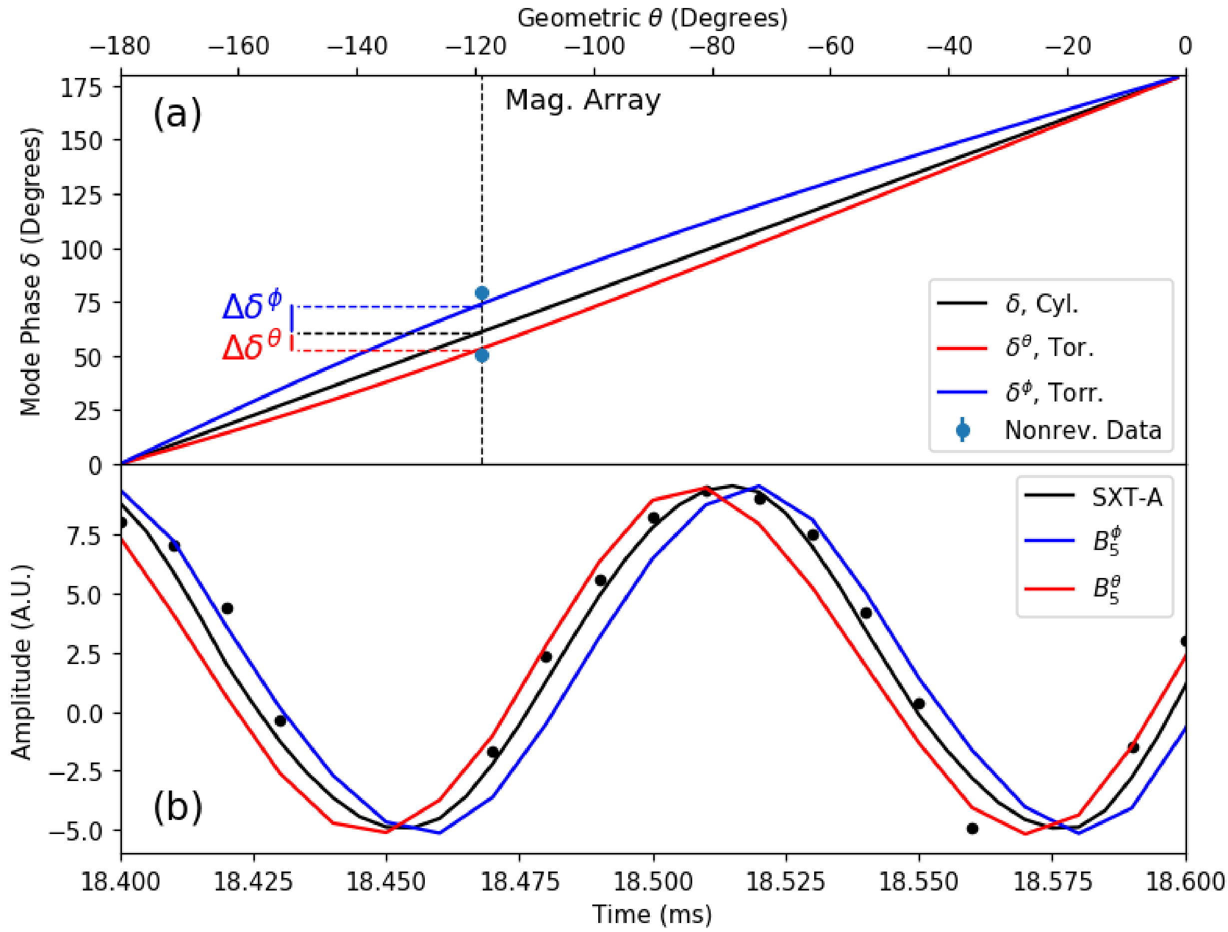

Figure D.8: (a) Phase of θ and ϕ component of n=6 eigenmode phase from a nonlinear NIMROD simulation of an MST plasma. Also plotted is the mode phase from a linear cylindrical calculation. This figure illustrates how $\Delta\delta_{\phi}$ and $\Delta\delta_{\theta}$ from the SXT data relate directly to the NIMROD calculations. The points along the dotted line represent the average $\Delta\delta_{\phi}$ and $\Delta\delta_{\theta}$ found from data. (b) Data from shot 1170313021 illustrating the observed phase shift between the SXT o-point oscillations in the core (cylindrical-like) and the two components of magnetic fluctuations data measured at the wall (toroidal-like). The black points are data while the black line is the p_{op} model. Magnetics data has been normalized to the offset and amplitude of the SXT data to aid in comparison.

Bibliography

[1] A. Ruzzon, P. Franz, and A. Fassina, "Analysis of the relation between SXR signals and the dominant magnetic mode phase in PPCD plasmas at MST," Consorzio RFX, Padova, Italy, Tech. Rep., 2013.

[2] C. Sovinec, A. Glasser, T. Gianakon, D. Barnes, R. Nebel, S. Kruger, D. Schnack, S. Plimpton, A. Tarditi, and M. Chu, "Nonlinear magnetohydrodynamics simulation using high-order finite elements," *Journal of Computational Physics*, vol. 195, pp. 355–386, 2004. [Online]. Available: https://doi.org/10.1016/j.jcp.2003.10.004